

Janice Eberly and James H. Stock, Editors

Brookings Papers

ON ECONOMIC ACTIVITY

FALL 2018

HARTMANN and SMETS

on the European Central Bank's Monetary Policy during Its First 20 Years

FARHI and GOURIO

on Accounting for Macro-Finance Trends: Market Power, Intangibles, and Risk Premia

BERNANKE

on the Real Effects of Disrupted Credit: Evidence from the Global Financial Crisis

COIBION, GORODNICHENKO, and ULATE

on the Cyclical Sensitivity in Estimates of Potential Output

FUHRER, OLIVEI, ROSENGREN, and TOOTELL

on Whether the Federal Reserve Should Regularly Evaluate Its Monetary Policy Framework

SYMPOSIUM

on Monetary Policy at the Effective Lower Bound

Brookings Papers

ON ECONOMIC ACTIVITY

BPEA

FALL
2018

BROOKINGS

BROOKINGS

www.brookings.edu



Brookings Papers

ON ECONOMIC ACTIVITY

FALL 2018

JANICE EBERLY
JAMES H. STOCK
Editors

BROOKINGS INSTITUTION PRESS
Washington, D.C.

Copyright © 2019 by

THE BROOKINGS INSTITUTION

1775 Massachusetts Avenue, N.W., Washington, D.C. 20036

ISSN 0007-2303

ISBN: Print 978-0-8157-3709-4; ebook: 978-0-8157-3710-0

Authorization to photocopy items for internal or personal use or the internal or personal use of specific clients is granted by the Brookings Institution for libraries and other users registered with the Copyright Clearance Center Transactional Reporting Service, provided that the basic fee is paid to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923. For more information, please contact CCC at (508) 750-8400. This authorization does not extend to other kinds of copying, such as copying for general distribution, for creating new collective works, or for sale. Specific written permission for such copying must be obtained from the Permissions Department, Brookings Institution Press, 1775 Massachusetts Avenue, N.W., Washington, D.C. 20036; fax (202) 536-3623; e-mail permissions@brookings.edu.

Brookings Papers

ON ECONOMIC ACTIVITY

FALL 2018

PHILIPP HARTMANN and FRANK SMETS

The European Central Bank's Monetary Policy during Its First 20 Years 1

Comments by Lorenzo Bini Smaghi and Lucrezia Reichlin 119

General Discussion 140

EMMANUEL FARHI and FRANÇOIS GOURIO

Accounting for Macro-Finance Trends: Market Power, Intangibles, and Risk Premia 147

Comments by Mark Gertler and Dimitris Papanikolaou 224

General Discussion 246

BEN S. BERNANKE

The Real Effects of Disrupted Credit: Evidence from the Global Financial Crisis 251

Comments by Olivier Blanchard and Raghuram Rajan 323

General Discussion 338

OLIVIER COIBION, YURIY GORODNICHENKO, and MAURICIO ULATE

The Cyclical Sensitivity in Estimates of Potential Output 343

Comments by Serena Ng and Valerie A. Ramey 412

General Discussion 431

JEFF FUHRER, GIOVANNI P. OLIVEI, ERIC S. ROSENGREN,
and GEOFFREY M.B. TOOTELL

Should the Federal Reserve Regularly Evaluate Its Monetary Policy Framework? 443

Comments by Jean Boivin and Stanley Fischer 498

General Discussion 513

SYMPOSIUM ON MONETARY POLICY AT THE EFFECTIVE LOWER BOUND

KRISTIN FORBES

*Monetary Policy at the Effective Lower Bound: Less Potent?
More International? More Sticky?*

521

JAMES D. HAMILTON

*The Efficacy of Large-Scale Asset Purchases When the Short-Term
Interest Rate Is at Its Effective Lower Bound*

543

ERIC T. SWANSON

*The Federal Reserve Is Not Very Constrained by the Lower Bound
on Nominal Interest Rates*

555

JANET YELLEN

Comments on Monetary Policy at the Effective Lower Bound

573

General Discussion 580

PURPOSE

The *Brookings Papers on Economic Activity* publishes research on current issues in macroeconomics, broadly defined. The journal emphasizes innovative analysis that has an empirical orientation, takes real-world institutions seriously, and is relevant to economic policy. Papers are presented and discussed at conferences held twice each year, and the papers and discussant remarks from each conference are published in the journal several months later. Research findings are described in a clear and accessible style to maximize their impact on economic understanding and economic policymaking; the intended audience includes analysts from universities, governments, and businesses. Topics covered by the journal include fiscal and monetary policy, consumption and saving behavior, business investment, housing, asset pricing, labor markets, wage- and price-setting, business cycles, long-run economic growth, the distribution of income and wealth, international capital flows and exchange rates, international trade and development, and the macroeconomic implications of health costs, energy supply and demand, environmental issues, and the education system.

We would like to thank the supporters of the *Brookings Papers on Economic Activity* conference and journal, including the Alfred P. Sloan Foundation; BlackRock Financial Management, Inc.; the National Science Foundation, under Grant No. 1756544; and State Farm Mutual Automobile Insurance Company.

The papers and discussant remarks reflect the views of the authors and not necessarily the views of the funding organizations or the staff members, officers, or trustees of the Brookings Institution.

CALL FOR PAPERS

Most papers that appear in the *Brookings Papers on Economic Activity* are solicited by the editors, but the editors also consider submitted proposals. Editorial decisions are generally made about a year in advance of each conference. Therefore, proposals should be received by September 1 for the following fall conference and by April 1 for the following spring conference. Proposals should be sent to bpea@brookings.edu.

ACCESSING THE JOURNAL

All past editions of the *Brookings Papers on Economic Activity*—along with appendix materials, data, and programs used to generate results—are made freely available for download at www.brookings.edu/past-bpea-editions. To purchase subscriptions or single copies, visit www.brookings.edu/press, or contact the Brookings Institution Press at (866) 698-0010 or P.O. Box 465, Hanover, PA 17331-0465. Brookings periodicals are available online through both the Online Computer Library Center (contact the OCLC subscription department at (800) 848-5878) and Project Muse (<http://muse.jhu.edu>). Archived issues of the *Brookings Papers on Economic Activity* are also available through JSTOR (www.jstor.org).

**EDITORS,
AUTHORS,
DISCUSSANTS,
ADVISERS,
PARTICIPANTS,
AND STAFF
FOR THE
ONE HUNDRED
AND SIXTH
CONFERENCE**

Stephanie Aaronson *Brookings Institution*
Randall Akee *Brookings Institution*
George Akerlof *Georgetown University*
Dany Bahar *Brookings Institution*
Martin Neil Baily *Brookings Institution*
Joe Beaulieu *Brevan Howard*
Sigridur Benediktssdottir *Yale University*
Ben S. Bernanke *Brookings Institution*
Lorenzo Bini Smaghi *Société Générale*
Olivier Blanchard *Peterson Institute for International Economics*
Alan Blinder *Princeton University*
Jean Boivin *BlackRock*
Lael Brainard *Federal Reserve Board*
William Brainard *Yale University*
Steven Braun *Council of Economic Advisers*
Jesse Bricker *Federal Reserve Board*
Jeremy Bulow *Stanford Business School*
Mark Calabria *The White House – Office of the Vice President*
Jeffrey Campbell *Federal Reserve Bank of Chicago*
Marcus Casey *Brookings Institution*
Tarun Chhabra *Brookings Institution*
Gerald Cohen *Beacon Macro*
Olivier Coibion *University of Texas at Austin*
Susan Collins *University of Michigan*
Charles Collyns *International Monetary Fund*
Diego Comin *Dartmouth College*
Richard Cooper *Harvard University*
Jason Cummins *Brevan Howard*
Troy Davig *Rokos Capital*
Steven J. Davis *University of Chicago*
John Driscoll *Federal Reserve Board*
Karen Dynan *Harvard University*
Janice Eberly *Northwestern University*
Wendy Edelberg *Congressional Budget Office*
Bruce Fallick *Federal Reserve Bank of Cleveland*
Emmanuel Farhi *Harvard University*
Martin Feldstein *Harvard University*
Stanley Fischer *Federal Reserve Board (ret.)*
Martin Fleming *IBM*
Kristin Forbes *Massachusetts Institute of Technology–
Sloan School of Management*
Benjamin Friedman *Harvard University*
Renee Fry-McKibbin *Australian National University*
Jeff Fuhrer *Federal Reserve Bank of Boston*
Jason Furman *Harvard Kennedy School*
Joseph Gagnon *Peterson Institute for International Economics*
William Gale *Brookings Institution*
Ted Gayer *Brookings Institution*
Mark Gertler *New York University*
Subir Gokarn *International Monetary Fund*
Robert J. Gordon *Northwestern University*
Yuriy Gorodnichenko *University of California, Berkeley*
François Gourio *Federal Reserve Bank of Chicago*

Megan Greene *John Hancock Financial Services /
Manulife Asset Management*
 Mark Griffiths *International Monetary Fund*
 Robert E. Hall *Stanford University*
 John Haltiwanger *University of Maryland*
 James D. Hamilton *University of California, San Diego*
 Philipp Hartmann *European Central Bank*
 Joshua Hausman *University of Michigan*
 Peter Henry *NYU Stern School of Business*
 Caroline Hoxby *Stanford University*
 Sebnem Kalemli-Ozcan *University of Maryland*
 Michael Kiley *Federal Reserve Board*
 Jeffrey Kling *Congressional Budget Office*
 Alan Krueger *Princeton University*
 Adriana Kugler *Georgetown University*
 Andreas Lehnert *Federal Reserve*
 Nellie Liang *Brookings Institution*
 John Lipsky *Johns Hopkins School of Advanced International Studies*
 Byron Lutz *Brookings Institution*
 Christos Makridis *Massachusetts Institute of Technology—Sloan
School of Management*
 N. Gregory Mankiw *Harvard University*
 Donald Marron *Urban Institute*
 Mark Mazur *Urban-Brookings Tax Policy Center*
 Warwick McKibbin *Australian National University*
 Karel Mertens *Federal Reserve Bank of Dallas*
 Frederic Mishkin *Columbia University*
 Adele Morris *Brookings Institution*
 Christopher Nekarda *Federal Reserve Board*
 Edward Nelson *Federal Reserve Board*
 Serena Ng *Columbia University*
 Giovanni P. Olivei *Federal Reserve Bank of Boston*
 Athanasios Orphanides *Massachusetts Institute of Technology*
 Dimitris Papanikolaou *Northwestern University*
 Karen Pence *Federal Reserve Board*
 George Perry *Brookings Institution*
 Jonathan Pingle *BlackRock*
 Adam Posen *Peterson Institute for International Economics*
 Eswar Prasad *Brookings Institution*
 Raghuram Rajan *University of Chicago Booth School of Business*
 Valerie Ramey *University of California, San Diego*
 Lucrezia Reichlin *London Business School*
 David Reifschneider *Colby College*
 Alice Rivlin *Brookings Institution*
 David Romer *University of California, Berkeley*
 Eric S. Rosengren *Federal Reserve Bank of Boston*
 Martha Ross *Brookings Institution*
 Glenn Rudebusch *Federal Reserve Bank of San Francisco*
 John Sabelhaus *Federal Reserve Board*
 Claudia Sahm *Federal Reserve Board*
 Brian Scholl *Securities and Exchange Commission*
 Hershil Shah *Council of Economic Advisers*
 Jay Shambaugh *Brookings Institution*

Louise Sheiner *Brookings Institution*
Tara Sinclair *George Washington University*
Frank Smets *European Central Bank*
Jón Steinsson *University of California, Berkeley*
James H. Stock *Harvard University*
Eric T. Swanson *University of California, Irvine*
Robert Tetlow *Board of Governors of the Federal Reserve*
Geoffrey M.B. Tootell *Federal Reserve Bank of Boston*
Mauricio Ulate *University of California, Berkeley*
Stan Veuger *American Enterprise Institute*
Alan Viard *American Enterprise Institute*
Ivan Vidangos *Federal Reserve Board*
Philip Wallach *R Street Institute*
David Wessel *Brookings Institution*
David Wilcox *Federal Reserve Board*
Jon Willis *Federal Reserve Bank of Kansas City*
Justin Wolfers *University of Michigan*
Janet Yellen *Brookings Institution*
Jeromin Zettelmeyer *Peterson Institute for International Economics*

Francisca P. Alba *Brookings Institution*
Alfred F. Imhoff *Brookings Institution*
Mary E. King *Brookings Institution*

PHILIPP HARTMANN

European Central Bank

FRANK SMETS

European Central Bank

The European Central Bank's Monetary Policy during Its First 20 Years

ABSTRACT On June 1, 2018, the European Central Bank (ECB) celebrated its 20th anniversary. This paper provides a comprehensive view of the ECB's monetary policy over these two decades. The first section gives a chronological account of the macroeconomic and monetary policy developments in the euro area since the adoption of the euro in 1999, going through four cyclical phases "conditioning" ECB monetary policy. We describe the monetary policy decisions from the ECB's perspective and against the background of its evolving monetary policy strategy and framework. We also highlight a number of the key, critical issues that were the subject of debate. The second section contains various assessments. We analyze the achievement of the price stability mandate and developments in the ECB's credibility, and we also investigate the ECB's interest rate decisions through the lens of a simple empirical interest rate reaction function. Finally, we present the ECB's framework for thinking about nonstandard monetary policy measures and review the evidence on their effectiveness. One of the main themes of the paper is how the ECB utilized its monetary policy to respond to the challenges posed by the European twin financial and sovereign debt crises and the subsequent slow economic recovery, making use of its relatively wide range of instruments, defining new ones where necessary, and developing the strategic underpinnings of its policy framework.

Conflict of Interest Disclosure: The authors are senior managers of the European Central Bank, the public authority whose monetary policy is discussed in this paper. They did not receive any financial support from any firm or person with a financial or political interest in the paper. No outside party had the right to review the paper before circulation.

European Economic and Monetary Union (EMU) is an unprecedented historical project, in which 11 European Union countries initially introduced a common currency—the euro—with a single central bank—the European Central Bank (ECB)—and a single monetary policy. By the time of writing, 19 quite diverse EU countries have joined the euro area, meaning that the ECB runs the monetary policy for about 341 million citizens (compared with about 326 million citizens for the U.S. Federal Reserve System) or an economic area that constitutes 11.6 percent of the world's GDP (compared with 15.3 percent for the U.S. or 18.2 percent for China, all in terms of purchasing power parity) (ECB 2018b). The motivation for this paper is that on June 1, 2018, the ECB celebrated its 20th anniversary. As two economists who have been on the staff of the ECB from the beginning, we take this opportunity to look back at the first two decades of our institution, describing and assessing its experience with monetary policy.

An important starting point is the statutory objectives of the ECB, as laid down in the Treaty on the Functioning of the European Union and the Treaty on European Union (EU 2012a, 2012b).¹ The ECB's primary objective is to maintain price stability. Without prejudice to the objective of price stability, the ECB also supports the general economic policies of the European Union, with a view to contributing to the achievement of its objectives. These (often called secondary) objectives include, for example, balanced economic growth and a highly competitive social market economy, aiming at full employment and social progress. This hierarchy of objectives is interpreted in lexicographic order (Driffill and Rotondi 2004; Artus and others 2008). Only to the extent that the primary objective is fulfilled can the ECB consider growth and employment. Such "single" central bank mandates, focusing on price stability as the primary objective, are quite common in advanced economies. For example, they apply to all the central banks of the Group of Seven, except the U.S. Federal Reserve.

From the euro's introduction in January 1999—the beginning of stage 3 of EMU—the ECB started with a strong and self-contained mandate to define and implement monetary policy for the euro area. For other tasks that central banks often fulfill, however, it had more indirect or

1. For simplicity, we are abstracting from the legally precise distinctions between the ECB, the Eurosystem (comprising the ECB and the national central banks of countries that have joined EMU), and the European System of Central Banks (comprising the Eurosystem and all other EU central banks). National central banks play an important role in ECB decisions, their preparation, and implementation; but unfortunately, we do not have the space in this paper to provide a proper account of these collective aspects of Eurosystem functioning.

contributing roles, notably in the prudential and financial stability arena.² (In November 2014, however, the ECB was given the role of banking supervisor for the countries that joined the European Banking Union—which is congruent with the euro area (EU 2013). It needs to conduct banking supervision and monetary policy separately.)

This paper focuses on the ECB's experience conducting monetary policy for the euro area.³ Our overall goals are to provide a rigorous and comprehensive “inside” view of what the ECB has been concerned with in this area, how its monetary policy has evolved during its first 20 years, and how it has performed in achieving its primary objective of maintaining price stability. Obviously, one main theme of the paper is how the ECB has responded to the enormous challenges posed by the European twin crises (in the European case, the financial crisis of 2007–9 morphed into the sovereign debt crisis of 2010–13) and the subsequent slow economic recovery, making use of its relatively wide range of instruments, defining new ones where necessary, and developing the strategic underpinnings of its policy framework. But given the main motivation for our paper, we should not limit the attention only to the second decade of the ECB's existence.

Before we delve deeper into the details of the ECB's monetary policy, we provide a perspective on the broader issues with which the ECB has been concerned during the last 20 years via the themes that ECB Executive Board members have addressed in their public communications.⁴ Figure 1 shows the number of public speeches Board members gave every year between 1999 and 2017. The figure's different shades and patterns refer to the shares of these speeches that were dedicated to any of nine different themes. We did not predetermine these themes. Instead, we applied a machine-learning approach to uncover them from the texts of the 1,892 board speeches displayed on the ECB's website for the period May 1998–April 2018. (As of 2014, the data set also began to include the speeches by the ECB's chair and vice chair, and the four ECB representatives from the Supervisory Board of the Single Supervisory Mechanism.)

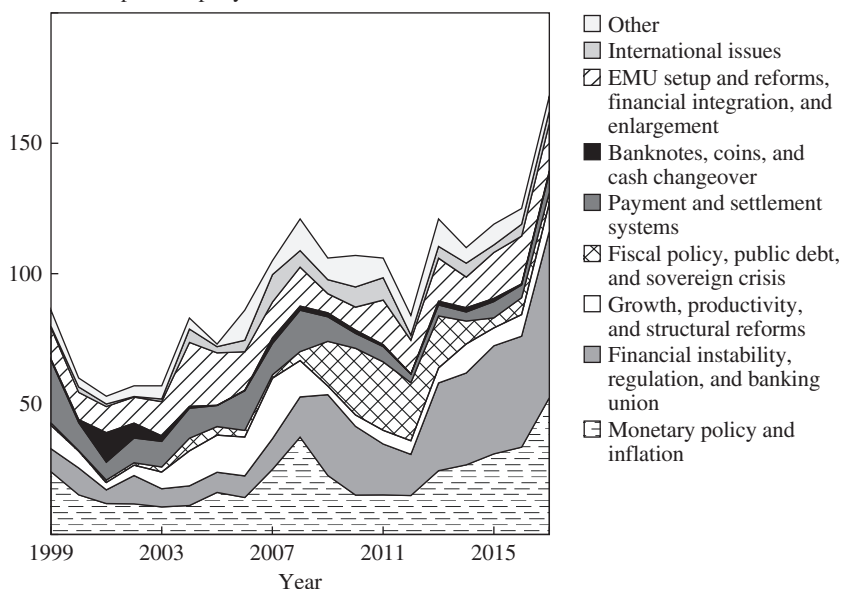
2. For complete lists of ECB tasks and functions, see EU (2012b, 2012c).

3. In a companion paper (forthcoming), we discuss the ECB's experience with financial stability.

4. The Board comprises the ECB president, vice president, and four further members, who are appointed by the European Council, usually for a term of eight years (EU 2012b, 2012c). They are collectively responsible for the current business of the ECB and play an important role in the Governing Council, the main decisionmaking body of the ECB and the Eurosystem. The other Governing Council members are the governors of the euro area national central banks.

Figure 1. Speeches by the ECB's Executive and Supervisory Board Members and Their Decomposition in General Themes, 1999–2017^a

Number of speeches per year



Sources: Authors' research; ECB data.

a. The figure is based on an application of the Latent Dirichlet Allocation methodology (Blei, Ng, and Jordan 2003) to identify the topics addressed in the public speeches given by Executive Board members of the European Central Bank between May 1998 and April 2018. All speeches on the ECB website section (<https://www.ecb.europa.eu/press/key/date/2017/html/index.en.html>) have been considered as documents. Since 2014, the speeches by the chair, vice chair, and ECB Supervisory Board members of the Single Supervisory Mechanism have also been included (<https://www.bankingsupervision.europa.eu/press/speeches/date/2017/html/index.en.html>).

Overall, the document set comprises 1,892 speeches. The figure shows results only for full years, that is, 1999–2017 (1,829 speeches). The upper line shows the total number of speeches per year. The shades and patterns of the areas underneath describe for a given year the shares of these speeches addressing nine general themes; see the legend. The themes have been derived by the authors grouping the topics found by the Latent Dirichlet Allocation machine-learning algorithm, which defines a topic as a set of words that occur together within documents and derives the probability that a given document addresses this topic. Applying the metric of Cao and others (2009), the total number of topics has been set at 50. A speech can address more than one topic. The full list of topics and their grouping in themes is available from the authors upon request.

Using the Latent Dirichlet Allocation method for textual analysis (Blei, Ng, and Jordan 2003) and the metric developed by Juan Cao and others (2009) for the optimal number of topics, we identify 50 specific topics that have been addressed in these speeches over time. For the purpose of the first general overview given in figure 1, we group this rather large number of topics into the 9 general themes displayed.

The results give a good impression of the breadth of issues that the ECB was concerned with (via the external communication of Board members) and how they changed over time. First, the core theme of “monetary policy and inflation” (the dashed area in figure 1) covered a sizable share most of the time, but it was particularly important at the time of the ECB’s inception in 1999; when the financial crisis hit, in the years 2007–9; and during the post-sovereign debt crisis, low-inflation recovery period, 2013–17. Clearly, these were three periods with increased needs for monetary policy communication. Second, financial stability and supervisory issues received particular attention when the financial crisis struck and after the 2012 agreement about the European Banking Union that granted supervisory responsibility to the ECB (the medium gray area, second from the bottom, of figure 1). Third, growth and productivity (the white area), fiscal matters (the checkered area) and international developments (the lighter gray area, second from the top), which all have implications for the conduct of monetary policy, received regular attention. But the attention paid to public debt and sovereign risk (part of the checkered area) was most pronounced when the financial crisis morphed into the European sovereign crisis in 2009–10. Structural reforms, productivity and competitiveness issues (part of the white area) were very much discussed before the start of the financial crisis.

Board members also addressed a number of other themes of great importance for the ECB that we do not touch upon at all in this paper. For example, one can see in figure 1 that in 2001—before the introduction of euro notes and coins in 2002—Board members prepared the public for the cash changeover (the black area of figure 1). In 2004 and 2005, at the time of the EU’s major eastern enlargement, they communicated more about accession and convergence issues (part of the diagonally striped area). Finally, payment and settlement issues (the dark gray area, fifth from the top) played a greater role in Board members’ external communications (in 2006), shortly before the ECB’s initial TARGET large-value payment system migrated to the single-platform TARGET2 system, and during 2008, when the TARGET2-Securities project was launched to establish a single, pan-European platform for securities settlement.


Returning to this paper's main focus, figure 2 shows only the speeches dealing with the core theme of "monetary policy and inflation" (the dashed area of figure 1), breaking it down in the shares of the 10 underlying topics that fall under this theme (out of 50 found by our machine-learning algorithm). In the beginning, the largest focus was on explaining the new institution's monetary policy strategy, including its monetary analysis aspect (the medium gray area, second from the bottom, of figure 2; see subsection I.A below). In the financial crisis period, the ECB's market operations and liquidity management (the white area) became an important focus, in line with the policy approach taken at the time (see subsection I.C below). In the last period, the focus of the ECB's Board members' public speeches moved to how the ECB used nonstandard monetary policy measures, such as large-scale asset purchases (the black area) and negative interest rates (the checkered area areas in figure 2), to strengthen the fragile recovery and ensure that inflation would return to the ECB's objective (the light gray area at the top) in an environment of interest rates close to their effective lower bound (see subsection I.C below).

The rest of the paper is organized in two main sections. Section I provides a chronological account of the macroeconomic, monetary, and financial developments in the euro area since the adoption of the euro, as well as of the ECB's monetary policy decisions. We divide the section into the four cyclical phases that "conditioned" ECB monetary policy between 1999 and 2018: the end of the technology cycle, the economic upturn with a buildup of imbalances, the "double-dip" recessions associated with the financial and sovereign debt crises, and the low-inflation recovery. Each of these four subsections in turn has three divisions: first, on developments in the ECB's monetary policy strategy and operational framework; second, on the conjuncture and actual decisions; and third, a discussion highlighting critical issues that were the subject of public debate.

In section II, we assess selected aspects of the ECB's monetary policy in the last 20 years. We first analyze the achievement of the price stability mandate and developments in the ECB's credibility and discuss possible implications for the ECB's inflation aim (subsection II.A). Next, we examine the ECB's interest rate decisions through the lens of a simple empirical interest rate reaction function (subsection II.B). This is appropriate until the ECB hits the zero lower bound in mid-2012. Finally, we present the ECB's framework for thinking about nonstandard monetary policy measures—many of which draw on its broad and flexible

Figure 2. Speeches by ECB Executive Board Members on Monetary Policy and Inflation and Their Decomposition in Topics, 1999–2017^a

Number of speeches per year




Sources: Authors' research; ECB data.

a. This figure is based on an application of the Latent Dirichlet Allocation methodology (Blei, Ng, and Jordan 2003) to identify the topics addressed in the public speeches given by Executive Board members of the European Central Bank between May 1998 and April 2018. All speeches on the ECB website section (<https://www.ecb.europa.eu/press/key/date/2017/html/index.en.html>) have been considered as documents. Since 2014, the speeches by the chair, vice chair, and ECB Supervisory Board members of the Single Supervisory Mechanism have also been included (<https://www.bankingsupervision.europa.eu/press/speeches/date/2017/html/index.en.html>).

Overall, the document set comprises 1,892 speeches. The figure shows results only for full years, that is, 1999–2017 (1,829 speeches). Applying the metric of Cao and others (2009), the total number of topics has been set to 50. The figure only refers to the 10 topics—see the legend—that can be grouped to a general theme denoted as “monetary policy and inflation” (see also figure 1). The upper line refers to the number of speeches per year addressing these 10 topics. The shades and patterns of the areas underneath describe the shares of these speeches addressing each topic for a given year. A speech can address more than one topic. The Latent Dirichlet Allocation machine-learning algorithm defines a topic as a set of words that occur together within documents and derives the probability that a given document addresses this topic. The descriptions of the topics shown in the legend have been formulated by the authors, based on the words included in the different topics and their reading of the speeches that addressed the topics with high likelihood.

Figure 3. Output Gap Estimates and the Unemployment Rate for the Euro Area, 1999–2017^a



Sources: European Commission (2018); International Monetary Fund (2018); Organization for Economic Cooperation and Development (2018); ECB data.

a. Yearly data and estimations are the latest available, and therefore are not in real time. The Unobserved Components Model (ECB) has been run following Szörfi and Tóth (2018), and the output gap estimates from it should not be regarded as official ECB output gaps. The uncertainty bands refer only to the Unobserved Components Model (ECB) output gap. Vertical dashed lines are indicative of business cycle troughs and peaks (also see figure 4). The most recent observations are for 2017.

framework for market operations—and we review the evidence on the effectiveness of the nonstandard instruments that have been used (subsection II.C). Section III offers conclusions, and considers how completing EMU could support the ECB’s monetary policy.

I. Two Decades of ECB Monetary Policy: From the Two Pillars to Quantitative Easing

This section discusses the ECB’s monetary policy during the past two decades. It gives a chronological overview of the main macroeconomic, monetary, and financial developments in the euro area since the euro’s adoption in January 1999 and how the ECB has responded to them in pursuit of its price stability mandate. Taking a business cycle perspective,

we use—among other indicators—the euro area output gap estimates and unemployment rate shown in figure 3 for identifying troughs and peaks (marked by dashed vertical lines). This leads to four episodes corresponding to cyclical downturns and upturns: the initial period of a growth slowdown following the collapse of the dot-com bubble accompanied by a weak euro exchange rate, 1999–2003; the boom period in money and credit growth accompanied by relatively stable inflation and accelerating growth, 2003–7; the subsequent double-dip recession due to the start of the U.S. financial crisis and the emergence of the euro area sovereign debt crisis, 2007–13; and, finally, the most recent low-inflation recovery period, 2013–18. Figure 4 (on the next two pages) provides an overview of the four periods—here marked with alternating gray and white areas whose transitions (business cycle troughs and peaks) are dated to the month—for main macroeconomic variables (on the next page) and two key monetary policy indicators (on the subsequent page). The figure also marks major ECB monetary policy actions (on the subsequent page) and other important events (on the next page) that were characteristic of the respective period.

1.A. The Beginning of ECB Monetary Policy toward the End of the Technology Cycle, January 1999–June 2003

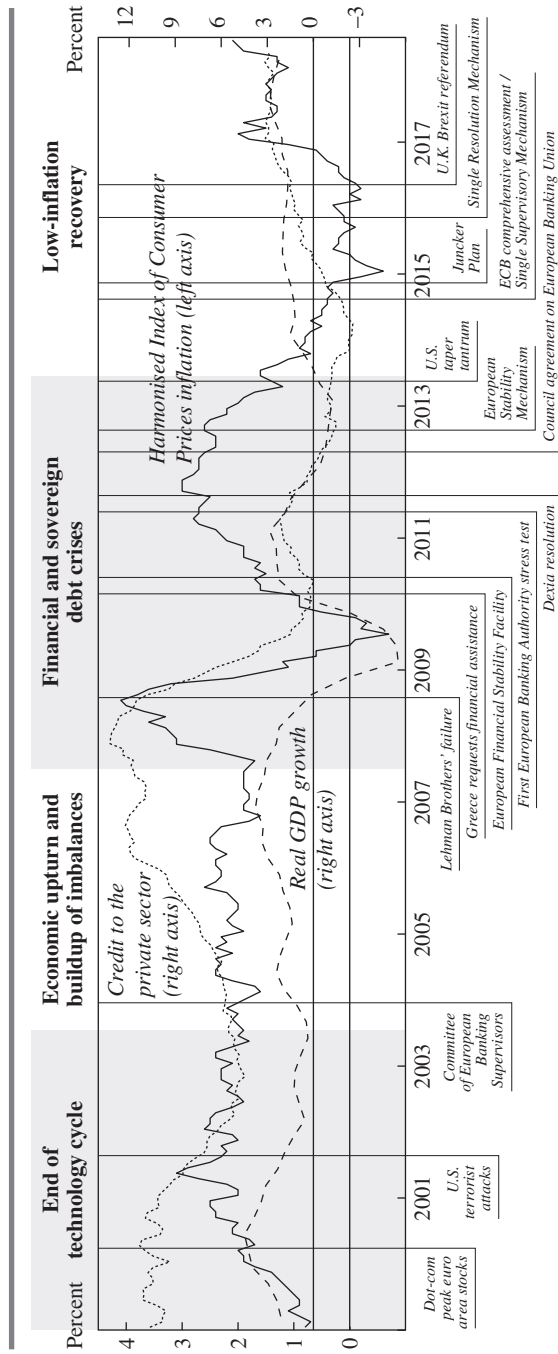
The beginning of the ECB's monetary policy was characterized by the first application of a new monetary policy strategy and framework. Challenges during the first cyclical period included a test of the ECB's anti-inflationary resolve related, among other things, to a protracted depreciation of the euro and a reversal of the perspective due to the collapse of the technology euphoria from the late 1990s.


A NEW STABILITY-ORIENTED MONETARY POLICY STRATEGY AND OPERATIONAL FRAMEWORK FOR THE ECB As discussed in the introduction, the treaty creating the European Economic and Monetary Union established price stability as the primary objective of monetary policy in the euro area. Under the leadership of Otmar Issing, its first chief economist, the ECB early on developed a monetary policy strategy that had the aim of providing a solid basis for the conduct and communication of monetary policy in pursuit of price stability.⁵ It also developed an elaborate operational framework for implementing monetary policy decisions.⁶

5. For an extensive presentation and justification of the original two-pillar, stability-oriented monetary policy strategy, see ECB (1999); and Issing and others (2001).

6. See ECB (2000a). The monetary policy strategy and operational framework were developed building on the extensive preparatory work carried out by the European Monetary Institute—the ECB's predecessor.

Figure 4. Four Cyclical Phases during the First 20 Years of the Euro: Key Macroeconomic and Monetary Policy Variables and Major Events, 1999–2018^a





Sources: Authors' research; ECB data.

a. Gray areas (downturns) and white areas (upturns) reflect four cyclical periods with troughs in about June–July 2003 and June–July 2013 and a peak in about July–August 2007 dated on the basis of a large number of indicators (also see figure 3). The panel on this page shows ECB monetary policy variables and events, and the panel on the previous page shows other euro area macroeconomic variables and events. ECB interest rate policy is proxied by the Main Refinancing Operations Rate before October 2008 and by the Deposit Facility Rate thereafter. The total amount of ECB monetary policy operations (the dashed line; see figure 16 for its components) is indicative of nonstandard monetary policy as of October 2008, and particularly when policy rates approached 0 or declined below. Harmonised Index of Consumer Prices inflation is the annual growth rate of these prices. Credit to the private sector is total loans and securities by monetary financial institutions vis-à-vis euro area nonmonetary financial institutions, firms, and households—that is, excluding general government.

In addition to being operationally ready from day one, there were two main and interrelated challenges. The first challenge was to establish, as quickly as possible, the credibility of the new institution for maintaining price stability. A high level of initial credibility would facilitate the transition to EMU and reduce the potential costs of having to build such credibility.⁷ The second challenge was to ensure a consistent and systematic approach to the conduct of monetary policy in an uncertain economic environment after a fundamental regime change and where the national central banks preceding the ECB/Eurosystem had different frameworks and traditions. Robustness in the face of pervasive uncertainty and country heterogeneity was seen as an important guiding principle for the design of the new strategy (Issing and others 2005; Issing 2008). In response to these two challenges, three main components were developed: first, a quantitative definition of the ECB's primary objective of price stability as a clear yardstick for accountability; second, a two-pillar framework as the organizing principle for the analysis underlying the assessment of the outlook for price developments and for a structured policy discussion; and third, an elaborate communication and accountability framework. Before describing the economic and monetary developments in this initial phase, we briefly describe these three elements. We also briefly characterize the initial operational framework. As we discuss in subsequent sections, elements of these building blocks have evolved in response to challenges over time (Constancio 2018).

The quantitative definition of price stability. In December 1998, the Governing Council of the ECB adopted a quantitative definition of price stability, which reads: "Price stability shall be defined as a year-on-year increase in the Harmonised Index of Consumer Prices (HICP) for the euro area of below 2 percent." Price stability "is to be maintained over the medium term" (ECB 1999, 46). This definition allows economic agents and observers to assess the ECB's performance at any time and over any horizon. It enhances the ECB's accountability by forcing the central bank to explain why inflation has at times deviated from its definition, and it thereby helps anchor medium- to long-term expectations. The definition focuses on the euro area as a whole, reflecting the fact that, within a monetary union, monetary policy cannot address country-specific inflation developments. It makes clear that medium-term inflation above 2 percent is not consistent with price stability. However, it also implies

7. Whether the ECB would have a deflationary bias in order to establish its anti-inflation credibility was hotly debated at the time. See, for example, Begg and others (1998, 1999).

that very low inflation rates, and especially deflation, are not consistent with price stability either. Following criticism of the perceived asymmetry of the quantitative definition, this was clarified—for example, by the ECB's president, Willem Duisenberg, in an early speech explaining the new strategy.⁸

Another important feature is the medium-term orientation of the ECB's strategy. Because monetary policy can affect price developments only with significant and variable time lags, and only to an uncertain extent, it is impossible to maintain a specific, predefined inflation rate at all times or to bring it back to a desired level within a very short period. Consequently, monetary policy needs to act in a forward-looking manner and focus on the medium term. This also helps to avoid excessive activism and the introduction of unnecessary volatility into the real economy, thereby contributing to the stabilization of output and employment. See, for example, the papers by Nicoletta Batini and Edward Nelson (2001) and Frank Smets (2003), who show the equivalence between the length of the policy horizon and the weight on output gap stabilization.

Against the background of the inflation forecast targeting strategies that were popular at the time, two aspects of the ECB's medium-term orientation are worth mentioning. First, the ECB has always emphasized that there is no fixed time horizon over which price stability needs to be reestablished, given that monetary policy should react differently to different sources of economic shocks (for example, demand versus supply shocks).⁹ Second, the medium-term orientation implies a lengthening of the monetary policy horizon beyond the usual two years typically associated with the horizon of inflation forecasts and the lags in monetary policy transmission. For example, then-ECB president Jean-Claude Trichet (2003b) states that “monetary policy needs to focus on the period covering the whole transmission process, bearing in mind that this may sometimes

8. As Duisenberg (1999) stated: “Some observers have criticised this strategy as ‘asymmetric.’ In other words, they argue that the Eurosystem is more concerned about inflation than it is about deflation. In their view, such asymmetry will impose a drag on the overall performance of the euro area economy as a whole because monetary policy will be overly restrictive on average, and risks triggering a damaging deflationary spiral in some circumstances. . . . I reject this criticism. The use of the word ‘increases’ in the definition imposes a floor of at least zero for the lower bound. . . . Let me state categorically, as I have often done in the past, that neither prolonged inflation nor prolonged deflation in the euro area would be deemed by the Governing Council to be consistent with the maintenance of price stability.”

9. This feature of the ECB's monetary policy strategy was eventually also adopted in inflation-targeting central banks, which have also recognised the need for a more flexible policy horizon. See, for example, Bean (2003).

span a protracted period of time.” As a result, the horizon for evaluating the credibility of the central bank should also extend beyond two years. In subsection II.A, we take an admittedly somewhat arbitrary five-year horizon, which typically should be enough to let the effects of the shocks that the central bank cannot control wash out.

The two-pillar framework. In the original formulation (ECB 1999), the “two pillars” of the ECB’s strategy were described as (1) a prominent role for money, as signaled by the announcement of a quantitative reference value for the growth rate of a broad monetary aggregate, known as M3; and (2) a broadly based assessment of the outlook for price developments and risks to price stability in the euro area as a whole, which includes the macroeconomic projections.¹⁰ The two-pillar framework was a unique feature of the ECB’s monetary policy strategy and was seen as a partial answer to the two challenges described above. First, the prominent role for money would help the ECB gain rapid credibility by borrowing some of the elements of the Deutsche Bundesbank’s stability-oriented monetary policy strategy. Second, the two-pillar framework would allow the ECB to bring different traditions under one roof and provide a robust framework in an environment of high uncertainty, pervasive structural change, cross-country heterogeneity, and convergence. It would also bring together perspectives from the two leading economic paradigms—Keynesianism and monetarism—that had very much shaped macroeconomic debates in the preceding decades, rather than focusing mainly on one of them.

The reference value for M3 growth (see the thick gray dashed line in figure 9 below) underlined both the relative importance of the role of money and the medium-term orientation of the ECB’s strategy. At a press conference on October 13, 1998, President Duisenberg was asked about the relative weight of the two pillars. He replied: “It is not a coincidence that I have used the words that money will play a prominent role. So if you call it the two pillars, one pillar is thicker than the other is, or stronger than the other, but how much I couldn’t tell you” (Duisenberg 1998). The choice of M3 was based on the evidence that this monetary aggregate exhibited a close relationship with the price level. At the same time, it was made clear from the very beginning that monetary policy would not react mechanically to deviations of M3 growth from the reference value; it was not a monetary growth target (ECB 1999). The

10. In fact, the internal briefing process supporting the Governing Council’s monetary policy decisions was, and still is at the time of writing, organized along the two pillars, later called economic and monetary analysis (see subsection I.B).

monetary pillar also involved an analysis of different monetary aggregates and the asset side of the banking system, in particular developments in credit to firms and households.

The reference value for M3 growth of 4.5 percent implicitly also revealed that the ECB was aiming at the upper half of the below-2-percent price stability definition. Because, in this period, the trend growth rate of GDP was assumed to lie in the range of 2 to 2.5 percent and the trend rate of decline in the velocity of circulation of M3 in the range of -0.5 to -1 percent, the arithmetic of the quantity equation for money suggests an operational inflation aim between 1 to 2 percent. Although it was never explicitly acknowledged, this was consistent with the emphasis on positive, but uncertain, measurement biases in HICP inflation (up to 1 percent), which was given as one of the explanations for why the ECB did not formulate a clear lower bound in the quantitative definition of price stability (Issing and others 2001).

Communication and accountability. Legally, the ECB is probably one of the most independent central banks in the world. Its independence is not simply a result of domestic law, but is based on the international Treaty on European Union (EU 2012a). Changing this treaty would require the agreement of every signing country.¹¹ Transparency and clear communication are a natural complement to strong independence, because it makes it easier to hold the central bank accountable, which in turn is a key element to maintain political support for the ECB's high degree of independence.¹² Clear communication is also important for effectively conducting monetary policy, as it helps anchor inflation expectations, reduce policy-induced uncertainty, and make the transmission process of policy decisions more effective.

From the very outset, the ECB put great emphasis on communicating its policy actions and the economic rationale underlying its decisions to financial market participants and the general public in a transparent and timely manner. Since the start, the main communication vehicle has been the monetary policy press conferences held by the president and the vice president

11. The treaty gives the ECB and the members of its decisionmaking bodies (the Governing Council) a very high degree of institutional (vis-à-vis Community institutions or bodies and any government of a member state), personal (relatively long fixed-term contracts), financial (own budget), and functional (exclusive competence for monetary policy in the euro area and prohibition of monetary financing) independence.

12. See Tucker (2018) for a recent in-depth discussion of the political economy of central bank independence.

immediately after each monetary policy Governing Council.¹³ On this occasion, the introductory statement is presented by the president on behalf of the Governing Council. It provides a summary of the policy-relevant assessment of economic and monetary developments, as well as the monetary policy stance, and it is structured along the lines of the ECB's monetary policy strategy. The press conference includes a question-and-answer session, which is attended by key media representatives from across the euro area and beyond. The press conference was seen as an effective means of presenting and explaining in a very timely manner the discussions in the Governing Council, and thus the monetary policy decisionmaking process. In the context of a global trend toward more detailed and transparent communications by central banks, this feature of the ECB's communication strategy has increasingly been adopted by other central banks (such as the Federal Reserve).¹⁴ Other important communication channels used by the ECB are the *Monthly Bulletin* (since January 2015, this has been called the *Economic Bulletin*, and it is published less frequently than monthly), which gives a detailed and comprehensive analysis of the economic environment and monetary developments, the quarterly appearances of the ECB president before the European Parliament's Committee on Economic and Monetary Affairs (Fraccaroli, Giovannini, and Jamet 2018), and a large number of public speeches (see figures 1 and 2) and interviews with media by members of the Executive Board.

The operational framework. The monetary policy stance decided on by the Governing Council is implemented through ECB market operations. As a matter of fact, the statute of the ECB delegated the conduct of these operations to the Executive Board from the start of the euro (see Article 12.1, second paragraph, in EU 2012c), creating some separation of the operational decisions from the general monetary policy debate. The operational decisions are further executed in a decentralized way between the counterparties and their respective national central banks. Originally, the market operations aimed primarily at keeping very-short-term money market rates close to the policy rate decided by the council. More goals

13. The frequency of the monetary policy Governing Council meetings was monthly (the first Governing Council meeting of the month) until December 2014, and was changed to eight times a year as of 2015 (a frequency very similar to the U.S. Federal Reserve's Federal Open Market Committee meetings). One reason for this was that after one month, often only a limited amount of new information was available but the fact that a new monetary policy decision had to be made could lead to some market volatility.

14. See, for example, Yellen (2012).

were only added much later, when the policy rate came close to its effective lower bound and other means than short-term interest rates needed to be used for easing monetary policy further (see subsection I.D). In designing its operational tools, the ECB prioritizes what is needed for the smooth implementation of its monetary policy. Next, it considers what is good for market functioning, neutrality, and risk control.

One important feature of the ECB's operational framework is its breadth, despite a focus on banks related to the predominant financial structure in the euro area (Hartmann, Maddaloni, and Manganelli 2003). Given the diversity of euro area countries' financial systems, a very broad framework was needed to meet the above-noted criteria. The very long experience of many euro area national central banks was particularly helpful in this regard. Until the present day, the ECB's operational framework has contained four types of instruments: (1) open market credit operations; (2) standing facilities; (3) minimum reserve requirements; and (4) outright asset purchases.¹⁵

Initially, only the first three instruments were actively used. Reserve requirements extended the liquidity deficit of the banking sector vis-à-vis the central bank that bank note issuance and government deposits create. Euro area banks need to hold a small share of their short-term liabilities (2 percent until January 2012, and 1 percent thereafter) on their Euro-system accounts, and these required reserves are remunerated at the rate set by the ECB for its Main Refinancing Operations (MROs), the MRO Rate (MROR). This needs to be the case only on average over a reserve maintenance period of a few weeks. Normally, the averaging procedure has a stabilizing effect, because it encourages liquidity planning and helps mitigate the effects of unexpected short-term liquidity shocks—the main purpose of the reserve requirements.

Open market operations allow ECB counterparties to acquire the liquidity needed to close the aggregate deficit, so that short-term money market rates stay close to the policy rate decided by the Governing Council. Before the European crises, the bulk of the liquidity was provided through MROs, so the MROR constituted a key policy rate for the Governing Council. MROs started as weekly tenders of two-week collateralized credit

15. The complete formal description of the framework is published in the *Official Journal of the European Union* as the “Guideline of the European Central Bank on the Implementation of the Eurosystem Monetary Policy Framework,” which originally was often called “General Documentation of Eurosystem Monetary Policy Instruments and Procedures.” An updated version can be found on the ECB website: https://www.ecb.europa.eu/ecb/legal/pdf/celex_02014o0060-20180416_en_txt.pdf.


operations (repurchase agreements) with Eurosystem counterparties, in which the ECB fixed the MROR, estimated the overall liquidity needed by the banking system, and allocated the amounts pro-rata according to the bids received. After some internal and external discussion about over-bidding and underbidding phenomena (Ayuso and Repullo 2001; Bindseil 2005; Ehrhart 2001; Nautz and Oechsler 2006), in June 2000 the ECB switched to variable rate tenders, with the minimum bid rate constituting the policy rate. In those, the ECB determined the total amount to be allotted and counterparty banks could bid for a larger or smaller share via the rates they were willing to pay at or above the minimum bid rate. In March 2004, the maturity of MROs was shortened to one week. A second type of open market operations from the start were Long-Term Refinancing Operations (LTROs), with a maturity of normally three months. In early times, a third type—fine-tuning operations—were used quite sparingly.

The ECB's two standing facilities create a corridor for very-short-term money market rates around the MROR. At the deposit facility, counterparties can “park” unused liquidity overnight, receiving an interest rate—the Deposit Facility Rate (DFR)—that is lower than the MROR. At the marginal lending facility, counterparties can borrow overnight (against eligible collateral) any liquidity that they are missing at the end of a day, paying a penalty rate—the Marginal Lending Facility Rate (MLFR), which is set above the MROR. Before the financial crisis, the corridor defined by the standing facilities was set most of the time symmetrically around the MROR, with a width of 200 basis points. Figure 5 shows the three policy rates—MROR, DFR, and MLFR—between January 1999 and August 2018.

The breadth of the ECB's operational framework is defined not only by the set of different instruments that can be used but also by the number of counterparties entitled to transact with the ECB and by the range of assets eligible as collateral. Any euro area credit institution that is financially sound, supervised in the EU (or under a comparable third country regime), and fulfills some operational criteria can become an ECB counterparty. The number of effective counterparties is about 2,000, which amounted to a quarter of all euro area banks during the early years of the euro and about a third of them more recently.

Its statutes stipulate that the ECB can lend to counterparties only against “adequate” collateral (EU 2012c, Article 18). Given the wide-ranging differences in EU member countries' banking and financial systems, the ECB decided from the start that a rather broad set of collateral assets need to be eligible for its operations. But they have to fulfill a number of criteria

Figure 5. The ECB's Policy Interest Rates and the Overnight Money Market Rate, 1999–2018^a



Source: ECB data.

a. The most recent observation is for October 10, 2018.

relating, among other things, to currency denomination; the location of the issuer, issuance, or any guarantor; and, notably, the risks involved. Risk control for collateral and counterparties is, of course, important for protecting the central bank from losses that could impair its credibility, hinder its operations, or even endanger its independence. It also shields euro area treasuries from reduced revenues originating from lower transfers of central bank monetary income (which ultimately means to protect taxpayers). Therefore, the ECB uses a risk management framework that has been adapted and improved over time, depending on new experiences. For example, like many other central banks, it applies haircuts to riskier assets and does not accept collateral below a certain quality in its credit operations (that is, not below a rating of A– before October 2008).¹⁶ Although the assets used as collateral for Eurosystem monetary policy operations changed over time, public sector debt securities, corporate bonds, asset-backed securities, and covered bonds, as well as various forms of credit claims, have always played significant roles.

16. This also applies to government bonds, because EMU does not include a fiscal union between member states.

During the global financial crisis starting in 2007, it turned out that the overall operational framework of the ECB was not only quite broad for dealing with the specific difficulties encountered but also quite flexible in adapting to new challenges.

THE ECB'S FIRST INTEREST RATE CYCLE Against this background, we next describe the first cyclical period experienced by the ECB.¹⁷ When describing economic developments, we take the ECB's perspective, as reflected in the introductory statements of its monthly press conferences and its *Monthly Bulletin*. The main macroeconomic, monetary, and financial developments to which we refer are depicted in figures 3 to 15 throughout the paper.


When stage 3 of EMU started in January 1999, the ripples of the financial crises in Asia in 1997 and Russia in August 1998, together with the near collapse of the Long-Term Capital Management hedge fund in September 1998, were still visible in the high volatility of financial markets. The high level of uncertainty clouded prospects for economic growth in the euro area. In a coordinated move on December 3, 1998, all the national central banks in the euro area had lowered their key central bank interest rates to 3 percent, which de facto determined the level of short-term interest rates with which the ECB started stage 3 of EMU.

In early 1999, it became increasingly clear that, on balance, the risks to price stability over the medium term were mainly on the downside. Inflation rates were very low by historical standards (below 1 percent; see figure 6) and were significantly below the ceiling of the ECB's definition of price stability amid emerging signs of a strong economic slowdown, which eventually did not materialize (figure 7). In spite of rising oil prices starting in mid-February 1999 (figure 8), a depreciating effective euro exchange rate, buoyant loan growth of about 10 percent, and headline M3 growth above the reference value (figure 9), the Governing Council reduced the policy rate by 50 basis points on April 8, 1999, from 3.0 to 2.5 percent (figure 5).

However, as sharp increases in oil prices and a general rise in import prices continued to exert upward pressure on prices in the short term in the context of robust economic growth, the risks of indirect and second-round effects on consumer price inflation via wage setting rose significantly in the course of 2000. These concerns were compounded by a trend depreciation of the euro exchange rate, especially in the second half of 2000, when it moved further out of line with the sound fundamentals of the euro area

17. For reviews of the first 10 years of the ECB, see ECB (2008a); Mackowiak and others (2008); Buti and others (2010); and Galí (2003).

Figure 6. The Euro Area’s Headline Inflation, Core Inflation, and Longer-Term Inflation Expectations, 1999–2018^a




Sources: ECB data; Survey of Professional Forecasters (SPF).

a. HICP = Harmonised Index of Consumer Prices; HICPX refers to HICP excluding energy and food. SPF five years ahead = the average HICP inflation rate expected by the respondents to the SPF. The most recent observation is for September 2018.


Figure 7. The Euro Area’s Real GDP Growth and Its Components, 1999–2018^a

Annual percentage changes
and percentage-point contributions



Source: ECB data.

a. This figure shows data for the 19 euro area countries (fixed composition). The most recent observation is for 2018:Q2.

Figure 8. Global Prices for Oil, Metals, and Food, 2000–2018^a


Sources: Bloomberg; Hamburgisches Weltwirtschaftsinstitut; ECB staff calculations.

a. Oil prices refer to the Daily Brent Oil spot prices per barrel in U.S. dollars. Food and metal prices are the respective subindexes of the Hamburgisches Weltwirtschaftsinstitut's total commodity price index, which are normalized to 100 in 2015. The most recent observation is for September 7, 2018.

(figure 10). Economic activity in the euro area expanded very rapidly in early 2000, heading above a 4 percent growth rate, and was set to continue along this path (figure 7), due to the strong dynamism of the world economy, especially in the sectors of the “new economy.” Also, the protracted monetary expansion above the reference value was increasingly pointing to upside risks to price stability at medium- to longer-term horizons over the course of 1999 and in early 2000 (figure 9). Against this background, the Governing Council raised the key ECB interest rates by a total of 225 basis points in a series of interest rate hikes between November 1999 and October 2000, bringing the main policy rate to a level of 4.75 percent in October 2000 (figure 5).

As of 2001, the prospects for economic growth deteriorated in the wake of severe shocks that hit the world economy and global financial markets, such as the collapse of the dot-com bubble and associated corporate scandals, the terrorist attacks in the United States on September 11, 2001, and the escalation of geopolitical tensions related to Iraq—all of which increased the degree of economic uncertainty and undermined confidence. Overall, economic growth in the euro area turned

Figure 9. Growth of M3 and Monetary Financial Institutions' Credit to the Private Sector for the Euro Area, 1999–2018^a



Source: ECB data.


a. Monetary financial institutions' credit to the private sector refers to total loans and securities vis-à-vis euro area nonmonetary financial institutions, firms, and households, excluding general government. The thick gray dashed line refers to the ECB's reference value of 4.5 percent for M3 growth, signaling a particularly prominent role of money until the ECB reviewed its monetary policy strategy in May 2003. The line is thinly dotted after the review, indicating that the annual review of the reference value was discontinued and the role of money diminished. The most recent observation is for August 2018.

rather weak in 2002, and this performance did not change fundamentally in 2003 (figure 7).¹⁸ Initially, annual HICP inflation rose further in 2000 and the first half of 2001, despite a marked fall in oil prices and a significant appreciation of the euro exchange rate against all major currencies after concerted foreign exchange interventions by the ECB, the Federal Reserve, and the Bank of Japan in September 2000.¹⁹ The

18. Note that in contrast to the National Bureau of Economic Research's business cycle dating committee for the United States, the Centre for Economic Policy Research's committee never called a recession in the euro area in the early years of the new millennium.

19. The ECB and several Eurosystem national central banks also intervened a number of times during the first half of November.

Figure 10. The Euro Exchange Rate Against the U.S. Dollar, and in Effective Terms, 1999–2018^a



Source: ECB data.

a. NEER = nominal effective exchange rate; REER = real effective exchange rate. The REER of the euro is calculated as the geometric weighted average of bilateral nominal exchange rates, which are deflated using relative price or cost measures; the weights used are the trade weights assigned to the currency of each trading partner. The 38 trading partners included in the NEER and REER are Algeria, Argentina, Australia, Brazil, Chile, Bulgaria, Canada, China, Croatia, the Czech Republic, Denmark, Hong Kong, Hungary, Iceland, India, Indonesia, Israel, Japan, Malaysia, Mexico, Morocco, New Zealand, Norway, the Philippines, Poland, Romania, Singapore, South Korea, Sweden, Switzerland, Russia, South Africa, Taiwan, Thailand, Turkey, the United Kingdom, the United States, and Venezuela. The most recent observation is for April 5, 2018.

concerns about second-round effects gradually dissipated over time as the outlook for the euro area economy continued to deteriorate. Average annual HICP inflation remained slightly above 2 percent from 2000 to the first half of 2003 (figure 6), but the subdued pace of economic activity and the significant appreciation of the euro after the spring of 2002 were expected to dampen inflationary pressures. Looking at the monetary developments, annual M3 growth accelerated strongly from mid-2001 onward (figure 9). However, this increase was not interpreted as implying risks to price stability at medium to longer horizons because it was mostly due to sizable shifts in private investors' portfolios from shares and other longer-term financial assets toward safe and more liquid monetary assets included in M3 in the aftermath of the global stock market correction and the terrorist attacks of September 11, 2001

(ECB 2008a, 44). This assessment was supported by the fact that annual growth of credit to the private sector continued to decline (figure 9), especially to nonfinancial corporations, in a context of rather subdued economic activity. In this period, the Governing Council lowered the key ECB interest rates by a total of 275 basis points (figure 5). This included a joint 50-basis-point cut coordinated with the Federal Reserve on September 17, 2001, in response to the adverse confidence effects of the terrorist attacks.²⁰ The policy rate reached a—at that time—historically low level of 2 percent in June 2003. At the same time, ECB policymakers saw the sustained growth in M3—correcting for the estimated impact of portfolio shifts—as an important indicator arguing against the emergence of deflationary risks for the euro area in 2002 and 2003 (ECB 2008a, 44).

DISCUSSION Overall, the ECB's first interest rate cycle contained a first test of the ECB's anti-inflation credibility as the euro exchange rate depreciated—and was only stopped by foreign exchange interventions—and annual headline inflation peaked at about 3 percent. The sources of the initial depreciation of the euro against the dollar (from a peak of 1.19 in January 1999 to a historic low of 0.83 in October 2000) were heavily discussed. As the main source, Giancarlo Corsetti and Paola Pesenti (1999) and Alberto Alesina and others (2001) pointed to fundamentals such as revisions in the forecasts of the output growth rate differential in the United States and in the euro area. In May 2000, President Duisenberg nevertheless issued a press release to EU citizens reassuring them of the euro's stability (ECB 2000b). And ultimately, the ECB intervened, together with the Federal Reserve and the Bank of Japan, based on a “shared concern about the potential implications of recent movements in the euro exchange rate for the world economy” (ECB 2000c). The underlying concern was that a disorderly depreciation process would add to the inflationary pressures in an environment of relatively high oil prices (figure 8), and affect its credibility (subsection II.A).

Once the cycle turned, after the bursting of the dot-com bubble in stock markets, the perspective reversed. As interest rates dropped to a historically low level in the euro area, and even more so in the United States, the

20. In the days after September 11, 2001, the ECB also undertook a series of crisis management operations to deal with the substantial effects of the severe damage to the U.S. financial market infrastructure and its effects on the euro area financial system. These included overnight fine-tuning operations and a swap line with the Federal Reserve Bank of New York that allowed Eurosystem national central banks to provide dollar liquidity to their banks.

policy and academic debate turned to the consequences of the lower-bound constraint on interest rates for the fulfillment of monetary policy objectives (Bernanke 2002).

The other feature of this period was the decoupling of money and credit growth (figure 9), which called into question the prominent role of money in the ECB's monetary policy strategy. From the start, money's prominent role was a controversial feature of the ECB's strategy. For example, Alesina and others (2001) thought the ECB should abandon the two pillars and adopt a flexible inflation-targeting strategy, which they regard as simpler. In their view, the M3 pillar stood in the way of effective communication. The ECB nevertheless used robust money growth to argue against further cuts in interest rates in 2003.²¹ Both issues featured in the review of the strategy in 2003, which we discuss in the next section.

Another discussion was related to the ECB's transparency and predictability. Although opinions differ about the ECB's degree of transparency (also compared with that of other central banks), the ECB generally scores quite high on this front, and over time it has also increased its transparency in response to demands from the European Parliament and other advocacy groups (Geraats 2002). For example, in December 2000 the ECB started to publish its macroeconomic projections (ECB 2013a). Nevertheless, two elements of criticism coming mostly from the inflation-targeting proponents were prominent in the early years. First, the ECB released neither the minutes of its policy deliberations nor the votes and their attribution to members of the Governing Council.²² It argued that the press conference gave a real-time account of the discussion and could therefore be seen as a substitute, and that publishing the minutes could expose the individual members of the Governing Council to pressure from their national constituencies and undermine the consensual nature of the ECB's decisionmaking and "one voice" communication strategy. As communication became more complex after the financial crisis, this was partly addressed in January 2015, when the Governing Council decided to publish an account of its monetary policy deliberations about four weeks after the meeting (Draghi 2014a).

The second criticism was that the ECB did not publish its own interest rate forecasts (Alesina and others 2001; Geraats, Giavazzi, and Wyplosz 2008). Instead, the ECB focused on trying to explain its reaction function. It argued that in view of the effects of various unexpected shocks that can

21. See, for example, the introductory statement of the monetary policy press conference in December 2003 (ECB 2003c).

22. See, for example, the debate between Buiter (1999) and Issing (1999).

hit the economy and the long and variable time lags with which monetary policy actions are transmitted to prices, the precise timing, and sometimes even the direction, of an interest rate decision is difficult to predict. Also, by publicly announcing its monetary policy strategy and communicating its regular assessment of economic developments in a transparent manner, it could clarify its reaction function, achieve a high degree of predictability, and thereby make monetary policy more effective (Blattner and others 2008). In fact, although the 50-basis-point sizes of the first and second interest rate decisions in April (a cut) and November 1999 (an increase) somewhat surprised market participants, various empirical studies showed that relatively early ECB interest decisions had already usually been predicted quite well by the market, at least as well as the decisions of the Federal Reserve or, for example, the Bank of England (Hartmann, Manna, and Manzanares 2001; Bernoth and von Hagen 2004; Wilhelmsen and Zaghini 2011). Still, the ECB often emphasized the need to maintain a full-information, state- and data-driven policy approach, and that it did not want to communicate or commit to future policy actions given the large uncertainties about the state of the economy in the future. This changed in 2013, when the ECB started giving forward guidance on its future policy actions (see subsection II.C).

1.B. Recovery and Growing Imbalances, July 2003–July 2007

At the transition between the first and second cyclical periods of the euro area, the ECB reviewed its monetary policy strategy against its experiences, and it clarified and amended some aspects. A little more than halfway through the period, the ECB started making a series of interest rate hikes in order to keep the inflationary pressures in check that emerged, among other things, from increasingly solidifying growth as well as increasingly vigorous money and credit dynamics. That the latter were a harbinger of a severe crisis only became clear during the next period (section II.C).

THE 2003 REVIEW OF THE STRATEGY In 2003, after about four years of experience with the ECB's new strategy, Otmar Issing initiated a review of it, which led to three main measures: (1) a clarification of the definition of price stability: the Governing Council would aim at a year-on-year HICP inflation rate of "below, *but close to* 2 percent over the medium term"; (2) the termination of the annual review of the reference value for M3 growth; and (3) a restructuring of the introductory statement of the president at the monthly monetary policy press conference, which now started with the economic analysis followed by the monetary analysis (ECB 2003a, 2003b).

The clarification of the price stability definition in the ECB's strategy was a response to the strengthened need to establish a sufficient inflation buffer as a discussion of deflation risks took place in 2002–3. Such a buffer was deemed to be necessary for two reasons. First, a small positive, steady state inflation rate would reduce the probability of hitting the lower bound on nominal interest rates. Second, a positive inflation rate also greases the wheels of the labor market, particularly in a monetary union with still segmented labor markets, because it reduces the need for wage deflation in the face of asymmetric economic developments. Such wage deflation was thought to be costly in the presence of widespread evidence of downward nominal wage rigidity in the euro area.²³ A number of studies had shown that an inflation buffer of close to 2 percent would significantly reduce the probability of hitting the zero lower bound on nominal interest rates or downward nominal wage rigidity constraints (Issing 2003a; Reifschneider and others 2000).

The specific formulation of the inflation aim of “below, but close to 2 percent over the medium term” was the result of a compromise that maximized the buffer, while remaining consistent with the definition of price stability and not giving a sense of unwarranted precision associated with inflation-targeting regimes. The sense of continuity was made clear by Issing at the press conference in May 2003 explaining the outcome of the review. When asked whether the aim of “below but close to 2 percent” is a change, he replied: “This ‘close to 2 percent’ is not a change, it is a clarification of what we have done so far, what we have achieved—namely, inflation expectations remaining in a narrow range of between roughly 1.7 and 1.9 percent—and what we intend to do in our forward-looking monetary policy” (ECB 2003b). Although all this should have removed (or very significantly reduced) the room for interpretation about how low the lower bound of the price stability definition was, the reformulation did not extinguish perceptions by some observers of an asymmetric inflation objective. Symmetry was seen as important by the proponents of inflation targeting (Bernanke and others 1999), but even German monetarists like Manfred Neumann (2010, 235) thought that “the lack of a lower bound as part of the definition was an unnecessary drawback.”

The second and third measures mentioned above *de facto* meant a downgrade of the prominent role of money in the ECB's strategy relative to the weight put on it, for example, by President Duisenberg (1999). This

23. See, for example, the findings of the Wage Dynamics Network, as given by ECB (2009b).

reflected the reality that, on a monthly basis, monetary policy decisions were mostly driven by the broadly based assessment of the outlook for price developments and the risks to price stability (the “economic analysis”), of which the ECB’s macroeconomic projections were an important part.²⁴ It also reflected emerging evidence on instability in money demand and the need to explain “distortions” or “portfolio adjustments” in M3 growth that were not linked to the medium-term risks to price stability as discussed above.²⁵ A revamped monetary analysis was now presented as a cross-check of the economic analysis from a medium- to long-term perspective, given the long-run monetary nature of inflation. It clarified that the main challenge facing monetary analysis is to see past inevitable short-term disturbances of the underlying relationship between money and prices so as to discern longer-term inflationary risks. This was also reflected in the changed structure of the introductory statement at the monetary policy press conferences, which now started with the economic analysis and ended with a cross-check from the monetary analysis.

The rearrangement of the pillars was applauded by academics favoring inflation targeting (Svensson 2003), while at the same time it was acknowledged that the money pillar had been useful during the first years of the ECB because it made it easier for it to gain credibility as a sign of “the new institution’s fidelity to principles stressed earlier by the Deutsche Bundesbank, which had in turn played a critical role as the anchor of the previous European Monetary System” (Woodford 2006, 87). But the debate on the role of monetary analysis and the need to have two separate pillars continued (Issing 2005). On November 9 and 10, 2006, the ECB held a symposium to discuss this from both academics’ and practitioners’ points of view (Beyer and Reichlin 2006). At the conference, Björn Fischer and others (2008) reviewed the actual ECB experience with its monetary analysis from 1999 through 2006 and emphasized the real-time and comprehensive nature of the monetary analysis that had been performed in the quarterly monetary assessments since December 1999. These authors described the tools that were used, making a distinction between money demand equations, judgmental analysis, and money-based inflation forecasts. They also assessed the forecasting performance of money-based tools and found that there was information value in addition to the

24. Also see the evidence on the ECB’s projections presented in subsection II.B.

25. For alternative views on money demand stability, see Alves, Robalo Marques, and Sousa (2007); and Bruggeman, Donati, and Warne (2003).

Broad Macroeconomic Projection Exercise forecasts. Finally, based on an in-depth analysis of the monetary analysis input, they concluded that the economic pillar prevailed in influencing the decision when the monetary pillar gave a blurred signal. This finding is confirmed below in the analysis of an interest rate rule in subsection II.B.

The broader discussion at the symposium pointed to two ongoing developments in the nature and role of monetary analysis. First, monetary analysis was evolving from a narrower perspective based on the quantity theory of money to a broader set of analyses that also include the role of financial frictions and financial intermediation in macroeconomic developments. This led to a revamping of the debate on why the two pieces of the analysis should be kept separate, given the intimate linkages between financial and real factors. At the symposium, ECB vice president Lucas Papademos conjectured that if “in the future, we will be in a position to develop and reliably estimate a single empirical approximation of a general theoretical framework in which money is of central importance. . . . It may be possible to merge the two pillars of our analysis into a single one. But this will be a larger pillar in which money will continue to play a prominent role in guiding our monetary policy decisionmaking” (Papademos 2006, 202). In 2007, the Governing Council endorsed a research program to further enhance monetary analysis, including by developing methodologies for cross-checking and building structural models that embody an active role for money and credit in the determination of inflation dynamics (Papademos and Stark 2010).

The second theme that received increasing attention during this period was the link between money and credit, asset price developments, and financial stability (for example, ECB 2005). Although this financial stability angle was not taken up as an explicit justification for the two-pillar approach in the 2003 review, the ECB paid increasing attention both in research and policy communication to this link and the associated view promoted by the Bank for International Settlements (Borio and Lowe 2002) that it may be necessary for monetary policy to lean against the wind of growing financial imbalances (Detken and Smets 2004; Issing 2003b). This also became part of the research program mentioned above (Detken, Gerdesmeier, and Roffia 2010). In a speech on asset price bubbles and monetary policy, then-ECB president Trichet (2005) conjectured that while “a leaning against the wind” approach is “compelling in many theoretical aspects, in practice. . . . It is likely that the circumstances will be rare that a policy maker will embark with confidence on an explicit leaning against the wind policy.” But he also argued that monetary

analysis helps in incorporating emerging financial stability risks with implications for price stability from a medium-term perspective: “The fact that our monetary analysis uses a comprehensive assessment of the liquidity situation that may, under certain circumstances provide early information on developing financial instability is an important element” (Trichet 2005). With particular reference to asset price bubbles and housing booms, this became part of the rationale for the monetary analysis (Issing 2005).

STABLE RATES, MONETARY TIGHTENING, AND NO ADDITIONAL “LEANING AGAINST THE WIND” Starting in June 2003, the ECB kept interest rates steady for almost two and a half years. So the previous decision to lower the MROR to a historically low level of 2.0 percent nurtured the economic recovery for quite a while. The overall picture of economic activity brightened during the second half of 2003, when the euro area’s exports increased significantly as a result of the renewed dynamism of the world economy. Also, domestic demand and investment picked up, not least in view of the low level of interest rates and the generally favorable financing conditions (figure 5; also see figures 14 and 15 below). The recovery in economic activity moderated somewhat in the second half of 2004 and the first half of 2005, partly on account of rising oil prices (figure 8), lower consumer confidence, a temporary deceleration of global economic growth, and the lagged effects of the past appreciation of the euro (figure 10). However, in the second half of 2005, the expansion of economic activity in the euro area regained momentum. On the price side, HICP inflation did not fall as swiftly and strongly as previously expected, largely due to adverse food price developments and oil prices that were higher than expected—although the latter were attenuated by the euro’s appreciation. Annual HICP inflation remained above 2 percent in 2005, but underlying domestic inflationary pressures were contained throughout 2004 and most of 2005 (figure 6), justifying the prolonged accommodative monetary policy stance.

As 2005 progressed, the ECB’s economic analysis suggested that upside risks were increasing, especially due to potential second-round effects in wage setting and price setting that stemmed from higher oil prices. But this time it was the monetary analysis that carried the day. As of mid-2004, robust credit and monetary expansion (see figure 9 below) reflected the stimulating effect of the then-prevailing very low level of interest rates in the euro area and, later on, renewed dynamism of the euro area’s economy, rather than portfolio shifts (as between 2001 and 2003), indicating increasing upside risks to price stability at medium- to longer-term horizons toward the end of 2005. In response, the ECB started raising its policy

rate as of December 2005, and eventually, by a total of 200 basis points, to a level of 4 percent by the end of June 2007 (figure 5).

The gradual withdrawal of monetary accommodation took place against the background of solid economic growth and a continued strong money and credit expansion in the euro area. The economic expansion gained momentum in the first half of 2006 and became gradually more broad-based and self-sustaining, with domestic demand as the main driver. Notwithstanding the impact of high and volatile oil prices, real GDP rose by about 3.3 percent in 2006, compared with about 1.7 percent in 2005 and about 2.1 percent in 2004, and continued to expand at a solid rate of about 3.1 percent in 2007 (see figure 7, which, however, shows annualized quarterly data). With regard to prices, average annual HICP inflation was slightly above 2 percent in 2006 and 2007, mainly driven by domestic demand, as underlying inflation developments were largely in line with the ECB's inflation aim (figure 6). Money and credit expansion became increasingly vigorous throughout this phase, supported by a persistently strong growth of bank loans to the private sector (figure 9).

DISCUSSION Overall, this second phase was characterized by an increasingly solid expansion of economic activity and increasingly vigorous money and credit growth (double the reference value toward the end), following a long period of low interest rates. Against the background of the discussion above on the approach of “leaning against the wind,” with the benefit of hindsight, the question emerges to what extent monetary analysis was used in guiding monetary policy in the face of growing financial imbalances. At the time, the ECB Board members warned of the potential for emerging misalignments in asset prices, notably in housing, due to strong money and credit growth.²⁶ Also, Trichet (2008) pointed to the December 2005 episode as one where the monetary pillar was crucial in driving the monetary policy decision. Indeed, based on a reading of the introductory statements at the end of 2005, Neumann (2010) argues that monetary analysis was one of the driving forces behind the decision to start raising interest rates in 2005. However, as we argue in subsection II.B below, it is difficult to detect significant deviations from the ECB's usual reaction to the outlook on growth and inflation in this period. This suggests that the tightening of policy rates in 2005 did not go beyond what would be indicated by the usual economic analysis and monetary cross-check, contrary

26. For example, Issing (2005): “Moreover, strong money and credit growth in a context of already ample liquidity in the euro area implies that asset price developments, particularly in housing markets, need to be monitored more closely, given the potential for misalignments to emerge.”

to what would have been the case under an active leaning-against-the wind approach trying to enhance financial stability through restrictive monetary policy action.

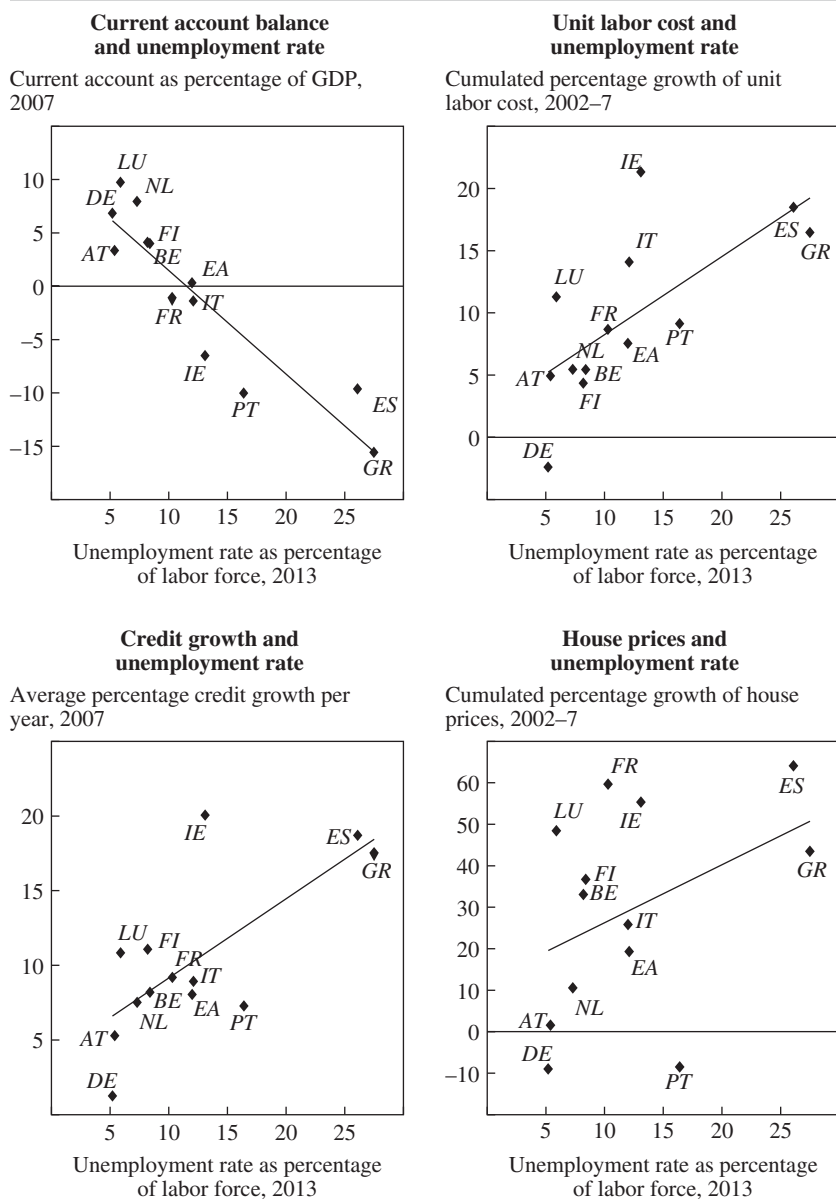
Although this does not prove conclusively that low monetary policy rates did not play any role in strong credit growth and bank risk-taking—in fact, to some extent they are a natural and desired effect of an expansionary monetary policy stance—the institutional setup for financial supervision in the euro area at the time located the primary responsibility for containing the buildup of financial risks with national prudential authorities. The ECB could only “contribute to the smooth conduct of policies pursued by the competent authorities relating to the prudential supervision of credit institutions and the stability of the financial system” as long as price stability was ensured (EU 2012b, Article 127.5). But not being a regulatory or supervisory authority itself, the ECB did not possess any prudential policy instruments that could address emerging financial imbalances. The available prudential instruments were all with national prudential authorities, subject to some cooperation through EU committees.²⁷ Moreover, before the climax of the financial crisis in 2008 the macroprudential approach to maintaining financial stability was not very well developed in the competent prudential authorities.²⁸

A related important feature of this period is that, underlying the aggregate euro area output and credit boom, there were diverging intra-euro area current account balances. These imbalances played an important part in the propagation of the subsequent twin financial and sovereign debt crises in the euro area, which we discuss in the next section. As shown in figure 11, the countries that, leading up to 2007, had accumulated large current account deficits along with high unit labor cost and credit and house price growth differentials relative to their euro area peers, were also among the ones that suffered the highest fallout from the financial crisis—for example, as measured by the subsequent level of the unemployment rate in 2013 (Constancio 2013; Smets 2014; Martin and

27. See our companion paper (forthcoming) for a description of the evolving prudential framework since the introduction of the euro and the ECB’s role in it. Some of the national central banks were banking supervisors but not as part of their Eurosystem roles. The European System of Central Banks’ Banking Supervision Committee brought all EU banking supervisors at one table. Though the ECB hosted its secretariat, it could not oblige the members to take any action.

28. The De Larosi re Report (High-Level Group on Financial Supervision 2009) led to the establishment of the European Systemic Risk Board in 2010, a macroprudential body whose secretariat is provided by the ECB but that can only make risk warnings or policy recommendations without having its own policy instruments.

Figure 11. The Euro Area Countries' Economic and Financial Imbalances Before the European Twin Crises and Unemployment Thereafter^a



Sources: Eurostat; authors' calculations.

a. These countries and regions are shown in this figure: DE = Germany, AT = Austria, LU = Luxembourg, NL = the Netherlands, BE = Belgium, FI = Finland, FR = France, EA = euro area, IT = Italy, IE = Ireland, PT = Portugal, GR = Greece, and ES = Spain.

Philippon 2017). Or, put differently, all the countries that ultimately ended up in macroeconomic adjustment programs—Greece, Ireland, Portugal, and Spain—ran substantial current account deficits in 2007. The ECB’s communication focused particularly on the need to address divergences in productivity and competitiveness across the various euro area countries (Trichet 2006; or the white area of figure 1).

Preparing the ground for subsection I.C, one narrative behind these boom-and-bust developments (put forward after the fact) runs as follows (Baldwin and Giavazzi 2015). Easy global financial conditions (partly driven by the global savings glut), as well as greater integration of wholesale financial markets within the monetary union (with disappearing risk premiums), encouraged cross-country capital flows from the EU’s “core” to its “periphery” (Blanchard and Giavazzi 2002; Lane 2015). Although the aggregate euro area current account was in balance throughout most of this period, large intra-euro area current account imbalances were building up, feeding nontradable sectors like government consumption and housing in the “periphery” countries, driving up wages and costs, and resulting in competitiveness losses that undermined the traded goods sectors and validated the current account deficits. With the exception of Greece, explicit public debt was not the first problem, according to this narrative, although from an *ex post* perspective, building up higher buffers may have been advisable, as shown by Philippe Martin and Thomas Philippon (2017). Instead, the private debt buildup was very significant, mimicking some of the developments in the U.S. and other countries that belong to the Organization for Economic Cooperation and Development.²⁹ Moreover, there was a mismatch between the longer-term loans to households and firms made by domestic banks and the short-term, cross-country interbank funding that financed this debt.

Other observers (Feld and others 2016) put more weight on the fiscal vulnerabilities of some euro area countries, even before the start of the financial crisis. The Stability and Growth Pact had been regularly broken by a variety of countries since the introduction of the euro. Between 1999 and 2008, the aggregate debt-to-GDP ratio of the euro area hovered close


29. Euro area countries with particularly high and increasing household debt levels in the years before the crisis included Cyprus, Ireland, the Netherlands, Portugal, and Spain. (Germany’s household debt was high in the early years of the euro but then consistently declined.) Countries with particularly high and increasing debt levels for nonfinancial corporations included Ireland, Luxembourg, the Netherlands (not increasing), Portugal, and Spain. Interestingly, neither Greece nor Italy had particularly high private debt levels, even though they increased in both cases. In many cases, the increases in private debt levels were part of a long-term trend, at least after the start of the euro.

to 70 percent, 10 percentage points above the Stability and Growth Pact's limit of 60 percent for individual countries. The countries that had entered the euro area with very high public debt levels (that is, significantly above 100 percent of GDP) were Belgium, Greece, and Italy. They all gradually reduced these levels in the early years, helped by strong nominal GDP growth and low interest rates; but because of rapidly eroding primary surpluses, this process stopped at levels of about 100 percent of GDP or slightly above, except for Belgium. In other words, the euro area entered the financial crisis with one large and one smaller fiscally vulnerable country.³⁰

In sum, among the countries that turned out to be stressed during the European twin crises (see the next section) beforehand, Cyprus, Ireland, Portugal, and Spain were more vulnerable in terms of private debt and Greece and Italy more in terms of public debt. Both groups together account for about a third of euro area GDP (roughly 39 percent of its total population), but the latter is a bit larger than the former. However, as we shall see further below, many more than these two important fundamental factors came together in determining the severity of the European twin crises and the obstacles that they implied for successful monetary policy.

The significant worsening of the financial crisis in key advanced countries in the course of 2008, which revealed the exposure of some European banks to toxic subprime mortgages in the U.S. and to other increasingly impaired credit instruments, and the revelation of the Greek deficit deceit in late 2009 were the triggers that led to a sudden stop of cross-country capital flows and exposed private and public debt overhangs in the respective euro area countries. Several negative propagation mechanisms then came into action. First, the need to backstop the weakened banks in the absence of a European resolution framework undermined the credit rating of a number of national governments. Second, weakened sovereigns and a faltering economy further increased the fragility and the undercapitalization of national banking sectors, leading to further deleveraging and "closing" the doom loop between national sovereign and banking instabilities. Third, the results of the Deauville Summit of France, Germany, and Russia in October 2010—which included a Franco-German agreement to promote "private sector involvement" in handling public debt overhangs—and associated discussions on a Greek debt restructuring that was only implemented

30. Other euro area countries whose public debt levels increased before the crisis and cut through the 60 percent limit included France, Germany, and Portugal. Austria fluctuated around 70 percent without a clear trend. Portugal had some vulnerabilities because of weak state-owned enterprises whose debt was not included in the Stability and Growth Pact's debt figures but migrated to them during the crisis.

Figure 12. Aggregate of the Euro Area Member Countries' Fiscal Policies, 1999–2018^a


Source: ECB data.

a. The euro area's fiscal stance is computed as the aggregate change in the cyclically adjusted primary balances of all member countries' governments. The cyclically adjusted primary balance corresponds to the budget balance, minus interest payments and adjusted for cyclical factors. The budget balance refers to the difference between total government revenues and expenditures. The most recent observation is for 2018, and was taken from the ECB's June 2018 Broad Macroeconomic Projection Exercise.

in 2012—though good for ex ante incentives of controlling public deficits—facilitated contagion toward other sovereigns ex post, because a euro area backstop for governments was lacking. Finally, weakened sovereigns also led to procyclical fiscal policy, which deepened the recession in the absence of fiscal policy coordination and a common budget. As shown in figure 12, the coordinated fiscal expansion of 2008–9 turned into a significant and protracted tightening of the fiscal policy stance from 2010 to 2013. As a result of the concurrence of all these factors, the sudden stop turned into a crisis and a prolonged double-dip recession (Corsetti 2015; Corsetti and Dedola 2016), to which we turn in the next section.

1.C. The Financial Crisis, the Sovereign Debt Crisis, and the Double-Dip Recession, August 2007–June 2013

The European sequence of crises starting in the summer of 2007 can be decomposed in the early turmoil in funding markets, the systemic banking crisis, and the sovereign debt crisis. The ECB managed the first phase with liquidity operations; it managed the second phase with decisive interest rate cuts, further enhanced with liquidity operations and a first asset

Figure 13. The Euro Area's Money and Government Bond Market Spreads, 1999–2018^a

Source: ECB data.

a. Euribor = Euro Interbank Offered Rate; € OIS = Euro Overnight Index Swap Rate. The euro area's 10-year yield is a GDP-weighted average of euro area member countries' government bond yields. The most recent observation is for September 2018.

purchase program for covered bonds; and it managed the third phase with a first asset purchase program for government bonds, very-long-term liquidity operations, and interest rate cuts to basically zero. The turning point toward recovery occurred when the EU's political leaders agreed on a series of EMU reforms in the summer of 2012 and President Draghi announced that the ECB would do “whatever it takes” to preserve the euro, backed up shortly afterward by the ECB's powerful Outright Monetary Transaction Programme.

THE ECB'S MONETARY POLICY MOVES TO CRISIS MANAGEMENT MODE In the early phase of the financial crisis, the ECB's operational framework took center stage.³¹ The reason was that wider problems first emerged in inter-bank and other short-term funding markets (figure 13), which could largely

31. For a recent description and chronology of the ECB's monetary policy responses since the onset of the crisis, see Camba-Méndez and Mongelli (2017).

be addressed with liquidity management tools. Moreover, it allowed the ECB to continue to follow the so-called separation principle, meaning that the conduct of monetary policy focused on setting policy rates for achieving price stability over the medium term and that market operations focused on ensuring that market turbulences would not disturb the transmission of the policy rates to the economy. Another way of saying this is that those operations acted as complements to conventional interest rate policy (and were not intended to act as substitutes). The separation principle was in line with the traditional analysis done by William Poole (1970), according to whom stabilizing the short-term interest rate in the face of purely financial shocks is the best way to insulate the real economy from the effects of those shocks (Fahr and others 2013; ECB 2008a). In the early phase, the operations were mainly focused on money and other bank funding markets; but when the sovereign debt crisis emerged in 2010, they also started to address malfunctioning government bond markets and extended liquidity beyond one year. These included asset purchases of both covered and government bonds (for an overview of the main ECB monetary policy measures during the two crises and the subsequent recovery, see figure 26 below).

As monetary policy moved into crisis management mode, the two-pillar approach took a backseat in communication. Although the introductory statements at the regular monetary policy press conferences remained structured along the two pillars, including a cross-checking section—also in line with how staff analyses still supported the Governing Council’s decisionmaking process—few speeches by Executive Board members dealt with the two-pillar structure of the ECB’s strategy (the medium gray area of figure 2 above becoming quite thin).

The crisis nevertheless had a big impact on the ECB’s monetary analysis. The focus turned to how to identify and address the impairments in the transmission mechanism of monetary policy. The financial crisis necessitated a further comprehensive broadening of the monetary analysis toward detailed macroeconomic and microeconomic analyses of the financial system and of the bank lending channel in particular (given the euro area’s financial structure). For example, the ECB’s Bank Lending Survey—already launched in 2003—became a prominent tool for understanding supply-side restrictions in bank credit markets (ECB 2018a). More generally, a deep analysis of the capital, leverage, and liquidity positions of banks became important, as well as a comprehensive and disaggregated look at both bank and nonbank financing conditions in the economy. This led to a thorough revamping of the ECB’s quarterly monetary assessments.

At the same time, given the intimate interaction between financial and real factors, it also blurred the distinction between the economic and the monetary analysis.

FROM LIQUIDITY OPERATIONS TO DECISIVE RATE CUTS AND EARLY ASSET PURCHASE PROGRAMS The third cyclical phase of the ECB's history can be divided into three subperiods. The first period, August 2007–September 2008, is often denoted as financial market turmoil (Evanoff, Hartmann, and Kaufman 2009). The collapse of the U.S. subprime mortgage market led to a general repricing of risk in the asset-backed securities (ABSs) and other structured credit markets of developed countries, which seriously impaired interbank and other short-term funding markets. The second period, October 2008–May 2010, covers the intense systemic financial crisis affecting many developed countries, after the failure of Lehman Brothers, the Great Recession, and the associated collapse of international trade. The third period, June 2010–June 2013, starts with the emergence of the sovereign debt crisis specific to the euro area, when the Greek fiscal situation deteriorated significantly and several other euro area countries subsequently became distressed.

Financial turmoil impairing money market functioning, August 2007–September 2008. Financial turmoil first erupted in Europe with the emergence of money market tensions on August 9, 2007, after the announcement that a number of investment funds had to close because they could no longer value their portfolios owing to the illiquidity of ABS markets. The uncertainty about the values of ABS and other structured credit products and the asymmetric information about their location among banks led to adverse selection, liquidity hoarding, and the freezing of interbank and other short-term funding markets (such as asset-backed commercial paper and repurchase agreements) (Cassola and others 2008; Gorton and Metrick 2012; Heider and others 2015). Despite these difficulties, large bank failures did not occur in the euro area during this period. Only a few mid-sized German banks, which had been particularly engaged in structured credit practices and wholesale funding, received public support. One indicator of the difficulties in bank funding markets (mixing credit and liquidity risks) is the spread between the unsecured interbank rate and the overnight swap rate, which is only subject to a minimum amount of counterparty risk (figure 13). After remaining very close to zero for years, this spread rose to about 60 basis points.

Reacting immediately on August 9 with a fixed-rate, overnight fine-tuning operation allotting the full demand of €95 billion to counterparties, the ECB was the first major central bank to respond to the turmoil. In the

following days, weeks, and months the series of operational measures addressing the euro money market disturbances included further fine-tuning operations, intra-maintenance period “front-loading” (meaning that the ECB provided very ample liquidity early in each reserve maintenance period, which then ran down until the end of each maintenance period), and a relative extension of the maturity profile of aggregate market operations (by running supplementary three-month LTROs) (ECB 2007a). In line with the separation principle, however, the measures were designed to keep the overall monetary policy stance unchanged. In the second half of December 2007, the ECB also joined forces with the Federal Reserve by providing U.S. dollar liquidity to Eurosystem counterparties through a swap arrangement. The Bank of Canada, the Bank of England, and the Swiss National Bank made parallel similar arrangements “to address elevated pressures in short-term funding markets” (ECB 2007b). Interestingly, none of these measures were mentioned in any of the introductory statements of the Governing Council’s press conference at the time, which only contained references to financial market volatility and reappraisals of risk and to the ECB paying great attention to them. They were announced in separate press releases, and were later summarized in the ECB’s *Monthly Bulletin*.

With the advent of financial turmoil, the outlook for future economic activity became clouded, and the balance of risks to the growth outlook tilted to the downside. Nevertheless, euro area growth remained above 2 percent for a while (figure 7), with corporate profitability sustained, employment growth strong, and the unemployment rate declining to 7.4 percent, a level not seen for 25 years (figure 3). At the same time, annual inflation rose sharply toward the end of 2007, reaching levels significantly above 2 percent (above 3 percent still in the same year, and above 4 percent in the summer of 2008; figure 6), driven largely by the very significantly increasing prices of commodities, including oil (figure 8). Although moderate wage developments and anchored medium- to longer-term inflation expectations helped to dampen inflationary pressures, the risks to price stability over the medium term were still judged to be on the upside. A cross-check with the ECB’s monetary analysis appeared to confirm this (figure 9). The ECB paid particular attention to monetary developments, also with a view to better understanding the shorter-term response of financial institutions, households, and firms to the financial market’s turmoil in the second half 2007. At the time, there was little evidence that the turmoil had strongly influenced the overall dynamics of money and credit expansion, also thanks to the effectiveness of ECB liquidity management,

which contained volatility in money market rates. Accordingly, the ECB decided to raise the MROR by 25 basis points in July 2008 to avert the risk of second-round effects on wages.

The financial crisis, the collapse of bank intermediation, and the Great Recession, October 2008–April 2010. This increase in the policy rate was quickly reversed when the financial turmoil escalated to a systemic financial crisis after the collapse of the U.S. investment bank Lehman Brothers on September 15, 2008.³² At that time, it became clear that even prominent and systemically important institutions could fail, and many more of them would have failed if they had not been taken over by other financial institutions or supported by the government.³³ So inter-bank and other financial markets seized up both internationally and within the euro area—for example, giving rise to large spreads between secured and unsecured money market rates (as shown in figure 13). Economic activity was disrupted, and many of the major economies were on the verge of collapse. Tensions spilled over from the financial sector into the real economy, leading to the Great Recession. The U.S. economy, which had slowed considerably when the financial turmoil first began, entered a severe recession in December 2007 and exited it in June 2009.

Owing to strong economic and financial ties, the crisis spread to the United States' main trade and financial partners, including the euro area countries. For example, a number of large euro area banks (compared with their home country) failed and/or were supported by their sovereigns—some more for their exposure to the collapse of the global credit trading system (triggered by the U.S. subprime mortgage crisis and the revelation of many toxic ABSs), and others more for their exposure to their tanking local economies and real estate markets.³⁴ On the back of confidence

32. The ECB's response to the financial crisis is described in detail by Pill and Reichlin (2014); and the crisis responses by the ECB and the euro area's fiscal and prudential authorities are described by Hartmann (2010).

33. In other words, the devastating systemic nature of the crisis was caused by a mixture of contagion among financial intermediaries and, notably, the unwinding of the widespread imbalances that had built up in the years before on financial institutions' balance sheets, particularly from the combination of originate-to-distribute behavior and the global trading of the resulting credit products (ABS, collateralized debt obligations, collateralized loan obligations, and the like). For the different forms of systemic risk, see ECB (2009a) and De Bandt, Hartmann, and Peydró-Alcade (2015).

34. See Hartmann (2015) for a discussion of different euro area countries' experiences with boom/bust cycles in residential real estate markets and problems with the associated prudential policies and frameworks.

effects and impaired trade finance, global trade plummeted by about 20 percent in both the fourth quarter of 2008 and the first quarter of 2009, respectively, which also transmitted economic instability to the countries whose financial intermediaries had not been engaged in unsound international credit trading practices or toxic investments.³⁵ The euro area experienced a “sudden stop” of capital flows across its member countries. Within a few months, it had entered its own severe recession, which lasted from the second quarter of 2008 until the second quarter of 2009. During this period, year-on-year GDP growth fell by more than 5 percent (figure 7), and headline HICP inflation was briefly negative in the summer of 2009 (figure 6), also on the back of falling oil prices (figure 8). Money and credit growth dropped to 0 percent at the beginning of 2010 (figure 9). The collapse of bank intermediation, which had gathered pace in the summer of 2008, amounted to about a 13-percentage-point reduction of credit to the private sector between December 2007 and January 2010.

In this period, standard and nonstandard monetary policy measures taken by the ECB worked in tandem, although the separation principle was maintained. After the internationally coordinated interest rate cut of October 8, 2008, by 50 basis points in response to the collapse in output and inflation, the ECB further decreased its key policy rates in six steps by 275 basis points, reaching a level of 1 percent for its main refinancing operation rate in May 2009, a new historical low (figure 5).

At the same time, the ECB took a number of nonstandard measures to satisfy the high demand for liquidity, foster an even transmission of monetary policy impulses across countries and banks, and help fend off risks of an even more dramatic financial meltdown. These measures drew on its broad and flexible operational framework (see subsection I.A), which turned out to be more readily employable for meeting the challenges of the crisis than was the case for some other major central banks. But they were still regarded as complements to interest rate decisions and not substitutes. Starting on October 15, 2008, the ECB’s MROs (and all its longer-term refinancing operations) were carried out through a fixed-rate tender procedure with full allotment at the interest rate on the main refinancing operation (MROR; ECB 2008b). The “excess liquidity” that this allowed in the banking sector moved overnight rates from close to the MROR down to close to the Deposit Facility Rate, which therefore became the effective

35. Calculated from the World Trade Organization’s international trade statistics and the International Monetary Fund’s Direction of Trade Statistics.

policy rate (figure 5).³⁶ In other words, the effective interest rate reduction after October 2008 added up to 400 basis points, from an MROR of 4.25 percent to a DFR of 0.25 percent. In addition, the ECB (2008b) expanded the list of marketable assets eligible as collateral in Eurosystem credit operations. Both measures stepped up significantly the ease with which counterparties could satisfy their liquidity demands. The ECB also reduced the minimum rating threshold for eligible collateral, from A– to BBB–, adjusting to the fact that the crisis had lowered the average credit quality of assets in the market. Furthermore, the ECB enhanced liquidity provision through longer-term refinancing (after having introduced six-month operations already in the preceding March), further lengthening the average maturity of its outstanding operations, and provided U.S. dollar liquidity through foreign exchange swaps (as was already the case in December 2007, January 2008, and March 2008). The former gave greater planning certainty to counterparties, and the latter helped manage dollar shortages in the euro area spilling over from instabilities in the U.S. (ECB 2014). Finally, the corridor of standing facilities was temporarily reduced from 200 to 100 basis points from October 2008 to January 2009 (figure 5), to further contain short-term money market rate volatility. As the ECB became the “market maker” in the money market, its balance sheet expanded significantly.


Additional nonstandard measures were adopted in May 2009—when the MRO rate reached the 1 percent level and the DFR reached the 0.25 percent level—to support the flow of credit to households and corporations.³⁷ These included announcements of the lengthening of the maximum maturity of refinancing operations to one year (one-year LTROs, starting in June) and the Covered Bonds Purchase Programme (CBPP, starting in July), the first outright purchase program carried out by the ECB with the aim of reviving the funding channel for banks and support for their credit intermediation. Together with those measures adopted in October 2008, these measures configured the ECB’s policy of “Enhanced Credit Support” in response to the financial crisis (Trichet 2009). Interestingly, the press conference after the Governing Council’s meeting on May 7, 2009, was the first time that

36. In figure 16 below, the difference between the ECB’s total net monetary policy operations, excluding recourse to standing facilities (the upper end of the figure, minus liquidity-absorbing operations) and the banking sector’s liquidity needs (thick black line) or the negative values for net recourse to the Deposit Facility and daily reserve surpluses illustrates this excess liquidity in the context of the Eurosystem’s balance sheet. Before the fall of 2008, there usually was no material excess liquidity. For the relationship between excess liquidity and money market rates, see figure 21.

37. A comprehensive description of ECB market operations between the first quarter of 2009 and the second quarter of 2012 is provided by Eser and others (2012).

Figure 14. Changes in the Euro Area's Bank Credit Standards, 2002–17^a

Net percentage of banks reporting tightening credit standards




Source: ECB Bank Lending Survey in July 2018.

a. NFCs = nonfinancial corporations. The solid line refers to changes in standards applied to the approval of loans or credit lines to NFCs. The dashed line refers to the standards applied to loans to households for house purchases. Net percentages are defined as the difference between the sum of the percentages of banks responding “tightened considerably” and “tightened somewhat” and the sum of the percentages of banks responding “eased somewhat” and “eased considerably.” “Actual” values are changes that the bank respondents to the survey report to have occurred, while “expected” values are changes anticipated by banks. The most recent observations are for 2018:Q2 for actual changes and for 2018:Q3 for expected changes.

some of these nonstandard measures were briefly included in the formal introductory statement by the president (and were only later detailed in separate press releases; ECB 2009c).

The combination of these standard and nonstandard monetary policy responses had a beneficial impact on interbank market spreads (figure 13) and on financing conditions more generally (figures 14 and 15). They contributed, together with expansionary fiscal policies (figure 12) and financial sector support measures, to the initial economic and financial recovery from the Great Recession.³⁸ For example, the cumulative government

38. In line with an agreement for strengthening growth reached at the first Group of Twenty's summit in Washington in November 2008, the European Commission combined national initiatives and a smaller share of EU funding to a \$200 billion concerted European Economic Recovery Plan to boost demand and stimulate confidence in the European Union (European Commission 2008). The total plan amounted to spending of about 1.5 percent of GDP, which was endorsed by the European Council in December 2008. For an analysis of the effects of this fiscal stimulus, see Coenen, Straub, and Trabandt (2012).

Figure 15. The Euro Area's Bank Lending Rates, 2003–18^a

Source: ECB data.

a. NFCs = nonfinancial corporations. The two indicators show the total cost of bank borrowing for NFCs (solid line) and for households financing house purchases (dashed line). They are calculated by aggregating short- and long-term rates using a 24-month moving average of new business volumes. The most recent observation is for August 2018.

support for euro area financial institutions in the form of commitments for capital injections, liability guarantees, or asset support between October 2008 and May 2010 has been estimated by Stéphanie Stolz and Michael Wedow (2010) at about 28 percent of GDP (although the effective amounts were only about half of this). Already at that time, however, bank stress tests did not have all the desirable effects. For example, not long after the first European coordinated tests of 22 major cross-border groups under the Committee of European Banking Supervisors (CEBS 2009)—which, however, was run without a minimum capital threshold—there were further bank failures in the euro area.³⁹

By the end of 2009, nevertheless, year-on-year real GDP growth turned positive again and continued to pick up in 2010 (figure 7). The fall in underlying inflation stopped at about 1 percent in late 2009 and early 2010;

39. For a comprehensive overview of national financial sector policies during the crisis, including national stress tests, see European Commission (2017).

and at about the same time, headline inflation rose quickly again, reaching about 1.7 percent in the second quarter of 2010 and, ultimately, levels above 2 percent (figure 6), as energy prices again increased (figure 8). In addition, a modest recovery in money and credit growth began in mid-2010 (figure 9). This led to an initial discussion about phasing out some of the exceptional monetary policies, which ex post proved to be premature.

The European sovereign debt crisis and the sovereign–bank nexus, redenomination risk, and the second recession, May 2010–June 2013. The financial crisis and the Great Recession had left their mark on public finances. Government bond yield spreads increased significantly in the euro area (figure 13), particularly in those countries whose deficits rose substantially owing to the impact of automatic stabilizers in the face of a deep recession, discretionary expansionary fiscal policy (figure 12), and, importantly, interventions to shore up the banking sector (Stolz and Wedow 2010; Domingues Semeano and Ferdinandusse 2018). For example, public debt in the euro area as a whole rose from about 65 percent of GDP in early 2008 to about 78 percent in early 2010, and to above 90 percent in 2013. Particularly large increases occurred—notably, in Cyprus, Greece, Ireland, Portugal, and Spain. At least to some extent, risk was transferred from the financial sector onto public sector balance sheets, leading to a deterioration of fiscal positions.

Moreover, after the Greek public debt deceit started to be revealed in October 2009—leading, among other things, to a large revision of the reported government deficit for 2009—in April 2010, the Greek sovereign debt market seized up and markets lost confidence in the authorities’ ability and willingness to address the large rise in Greek government debt. Despite European governments putting together a rescue package and associated adjustment program for Greece and establishing the European Financial Stability Facility in June 2010—a (temporary) backstop vehicle for future crisis incidents—other “peripheral” countries faced their own crises in the following two years. Against the background of the discussion on countries with (private and public) debt overhang problems in subsection I.B and the further public debt increases referred to above, it is instructive to note that the affected countries, which needed some form of adjustment program with financial assistance, were precisely Ireland (2010), Portugal and Cyprus (both 2011), and Spain (2012). Over time, Italy also became seriously distressed, but never to the point that it had to start a rescue program.

The ECB’s policy response continued to abide by the separation principle. On one hand, in order to ensure depth and liquidity in the sovereign

bond markets of distressed countries and to restore the appropriate functioning of the monetary policy transmission mechanism, in May 2010 the ECB introduced its first sovereign bond purchase program, the Securities Markets Programme (SMP). To signal that the SMP was not designed to alter the stance of monetary policy, the SMP's purchases of debt securities were sterilized.

On the other hand, as both GDP growth and headline inflation picked up, the ECB raised its key policy rates by 25 basis points in April 2011 and again in July 2011, after almost two years of no change (figure 5). The euro area economy had grown at a quarterly rate of 0.8 percent in the first quarter of 2011, and the economic analysis revealed some upside risks to price stability. In fact, inflation had risen to 2.6 percent in March 2011 (and actually reached 3 percent toward the end of 2011, way above the medium-term objective below but close to 2 percent). There were concerns about second-round effects in the setting of prices and wages, and a risk of inflation expectations becoming unanchored from the ECB's definition of price stability.

Although the econometric evidence reviewed in subsection II.C finds that the SMP's interventions did put downward pressures on and lowered the volatility of sovereign yields for most countries, they did not stop the rise in sovereign spreads. By mid-July 2011, financial tensions intensified again due to the worsening of public finances in several euro area countries and contagion from the agreement to restructure Greek sovereign debt (which was, however, not implemented before 2012).⁴⁰ After bank bailouts had weakened sovereigns, the sovereign–bank nexus closed because the weakened sovereigns implied mark-to-market losses on banks' government debt holdings and an erosion of public guarantees (Acharya and others 2014). The sovereign debt crisis increasingly turned into a twin sovereign debt and banking crisis. Further negative feedback loops between vulnerable banks, indebted sovereigns, and weak economies took hold in several

40. The restructuring of Greek debt reflected one aspect of the Deauville agreement between Chancellor Merkel and President Sarkozy in October 2010, to promote “private sector involvement” in handling a public debt overhang. The lingering question about its application to the Greek case after October 2010, the decision in July 2011 to restructure Greek debt and the delay in its implementation until the spring of 2012 implied an ongoing source of uncertainty and volatility over an extended period of time. It should not be forgotten, however, that the Franco-German Deauville agreement constituted a much more wide-ranging public debt crisis resolution framework for Europe, which included—among other things—the establishment of a permanent rescue facility when the European Financial Stability Facility would expire in 2013. See Zettelmeyer, Trebesch, and Gulati (2013) for a detailed history of the Greek debt restructuring and events around it.

countries and led to acute financial fragmentation along country borders (Shambaugh 2012; ECB 2012a, 2013b). The most affected countries lost market access and entered adjustment programs (see above), contributing to a period of procyclical fiscal consolidation (figure 12) and stabilization slowdowns.

In addition to the fiscal woes and associated high and diverse sovereign yields across the euro area, monetary transmission remained severely hampered by lingering bank instabilities, which constrained the flow of credit to the economy and imposed significant obstacles to the ECB's Enhanced Credit Support. The reason was that in many countries, particularly the fiscally weak countries, bank recapitalization or resolution processes progressed only very slowly. Cases in point are the two EU-wide coordinated stress-testing exercises in 2010 and 2011. Although 7 euro area banks out of 91 EU banks (a much larger set than in the first, 2009 CEBS exercise) failed the July 2010 tests and had to raise new capital (CEBS 2010), their potentially beneficial effects on confidence in European banks were soon undermined by the fact that the two largest Irish banks, which had passed the test like many others, needed to be bailed out only a few months later. Similarly, 8 euro area banks out of 90 EU banks tested failed to meet the minimum threshold in the 2011 exercise, which was coordinated for the first time by the new European Banking Authority (2011), and were asked to replenish their capital. In October 2011, however, the large Franco-Belgian group Dexia, which had passed the test by a wide margin, went into resolution.⁴¹ At this point, the credibility of prudential and fiscal authorities' ability to solve Europe's banking problems was in a sorry state. Bank fragility and fragmentation remained a serious impediment to an effective monetary policy for this whole period.

As the financial tensions intensified and fiscal consolidation took hold, economic confidence fell, the economy slowed down rapidly and the euro area entered a double-dip recession in the last quarter of 2011 (figure 7). An important contributing factor was banks' deleveraging needs and the associated tightening of bank lending standards and further reductions in money and credit growth (figures 14, 15, and 9).

In response, the ECB entered a new monetary policy easing phase, during which—in November 2011—Mario Draghi also succeeded Jean-Claude Trichet as ECB president. On August 7, 2011, Trichet made a statement on Italy and Spain and announced that the ECB would reactivate its SMP (ECB 2011b). Toward the end of 2011, the

41. Spain's Bankia collapsed in April 2012.

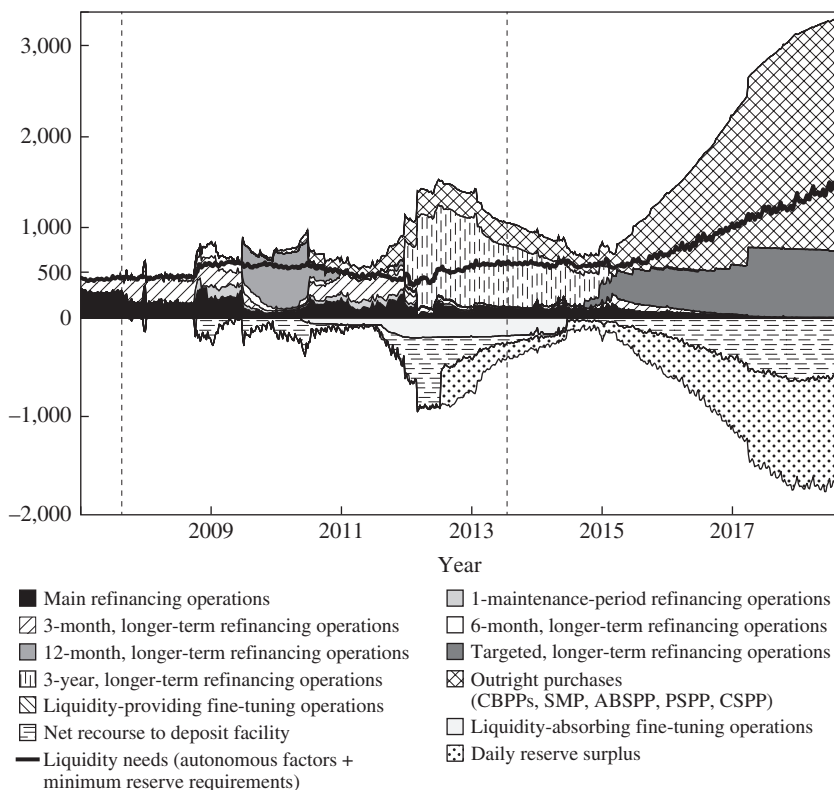
ECB introduced several new, nonstandard measures. Two LTROs of 12 and 13 months were announced on October 6, 2011, as well as a second covered bond purchase program (CBPP2) for an intended amount of €40 billion. Then the ECB reversed the interest rate hikes of April and July 2011 by cutting key policy rates in November and December 2011 by a total of 50 basis points. Moreover, in December 2011 and February 2012, two three-year Very-Long-Term Refinancing Operations (VLTROs), with the option of early repayment after one year, were conducted, with a combined gross amount of more than €1 trillion. The vertically dashed area in figure 16 shows what a large share these VLTROs assumed in total ECB monetary policy operations—for example, compared with the 1-year LTROs a few years earlier (the medium gray area of the figure). They gave banks funding certainty, eased redemption of maturing bonds, and helped them sustain credit lines with private customers. Finally, on December 8, 2011, the ECB also decided to again enlarge the collateral list via a reduction of the rating threshold for certain asset-backed securities and reduced the minimum reserve ratio from 2 to 1 percent (see subsection I.A).

These measures brought much needed relief for banks' funding, but—by definition of central bank liquidity operations—could not ensure the much needed balance sheet repair of many euro area banks. Also the need for fiscal consolidation lingered on. In early 2012, weak growth and news of fiscal slippages in several countries once more strained financial markets, and financial tensions rose again. Over the course of the sovereign debt crisis, a new phenomenon had slowly emerged, redenomination risk—the risk that euro assets could be redenominated in legacy currencies (De Santis, forthcoming). In other words, some premiums priced into the government bond yields of a few countries reflected increasing market-derived probabilities that those countries could leave the euro. In 2011 and, particularly, in 2012 some of them reached new heights (De Santis, forthcoming), increasing the cost of funding for several stressed euro area countries and seriously hampering the transmission of the ECB's policy stance to the real economy in those countries. Preserving the unity of the euro area became the defining challenge of the crisis.

This was the context in which, finally, decisive steps were also taken at the political level. For example, the “Fiscal Compact” was signed in March 2012, involving—among other things—a balanced-budget rule. More important, at a key European summit on June 28–29, 2012, the president of the European Council proposed significant reforms to EMU's financial, budget, and economic policy frameworks, notably the establishment

Figure 16. Quantities of the ECB's Market Operations from a Balance Sheet Perspective, 2007–18^a

Billions of euros



Source: ECB data.

a. CBPPs = Covered Bond Purchases Programmes; SMP = Securities Markets Programme; ABSPP = Asset-Backed Securities Purchase Programme; PSPP = Public Sector Securities Programme; CSPP = Corporate Sector Purchase Programme. This figure shows monetary policy items on the Eurosystem's balance sheet. Total ECB monetary policy operations (see also figure 4) are equal to the sum of Main Refinancing Operations, Longer-term Refinancing Operations (3-month, 6-month, 12-month, and 3-year), and outright purchase programs (CBPPs, SMP, ABSPP, PSPP, and CSPP), but excluding the net recourse to standing facilities (Marginal Lending and Deposit Facility). Liquidity needs are the sum of autonomous factors and minimum reserve requirements. Autonomous factors are factors, like banknotes in circulation and government deposits, that affect the liquidity needs of the banking system but are outside of the control of the central bank. Daily reserve surplus refers to the difference between banks' current account balances held with the central bank and banks' minimum reserve requirements. Excess liquidity can be approximated by adding daily reserve surplus and net recourse to deposit facility. Net recourse to deposit facility is the difference between recourse to the marginal lending facility and recourse to the deposit facility. The most recent observation is for July 17, 2018.

of the main elements of the European Banking Union—single supervision, resolution, and deposit insurance (European Council 2012a, 2012c).⁴² Making explicit reference to the need for breaking the sovereign–bank nexus, the euro area countries agreed to start with a Single Supervisory Mechanism at the ECB (European Council 2012a). Earlier, it had been agreed that the temporary European Financial Stability Facility (EFSF) would be replaced in October 2012 by the European Stability Mechanism (ESM), an intergovernmental organization to safeguard the financial stability of the euro area through financial assistance against strict conditionality to member states with severe financing problems. The ESM has a lending capacity of €500 billion, and it later also assumed the possibility of direct bank recapitalizations (European Council 2012a). Details of the reforms were worked out in the Four Presidents’ Road Map toward a genuine EMU, published in December 2012 (European Council 2012b), and in subsequent legislation.⁴³

In this new context of a much clearer path for fixing some of EMU’s most important financial and fiscal weaknesses, on July 11, 2012, the ECB lowered rates by 25 basis points, bringing the Deposit Facility Rate to 0 percent (which was then left unchanged for almost two years; figure 5). More important, on July 26, 2012, ECB president Draghi (2012) delivered a speech in London in which he gave the assurance that “within our mandate, the ECB is ready to do whatever it takes to preserve the euro. And believe me, it will be enough.” Several days later, on August 2, 2012, the ECB’s Governing Council announced it would introduce the Outright Monetary Transactions (OMT) Programme—which consists of purchasing sovereign bonds in secondary markets under strict conditions, with the aim of “safeguarding an appropriate monetary policy transmission and the singleness of monetary policy” (ECB 2012b) in the face of potentially self-fulfilling redenomination risks. The technical framework of the OMT was announced on September 6, 2012, and on the same day, the SMP was terminated. A necessary requirement for the OMT was strict and effective conditionality attached to an appropriate EFSF/ESM program (including a precautionary program). The OMT backstop was seen as credible, supported by the political agreements at the June Summit and the imminent start of the permanent ESM, and led to an immediate contraction of

42. In June 2012, the European Commission had also presented a first draft of the European Bank Recovery and Resolution Directive.

43. Also, Greek debt restructuring had finally taken place in March and April 2012, although the agreed-on bond exchange already had to be complemented with an EFSF buyback of newly issued debt in December (Zettelmeyer, Trebesch, and Gulati 2013).

sovereign bond spreads, which rapidly declined to more sustainable levels (figure 13).⁴⁴

On May 8, 2013, the ECB lowered the MRO rate by 25 basis points and the Marginal Lending Facility Rate by 50 basis points, further narrowing the interest rate corridor (figure 5). With the DFR already at 0 percent, room for further cuts in interest rates was increasingly limited. In response to the partial normalization of financial tensions, growth slowly picked up in the course of 2013.

DISCUSSION Overall, in the period between August 2007 and June 2013, the ECB entered the uncharted territory of nonstandard monetary policy measures.⁴⁵ At first, the ECB's operational framework was well suited to address impairments in the interbank market by providing ample liquidity for its wide set of counterparties and against a wide variety of collateral (Cassola, Durré, and Holthausen 2011; Eser and others 2012). The ECB particularly "lent to the market" like a traditional lender of last resort for the banking system.⁴⁶ In so doing, it relied on the separation principle to distinguish very generous liquidity provision from setting the monetary policy stance.

One question in this regard is whether (with the benefit of hindsight) the ECB was too optimistic about its (or other policy branches') ability to contain those impairments—notably, the later and more severe ones (see the next paragraph)—and their macroeconomic effects. This question has become subject to debate, in particular with respect to the short-lived tightening of standard monetary policy in 2008 and 2011 in parallel with continued easy liquidity provision. The reaction function analysis given in subsection II.B, which adopts the adequate real-time perspective, suggests that the July 2008 interest rate increase, although quite short-lived, was not fully in line with the ECB's own falling growth and inflation forecasts. The interest rate increases in 2011 were more in line with the strong growth and inflation forecasts in early 2011, though somewhat delayed. As nonstandard monetary policy measures became

44. For example, the Commission tabled a proposal for the Single Supervisory Mechanism in September 2012.

45. For various studies discussing how monetary and other central bank policies have changed over the last decade and how this is affecting central banks' roles more broadly, see Hartmann, Huang, and Schoenmaker (2018).

46. See, for example, Garcia-de-Andoain and others (2016) for an in-depth analysis of this "lending to the market" between 2008 and 2014. Emergency liquidity assistance to individual banks was undertaken, where needed, by euro area national central banks outside their Eurosystem responsibilities. But banks with sufficient Eurosystem-eligible collateral could also tap the ECB's marginal lending facility.

more forceful, the distinction between monetary policy stance and market operations started to soften. For example, as mentioned above, with the introduction of the fixed-rate/full-allotment credit operations in October 2008, the excess liquidity that started to build up in the banking system (see figure 16) pushed the overnight rate in the money market from the middle of the ECB's interest corridor to the bottom, making the DFR the effective policy rate. Also, nonstandard measures based on the ECB's market operations sometimes started to be mentioned in the introductory statement at the Governing Council's press conference.

However, as first the financial crisis and then the sovereign debt crisis took hold and the underlying solvency problems of both banks and sovereigns lingered on and reinforced each other, the incompleteness of EMU in the banking and fiscal areas became increasingly obvious (see also European Commission 2015; and Baldwin and Giavazzi 2015, 2016) and undermined the effectiveness of the ECB's monetary policy. The imperfect ways in which major financial and fiscal instabilities were addressed by the competent authorities, and the absence of sufficient institutions and tools for solving the related collective action problems in a highly integrated monetary union of sovereign states with primarily national fiscal and supervisory policies, posed formidable challenges for the ECB's monetary policy. An early indication of this was that in spite of very early generous liquidity provision, the ECB did not succeed in pushing interbank market rates all the way back down close to precrisis levels, as shown in figure 13. One plausible explanation is that these spreads contained a significant credit risk component and that credit risks and liquidity risks were strongly intertwined (Eisenschmidt and Taping 2009; Angelini, Nobili, and Piscillo 2011). Relatedly, the pass-through of the lower policy rates to bank lending rates became very uneven across countries over time as financial fragmentation took hold, again undermining the effectiveness of monetary policy (ECB 2015b).

ECB monetary policy itself could not address the underlying solvency issues of either banks or governments. In fact, the prohibition of monetary financing laid down in the EU treaty forbids the ECB from directly financing governments or government tasks, such as the recapitalization of banks.⁴⁷ It provides an important protection of the ECB from fiscal

47. Article 123 of the Treaty on the Functioning of the European Union prohibits overdraft facilities or any other type of credit facility for governments or government institutions with the European Central Bank or the Eurosystem, as well as the direct purchase from them of debt instruments.

dominance over its monetary policy, thereby supporting the achievement of price stability in the medium to long term. Instead, such solvency issues can only be effectively addressed by prudential and fiscal authorities. Unfortunately, major progress in addressing the institutional limitations in the field of supervision and resolution was only achieved toward the end of this period, as political agreements were reached to build a banking union—with the setting up of the Single Supervisory Mechanism at the ECB and the Single Resolution Mechanism—and to strengthen the backstop for governments through the permanent ESM.

Against this background, the ECB's actions had to balance the need to address impairments in the transmission of monetary policy due to malfunctioning financial markets and self-fulfilling market dynamics with the prohibition of monetary financing. This partly explains what some observers regard as the initial timid interventions in the government bond market through the SMP based on implicit conditionality.⁴⁸ Leading up to the June 2012 European Summit, however, the necessary institutions and reforms to improve on the main weaknesses of EMU in the prudential and fiscal fields were put on a credible path. In this context, the ECB stepped up its nonstandard tool kit to the next level, starting with President Draghi's "whatever it takes" speech and the powerful OMT program, based on the explicit conditionality of an adequate EFSF/ESM program.

1.D. Deflation Risks and Low-Inflation Recovery, June 2013–June 2018

The last cyclical period of the euro area that we cover in this paper concerns the slow recovery after the crises. The protracted low-inflation fallout of the sovereign debt crisis and risks of deanchoring inflation expectations led the ECB to further extend its nonstandard monetary policy and communications tool kit. Although this reinforced discussions about the benefits and risks of such policies, in various dimensions it made the ECB more similar to its main peers.

ADDRESSING THE LOWER BOUND ON INTEREST RATES The fourth and most recent episode was characterized by the ECB's actions to overcome the zero lower bound on interest rates in its attempt to address deflation risks and bring inflation back to levels close to 2 percent. In doing so, the ECB turned to policies such as quantitative easing, funding for lending, and explicit forward guidance that had been used before by other central banks,

48. As discussed in subsection II.C, the ECB characterized the SMP's interventions as limited and temporary, leading markets to doubt that it was prepared to offer a full backstop.

such as the U.S. Federal Reserve and the Bank of England. The ECB was, however, the first major central bank to also go into negative interest rate territory. We review existing evidence on the effectiveness of these programs in subsection II.C.

As monetary policy became much more complex, there was an increased need for communication. As part of the efforts to provide enhanced communication in a more complex environment, in January 2015 the Governing Council decided to release the accounts of its meetings on monetary policy, about four weeks after each meeting (Draghi 2014a). At the same time, the frequency of these meetings was changed from monthly to eight times a year, in order to better align them with the arrival of sufficient new information and to reduce the number of instances when expectations could cause market volatility. Unlike previous communications, in which the ECB had stated that it would not precommit on monetary policy decisions, it also turned to forward guidance in this period (see figure 26 below). Following the taper tantrum in the U.S., which led to significant undesired interest rate spillovers to the euro area, the ECB introduced explicit forward guidance about the future path of key interest rates in July 2013. As in other central banks, the precise formulation of the forward guidance evolved over time, as we describe later in this subsection.

NEGATIVE RATES, TARGETED LENDING, AND QUANTITATIVE EASING The sovereign debt crisis abated, and the recovery started to take hold, as some of its underlying causes were addressed by the various country adjustment programs, the creation of a banking union with common supervision and resolution, and the establishment of a backstop for governments via the ESM and the ECB's OMT program. However, the damage of high unemployment and negative output gaps in 2012 and 2013 was done (figure 3). Toward the second half of 2013, both headline and core inflation dropped below 1 percent, and headline inflation became negative in the course of 2015 (to a minimum of -0.7 percent in January 2015; see figure 6), largely on account of falling energy prices (figure 8). Inflation expectations, which up until then had remained well anchored, started to decline and to exhibit a significant downward skewing (subsection II.A).⁴⁹ Concerns grew about deflation risks and a prolonged period of low inflation. Moreover, it became increasingly clear that the transmission of the easing of ECB key policy rates had remained impaired and uneven. In particular,

49. For example, in his speech at the Federal Reserve Bank of Kansas City's Jackson Hole Symposium in August 2014, President Draghi (2014b) digressed from his main topic of euro area unemployment to point out that inflation expectations were declining significantly at all horizons (see figures 18 and 20 below).

the cumulative reduction of 125 basis points in the MROR (75 basis points in the DFR) from November 2011 to November 2013 had not yet been transmitted to households and firms in the stressed euro area countries. Over time, as the medium-term outlook for inflation continued to worsen, the ECB's balance sheet shrank; credit growth remained negative, reflecting ongoing deleveraging; and until March 2014, the exchange rate of the euro strengthened (figures 16, 9, and 10).

To stave off emerging deflation risks and address the impairment of the bank lending channel, the ECB embarked on a three-pronged, comprehensive monetary policy easing strategy starting in June 2014. This strategy was foreshadowed in a speech by President Draghi (2014a), in which he laid out the conditions for the three elements of the easing strategy. A first measure was to go into negative interest rate territory. In June 2014, and again in September 2014, the ECB lowered the DFR by 10 basis points, to -0.2 percent. Second, to revive the provision of credit and address the fragmented policy transmission, it announced a renewed round of credit easing measures with a series of Targeted Longer-Term Refinancing Operations (TLTROs) fixed at the MRO rate plus 10 basis points. The surcharge was abolished in January 2015. The maximum maturity for the TLTROs was set to September 2018, and the initial allowance for the maximum amount of borrowing under the TLTRO program amounted to 7 percent of outstanding loans to the euro area's nonfinancial private sector. The maturity of the loans was conditional on banks exceeding certain lending thresholds for the corporate sector. These credit-easing measures were complemented by an asset-backed securities purchase program, and a third covered bond purchase program in September 2014. Third, to provide additional stimulus in an environment where further cuts in short-term rates were constrained, in January 2015 the ECB announced an expanded Asset Purchase Programme (APP), with average monthly purchases of public and private sector securities of €60 billion. Through the portfolio rebalancing and signaling channels, this put further downward pressure on long-term interest rates and flattened the slope of the yield curve (Coeuré 2015). At the same time, it led to a big expansion of the ECB's balance sheet (figure 16).

The combined impact of these measures was to lower market- and bank-based financing costs and ease financial conditions more broadly (subsection II.C). Figure 15 shows that the composite indicator of the cost of borrowing for nonfinancial corporations (NFCs) and households fell from 3 percent to close to 2 percent at the end of 2015, and bank lending rates started to converge in the wake of the earlier fragmentation. At the same time, banks started easing their lending standards, and credit growth to the

private sector gradually started recovering (figures 9 and 14). By the summer of 2015, GDP growth had picked up to close to 2 percent (figure 7), and both headline and underlying inflation had stabilized, but at relatively low levels of 0 and 1 percent, respectively (figure 6).

Against this background of still uneven and fragile growth and low inflation, the growth in global demand faltered in the summer of 2015, as a stock market collapse in China and an unexpected depreciation of the renminbi caused financial turbulence in emerging market economies. In order to avoid a renewed increase in deflation risk and to continue to support the gradual recovery of the euro area's economy, the three-pronged package of measures was recalibrated again in December 2015 and March 2016 with a view to adding further monetary policy stimulus. On December 3, 2015, the ECB lowered interest rates further by 10 basis points and announced a recalibration of the APP, prolonging the program until March 2017, or beyond if necessary, to ensure a sustained adjustment of inflation toward the aim of being below, but close to, 2 percent (ECB 2015d). At the same time, the ECB announced that it would reinvest the principal repayments, keeping the stock of the APP portfolio constant after the end of the net purchases for as long as necessary, and extended the list of APP-eligible assets to include securities issued by regional and local governments. On March 10, 2016, the ECB decided again to lower rates, with effect from March 16, bringing the interest rate corridor down to 65 basis points and lowering the DFR by 10 basis points, to -0.4 percent (ECB 2016b, 2016c). At the same time, a considerable expansion of the APP was announced, with average monthly purchases being increased to €80 billion. The ECB also launched the Corporate Sector Purchase Programme as an integral part of the APP. Finally, four new TLTROs, known as TLTRO-II, each with a maturity of four years, were announced, starting in June 2016 and running until March 2017. The rates on these operations could be as low as the negative DFR, if banks exceeded certain lending benchmarks. The combined effect of these additional easing measures was to further improve financing conditions. This stimulated domestic demand and turned a fragile and uneven recovery into a solid and broad-based expansion, in spite of the temporary weakness of the international economy. Accordingly, but also partly due to rising commodity prices, inflation picked up as of 2016:Q3, reaching almost 1.8 percent in January 2017 (figures 6 and 8).

As the euro area's economy strengthened, but underlying inflation remained subdued, further monetary stimulus was deemed appropriate, but the intensity of the stimulus was gradually adjusted. On December 8, 2016, the Governing Council decided to extend the net APP until the

end of 2017, while at the same time reducing its monthly pace as of April 2017 to €60 billion. As of 2017:Q2, growth further increased, peaking at 2.8 percent in September (figure 7). On October 26, 2017, it was decided to further recalibrate the APP, and the program was extended until September 2018, with a reduced monthly pace of purchases of €30 billion starting in January 2018. Finally, on June 14, 2018, the Governing Council announced an anticipated extension of the net APP until the end of 2018 at a reduced pace of €15 billion, after which the net APP was expected to end. At the same time, it enhanced its forward guidance on policy rates by stating that it expects policy rates to remain at their present levels at least through the summer of 2019 and, in any case, for as long as necessary to ensure the sustained convergence of inflation to levels that are below, but close to, 2 percent. In fact, headline inflation stabilized at close to 2 percent during the summer of 2018, whereas core inflation continued to “creep up” only very slowly.

DISCUSSION Overall, the fourth episode was characterized by the ECB’s actions to overcome the zero lower bound on interest rates in its attempt to address deflation risks and bring inflation back to levels close to 2 percent. In doing so, the ECB turned to policies such as negative interest rates, quantitative easing, funding for lending, and explicit forward guidance; and in this respect, it started to more closely resemble many of its peers.

Most of the debates in this period related to the rationale, the sequencing, and the costs and benefits of the new nonstandard measures. We review the rationale for these various measures and the evidence of their effectiveness in subsection II.C. Here, it is important to realize that, as the ECB ventured into uncharted territory, it learned from its own and other central banks’ experience. A prominent example is the introduction of a negative DFR, which was introduced in small steps of 10 basis points and followed the positive experience with negative rates in a number of smaller countries, such as Denmark and Switzerland (Jackson 2015; Martínez Pagés and Millaruelo 2016).

Also, the ECB’s forward guidance evolved in this period (see figure 26 below). After the taper tantrum in the United States, the ECB announced that the policy rates were expected “to remain at present or lower levels for an extended period of time” and that this expectation was “based on the overall subdued outlook for inflation extending into the medium term, given the broad-based weakness in the real economy and subdued monetary dynamics” (ECB 2013c). The aim was to anchor policy expectations and maintain an accommodative level of long-term interest rates in the face of rising bond yields in the global market and a still very

subdued and fragile euro area recovery. As explained by Peter Praet (2013), the forward guidance on interest rates was meant to clarify the ECB's reaction function. As in other central banks, the ECB's forward guidance framework subsequently evolved. It took on a more complex and time- and state-dependent form when the expanded APP was announced in January 2015. On this occasion, the ECB also gave forward guidance on the net asset purchases and announced that they "are intended to be carried out until end-September 2016 and . . . in any case . . . until [it sees] a sustained adjustment in the path of inflation which is consistent with its aim of achieving inflation rates below, but close to, 2 percent over the medium term" (ECB 2015c). This forward guidance therefore had both time- and state-dependent conditioning elements. The former underscored the commitment made by the Governing Council, whereas the latter made the state-dependent nature of the forward guidance clear. A direct link with the ultimate objective was seen as more appropriate than alternative intermediate targets, also in light of the mixed experience with conditioning variables, such as unemployment in the United States and the United Kingdom. The APP was subsequently extended in December 2015, in December 2016, and in October 2017, maintaining a similar formulation.

In March 2016, when the APP's monthly purchases were increased from €60 billion to €80 billion, the ECB also for the first time linked forward guidance on interest rates to that on the APP, by stating that the "key interest rates would remain at present or lower levels for an extended period of time and well past the end of the net asset purchases" (ECB 2016c). This helped to secure the credibility of the interest rate forward guidance (Coenen and others 2017), thereby reinforcing both parts of the easing program, and it also provided clarity on the sequencing in the normalization of the various elements of the easing measures (Praet 2018). In June 2017, the reference to lower interest rates (the "easing bias") was dropped (ECB 2017d). And in June 2018, when the anticipation of the end of the net asset purchases by the end of 2018 was announced, forward guidance on interest rates was delinked from the APP, and it was stated that "the Governing Council expects key interest rates to remain at the present level at least through the summer of 2019 and, in any case, as long as necessary to ensure . . . the continued sustained convergence of inflation to levels that are below, but close to, 2 percent over the medium term" (ECB 2018c). A time- and state-based element is now attached to the liftoff of policy rates.

Controversy about the ECB's policy decisions in this period focused mostly on the Public Sector Purchase Programme (PSPP), which constituted the largest part of the APP. Despite an observable slide in inflation,

there was some opposition to a large-scale bond purchase program because of concerns about potential monetary financing (Article 123 of the Treaty on the Functioning of the European Monetary Union, EU 2012b), additional central bank balance sheet risks, independence in making interest rate decisions with a large government bond portfolio, and possible effects on governments' willingness to pursue debt consolidation and enact reforms (Weidmann 2015). The fact that deploying asset purchases was fully in line with the ECB's mandate was confirmed by the European Court of Justice. In its judgment on the OMT program, it ruled that purchases of government bonds are legal under the ECB's statute and are a legitimate tool of monetary policy (Court of Justice 2015). To ensure that secondary market purchases of government bonds cannot be assimilated to primary market purchases that are forbidden under the monetary financing prohibition, it is, however, also important to ensure that the program is consistent with the ultimate objectives of Article 123, namely, safeguarding (1) the primary objective of price stability, (2) the central bank's independence, and (3) the fiscal discipline of a member state. To this effect, the ECB built sufficient safeguards into the PSPP. First, PSPP purchases adhere to a blackout period; that is, the Eurosystem does not buy near the date of a new issuance, which facilitates the formation of market prices for PSPP-eligible securities (ECB 2015a). Furthermore, the relevant securities are subject to an issue share limit and an issuer limit, which preserve market functioning. Finally, to avoid free-riding by national governments, risk-sharing of the PSPP was limited to 20 percent of the portfolio and the portfolio weights were guided by the capital key—that is, the share of each national central bank in the ECB's capital (ECB 2015e). All these safeguards were designed to ensure that PSPP purchases stay well clear of monetary financing.

II. Assessing the ECB's Monetary Policy


After the chronological part in the previous section, we now move to assessing key aspects of the ECB's monetary policy during its first 20 years. In turn, we cover the achievement of price stability, the ECB's primary objective, the standard interest rate decisions, and the more recent nonstandard monetary policy measures.

II.A. The Objective of Price Stability: Performance, Credibility, and Challenges

Let us start by analyzing the performance and credibility regarding the ECB's primary mandate of medium-term price stability for the euro area.

Figure 17. The Euro Area's Headline Inflation and a Five-Year Moving Average, 1999–2018^a

Year-on-year percentage change



Source: ECB data.

a. The long-term average of inflation according to the Harmonised Index of Consumer Prices since 1999 is 1.7 percent. The moving average is centered, and its values toward the end of the period are calculated using the ECB/Eurosysteem inflation projections. The most recent observation is for September 2018.

The key question is to what extent the ECB managed to anchor medium-term inflation expectations in a way that is consistent with its mandate, particularly in the aftermath of the twin crises. Later in this subsection, we discuss implications for the definition of the ECB's inflation aim.

HOW WELL ANCHORED ARE INFLATION EXPECTATIONS IN THE EURO AREA?

Figure 17 shows that over the past two decades, average euro area inflation has been about 1.7 percent. This average outcome is consistent with, but on the low side of, Issing's indication of an inflation aim between 1.7 and 1.9 percent. Over this period, annual HICP inflation has roughly fluctuated between 0 and 4 percent, mostly reflecting the impact of volatile energy and food price inflation. The range for core inflation (that is, HICP inflation, excluding energy and food) is smaller, and lies between 0.6 and 2.6 percent, reflecting its more sluggish nature (figure 6). Figure 17 also depicts a five-year centered moving average of HICP inflation, which may capture a more appropriate medium-term horizon for assessing the ECB's performance. This moving average fluctuated closely around 2 percent until the sovereign debt crisis, but started to decline below its previous range in the second half of 2012 and fell to a historic low of about


0.7 percent at the beginning of 2016, and since then has been expected to slowly recover.⁵⁰

Given the imperfect short-term control of inflation by the central bank, it is also useful to examine the stability of medium- to longer-term inflation expectations. The anchoring of longer-term inflation expectations to the ECB's inflation aim is a good measure of the ECB's credibility for maintaining price stability over the medium term. The empirical literature has shown that the degree to which inflation expectations are anchored has been dispersed across countries and time, and appears to co-move with the degree of credibility of monetary policy. The tendency toward better-anchored expectations was typically stronger in countries with official inflation targets, suggesting that agents use inflation targets as focal points when forming longer-term inflation expectations (Demertzis, Marcelino, and Viegi 2009; Gürkaynak, Levin, and Swanson 2010).

A study focusing on the earlier part of the EMU period (Beechey, Johannsen, and Levin 2011) showed that, on average, the euro area's long-run inflation expectations were more firmly anchored than those in the United States.⁵¹ In this subsection, we follow Jonas Dovern and Geoff Kenny (2017), and use data from the ECB's Survey of Professional Forecasters (SPF) to examine how the various moments of longer-term inflation expectations in the euro area have evolved over the past two decades. Figure 18 shows the evolution of two measures of average 5-year-ahead inflation expectations taken from the SPF (together with two measures of market-based inflation expectations derived from swap rates between inflation-adjusted and nominal government bonds). The average point forecast (the dotted line, as also included in figure 6) stayed close to 2.0 percent over the full EMU period, roughly fluctuating between 1.8 and 2.0 percent. The average mean of the individual forecasters' distributions (dashed line) has fluctuated a bit more, and reached a minimum of 1.65 percent at the beginning of 2016. As shown by Tomasz Lyziak and Maritta Paloviita (2017), there is some dependence of these average forecasts on a moving average of actual inflation, but overall these movements have been very contained. Using more formal tests for breaks in mean longer-term inflation expectations, Dovern and Kenny (2017) find two significant breaks in 2005:Q2 and then again in 2013:Q2. In 2005:Q2, the mean expectation shifted upward, from an estimated 1.85 percent to 1.92 percent. Arguably, this may be due

50. The ECB/Eurosystem's inflation projections are used to calculate the 5-year moving average toward the end of the period.

51. See also Ehrmann and others (2011).

Figure 18. Survey and Market-Based Inflation Expectations in the Euro Area, 2005–18^a

Sources: ECB data; Survey of Professional Forecasters.


a. This figure shows different measures of the first moment of inflation expectations. The average point estimate refers to the average of 5-year-ahead point forecasts for inflation per the Harmonised Index of Consumer Prices (HICP) across contributors to the ECB Survey of Professional Forecasters. The average distributional mean refers to the mean of the aggregate 5-year-ahead forecast distribution for HICP inflation across contributors to the SPF. For further explanations, see ECB (2017a). The most recent observation is for July 2, 2018.

to the clarification of the definition of price stability as below, but close to, 2 percent in 2003. This upward movement in expectations was, however, more than reversed in 2013:Q2, when the mean inflation expectation dropped back to about 1.8 percent, partly in response to the persistently low level of inflation after the sovereign debt crisis in the euro area.

It is also instructive to look at the second moment of the longer-term forecast distribution. Figure 19 shows three measures of longer-term inflation uncertainty based on the SPF. First, it shows a measure of disagreement among professional forecasters, that is, the standard deviation of individual forecasters' point forecasts (the solid line). Disagreement fell significantly in the first decade of EMU, from 0.4 to 0.1 percentage point, suggesting that the ECB's extensive communication about its stability-oriented monetary policy strategy (see figure 2, the medium gray area) and the quantitative definition as well as the consistent and transparent conduct of monetary policy

Figure 19. Survey-Based Longer-Term Inflation Uncertainty in the Euro Area, 1999–2018^a

Standard deviation




Sources: ECB data; ECB (2017a).

a. This figure shows different measures of the second moment of 5-year-ahead inflation expectations. Disagreement refers to the standard deviation of point inflation forecasts per the Harmonised Index of Consumer Prices (HICP) made by contributors to the ECB Survey of Professional Forecasters. Individual uncertainty refers to the average of individual forecasters' standard deviations for HICP inflation. Aggregate uncertainty refers to the standard deviation of the aggregate distribution of HICP inflation forecasts contributed by participants in the survey. For further explanations, see ECB (2017a). The most recent observation is for July 1, 2018.

were effective in aligning longer-term expectations across forecasters. Although disagreement rose significantly after the start of the Great Recession, it has fallen back, reaching levels close to 0.15 percentage point since then. The other two measures shown in figure 19 take into account the individual forecast uncertainty. After the financial crisis, longer-term inflation forecast uncertainty has clearly increased, also reflecting the higher variance of actual HICP inflation after 2007 (Dovern and Kenny 2017). There is no evidence that this measure of uncertainty has so far significantly reverted to its precrisis level.

Finally, one can also analyze the balance of longer-term inflation risks as captured by the SPF expected distributions. Figure 20 gives the range of a number of such measures, as well as their average. It shows that before the financial crisis, the risks around the longer-term inflation forecast were roughly balanced. Interestingly, a slight negative skewing emerged in about 2003–4, when, as discussed above, there was a debate about the impact of

Figure 20. The Survey-Based Balance of Longer-Term Inflation Risks and the Inflation Risk Premium in the Euro Area, 1999–2018^a



Sources: ECB data; Joslin, Singleton, and Zhu (2011); Camba-Méndez and Werner (2017); ECB (2017a).

a. This figure shows the average and the range of 12 different measures of the third moment (skew) of 5-year-ahead inflation expectations from the ECB Survey of Professional Forecasters. The Inflation Risk Premium decomposition is based on an affine term structure model and fitted to the euro area zero-coupon, inflation-linked swap curve. The estimation method follows Joslin, Singleton, and Zhu (2011); for details, see Camba-Méndez and Werner (2017). For further information, see ECB (2017a). The most recent observation is for August 1, 2018.

the zero lower bound on optimal inflation targets and the ECB's inflation aim was clarified. However, the skewing became persistently negative after the beginning of the financial crisis, and in particular after the sovereign debt crisis. Most recently, there has been a return toward more balanced risks. This is consistent with recent evidence given in a paper by Olesya Grishchenko, Sarah Mouabbi, and Jean-Paul Renne (2017).

Figure 20 also shows that the negative skewing is highly correlated with model-based estimates of the inflation risk premium in inflation-indexed bonds and can explain why market-based, 5-year-ahead, 5-year-forward inflation rates have been more responsive to actual headline inflation than the average survey expectations (figure 18). As the probability of getting trapped in a low inflation or deflation regime increases, the demand for deflation protection rises, affecting inflation risk premiums.

A second approach for analyzing the anchoring of inflation expectations is to investigate the sensitivity of longer-term inflation expectations to short-term macroeconomic news and inflation developments, as reviewed by Matteo Ciccarelli and Chiara Osbat (2017). Although the reviewed studies differ in the details of their respective methodologies, there are some common findings that are consistent with the evidence noted above. Before the financial crisis, no significant pass-through effects were recorded. But the overall picture is less clear after the start of the crisis.⁵² However, after the negative oil price shock of mid-2014, three out of four pass-through measures identified increasing risks of a deanchoring of longer-term inflation expectations. In 2015, the announcement and subsequent implementation of the APP seem to have softened these risks, and some studies suggest that the pass-through signal has become insignificant.

Overall, this review of the evidence suggests that in contrast to some early fears, the ECB was effective in anchoring medium- to longer-term inflation expectations to its inflation aim early on (Smets 2010). Moreover, modal expectations remained anchored below, but close to, 2 percent throughout the financial and sovereign debt crises. However, the higher uncertainty about the expected longer-term inflation forecast and the emergence of a significant negative skewing in the balance of risks after the beginning of the sovereign debt crisis in particular suggests that the ECB was not able to fully dispel the probability of ending up in a low inflation / deflation regime (as had happened in Japan). This may not necessarily be related to the credibility of the ECB's commitment to maintain price stability, but it may be due to doubts that the ECB had the necessary tools to fight deflation in an environment of low interest rates. Not the willingness of the central bank, but its ability, may have been put in doubt, as the ECB was relatively slow in applying large-scale purchases of government bonds as a monetary policy tool—particularly as compared with other major central banks, such as the U.S. Federal Reserve System, the Bank of England, and the Bank of Japan. This explanation is also borne out by some evidence of asymmetry between the response of longer-term market inflation compensation measures to inflationary and deflationary shocks (Natoli and Sigalotti 2018). Though this may have been more important for the ECB, where a discussion on the

52. For the United States, a number of studies have shown that longer-term mean inflation expectations started to react more strongly to macroeconomic news after the financial turmoil of 2008; see Galati, Poelhekke, and Zhou (2011); and Autrup and Grothe (2014).

use of quantitative easing was more intense and may explain its delayed implementation, the fact that this feature has to some extent also been observed in other jurisdictions with a single fiscal authority suggests that it may be a more general phenomenon related to the risk that one can get trapped in a deflation regime once inflation expectations adjust (Benhabib, Schmitt-Grohé, and Uribe 2001). As discussed above, the emergence of medium-term deflation risks eventually led the ECB to embark on a comprehensive, unconventional easing program, which helped to remove deflation risks (Andrade and others 2016).

POSSIBLE IMPLICATIONS FOR THE DEFINITION OF PRICE STABILITY A review of the ECB's credibility highlights the reality that over the past two decades, the ECB's initial concerns that it may not have had the same anti-inflation credibility as some of its predecessors, such as the Deutsche Bundesbank, later were turned around into the opposite concern that it may not be sufficiently equipped to avoid a low inflation or deflation equilibrium. In this light, it is worthwhile to review some of the elements of the ECB's definition of price stability.

One issue is whether the excess sensitivity of longer-term inflation expectations to low inflation is partly due to a persistent perception of a lack of *symmetry* in the ECB's inflation objective. Due to the formulation of the inflation aim ("below, but close to"), many observers continue to think that the ECB's tolerance for lower inflation is higher than its tolerance for higher inflation, although ECB policymakers have continuously stressed the importance of symmetry.⁵³ The question of symmetry can be addressed within the literature on the policy reaction function. Maritta Paloviita and others (2017) find no evidence of asymmetry if the inflation target is assumed to be 1.7 percent, but some evidence of asymmetry if the target is assumed to be 2 percent.⁵⁴ In subsection II.B, we test for asymmetry in a simple policy reaction function setup, and find little evidence of a stronger response to positive deviations of inflation than to negative deviations from the ECB's inflation aim.

A related question is whether the *precision* of the medium-term inflation objective matters. As mentioned above, empirical evidence suggests that a point target helps agents to focus when forming inflation expectations

53. For example, in a recent speech, President Draghi (2016) emphasized the importance of pursuing the price stability objective symmetrically, particularly in a zero-lower-bound and high-debt environment. This criticism was around from day one, as discussed above.

54. Examples of other earlier studies of possible asymmetries in the ECB's monetary policy include those by Aguiar and Martins (2008, 1651), who find a "precautionary demand for price stability"; and Surico (2007).

and contributes to the anchoring of those expectations. This is why many academics were originally in favor of a point target (Bernanke and others 1999). It also explains why most inflation targeting by central banks has a clear focal point, even if this is often embedded within a target range to underline that a central bank cannot precisely pin down inflation at all times.⁵⁵

But what is the optimal focal point for inflation? In the advanced economies, there has been a convergence of inflation targets to 2 percent since the start of inflation-targeting regimes in New Zealand in 1989. Recent examples are the U.S. Federal Reserve in 2012, the Bank of Japan in 2013, and the Norges Bank in 2018. One argument against being very precise is that there is uncertainty surrounding *the optimal medium-run inflation objective* and that it may change over time. In the academic literature, estimates of the optimal inflation target vary from mild deflation to 4 percent and higher. The recent experience of higher macroeconomic volatility and a lower equilibrium real interest rate have led some macroeconomists to argue for higher inflation targets of 4 percent (Blanchard, Dell’Ariccia, and Mauro 2010; Ball 2014; Krugman 2014).⁵⁶ The higher probability of hitting the zero lower bound in an environment of low interest rates is also brought out in quantitative simulation studies like those by Michael Kiley and John Roberts (2017). At the same time, central banks, including the ECB, have gained much positive experience with the use of unconventional policy measures to circumvent the effective lower bound on short-term interest rates. Recent empirical research suggests that these tools may have been just as effective as the more standard short-term interest rate tools in steering the economy (as discussed in subsection II.C), although they may come with additional side effects.⁵⁷ And changing inflation objectives always runs the risk of undermining the central bank’s credibility and increasing uncertainty and the inflation risk premium.⁵⁸

55. A precise numerical target also helps in communication. In the words of Stephen Nickell (2006, 252), former member of the Bank of England’s Monetary Policy Committee: “In my own experience, I find being provided with a precise numerical inflation target enormously helpful, since I can then explain my own policy decisions very simply in terms of avoiding an undershoot or overshoot of this target.”

56. Early on, Wyplosz (2001) argued for a higher inflation target of 4 to 10 percent for the euro area on the basis of the presence of more significant downward nominal wage rigidities.

57. See, for example, Swanson (2018)—but for an opposite view, see Hamilton (2018).

58. For example, raising an inflation objective could increase the risk that inflation expectations could become unanchored (Ascari, Florio, and Gobbi 2017; Deutsche Bundesbank 2018) or be “too blunt an instrument” compared with alternative options (Coibion, Gorodnichenko, and Wieland 2012, 1371).

A suggested compromise has therefore been to keep the 2 percent focal point, but to strengthen the role of inflation expectations as an automatic stabilization mechanism to further alleviate the zero lower bound on interest rates. This can, for example, be done by average inflation targeting (Svensson 1999b; Nessen and Vestin 2005). Vitor Gaspar, Smets, and David Vestin (2010) show that the benefits of such an approach continue to exist even in the absence of rational expectations, as long as the agents learn and adapt their expectation formation to changes in the regime.⁵⁹

II.B. The Conduct of Monetary Policy: The ECB's Interest Rate Decisions

This subsection analyzes the ECB's interest rate decisions through the lens of an empirical interest rate reaction function. This is particularly appropriate until the ECB hits the zero lower bound in July 2012.


We explained in subsection I.A above how the ECB used its operational framework to steer short-term money market rates close to the MROR, during the first decade the main monetary policy rate agreed by the Governing Council. Figure 5 shows the developments of the main policy-controlled interest rates since the start of EMU and how during the first decade of EMU, the euro overnight interest rate—measured by the Euro Overnight Index Average (EONIA)—fluctuated within the corridor given by the MLFR and the DFR. The EONIA stayed relatively close to the MROR—that is, at the midpoint of the corridor—with exceptions at the end of the maintenance period, when unexpected liquidity shortages or surpluses can lead to sharp deviations within the corridor.

Although a number of refinements were made to the ECB's operational framework during its first decade, as we described in subsection I.C the big changes came with the severe worsening of the financial crisis in October 2008 (ECB 2011a). A key one was the switch to fixed-rate/full allotment tenders, as it led the DFR to become the effective monetary policy rate (and not any longer the MRO minimum bid rate). It triggered increasing excess liquidity, which made the EONIA drop below the MROR and toward the bottom of the corridor given by the DFR. The distance of the EONIA rate from the DFR is a (nonlinear) function of the amount of excess liquidity in the banking system, as illustrated in figure 21 (covering data between

59. An argument against average inflation targeting is that it may require short periods of deflation after periods of inflation. This is addressed in the proposal by Bernanke (2017) to install a price-level target only after periods in which the lower bound has been binding.

Figure 21. The Euro Area Banks' Excess Liquidity and the EONIA–DFR Spread^a

EONIA–DFR spread (percentage points)



Source: ECB data.

a. EONIA = Euro Overnight Index Average; DFR = Deposit Facility Rate; LTRO = Longer-Term Refinancing Operation; APP = Asset Purchase Programme. Each sign (dot, rhombus, or triangle) corresponds to the average spread between the EONIA and the DFR for a specific reserve maintenance period. Low excess liquidity levels refer to excess liquidity below €200 billion and correspond to the period before December 2011 and between the end of the 3-year LTROs and the start of the APP (about the end of 2013 and 2014). Medium levels refer to excess liquidity between €200 billion and €400 billion, and high levels refer to excess liquidity above €400 billion. The sample period is from January 20, 2010, to May 2, 2018.

2010 and 2018). For example, two periods were characterized by a rising EONIA relative to the deposit rate (the gray dots in the figure). The first one was 2011, when the macroeconomic picture improved, the ECB raised rates twice, and excess liquidity dropped to very low levels. The second period was toward the end of 2013 and 2014, when excess liquidity again fell to low levels as banks started repaying VLTROs. In figure 21, medium and high levels of excess liquidity are marked, respectively, with black rhombuses and light gray triangles.

In the rest of this section, we analyze the setting of the main policy rate through the lens of a simple, but robust, first-difference policy rule originally proposed by Athanasios Orphanides (2003). This rule links the change in the main policy rate of the ECB (the minimum bid rate in MROs before October 2008, and the DFR after October 2008) to deviations of


the one-year-ahead inflation forecast from the ECB's inflation aim and deviations of the one-year-ahead real GDP growth forecast from potential output growth:

$$(1) \quad \Delta i = 0.5(E\pi_{t+1} - \bar{\pi}) + 0.5(E\Delta y_{t+1} - \Delta \bar{y}).$$

Orphanides (2003) showed that this rule describes quite well the behavior of U.S. policy rates during the Volcker–Greenspan period. As discussed by Orphanides (2006), one of the advantages of this simple rule is that it avoids having to rely on unobservable concepts such as the output gap and the natural real interest rate, which are subject to considerable uncertainty. Moreover, the first-difference rule has been shown to be robust in a variety of models, reflecting a wide range of data, parameter, and model uncertainties (Orphanides and Williams 2005, 2008). Finally, because the rule can be implemented on the basis of short-term forecasts for growth and inflation that were available at the time of the policy decision, it is an easy way of constructing a real-time policy benchmark that is not contaminated by ex-post information. This rule has been applied to the euro area by, among others, Smets (2010); Orphanides and Volker Wieland (2013); and Tilman Bletzinger and Wieland (2016).

Figure 22 replicates and extends the rule given by Orphanides and Wieland (2013). The dotted line depicts the changes in the relevant policy-controlled interest rate.⁶⁰ The shaded area shows the predictions of the Orphanides rule, where we use the one-year-ahead forecasts for inflation and growth from the SPF and the European Commission's real-time estimate of potential GDP growth as input variables. The upper and lower limits of the shaded area correspond to a range for the inflation aim between 1.5 and 2.0 percent. As also shown by Smets (2010) and Orphanides and Wieland (2013), this simple rule captures the changes in the ECB's policy rate very well (until it becomes zero in July 2012). If we impose the condition that the average error between the actual and predicted interest rate changes is zero (as in a regression analysis), then we can use this rule to calculate the ECB's implied inflation aim, which is 1.76 percent, very close to the midpoint between 1.5 and 2.0 percent and consistent with the range highlighted by Issing at the May 2003 press conference on the

60. Note that the changes in the policy rate are quarterly changes to align it with the quarterly frequency of the SPF forecasts, whereas the policy decisions are monthly through most of the period. We take the policy rate set in the middle of the quarter to align it with the time when the SPF forecasts are first available to the Governing Council.

Figure 22. The Orphanides Rule for the Euro Area, 1999–2018^a

Sources: Authors' research; ECB data; ECB Survey of Professional Forecasters (SPF); European Commission.

a. This estimation uses the SPF findings, as given by Orphanides and Wieland (2013). The short rate changes combine the time series of the changes in the main refinancing operations rate up to 2008:Q3 with the changes in the deposit facility rate from 2008:Q4 onward. Changes are mid-quarter-on-quarter changes. The most recent observation is for 2018:Q1.

occasion of the announcement of the results of the ECB's monetary policy strategy review (ECB 2003b).

In the rest of this subsection, we go beyond the previous papers by using the ECB's own growth and inflation projections to derive the rule. Since their start in 1998, the ECB and the Eurosystem have produced quarterly macroeconomic projections, which typically are presented to the Governing Council in the first meeting in March, June, September, and December of each year as part of the economic analysis.⁶¹ Kontogeorgos


61. The June and December projections are called the Broad Macroeconomic Projection Exercise because it is a Eurosystem exercise involving the contributions of all the national central banks of the euro area, whereas the March and September Macroeconomic Projection Exercises are intermediate updates of the December and June Broad Macroeconomic Projection Exercises produced by ECB staff. Note that the ECB/Eurosystem's projections at first were based on a constant interest rate assumption; but since the June 2006 projection exercise, they have been based on market expectations of short- and long-term interest rates. Differences in technical assumptions for the oil prices or the exchange rate may explain part of the differences between SPF and ECB/Eurosystem projections. For a comprehensive description of the exercises, see ECB (2016a).

and Lambrias (2018) recently investigated some of the properties of the ECB/Eurosystem staff projections for GDP growth and HICP inflation, and they find that they satisfy the properties of optimal forecasts. They are generally unbiased; errors are not correlated beyond what one theoretically could expect; and the uncertainty in the forecasted increases with the horizon. They outperform simple benchmarks—such as the Random Walk and an Autoregressive Model of Order 1 [AR(1)]—and, in the case of inflation, are rational.⁶² Unfortunately, a direct comparison with the SPF forecasts is difficult because the professional forecasters use different information sets and different technical assumptions. Paloviita and others (2017) focus on the properties of the ECB/Eurosystem staff forecasts over the projection horizon. They find that the ECB/Eurosystem staff projections exhibit stronger and faster mean reversion than are implied by persistence in the actual data. After about six quarters, the median inflation projections are already in the proximity of their levels at the end of the forecast horizon. They also find that inflation forecasts are too often close to the mean, and that three to four quarters out, the inflation and growth forecasts are not correlated with the actual outcomes. Some of these findings are illustrated in figure 23, which shows the mean, maximum/minimum, and 25th/75th percentiles of the ECB/Eurosystem staff projections of year-on-year inflation and real GDP growth for different horizons.⁶³

Figure 24 shows the outcome of applying the Orphanides rule to the ECB/Eurosystem staff projections. In order to align the interest rate decisions with the ECB/Eurosystem projections, we take the policy rate set when the projections are presented (that is, in the last month of the quarter). This explains the slightly different pattern of interest rate changes compared with figure 22. The conclusions remain, however, roughly the same. The simple policy rule captures the ECB's policy decisions quite well. The increase in rates in 1999 and 2000 and the subsequent fall, the pause in 2004–5, the rise starting in 2006, the sharp fall in 2008 and 2009, and the slight increase in 2011, as well as the fall in 2012, are all captured fairly well by a simple response to deviations of the one-year-ahead inflation projection from the inflation aim and the deviations of the one-year-ahead growth projection from estimated potential output growth. Not surprisingly, the correspondence is less striking as of July 2012, when the deposit rate is constrained by reaching zero (see subsection I.C), and

62. See also ECB (2013a); and Alessi and others (2014).

63. Paloviita and others (2017) show a similar figure.

Figure 23. ECB/Eurosystem Staff Projections for Year-on-Year HICP Inflation and Real GDP Growth^a

Sources: ECB data; ECB staff projections.


a. HICP = Harmonised Index of Consumer Prices. F0–F8 refer to the 0–8 quarters-ahead forecast horizons, where the 0 quarter is the forecast for the current quarter. For each forecasted horizon, we show the minimum, maximum, 25th percentile, 75th percentile, median, and unconditional mean of forecasts over the sample period. The HICP inflation sample period is 1999:Q1–2018:Q1, and the real GDP growth sample period is 2000:Q1–2018:Q1.

only relatively small further changes into negative territory were feasible. Some of it can be reestablished with the help of “shadow interest rates” (the dashed line in figure 24), which we discuss in subsection II.C when we assess the ability of ECB nonstandard monetary policy measures to provide additional stimulus at the lower bound of interest rates.

Table 1 shows the results from estimating this rule.⁶⁴ The estimated coefficients are somewhat smaller than, but not significantly different from, 0.5. The ECB’s implicit inflation aim, which can be deduced from the estimated constant, is 1.81 percent. The R^2 is higher than 0.5, which is

64. Other studies that have estimated policy reaction functions for the ECB include Gerdesmeier and Roffia (2003); Gorter, Jacobs, and de Haan (2008); Gerlach and Lewis (2014); and Paloviita and others (2017). Paloviita and others (2017) find support for monetary policy reaction functions with very-short-run (one quarter ahead) GDP growth projections; somewhat longer (one-year-ahead) inflation projections; and a proxy for the natural rate of interest.

Figure 24. The Orphanides Rule for the Euro Area, with Forecasts Based on ECB/Eurosystem Staff Projections, 1999–2018^a



Sources: Authors' research; ECB data; ECB staff projections; European Commission. The shadow rates come from Krippner (2015), Kortela (2016), Lemke and Vladu (2017), and Wu and Xia (2017).

a. The short rate changes (the dotted line) combine the time series of the changes in the main refinancing operations rate up to 2008:Q3, with the time series of the changes in the deposit facility rate from 2008:Q4 onward. The shadow rate changes (the dashed line) are based on a shadow short-term interest rate, for which the zero lower bound is not binding and that therefore also captures the impact of unconventional monetary policy tools. It is calculated as the first principal component of the five shadow rates in figure 27 from 2012:Q3 (the time when the ECB's deposit facility rate reached zero) onward. Changes are end of quarter-on-quarter changes. The most recent observation is for 2018:Q1.

quite high, given that the variable we are trying to explain is expressed in first differences. Shortening the sample until the second quarter of 2012, when the ECB reached the zero lower bound on its deposit rate (see the second column of table 1), does not significantly change these results. In these regressions, we chose the horizon for year-on-year GDP growth to be $t + 3$ quarters, reflecting the fact that at the time of the interest rate decisions in the last month of the quarter, the current quarter is not yet known, while the previous quarter is known, whereas for inflation we have $t + 11$ months, reflecting the fact that inflation in the previous month is known. We tested for different forecasted horizons and found that for both GDP growth and inflation, the one-year-ahead projections are the most informative for policy decisions (the highest R^2).


Figure 25 shows the cumulated errors of both the calibrated and estimated rules. Using this set of benchmarks suggests that interest rate policy

Table 1. Regression Results for the Orphanides Rule^a

<i>Y = short rate changes</i>	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ECB inflation projections	0.34*** (0.09)	0.30*** (0.13)	0.20 (0.19)	0.37*** (0.14)	0.33** (0.14)	0.17 (0.13)	0.36*** (0.09)
ECB growth projections	0.37*** (0.08)	0.43*** (0.09)	0.54*** (0.20)	0.36*** (0.08)	0.37*** (0.08)	0.52*** (0.14)	0.40*** (0.09)
SPF inflation projections			0.27 (0.22)				
SPF growth projections			-0.17 (0.19)				
ECB core inflation projections				-0.073 (0.16)			
Interaction positive inflation deviation					0.27 (0.17)	0.50*** (0.18)	
Interaction positive growth deviation						-0.46*** (0.17)	
Change in credit growth							-0.07 (0.07)
Positive inflation deviation					-0.10 (0.08)	-0.15** (0.07)	
Positive growth deviation						0.02 (0.02)	
Constant	-0.62*** (0.164)	-0.52** (0.26)	-0.85*** (0.14)	-0.56*** (0.16)	-0.60** (0.22)	-0.22 (0.20)	-0.65*** (0.16)
Inflation target	1.81	1.75	1.78	1.85			1.82
Observations	77	54	77	77	77	77	77
Adjusted <i>R</i> ²	0.52	0.58	0.54	0.52	0.53	0.57	0.53

Sources: Authors' research; ECB and European Commission data.

a. Robust standard errors are in parentheses. Statistical significance is shown for *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$. Coefficients are rounded to the second decimal. The dependent variable, short rate changes, combines the time series of the changes in the Main Refinancing Operations Rate up to 2008:Q3 with the changes in the Deposit Facility Rate from 2008:Q4 onward. Changes are quarter-on-quarter changes of the implementation dates of monetary policy decisions falling in the last month of the quarter. The sample period is 1999:Q1 to 2018:Q1. For column 2, the sample is shortened to end in 2012:Q2. The variable Inflation target is equal to minus the constant over the estimated inflation coefficient. The variables ECB inflation projections, ECB core inflation projections, and ECB growth projections refer to 1-year-ahead inflation projections for the Harmonised Index of Consumer Prices (HICP), the 1-year-ahead HICP inflation excluding food and energy projections, and the 1-year-ahead GDP growth projections in deviation from potential growth of the ECB/Eurosystem staff projections, respectively. The variables SPF inflation projections and SPF growth projections refer to 1-year-ahead HICP inflation projections and 1-year-ahead GDP growth projections in deviation from potential growth of the ECB Survey of Professional Forecasters, respectively. Positive inflation deviation and positive growth deviation refer to dummies equal to 1 if projected inflation is greater than the predicted inflation target of 1.81, and if projected GDP growth is greater than potential output, respectively. Interaction positive inflation deviation and interaction positive growth deviation are terms interacting ECB inflation projections and ECB growth projections with the dummy variables positive inflation deviation and positive growth deviations, respectively. Change in credit growth is the change in the year-on-year growth rate of credit to households and nonfinancial corporations.

Figure 25. Cumulative Errors from the Orphanides Rule for the Euro Area, 1999–2018^a

Sources: Authors' research; ECB data; ECB staff projections; ECB Survey of Professional Forecasters; European Commission data.

a. The cumulated error is calculated as the cumulated difference between the change in the short rate and the Orphanides rule, using the corresponding inflation targets—e.g., 1.73 percent, 1.76 percent, and 1.81 percent (for the predicted rule). The short rate combines the time series of the changes in the main refinancing operations rate up to 2008:Q3, with the time series of the changes in the deposit facility rate from 2008:Q4 onward. The most recent observation is for 2018:Q1.

may have been somewhat too loose in 2002 and too tight in 2009 and 2013.⁶⁵ This finding is consistent with more elaborate thick-modeling exercises by ECB staff, which identify both 2009 and 2012–13 as periods in which the actual interest rate is above what a range of Taylor-type rules estimated before 2008 would have suggested. Of course, the latter periods are also when the ECB implemented a range of unconventional measures, as we discuss below. The interest rate increase in July 2008 does not appear to be justified by the ECB's own outlook for growth and inflation,

65. The finding of too-tight policy in 2009 is somewhat at odds with the findings of Giannone and others (2012) and Pill and Smets (2013), who show that by the end of 2009 and until 2012, the actual path of 3-month Euribor was below the counterfactual one based on the historical ECB monetary policy rule. Pill and Reichlin (2015) argue that the euro area experience contrasts with evidence from the United States, where the zero lower bound appears to have been a binding constraint on rate setting throughout the crisis period.

but was quickly reversed. The interest rate increases in 2011 do not show up as a major policy mistake, but seem delayed as the inflation and growth projections suggested an earlier tightening move. Of course, this does not exclude the possibility that the ECB underestimated the impact of the financial and sovereign debt crises on economic activity and inflation; but similar results using SPF forecasts suggest that the ECB was not the only institution to do so. Finally, these benchmarks do not suggest that monetary policy was too loose during the time before the crisis, as suggested by Taylor (2007) for the United States.

In table 1, we also test a number of alternative specifications. First, the third column shows that the ECB/Eurosystem staff projections outperform the SPF forecasts in explaining the ECB's interest rate decisions. This is not surprising, given that the SPF forecasts are collected one to two months earlier than the ECB/Eurosystem staff projections and therefore do not incorporate the latest data available at the time of the interest rate decisions. Second, we test whether the projections for HICP inflation excluding food and energy add value in explaining the interest rate decisions (the fourth column). The estimated coefficient on the projection for core inflation is negative, but insignificant. This is consistent with empirical findings for the euro area that headline inflation leads core inflation and not the other way around and with the descriptive analysis in section I, which points to the fact that on a number of occasions the ECB was worried about second-round effects of changes in headline inflation driven by rising oil prices on wages and underlying inflation. This was, for example, the case for the interest rate increases in 2008 and 2011.

Next, we test whether the ECB responded more aggressively to positive deviations of projected inflation from its inflation aim than to negative deviations. The fifth column of table 1 shows that the relevant coefficient is not significantly different from zero. However, when we interact both inflation and output terms with a dummy when positive, we get the interesting finding that the coefficient is large, positive, and significant when inflation is above target, but otherwise is insignificant. However, we get the opposite finding for growth: It is large and significant when growth is below potential and insignificant when growth is above potential. Thus, over the sample period, the ECB seems to ease policy mainly in response to expected growth slowdowns and tighten policy mainly in response to expected inflation above its inflation aim.

Finally, we also tested whether indicators coming from the ECB's monetary analysis have additional explanatory information value for its interest rate decisions. Fischer and others (2008) and Smets (2010) do not

find additional explanatory power coming from monetary analysis. This is consistent with the idea of monetary analysis being a cross-check. It is also consistent with the argument by Orphanides (2006) that the simple policy rule can also be derived from the combination of the quantity theory of money and a money demand function, and therefore already embeds an implicit role for money. The last column of table 1 includes changes in annual credit growth as an additional explanatory variable in the interest rate rule. The related coefficient is not significant and, if anything, is negative. Similar results are obtained with M3 growth or other money and credit growth indicators. Of course, this does not exclude the usefulness of monetary analysis as a cross-checking device (Beck and Wieland 2008, 2010).

II.C. Reviewing the ECB's Nonstandard Monetary Policy Measures

Understanding the working of and assessing nonstandard monetary policy are the subjects of an evolving literature. In this subsection, we first offer a conceptual framework for how to think about nonstandard policies from an ECB perspective. Subsequently, we review the literature about the success with which the ECB has used these policies to repair the monetary transmission mechanism (the complement of standard interest rate policy) and about the effectiveness with which the ECB has provided additional monetary stimuli with these measures (substitute for standard rate policy).

CLASSIFYING THE ECB'S NONSTANDARD POLICY MEASURES Figure 26 gives an overview of the nonstandard monetary policy measures the ECB has taken since 2007, reflecting the different crisis phases in the columns.⁶⁶ These measures can be divided into four categories, as shown in the four rows of the figure: (1) credit operations with the ECB's counterparties, that is, euro area monetary and financial institutions; (2) outright asset purchases of both private and public sector securities; (3) negative interest rates; and (4) forward guidance, that is, enhanced communication about future policy actions.⁶⁷

66. Most of these measures were using or amending the Eurosystem's operational framework. For detailed and comprehensive descriptions of this framework and the ECB's monetary policy instruments since the start of the financial crisis, see Eser and others (2012); Alvarez and others (2017); Task Force on the Use of Monetary Policy Instruments (2018); and Bindseil and others (2017).

67. The extent to which these measures can be classified as nonstandard is of course debatable. For example, in the early period of the financial crisis, the ECB primarily adjusted the conditions and features of its credit operations, which are standard instruments of the ECB's monetary policy operational framework. Similarly, negative interest rates and forward guidance can be seen as variants of the standard setting of policy-controlled interest rates and their communication.

Figure 26. Timeline of the ECB's Monetary Policy Measures since the Start of the Financial Crisis, August 2007–June 2018^a

	Financial crisis		Sovereign debt crisis		Low-inflation recovery (with a lower-bound constraint)	
Interest rate policy	+25 basis points MROR: 4.25%	+400 basis points DFR: 0.25%	+50 basis points DFR: 0.75%	-75 basis points DFR: 0%	-20 basis points DFR: -0.20%	-20 basis points DFR: -0.40%
Credit operations	Overnight FTOs "Front-loading" Maturity extension Dec. 2007 Dollar swaps	Oct. 2008 FRFA Expanded collateral LTROs (6 months) Dollar swaps May 2009 LTROs (1 year)	Oct. 2011 LTROs Dec. 2011 VLTRO I (3 years) Feb. 2012 VLTRO II (3 years)	Oct. 2011 LTROs Dec. 2011 VLTRO I (3 years) Feb. 2012 VLTRO II (3 years)	June 2014 TLTRO I	March 2016 TLTRO II
Asset purchases		May 2009 CBPP I	May 2010 SMP I	Aug. 2011 SMP II Oct. 2011 CBPP II Sept. 2012 OMT	June 2014 ABSPP CBPP III Jan. 2015 PSPP	Dec. 2015 APP I March 2016 CSPP APP II (80 billion)
Forward guidance					July 2013 FG I: Policy rate extended period Jan. 2015 FG II: APP date and SAPI	Dec. 2016 APP III (60 billion) Oct. 2017 APP IV (30 billion) June 2018 APP V (15 billion)

Aug. 2007 Sept. 2008 May 2010 Aug. 2011 June 2013 Aug. 2015 Dec. 2016 June 2018

☐ Standard interest rate policies
☐ Negative Deposit Facility Rate
☒ Nonstandard policies to address lower bound of rates
☐ Impaired interbank and bank funding markets, and later also bank-lending channel
☐ Sovereign bank nexus and redenomination risk
☒ Heterogeneous pass-through in bank-lending markets

Sources: Authors' research; ECB data.

a. MROR = Main Refinancing Operations Rate; DFR = Deposit Facility Rate; FTOs = fine-tuning operations; FRFA = fixed-rate/full-allotment credit operations; LTROs = Longer-Term Refinancing Operations; VLTROs = Very-Long-Term Refinancing Operations; TLTROs = Targeted Longer-Term Refinancing Operations; CBPP = Covered Bond Purchase Programme; SMPs = Securities Markets Programmes; OMTs = Outright Monetary Transactions; ABSPP = Asset-Backed Securities Purchase Programme; PSPP = Public Sector Purchase Programme; APPs = Asset Purchase Programmes; FG = forward guidance; SAPI = Sustained Adjustment in the Path of Inflation.

Broadly speaking, the use of the nonstandard measures served two purposes. First, some of the measures complemented standard reductions in policy-controlled interest rates in the presence of impairments in monetary policy transmission. In a financial crisis, it may be optimal to address the rise in funding and financing costs arising from malfunctioning financial markets through direct market interventions such as asset purchases or through lending operations, rather than try to offset them through a reduction in policy-controlled interest rates. Second, other measures were substitutes for standard policy: They provided additional stimulus in the presence of limited room for further standard interest rate easing close to the zero lower bound.⁶⁸ The two purposes are marked with graphical patterns in the different cells of figure 26. Measures complementing standard policy are indicated with stripes, where the different shades of gray (and the thickness of stripes) indicate the different types of impairments in the monetary transmission mechanism addressed. Standard interest policy (the light gray cells) and its various substitutes for providing an additional stimulus (the negative policy rates in medium gray, and asset purchases and forward guidance in black) are indicated with gray/black shades without stripes.

A number of observations are worth making regarding these two purposes. First, in the early stages of the financial crisis, when short-term interest rates were not yet constrained by the zero lower bound, the ECB in its communication made a clear distinction—through the so-called separation principle—between standard policy, which was geared at maintaining price stability, and nonstandard measures that were focused on addressing malfunctioning financial markets and impairments in policy transmission. In practice, the two policies of course interact and together determine the monetary policy stance; but arguably, highlighting this distinction allowed the ECB to more easily take different directions in its standard and nonstandard monetary policy. This was, for example, the case in 2008 and 2011, when the ECB tightened standard monetary policy while nonstandard measures were still in place. One signal of the separation principle during the sovereign debt crisis was the decision to sterilize the SMP and potential OMT interventions.⁶⁹

Second, the nature of the nonstandard measures depends on which impairments are being addressed. As discussed in section I and shown in figure 26, three stages can be distinguished. In the early stages of

68. For a discussion of the motivations, effectiveness, and risks of the ECB's nonstandard measures, also see Neri and Siviero (2019).

69. The ECB conducted regular one-week FTOs between May 2010 and June 2014 to absorb the liquidity effect of the SMP initiated on May 10, 2010.

the financial crisis, the focus was primarily on banks' funding markets, in particular the money market and the covered bond market, but later also on bank lending (the striped cells in light gray). In the second stage, the financial crisis turned into a sovereign debt crisis with repercussions for bank funding markets (the so-called sovereign–bank nexus) and the emergence of self-fulfilling redenomination risk (the striped cells in medium gray). The last stage focused on the heterogeneous transmission in bank lending markets and involved funding for lending operations (TLTROs; the striped cells in dark gray). An evaluation of these different nonstandard measures therefore involves an assessment of whether the specific impairments were addressed.

Third, the nonstandard measures geared at addressing impairments in the monetary transmission process are akin to classical lender-of-last-resort policies, whereby the central bank steps in to provide liquidity and avoid having market runs and self-fulfilling speculative attacks turn into solvency issues. One issue with these policies is that it is often not easy to distinguish between liquidity and solvency problems. In lending operations to multinational financial institutions, this is solved by requiring collateral, which are often government bonds. However, in a monetary union with national fiscal policies, sovereign risks may undermine the safety of such collateral and may make direct interventions in sovereign bond markets more problematic. This explains why nonstandard measures to address illiquidity and self-fulfilling redenomination risks in sovereign bond markets (SMP and OMT) required conditionality to ensure the soundness and sustainability of the underlying fiscal policies.

Finally, from figure 26, it is also clear that over time, as the euro area economy fell in a double-dip recession, more of the measures—in particular, the negative DFR, the large-scale APP, and enhanced forward guidance—served the second purpose of easing policy close to the zero lower bound. In line with this distinction, we next review the evidence on the effectiveness of nonstandard measures.

ADDRESSING IMPAIRMENTS IN THE MONETARY POLICY TRANSMISSION PROCESS
During the early stages after the start of the financial crisis, nonstandard measures mostly focused on bank funding markets. Due to the fixed-rate/full-allotment procedure, liquidity provision was primarily demand-determined during that period. The Enhanced Credit Support (the right column with the striped cells in light gray in figure 26) program helped ease tensions in the money market, as indicated by the reduction in the Euribor–OIS spreads at various maturities (figure 13). Lucrezia Reichlin (2014, 388) and Huw Pill and Reichlin (2015) describe this period as the

ECB taking a “market operation approach” to its role as lender of last resort (see also Garcia-de-Andoain and others 2016), and conclude that it contributed to the recovery of economic activity, which started in 2009:Q3. Michele Lenza, Pill, and Reichlin (2010); Domenico Giannone and others (2012); Gert Peersman (2011); and Seth Carpenter, Selva Demiralp, and Jens Eisenschmidt (2013) use a variety of counterfactual exercises to conclude that in this period, the effectiveness of the ECB’s actions was not constrained by the zero lower bound and that these measures were supportive of economic activity, largely by preventing a more discontinuous and dramatic curtailment of credit provision to the real economy. See also Jef Boeckx, Maarten Dossche, and Peersman (2017). A model-based analysis is done by Christophe Cahn, Julien Matheron, and Jean-Guillaume Sahuc (2017). Using an estimated dynamic stochastic general equilibrium model with a frictional banking sector, they find that liquidity injections have played a key role in averting a major credit crunch. A counterfactual analysis suggests that during 2009, absent these nonstandard measures, output, consumption, investment, and the GDP deflator on average would, respectively, have been 2.5 percent, 0.5 percent, 9.7 percent, and 0.5 percent lower. For a similar analysis, also see the work of Dominic Quint and Oreste Tristani (2018).

Part of the Enhanced Credit Support policy was the first Covered Bond Purchases Programme (CBPP1). Purchases of €60 billion were made from July 2009 through June 2010, distributed across the euro area in both primary and secondary markets. John Beirne and others (2011) discuss the modalities and the impact of the CBPP1 and find that it has contributed to (1) a decline in money market term rates, (2) an easing of funding conditions for credit institutions and enterprises, (3) encouraging credit institutions to maintain and expand their lending to clients, and (4) improving market liquidity. Second and third installments of the CBPP were decided on, respectively, in October 2011, in the context of the intensification of the sovereign debt crisis, which again affected the bank’s funding conditions; and in September 2014, as part of the comprehensive easing package to fight risks of deflation starting in June 2014.

The SMP was introduced to address malfunctioning sovereign bond markets after the start of the sovereign debt crisis, in particular in Greece, Portugal, and Ireland, which suffered from illiquidity and which were deemed to threaten monetary policy transmission. Interventions faded out in the relatively stable first half of 2011; but as the sovereign debt crisis negatively affected Italy and Spain in July 2011, a reactivation of the SMP was announced on August 7, 2011. The SMP ran until the end

of December 2012 and reached an outstanding nominal amount of about €218 billion, although the volumes were not announced *ex ante*.

Various authors have assessed the impact of the SMP on sovereign bond yields. The SMP interventions succeeded in reducing yields and volatility of government bond segments of the countries under the program. Using a counterfactual exercise, Eric Ghysels and others (2017) find that purchases of Italian and Spanish bonds lowered two-year yields by 320 and 180 basis points, respectively, and 10-year yields by 230 basis points for both countries. Similarly, Fabian Eser and Bernd Schwaab (2016) find a significant impact of the SMP on the yields of those securities that were purchased. Their baseline model suggests that, on average, a daily SMP intervention of €100 million lowered yields by 0.1 to 2.0 basis points. This impact is stronger in markets that are smaller and less liquid, and where risk premiums are higher. (Also see Trebesch and Zettelmeyer 2018; and De Pooter, Martin, and Pruitt 2018.)

Nevertheless, the SMP was not able to stem the rising redenomination risk. Pill and Reichlin (2015) point to three reasons why the SMP did not succeed in stemming the rise in sovereign spreads. First, the SMP actions were characterized as limited and temporary, which undermined market confidence that the ECB was prepared to offer a full backstop. Second, the ECB had conditioned its provision to Italy and Spain on certain policy commitments that threatened the political feasibility of the support. Third, there were concerns about the subordination of private sector bond holders.

As discussed in subsection I.C, bolder ECB action became possible after European governments had started to strengthen fiscal governance, provided a backstop for governments in the form of the ESM, and decided to create a banking union with common supervision and resolution. After the famous “whatever it takes” speech of President Draghi in July 2012, the ECB announced its readiness to undertake *ex-ante* unlimited OMTs in euro area secondary sovereign bond markets, subject to countries complying with conditionality.⁷⁰ Although, so far, OMTs have not been activated, the announcement was instrumental in addressing excessive risk premiums and improving financial market confidence, as shown in figure 13 above. The success of the OMT was dependent on a number of features: a strict and effective conditionality attached to an appropriate EFSF/ESM program, a focus on the shorter segment of the yield curve, no *ex-ante* quantitative limits on size, and *pari-passu* treatment. The conditionality was key

70. The technical features of the OMTs are given in ECB (2012b).

for preserving the appropriate incentives for fiscal discipline and monetary dominance as well as to ensure proper risk management by the central bank. Using high-frequency data, Carlo Altavilla, Domenico Giannone, and Michele Lenza (2016) find that OMT announcements decreased the Italian and Spanish two-year government bond yields by about 2 percentage points, while leaving the bond yields in Germany and France unchanged. Using a multicountry vector autoregression model, they also find that the reduction in bond yields due to the OMT was associated with a significant increase in real activity, credit, and prices in Italy and Spain, with some positive spillovers in France and Germany. (For additional evidence on the financial market effects, see the papers by Krishnamurthy, Nagel, and Vissing-Jorgensen 2018; Szczerbowicz 2015; and De Santis 2016, 2018, forthcoming.) Philippe Aghion, Emmanuel Farhi, and Enisse Kharroubi (2017) find that growth effects worked particularly through highly indebted corporate sectors, notably via more easily adjustable short-term debt, but only if they were located in countries with relatively less regulated product markets. This bolsters the view that demand policies are more effective when accompanied by adequate supply policies. Using evidence from the ECB's Survey on the Access to Finance of Enterprises, Annalisa Ferrando, Alexander Popov, and Gregory Udell (2015) find that the ECB's OMT announcement was followed by an immediate decline in the share of credit-rated firms and of firms discouraged from applying for loans. Firms with an improved outlook and credit history were particularly likely to benefit from easier credit access. Viral Acharya and others (2017, 2) also find positive effects of the revaluation of sovereign bond portfolios due to OMT on bank lending. They argue though that a significant fraction of this lending went to "zombie firms."

As part of the attempt to stop the doom loop, the ECB (2011c) also conducted two three-year VLTROs in December 2011 and February 2012. A combined gross amount of more than €1 trillion was allotted (see the vertically dashed area of figure 16), giving banks funding certainty, easing redemption of maturing bonds, and helping to sustain credit lines with households and firms. Matthieu Darracq-Paries and Roberto De Santis (2015) show that VLTROs increased real output and lending to NFCs over a two- to three-year horizon.⁷¹ Martina Jasova, Caterina Mendicino, and Dominik Supera (2018) use microeconomic bank-firm level data for Portugal to show that the lengthening of bank debt maturity with the ECB

71. For evidence on Spain, see Garcia-Posada and Marchetti (2016); and Szczerbowicz (2015).

(reduction of rollover risk) had a positive and economically sizable impact on bank lending. Banks with a 1-standard-deviation-greater ability to draw on the VLTROs (for example, due to more available collateral) increased both existing and new lending by 5.3 percent. The effects are stronger on the supply of credit to smaller, younger, and riskier firms. However, they also show that unrestricted liquidity provision incentivized banks to purchase more government securities, partly offsetting the positive effects on lending. Matteo Crosignani, Miguel Faria-e-Castro, and Luis Fonseca (2018) find that VLTROs induced Portuguese banks to purchase short-term domestic government bonds and pledge them to obtain central bank liquidity.⁷²

Turning to the funding for lending policies (see the striped cells in dark gray in figure 26), it is difficult to disentangle the effects of TLTROs from the other measures that were part of the comprehensive easing package that started in June 2014 and that also included negative rates and asset purchases.⁷³ The ECB (2015b) shows that the rates on loans to NFCs declined markedly immediately after the announcement of the first series of TLTROs. The declines were sharper in countries where the composite lending rates to NFCs had been more elevated. Moreover, in vulnerable countries, banks that borrowed under TLTROs reduced their rates by more than banks that abstained from bidding. Altavilla, Fabio Canova, and Matteo Ciccarelli (2016) explicitly analyze developments over time in the pass-through of monetary policy measures on bank lending rates and find that, after 2014, nonstandard policy measures (including the TLTROs) significantly normalized the capacity of banks to grant loans and reduced the cross-sectional dispersion of interest rate pass-throughs.

PROVIDING ADDITIONAL STIMULUS AT THE EFFECTIVE LOWER BOUND As policy-controlled interest rates were increasingly constrained by the effective lower bound in 2013, the ECB took a number of additional nonstandard measures, such as the expanded APP and forward guidance, with the aim of further lowering medium- to long-term interest rates through portfolio rebalancing and signaling channels (see the black cells in figure 26). One way of capturing the impact of these unconventional measures is to calculate a shadow short-term interest rate, as proposed by Leo Krippner (2015). A shadow rate is the shortest maturity rate extracted from a term structure model that would generate the observed yield curve in the absence of a lower bound. It coincides with the policy rate in normal times, and is free to go into negative territory when the policy rate is stuck at the lower bound.

72. See also Acharya and Steffen (2015); and Van Bakkum, Gabarro, and Irani (2018).

73. The ECB (2017b) explains the features of the two TLTRO programs, as well as their impact on bank lending; also see ECB (2017c).


Various researchers have shown that the shadow rate captures the stance of monetary policy during lower-bound periods in the same way the policy rate does in normal times (Claus, Claus, and Krippner 2016; Francis, Jackson, and Owyang 2017; Van Zandweghe 2015). They show (1) that the shadow rate captures the impact of both conventional and nonconventional policy measures, such as asset purchase programs, forward guidance on interest rates, and long-term refinancing operations; and (2) that the dynamic interactions between macroeconomic activity and the short-term rate are preserved through the shadow rate. The latter is consistent with the results of David Debortoli, Jordi Galí, and Luca Gambetti (2018), who find that there has been no structural break in the macroeconomic relations since the use of nonstandard measures. These researchers conclude that nonconventional tools must have had a similar impact on the macroeconomy as conventional interest rate policy. Similarly, a number of vector autoregression exercises, where unconventional monetary policies are identified through the term structure changes during a narrow window around monetary policy decisions, have shown that quantitative easing has very similar effects on the economy.⁷⁴ Finally, Jing Wu and Ji Zhang (2017) show that in a New Keynesian model for the United States, the negative shadow rates are a useful summary statistic to capture the impact of unconventional policies, especially quantitative easing and lending facilities.

At the same time, estimates of shadow rates are quite sensitive to differences in term structure models, and in particular to the assumptions made about where the effective lower bound on interest rates lies. This may particularly be an issue for the euro area, where the perceived effective lower bound has changed over time as interest rates have gone into negative territory. Figure 27 plots several shadow rate estimates for the euro area, together with the EONIA. It shows, generally speaking, that the shadow rates are close to the EONIA before 2012 and that nonstandard measures have had an easing impact on the yield curve since 2012. Although there is considerable co-movement, the levels of the shadow rates are however very diverse.

We therefore follow Mouabbi and Sahuc (2017) and use a common factor of five alternative shadow rate models for the euro area as a summary statistic for the stance of monetary policy in the euro area after the second quarter of 2012. The results of this exercise are also shown in figure 24, which compares changes in the shadow rate (the dashed line) with the outcome of the Orphanides rule (the dotted line) after the DFR reached

74. See Bundick and Smith (2016); Swanson (2017); and Inoue and Rossi (2018).

Figure 27. Estimated Shadow Rates for the Euro Area and the ECB's Effective Policy Rate, 1999–2018^a



Sources: ECB data; Kortela (2016); Krippner (2015); Lemke and Vladu (2017); Wu and Xia (2017).

a. MROR = Main Refinancing Operations Rate; DFR = Deposit Facility Rate. The shadow rate refers to a shadow short-term interest rate, for which the zero lower bound is not binding and that therefore also captures the impact of nonstandard monetary policy tools (see references in the sources). Lemke and Vladu LB-Adaptive and LB-Monotonic stand for two shadow rate versions based on different specifications of the lower bound. The version of LB-Adaptive sets the lower bound according to the minimum of forward rates observed at that point in time. LB-Monotonic also follows the same minimum rule, but the lower bound is never allowed to go up again. MROR and DFR combined refer to the effective policy rate, which is the main refinancing operations rate up to 2008:Q3 and the deposit facility rate from 2008:Q4 onward. The most recent observation is for March 2018.

0 percent in July 2012. As expected, the common component of estimated shadow rates tracks the range predicted by the policy rule for most times of this period better than the DFR. Broadly speaking, changes in the shadow rate capture the two periods that correspond to a slowdown in expected growth and inflation and the resulting intensification of nonstandard measures taken by the ECB, as discussed in subsection I.D and also reflected in figure 26 (three-pronged easing as of mid-2014 and its recalibration at the end of 2015 and in early 2016). In 2017, increases in the shadow rate reflect a relative tightening of monetary policy in line with the prescription of the rule (gray range). However, movements in the shadow rate in late 2012 and early 2013 do not capture the need for additional easing at that time. One issue here is that the powerful OMT announcement is not picked up well


by the dashed line, as it happened quite close to the start of the calculation of the common component of the shadow rates. Further research seems to be needed in this area.

Mouabbi and Sahuc (2017) use the shadow rates to capture monetary policy after 2013 in an estimated dynamic stochastic general equilibrium model for the euro area, and they find that without the implemented non-standard measures, year-on-year inflation and GDP growth would have been lower by 0.66 percent and 0.99 percent, respectively, over the period 2014:Q1–2017:Q2.

Overall, these estimates are in the same ballpark as estimates by the ECB that are based on a variety of methods (Draghi 2017; Praet 2017; Hutchinson and Smets 2017).⁷⁵ ECB staff estimates indicate that the monetary policy contribution of the easing package since 2014 to euro area GDP has been about 1.8 percentage points, cumulatively over the period 2016–19 (see, for example, Hammermann and others 2019). About one-third of the 5-percentage-point increase in the employment rate observed in the euro area as a whole since mid-2014 is estimated to be due to the ECB's measures. This roughly corresponds to 2 to 3 million more jobs. Absent the ECB's policy package, inflation would on average be about 45 basis points lower than what is realized or currently projected for each year over the 2016–19 period.

The main transmission channel is through the easing of financial conditions and financing costs. Counterfactual simulations by ECB staff estimate that the 2014 policy package has had a considerable impact on euro area financing conditions. Figure 28 shows some of the results. For example, without the ECB's measures, the 10-year sovereign yield for a euro area GDP-weighted aggregate would be about 150 basis points higher and lending rates to euro area NFCs would be about 70 basis points higher. The ECB's measures have also had a sizable impact on the nominal euro effective exchange rate, which would have been about 13 percent higher without the measures (Altavilla, Carboni, and Motto 2015; Ambler and Rumler

75. The approaches can be categorized into two groups: a "direct" and an indirect, or "two-step," approach. In the direct approach, models tend to be fully specified structural models, such as dynamic stochastic general equilibrium models, which incorporate mechanisms to directly allow for asset purchases to affect economic activity and inflation. Typically, these models extend the workhorse New Keynesian model by including financial frictions so that central bank asset purchases have an impact on the economy. In the two-step approach, the first step involves estimating, off model, the impact of asset purchases on long-term yields and other financial prices. In the second step, this is fed into a macro model, which then estimates the impact on activity and inflation.

Figure 28. Changes in Key Euro Area Financial Indicators since June 2014, and the Impact of the ECB's Policy Measures^a

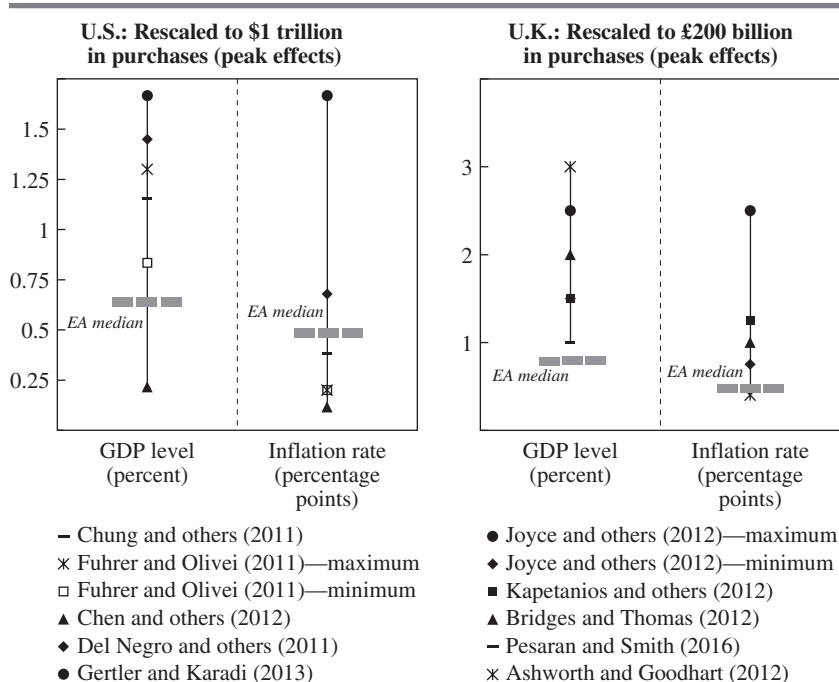
Sources: Bloomberg; ECB data; ECB staff calculations.

a. OIS = Euro Overnight Index Swap Rate; NFC = nonfinancial corporations; NEER = Nominal Effective Exchange Rate; APP = Asset Purchase Programme; DFR = Deposit Facility Rate. The impact of credit easing is estimated on the basis of an event-study methodology that focuses on the announcement effects of the June–September 2014 package; see ECB (2015b). The effects of the DFR cuts rest on the announcement effects of the September 2014 DFR cut. The APP encompasses the effects of measures taken in January 2015, December 2015, March 2016, December 2016, and October 2017. The January 2015 APP impact is estimated on the basis of two event-study exercises by considering a broad set of events that, starting in September 2014, have affected market expectations about the program; see Altavilla, Carboni, and Motto (2015) and De Santis (2016). The quantification of the impact of the December 2015 policy package on asset prices rests on a broad-based assessment comprising event studies and model-based counterfactual exercises. The impact of the March 2016 measures, the impact of the December 2016 measures, and the impact of the October 2017 measures are assessed via model-based counterfactual exercises. Changes in lending rates are based on monthly data, the reference period for which is June 2014 to April 2017. The most recent observation is for November 27, 2017.

2017; De Santis 2016). There is also emerging evidence of the portfolio rebalancing effects of the APP (Paludkiewicz 2018).

To put all this in perspective, in figure 29 we compare estimates of GDP and inflation effects of central bank asset purchases from a selection of studies for the U.S. and U.K. with those for the euro area in a standardized format. Median ECB staff euro area estimates (the horizontal dashed lines) are based on a suite of models (encompassing both direct and two-step

Figure 29. Comparison of the Effectiveness of Asset Purchases in the Euro Area, the United States, and the United Kingdom^a



Sources: ECB calculations; Ashworth and Goodhart (2012); Bridges and Thomas (2012); Chen, Vasco-Cúrida, and Ferrero (2012); Chung and others (2011); Del Negro and others (2016); Fuhrer and Olivei (2011); Gertler and Karadi (2013); Joyce and others (2012); Kapetanios and others (2012); Pesaran and Smith (2016).

a. For the United States, the macroeconomic impact is scaled to \$1 trillion in asset purchases to allow for comparison across studies. Some of the studies provide the impact only for real GDP. The euro area median for GDP refers to the median of the cumulated impact over 2015, 2016, and 2017 of a range of models: vector autoregression; National Institute Global Econometric; Christiano, Motto, and Rostagno (2014); Darracq-Paries, Kok Sorensen, and Rodriguez-Palenzuela (2011); and dynamic stochastic general equilibrium. The euro area median for the inflation rate refers to the median of the peak impact of 2015–17 for the same models.

approaches). Though the euro area GDP estimates are in the lower mid-range of the U.S. estimates, they are below the estimates for U.K. GDP and are in the lower part of the range for inflation. Also see the papers by Philippe Andrade and others (2016) and the ECB (2017b).

Finally, as part of the comprehensive easing program, the ECB also lowered the DFR into negative territory, a move that before the ECB only central banks of smaller jurisdictions had dared. Massimo Rostagno and others (2016) and Hartmann (2018) show that this shifted the yield curve down

and twisted it (long rates coming down more), as one would expect from term-structure models allowing for negative rates (Lemke and Vladu 2017). In other words, the interest rate channel of monetary policy was extended. Moreover, contrary to the concerns of some skeptics, in the euro area case, it did not seem to hinder the bank-lending channel—quite the contrary. Florian Heider, Farzad Saidi, and Glenn Schepens (2018) find enhanced lending of banks with small retail depositor bases relative to banks with large retail deposits (which would suffer more from not being able to pass negative rates on liabilities on to households). Demiralp, Eisenschmidt, and Thomas Vlassopoulos (2019) find evidence that this also amounted to an aggregate lending effect. Eisenschmidt and Smets (2018) review the euro area's monetary policy experience with negative rates and the related literature further. They document the pass-through of negative policy rates on bank deposit and lending rates as well as on loan volumes in the euro area. They confirm that the zero-lower-bound constraint is binding for interest rates on household deposits held at banks. Nevertheless, the pass-through on loan rates is broadly unchanged in their analysis, even for banks with a high reliance on household deposit funding. The negative effect on the interest rate margin and profitability is generally offset by the positive impact of lower market rates on asset values and loan loss provisions (Altavilla and others 2019). Or, in other words, the “reversal rate” below which bank lending could be hurt does not seem to have been reached so far (Brunnermeier and Kobe 2018). At the same time, it needs to be acknowledged that the effects of negative policy rates cannot be perfectly disentangled from other nonstandard monetary policy measures active at the time. For example, in April 2014 and March 2016, TLTROs helped to reduce funding rates into negative territory for banks that exceeded certain lending targets (Rostagno and others 2016).

Overall, the research evidence on the effectiveness of the ECB's non-standard measures in easing financial conditions, stimulating the economy, and bringing inflation back to the ECB's inflation aim is quite encouraging for the time period covered in this paper. It suggests that concerns that central banks may be powerless when interest rates hit the zero lower bound may be excessive (Swanson 2017).

III. Conclusion

In this paper, we have reviewed the ECB's monetary policy during its first 20 years of existence. Overall, the ECB has delivered on its price stability mandate, despite the very challenging crisis times of the last decade.

Average inflation over this period has been 1.7 percent, which is in line with the ECB's aim of maintaining inflation below, but close to, 2 percent over the medium term. However, this average number masks quite stable inflation of about 2 percent before the start of the financial and sovereign debt crises and a much more volatile and, on average, lower inflation rate of about 1.5 percent thereafter. Throughout the whole 20 years, average five-year-ahead inflation expectations, as captured by the Survey of Professional Forecasters, have remained stable within a narrow range between 1.8 and 2.0 percent, underlining the ECB's credibility. But after the sovereign debt crisis, when headline inflation and various core inflation measures declined significantly below 1 percent, a series of indicators pointed to the emergence of tangible risks of deanchoring inflation expectations and even deflation risks. They only disappeared after the ECB initiated a comprehensive easing package starting in June 2014—including quantitative easing, targeted credit operations, and negative policy rates—and thereby dispelled doubts about whether it had an effective tool kit to address those risks in an environment of close-to-zero interest rates. Headline inflation is currently about 2 percent (August 2018); and underlying inflation, though still subdued, is slowly increasing toward values close to 2 percent.

One issue that has been debated regarding this price stability track record is whether the ECB could have been more proactive in responding to the fallout from the sovereign debt crisis from mid-2010 to mid-2012. A fair assessment requires a real-time and not an ex-post perspective. The simple real-time policy reaction function used in this paper arguably suggests that both the policy rate tightening in 2011 and the subsequent easing were broadly in line with the ECB's own and other professional forecasters' growth and inflation projections at the time. Moreover, this period was increasingly characterized by solvency issues in both banking and government finances, which lingered for too long and reinforced each other in the absence of sufficient institutions and tools for solving the related collective action problems in a highly integrated monetary union of sovereign states with primarily national fiscal and supervisory policies. The unresolved public and private balance sheet problems and the resulting financial fragmentation in the euro area imposed tremendous obstacles on the effectiveness of the ECB's monetary policy.

At the same time, monetary policy cannot directly address such solvency issues. In fact, the prohibition of monetary financing (Article 123 of the Treaty on the Functioning of the European Union) forbids the ECB from directly financing governments or government tasks such as the recapitalization of banks. Against this background, the ECB's actions had to

balance the need to address impairments in the transmission of monetary policy due to malfunctioning financial markets and self-fulfilling market dynamics with the prohibition of monetary financing. This may explain, in part, what some observers regard as initially timid interventions in the government bond market through the SMP in 2010 and 2011. Leading up to the key June 2012 European Summit, necessary institutions and reforms to improve on the main weaknesses of EMU in the prudential and fiscal fields were put on a credible path. In this new context, the ECB stepped up its nonstandard tool kit to the next level, starting with President Draghi's "whatever it takes" speech and the powerful OMT program in the summer of 2012, as well as the comprehensive easing package (mentioned above) later on, in June 2014.

Overall, the main building blocks of the ECB's original monetary policy strategy and framework—its quantitative definition of price stability, the two pillars of economic and monetary analysis, the communication and accountability framework, and the broad-based and flexible operational framework—have served the ECB well during the past 20 years. However, as described in this paper, it was important that they evolved in response to challenges over time.

For example, as initial doubts by some observers about the ECB's anti-inflation credibility during the early years turned into concerns about its ability to address downward risks to price stability in a low-interest-rate environment, the quantitative inflation aim was clarified as being close to 2 percent, providing a buffer against the zero lower bound. Our analysis of the ECB's interest rate reaction function in subsection II.B suggests that the ECB pursued this inflation aim symmetrically. Moreover, this analysis indicates that the ECB's economic analysis and its quarterly macroeconomic projections formed the main basis for its monthly monetary policy decisions. At the same time, its monetary analysis provided a cross-check. It evolved from a narrower focus—with an emphasis on a reference value for M3 growth based on the quantity theory of money, which was useful in the first years to borrow the Deutsche Bundesbank's credibility—to a broad-based assessment of monetary developments and the state of financial intermediation and bank lending in the euro area economy. Before the crisis, this broad-based analysis was useful for considering the buildup of financial imbalances, though our interest-rate analysis does not show evidence that the ECB pursued a leaning-against-the-wind monetary policy approach. At the time, the ECB had neither a microprudential nor a macroprudential policy mandate and the related tools to address the financial imbalances at the source. Only with the advent of Banking

Union did the ECB acquire an important banking supervisory role as of November 2014, which implied comprehensive microprudential and some limited macroprudential responsibilities. Following the start of the financial crisis, the broadened monetary analysis was increasingly helpful in assessing fragilities in the banking sector and how they influence bank lending and the monetary policy transmission mechanism, as well as the effectiveness of some of the nonstandard measures.

Moreover, the ECB's communication and accountability framework was adjusted, as the need for additional communication in a complex (nonstandard) policy environment arose and forward guidance became an essential tool for easing policy in a low-interest-rate context. Finally, the ECB's operational framework was well suited to provide ample liquidity to its wide range of counterparties and quickly against a wide set of collateral when the money market froze. This helped address impairments in the early steps of the monetary transmission mechanism and also contributed to financial stability. Moreover, when the zero lower bound became more and more a constraint after the sovereign debt crisis, the operational framework proved broad and flexible enough to allow the ECB to expand its tool set with other nonstandard policy measures. A review of the available research on the effectiveness of the ECB's nonstandard measures for easing financial conditions, stimulating the economy, and bringing inflation back to its inflation aim—also in comparison with the evidence from other constituencies having used similar instruments, such as the U.S. and the U.K.—is quite encouraging and suggests that concerns that central banks may be powerless when interest rates hit the zero lower bound may be excessive.

All in all, the ECB has adjusted its monetary policy to changing and challenging circumstances over time, making effective use of its strategy and framework and maintaining a clear focus on its primary mandate of price stability in the medium term. As it has broadened its tools over time, it has become more similar to many of its peers as well. At the same time, some elements of its policy framework seem to have inspired changes in other central banks' frameworks—including the medium-term orientation of its price-stability objective, the transparency and accountability associated with the press conferences conducted by its president and vice president soon after its formal monetary policy meetings, and its broad and flexible operational framework.

A series of important reforms after the crises—in particular, the establishment of the European Stability Mechanism; the implementation of the first two legs of the Banking Union, the Single Supervisory and Resolution

Mechanisms; the signing of the Fiscal Compact; and the introduction of the European Semester, with the Macroeconomic Imbalance Procedure—have addressed some aspects of EMU’s incompleteness that complicated the ECB’s mission to maintain price stability over the past decade. In future years, the ECB’s monetary policy will benefit tremendously from the thorough implementation of these reforms, compliance with their objectives and rules, and further progress toward completing European Economic and Monetary Union along the lines of the 2015 Five Presidents’ Report (European Commission 2015).

ACKNOWLEDGMENTS Any views expressed are only the authors’ own and should not be interpreted as views of the European Central Bank or the Eurosystem. The previous working paper version of this paper was circulated with the title “The First Twenty Years of the European Central Bank: Monetary Policy.” We are beholden to numerous current and former ECB colleagues for comments—notably, Vítor Constâncio, Vítor Gaspar, Otmar Issing, and Jean-Pierre Vidal—and for generous help in compiling the information used for writing this paper. We give special thanks to Rupert de Vincent-Humphreys, Philipp Hochmuth, Christophe Kamps, Francesco Mongelli, Eric Persson, Benjamin Sahel, and David Sondermann. Also very helpful for improving the paper was the feedback that we received from participants in the Brookings Panel on Economic Activity, from the Hong Kong Monetary Authority–Federal Reserve Bank of New York conference on “Monetary Policy and Heterogeneity,” from the Federal Reserve Bank of Cleveland conference on “Inflation: Drivers and Dynamics,” and from the Eurosystem monetary policy committee—as well as from lectures at the ECB, at the SAFE Institute at the University of Frankfurt, and at the European School of Management and Technology in Berlin, and from a roundtable at the Banque de France and Paris School of Economics. Alexandre Mendonça, Adrian Monninger, Lea Steininger, Desislava Tartova, Hannes Tuieling, and Alexia Ventula Veghazy provided excellent research assistance. The remaining errors or omissions are our own.

References

- Acharya, Viral V., Itamar Drechsler, and Phillip Schnabl. 2014. "A Pyrrhic Victory? Bank Bailouts and Sovereign Credit Risk." *Journal of Finance* 69, no. 6: 2689–2739.
- Acharya, Viral V., Tim Eisert, Christian W. Eufinger, and Christian Hirsch. 2017. "Whatever It Takes: The Real Effects of Unconventional Monetary Policy." Working Paper 152. Frankfurt: Sustainable Architecture for Finance in Europe. <http://hdl.handle.net/10419/158505>.
- Acharya, Viral V., and Sascha Steffen. 2015. "The 'Greatest' Carry Trade Ever? Understanding Eurozone Bank Risks." *Journal of Financial Economics* 115, no. 2: 215–36.
- Aghion, Philippe, Emmanuel Farhi, and Enisse Kharroubi. 2017. "On the Interaction between Monetary Policy, Corporate Balance Sheets and Structural Reforms." In *Proceedings of the ECB Forum on Central Banking: Investment and Growth in Advanced Economies*. Frankfurt: European Central Bank.
- Aguir, Alvaro and Manuel M. F. Martins. 2008. "Testing for Asymmetries in the Preferences of the Euro Area Policy-Maker." *Applied Economics* 40, no. 13: 1651–67.
- Alesina, Alberto, Olivier Blanchard, Jordi Galí, Francesco Giavazzi, and Harald Uhlig. 2001. *Defining a Macroeconomic Framework for the Euro Area: Monitoring the European Central Bank 3*. London: Centre for Economic Policy Research.
- Alessi, Lucia, Eric Ghysels, Luca Onorante, Richard Peach, and Simon Potter. 2014. "Central Bank Macroeconomic Forecasting during the Global Financial Crisis: The European Central Bank and Federal Reserve Bank of New York Experiences." *Journal of Business & Economic Statistics* 32, no. 4: 483–500.
- Altavilla, Carlo, Desislava Andreeva, Miguel Boucinha, and Sarah Holton. 2019. "Monetary Policy, Credit Institutions and the Bank Lending Channel in the Euro Area." ECB Occasional Paper 222. Frankfurt: European Central Bank.
- Altavilla, Carlo, Fabio Canova, and Matteo Ciccarelli. 2016. "Mending the Broken Link: Heterogeneous Bank Lending and Monetary Policy Pass-Through." Working Paper 1987. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp1978.en.pdf>.
- Altavilla, Carlo, Giacomo Carboni, and Roberto Motto. 2015. "Asset Purchase Programmes and Financial Markets: Lessons from the Euro Area." Working Paper 1864. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp1864.en.pdf>.
- Altavilla, Carlo, Domenico Giannone, and Michele Lenza. 2016. "The Financial and Macroeconomic Effects of OMT Announcements." *International Journal of Central Banking* 12, no. 3: 29–57.
- Alvarez, Inmaculada, Fabio Casavecchia, Marino De Luca, Alexander Duering, Fabian Eser, Caspar Helmus, Christophe Hemous, Niko Herrala, Julija Jakovicka, Michelina Lo Russo, Filippo Pasqualone, Marc Rubens, Rita Soares, and

- Fabrizio Zennaro. 2017. "The Use of the Eurosystem's Monetary Policy Instruments and Operational Framework since 2012." Occasional Paper 188. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/scpops/ecb.op188.en.pdf>.
- Alves, Nuno, Carlos Robalo Marques, and João Sousa. 2007. "Is the Euro Area M3 Abandoning Us?" Working Paper 20-2007. Lisbon: Banco de Portugal. <https://www.bportugal.pt/sites/default/files/anexos/papers/wp200720.pdf>.
- Ambler, Steve, and Fabio Rumler. 2017. "The Effectiveness of Unconventional Monetary Policy Announcements in the Euro Area: An Event and Econometric Study." Working Paper 212. Vienna: Oesterreichische Nationalbank. <https://www.oenb.at/dam/jcr:774a56cf-aff2-4064-b0c3-4b3a55bae7c4/Working-Paper-212.pdf>.
- Andrade, Philippe, Johannes Breckenfelder, Fiorella De Fiore, Peter Karadi, and Oreste Tristani. 2016. "The ECB's Asset Purchase Programme: An Early Assessment." Working Paper 1956. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp1956.en.pdf>.
- Angelini, Paolo, Andrea Nobili, and Cristina Piscillo. 2011. "The Interbank Market after August 2007: What Has Changed, and Why?" *Journal of Money, Credit and Banking* 43, no. 5: 923–58.
- Artus, Patrick, Alan S. Blinder, Willem Buiter, Otmar Issing, and Robert M. Solow. 2008. "Round Table Discussion: Monetary Policy in the New International Environment." In *Central Banks as Economic Institutions*, edited by Jean-Philippe Touffut. Cheltenham, U.K.: Edward Elgar.
- Ascari, Guido, Anna Florio, and Alessandro Gobbi. 2017. "Transparency, Expectations Anchoring and Inflation Targeting." *European Economic Review* 91: 261–73.
- Ashworth, Jonathan P., and Charles A. E. Goodhart. 2012. "QE: A Successful Start May Be Running into Diminishing Returns." *Oxford Review of Economic Policy* 28, no. 4: 640–70.
- Autrup, Søren Lejsgaard, and Magdalena Grothe. 2014. "Economic Surprises and Inflation Expectations: Has Anchoring of Expectations Survived the Crisis?" Working Paper 1671. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp1671.pdf>.
- Ayuso, Juan, and Rafael Repullo. 2001. "Why Did the Banks Overbid? An Empirical Model of the Fixed Rate Tenders of the European Central Bank." *Journal of International Money and Finance* 20, no. 6: 857–70.
- Baldwin, R., and F. Giavazzi. 2015. *The Euro Zone Crisis: A Consensus View of the Causes and a Few Possible Solutions*. VoxEU e-book. London: Centre for Economic Policy Research. https://voxeu.org/sites/default/files/file/reboot_upload_0.pdf.
- . 2016. *How to Fix the Euro Zone: Views from Leading Economists*. VoxEU e-book. London: Centre for Economic Policy Research.
- Ball, Laurence. 2014. "The Case for a Long-Run Inflation Target of Four Percent." Working Paper 14/92. Washington: International Monetary Fund.

- Batini, Nicoletta, and Edward Nelson. 2001. "Optimal Horizons for Inflation Targeting." *Journal of Economic Dynamics and Control* 25, nos. 6–7: 891–910.
- Bean, Charles. 2003. "Inflation Targeting: The U.K. Experience." Speech given at Annual Congress of German Economic Association, Zurich, October 1.
- Beck, Guenter W., and Volker Wieland. 2008. "Central Bank Misperceptions and the Role of Money in Interest Rate Rules." *Journal of Monetary Economics* 55 (supplement): S1–S17.
- . 2010. "Money in Monetary Policy Design: Monetary Cross-Checking in the New Keynesian Model." Working Paper 1191. Frankfurt: European Central Bank.
- Beechey, Meredith J., Benjamin K. Johansson, and Andrew T. Levin. 2011. "Are Long-Run Inflation Expectations Anchored More Firmly in the Euro Area Than in the United States?" *American Economic Journal: Macroeconomics* 3, no. 2: 104–29.
- Begg, David, Paul De Grauwe, Francesco Giavazzi, Harold Uhlig, and Charles Wyplosz. 1998. *The ECB: Safe at Any Speed?* London: Centre for Economic Policy Research. https://cepr.org/sites/default/files/geneva_reports/GenevaP116.pdf.
- . 1999. *MECB Update: May 1999*. London: Centre for Economic Policy Research.
- Beirne, John, Lars Dalitz, Jacob Ejsing, Magdalena Grothe, Simone Manganelli, Fernando Monar, Benjamin Sahel, Matjaž Sušec, Jens Tapking, and Tana Vong. 2011. "The Impact of the Eurosystem's Covered Bond Purchase Programme on the Primary and Secondary Markets." Occasional Paper 122. Frankfurt: European Central Bank.
- Benhabib, Jess, Stephanie Schmitt-Grohé, and Martín Uribe. 2001. "Monetary Policy and Multiple Equilibria." *American Economic Review* 91, no. 1: 167–86.
- Bernanke, Ben S. 2002. "Deflation: Making Sure 'It' Doesn't Happen Here." Speech given at National Economists Club, Washington, November 21.
- . 2017. "Monetary Policy in a New Era." In *Proceedings of the Rethinking Macroeconomic Policy Conference*. Washington: Peterson Institute for International Economics.
- Bernanke, Ben S., Thomas Laubach, F. Mishkin, and A. Posen. 1999. "Inflation Targeting for the United States and the European Monetary Union." In *Inflation Targeting: Lessons from the International Experience*. Princeton University Press.
- Bernoth, Kerstin, and Jürgen von Hagen. 2004. "The Euribor Futures Market: Efficiency and the Impact of ECB Policy Announcements." *International Finance* 7, no. 1: 1–24.
- Beyer, Andreas, and Lucrezia Reichlin, eds. 2006. *Proceedings of the Fourth ECB Central Banking Conference: The Role of Money—Money and Monetary Policy in the Twenty-First Century*. Frankfurt: European Central Bank.

- <https://www.ecb.europa.eu/pub/pdf/other/roleofmoneyen2008en.pdf?646922df8deac54767d75918036d8dfd>.
- Bindseil, Ulrich. 2005. "Over- and Underbidding in Central Bank Open Market Operations Conducted as Fixed-Rate Tender." *German Economic Review* 6, no. 1: 95–130.
- Bindseil, Ulrich, Marco Corsi, Benjamin Sahel, and Ad Visser. 2017. "The Euro-system Collateral Framework Explained." Occasional Paper 189. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/scpops/ecb.op189.en.pdf>.
- Blanchard, Olivier, and Francesco Giavazzi. 2002. "Current Account Deficits in the Euro Area: The End of the Feldstein-Horioka Puzzle?" *Brookings Papers on Economic Activity* 2: 147–210. https://www.brookings.edu/wp-content/uploads/2002/06/2002b_bpea_blanchar.pdf.
- Blanchard, Olivier, Giovanni Dell'Ariccia, and Paolo Mauro. 2010. "Rethinking Macro Policy." VoxEU.org, February 16. <https://voxeu.org/article/rethinking-macro-policy>.
- Blattner, Tobias, Marco Catenaro, Michael Ehrmann, Rolf Strauch, and Jarkko Turunen. 2008. "The Predictability of Monetary Policy." Occasional Paper 83. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/scpops/ecbocp83.pdf>.
- Blei, M. David, Andrew Y. Ng, and Michael I. Jordan. 2003. "Latent Dirichlet Allocation." *Journal of Machine Learning Research* 3: 993–1022.
- Bletzinger Tilman, and Volker Wieland. 2016. "Forward Guidance and 'Lower for Longer': The Case of the ECB." Working Paper 102. Frankfurt: Institute for Monetary and Financial Stability. <http://publikationen.ub.uni-frankfurt.de/frontdoor/index/index/year/2016/docId/39578>.
- Boeckx, Jef, Maarten Dossche, and Gert Peersman. 2017. "Effectiveness and Transmission of the ECB's Balance Sheet Policies." *International Journal of Central Banking* 13, no. 1: 297–333.
- Borio, Claudio, and Philip Lowe. 2002. "Asset Prices, Financial and Monetary Stability: Exploring the Nexus." BIS Working Paper 114. Basel: Bank for International Settlements.
- Bridges, Jonathan, and Ryland Thomas. 2012. "The Impact of QE on the UK Economy: Some Supportive Monetarist Arithmetic." Working Paper 442. London: Bank of England. <https://www.bankofengland.co.uk/working-paper/2012/the-impact-of-qe-on-the-uk-economy-some-supportive-monetarist-arithmetic>.
- Bruggeman, Annick, Paola Donati, and Anders Warne. 2003. "Is the Demand for Euro Area M3 Stable?" Working Paper 255. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp006.pdf?14e04cc5febe813648551a4836f7ead1>.
- Brunnermeier, Markus, and Yann Koby. 2018. "The Reversal Interest Rate." Working paper. https://scholar.princeton.edu/sites/default/files/markus/files/20p_reversalrate.pdf.

- Buiter, Willem. 1999. "Alice in Euroland." *Journal of Common Market Studies* 37, no. 2: 181–210.
- Bundick, Brent, and A. Lee Smith. 2016. "The Dynamic Effects of Forward Guidance Shocks." Research Working Paper 16-02. Kansas City: Federal Reserve Bank of Kansas City. <https://www.kansascityfed.org/~media/files/publicat/reswkpap/pdf/rwp16-02.pdf?la=en>.
- Buti, M., S. Deroose, V. Gaspar, and J. Nogueira Martins, eds. 2010. *The Euro: The First Decade*. Cambridge University Press.
- Cahn, Christophe, Julien Matheron, and Jean-Guillaume Sahuc. 2017. "Assessing the Macroeconomic Effects of LTROs during the Great Recession." *Journal of Money, Credit and Banking* 49, no. 7: 1443–82.
- Camba-Méndez, Gonzalo, and Francesco Mongelli. 2017. "Policy Responses in Europe since the Onset of the Crisis (2007–Mid 2017)." Paper, ECB, October 26. https://www.researchgate.net/publication/321637875_Policy_Responses_in_Europe_Since_the_Onset_of_the_Crisis_2007-Mid_2017.
- Camba-Méndez, Gonzalo, and Thomas Werner. 2017. "The Inflation Risk Premium in the Post-Lehman Period." Working Paper 2033. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp2033.en.pdf>.
- Cao, Juan, Tian Xia, Jintao Li, Yongdong Zhang, and Sheng Tang. 2009. "A Density-Based Method for Adaptive LDA Model Selection." *Neurocomputing* 72, nos. 7–9: 1775–81.
- Carpenter, Seth, Selva Demiralp, and Jens Eisenschmidt. 2013. "The Effectiveness of the Non-Standard Policy Measures during the Financial Crises: The Experiences of the Federal Reserve and the European Central Bank." Working Paper 1562. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp1562.pdf>.
- Cassola, Nuno, Mathias Drehmann, Philipp Hartmann, Marco Lo Duca, and Martin Scheicher. 2008. "A Research Perspective on the Propagation of the Credit Market Turmoil." Research Bulletin 7. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/other/researchbulletin07en.pdf>.
- Cassola, Nuno, Alain Durré, and Cornelia Holthausen. 2011. "Implementing Monetary Policy in Crisis Times: The Case of the ECB." In *Proceedings of the Sixth ECB Central Banking Conference: Approaches to Monetary Policy Revisited—Lessons from the Crisis*, edited by M. Jarocinski, F. Smets, and C. Thimann. Frankfurt: European Central Bank.
- CEBS (Committee of European Banking Supervisors). 2009. "CEBS's Press Release on the Results of the EU-wide Stress Testing Exercise." London, October 1. <https://eba.europa.eu/documents/10180/15977/CEBS-2009-180-Annex-2-%28Press-release-from-CEBS%29.pdf>.
- . 2010. "Aggregate Outcome of the 2010 EU-wide Stress Test Exercise Coordinated by CEBS in Cooperation with the ECB." London, July 23. <https://eba.europa.eu/documents/10180/15938/Summaryreport.pdf/95030af2-7b52-4530-afe1-f067a895d163>.

- Chen, Han, Vasco Cúrida, and Andrea Ferrero. 2012. "The Macroeconomic Effects of Large-Scale Asset Purchase Programmes." *Economic Journal* 122: F289–F315.
- Christiano, Lawrence J., Roberto Motto, and Massimo Rostagno. 2014. "Risk Shocks." *American Economic Review* 104, no. 1: 27–65.
- Chung, Hess, Jean-Phillippe Laforge, David Reifschneider, and John C. Williams. 2011. "Estimating the Macroeconomic Effects of the Fed's Asset Purchases." *Economic Letter* 2011-03. San Francisco: Federal Reserve Bank of San Francisco. <https://www.frbsf.org/economic-research/files/el2011-03.pdf>.
- Ciccarelli, Matteo, and Chiara Osbat, eds. 2017. "Low Inflation in the Euro Area: Causes and Consequences." Occasional Paper 181. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/scpops/ecbop181.en.pdf>.
- Claus, Edda, Iris Claus, and Leo Krippner. 2016. "Asset Markets and Monetary Policy Shocks at the Zero Lower Bound." Discussion Paper 2014/03a. Wellington: Reserve Bank of New Zealand.
- Coenen, Günter, Michael Ehrmann, Gaetano Gaballo, Peter Hoffmann, Anton Nakov, Stefano Nardelli, Eric Persson, and Georg Strasser. 2017. "Communication of Monetary Policy in Unconventional Times." Working Paper 2080. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/scpwp/ceb.wp2080.en.pdf>.
- Coenen, Günter, Roland Straub, and Mathias Trabandt. 2012. "Fiscal Policy and the Great Recession in the Euro Area." *American Economic Review: Papers & Proceedings* 102, no. 3: 71–76.
- Cœuré, B. 2015. "Embarking on Public Sector Asset Purchases." Speech given at Second International Conference on Sovereign Bond Markets, March 10.
- Coibion, Olivier, Yuriy Gorodnichenko, and Johannes Wieland. 2012. "The Optimal Inflation Rate in New Keynesian Models: Should Central Banks Raise Their Inflation Targets in Light of the Zero Lower Bound?" *Review of Economic Studies* 79: 1371–1406.
- Constâncio, Vítor. 2013. "The European Crisis and the Role of the Financial System." Speech given at Bank of Greece Conference on the Crisis in the Euro Area, Athens, May, 23.
- . 2018. "Past and Future of the ECB's Monetary Policy." Speech at Conference on Central Banks in Historical Perspective: What Changed after the Financial Crisis, Central Bank of Malta, Valletta, May 4.
- Corsetti, Giancarlo. 2015. "Roots of the EZ Crisis: Incomplete Development and Imperfect Credibility of Institutions." VoxEU article, Centre for European Policy Research, September 7. <https://voxeu.org/article/roots-ez-crisis>.
- Corsetti, Giancarlo, and Luca Dedola. 2016. "The Mystery of the Printing Press: Monetary Policy and Self-Fulfilling Debt Crises." *Journal of the European Economic Association* 14, no. 6: 1329–71.
- Corsetti, Giancarlo, and Paolo Pesenti. 1999. "Stability, Asymmetry, and Discontinuity: The Launch of European Monetary Union." *Brookings Papers on Economic Activity*, no. 2: 295–372.

- Court of Justice. 2015. "Judgment of the Court of Justice of the European Union (Grand Chamber) of June 16, 2015, *Peter Gauweiler and Others v. Deutscher Bundestag*." <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:62014CJ0062>.
- Crosgnani, Matteo, Miguel Faria-e-Castro, and Luís Fonseca. 2018. "The (Unintended) Consequence of the Largest Liquidity Injection Ever." Working Paper 2017-039C. Saint Louis: Federal Reserve Bank of Saint Louis. <https://research.stlouisfed.org/wp/more/2017-039>.
- Darracq-Paries, Matthieu, Christoffer Kok Sorensen, and Diego Rodriguez-Palenzuela. 2011. "Macroeconomic Propagation Across Different Regulatory Regimes: Evidence from an Estimated DSGE Model for the Euro Area." *International Journal of Central Banking* 7, no. 4: 49–113.
- Darracq-Paries, Matthieu, and Roberto A. De Santis. 2015. "A Non-Standard Monetary Policy Shock: The ECB's 3-Year LTROs and the Shift in Credit Supply." *Journal of International Money and Finance* 54: 1–34.
- De Bandt, Olivier, Philipp Hartmann, and José Luis Peydró-Alcalde. 2015. "Systemic Risk in Banking after the Great Financial Crisis." In *The Oxford Handbook of Banking*, edited by Allen N. Berger, Philip Molyneux, and John O. S. Wilson. Oxford University Press.
- Debortoli, Davide, Jordi Galf, and Luca Gambetti. 2018. "On the Empirical (Ir)relevance of the Zero Lower Bound Constraint." Working Paper 1594. Barcelona: Universitat Pompeu Fabra. <https://econ-papers.upf.edu/papers/1594.pdf>.
- Del Negro, Marco, Gauti Eggertsson, Andrea Ferrero, and Nobuhiro Kiyotaki. 2016. "The Great Escape? A Quantitative Evaluation of the Fed's Liquidity Facilities." Staff Report 520. New York: Federal Reserve Bank of New York. https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr520.pdf.
- Demertzis, Maria, Massimiliano Marcellino, and Nicola Viegi. 2009. "Anchors for Inflation Expectations." Working Paper 229. Amsterdam: De Nederlandsche Bank. https://www.dnb.nl/en/binaries/229%20Anchors%20for%20Inflation%20Expectations_tcm47-225545.pdf.
- Demiralp, Selva, Jens Eisenschmidt, and Thomas Vlassopoulos. 2019. "Negative Interest Rates, Excess Liquidity and Retail Deposits: Banks' Reaction to Unconventional Monetary Policy in the Euro Area." ECB Working Paper 2283. Frankfurt: European Central Bank.
- De Pooter, Michiel, Robert F. Martin, and Seth Pruitt. 2018. "The Liquidity Effects of Official Bond Market Intervention." *Journal of Financial and Quantitative Analysis* 53, no. 1: 243–68.
- De Santis, Roberto A. 2016. "Impact of the Asset Purchase Programme on Euro Area Government Bond Yields Using Market News." Working Paper 1939. Frankfurt: European Central Bank.
- . 2018. "Unobservable Country Bond Premia and Fragmentation." *Journal of International Money and Finance* 82: 1–25.
- . Forthcoming. "Redenomination Risk." *Journal of Money, Credit and Banking*. <https://onlinelibrary.wiley.com/doi/10.1111/jmcb.12582>.

- Detken, Carsten, Dieter Gerdesmeier, and Barbara Roffina. 2010. "Interlinkages between Money, Credit and Asset Prices and Their Implications for Consumer Price Inflation: Recent Empirical Work." In *Enhancing Monetary Analysis*, edited by Lucas Papademos and Jürgen Stark. <https://www.ecb.europa.eu/pub/pdf/other/enhancingmonetaryanalysis2010en.pdf?13ffbb3e72c01de37830173b8a5b8217>.
- Detken, Carsten, and Frank Smets. 2004. "Asset Price Booms and Monetary Policy." In *Macroeconomic Policies in the World Economy*, edited by Horst Siebert. Berlin: Springer-Verlag.
- Deutsche Bundesbank. 2018. "Lower Bound, Inflation Target and the Anchoring of Inflation Expectations." *Monthly Report*, June, 31–50. <https://www.bundesbank.de/resource/blob/749842/98826deabb00741b7796b903867f5c40/mL/2018-06-zinsuntergrenze-data.pdf>.
- Domingues Semeano, João, and Marien Ferdinandusse. 2018. "The Fiscal Impact of Financial Sector Support Measures: Where Do We Stand a Decade on from the Financial Crisis?" *ECB Economic Bulletin*, June. https://www.ecb.europa.eu/pub/economic-bulletin/focus/2018/html/ecb.ebbox201806_04.en.html.
- Dovern, Jonas, and Geoff Kenny. 2017. "The Long-Term Distribution of Expected Inflation in the Euro Area: What Has Changed since the Great Recession?" Working Paper 1999. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/scpwp/scpwp1999.en.pdf>.
- Draghi, Mario. 2012. "Verbatim of the remarks made by Mario Draghi, Speech at the Global Investment Conference, London, 26 July." <https://www.ecb.europa.eu/press/key/date/2012/html/sp120726.en.html>.
- . 2014a. "Monetary Policy Communication in Turbulent Times." Speech at conference "De Nederlandsche Bank 200 Years: Central Banking in the Next Two Decades," Amsterdam, April 24. <https://www.ecb.europa.eu/press/key/date/2014/html/sp140424.en.html>.
- . 2014b. "Unemployment in the Euro Area." Speech given at annual central bank symposium in Jackson Hole, Wyo., August 22.
- . 2016. "Delivering a Symmetric Mandate with Asymmetric Tools: Monetary Policy in a Context of Low Interest Rates." Speech given at ceremony to mark 200th anniversary of Oesterreichische Nationalbank, Vienna, June 2.
- . 2017. "Monetary Policy and the Economic Recovery in the Euro Area." Speech given at ECB and its Watchers XVIII Conference, Frankfurt, April 6.
- Driffill, John, and Zeno Rotondi. 2004. "Monetary Policy and Lexicographic Preference Ordering." CEPR Discussion Paper 4247. https://cepr.org/active/publications/discussion_papers/dp.php?dpno=4247.
- Duisenberg, Willem F. 1998. "Introductory Statement and Transcript of Questions and Answers." ECB press conference, Frankfurt, October 13. <https://www.ecb.europa.eu/press/pressconf/1998/html/is981013.en.html>.
- . 1999. "The Eurosystem's Monetary Policy Strategy: The First Year's Experience." Speech given at joint congress of the federations EUROFINAS and LEASEUROPE, Paris, October 11. <https://www.ecb.europa.eu/press/key/date/1999/html/sp991011.en.html>.

- ECB (European Central Bank). 1999. "The Stability-Oriented Monetary Policy Strategy of the Eurosystem." *ECB Monthly Bulletin*, January, 39–50. <https://www.ecb.europa.eu/pub/pdf/mobu/mb199901en.pdf>.
- . 2000a. "Guideline of the European Central Bank of 31 August 2000 on Monetary Policy Instruments and Procedures of the Eurosystem." *Official Journal of the European Communities*, L 310/1, December 11. https://www.ecb.europa.eu/ecb/legal/pdf/01_a_en_ecb_2000_7.pdf.
- . 2000b. "Statement on the Euro by Dr. Willem F. Duisenberg, President of the European Central Bank." Press release, Frankfurt, May 5. <https://www.ecb.europa.eu/press/pr/date/2000/html/pr000505.en.html>.
- . 2000c. "The ECB Announces Joint Intervention in the Exchange Markets." Press release, September 22. <https://www.ecb.europa.eu/press/pr/date/2000/html/pr000922.en.html>.
- . 2003a. "The Outcome of the ECB's Evaluation of Its Monetary Strategy." *ECB Monthly Bulletin*, June, 79–92. https://www.ecb.europa.eu/pub/pdf/other/pp79_92_mb200306en.pdf?e366490ebb678c0d82b6d62a85e10d44.
- . 2003b. "Press Seminar on the Evaluation of the ECB's Monetary Policy Strategy." Press release, May 8. https://www.ecb.europa.eu/press/pressconf/2003/html/is030508_1.en.html.
- . 2003c. "Introductory Statement by Jean Claude Trichet." December 4. <https://www.ecb.europa.eu/press/pressconf/2003/html/is031204.en.html>.
- . 2005. "Asset Price Bubbles and Monetary Policy." *Monthly Bulletin*, April, 47–60.
- . 2007a. "The ECB's Additional Open Market Operations in the Period from 8 August to 5 September 2007." *Monthly Bulletin*, September, 30–34. <https://www.ecb.europa.eu/pub/pdf/mobu/mb200709en.pdf>.
- . 2007b. "Measures Designed to Address Elevated Pressures in Short-Term Funding Markets." Press release, December 12. <https://www.ecb.europa.eu/press/pr/date/2007/html/pr071212.en.html>.
- . 2008a. "The 10th Anniversary of the European Central Bank." *Monthly Bulletin*, June. <https://www.ecb.europa.eu/pub/pdf/other/10thanniversaryoftheecbmb200806en.pdf>.
- . 2008b. "Measures to Further Expand the Collateral Framework and Enhance the Provision of Liquidity." Press release, October 15. <https://www.ecb.europa.eu/press/pr/date/2008/html/pr081015.en.html>.
- . 2009a. "The Concept of Systemic Risk." *Financial Stability Review*, December, 134–42. <https://www.ecb.europa.eu/pub/pdf/fsr/financialstabilityreview200912en.pdf?10adf78b3dc425f72013840711c53aa7>.
- . 2009b. "Wage Dynamics in Europe: Final Report of the Wage Dynamics Network (WDN)." December 4. https://www.ecb.europa.eu/home/pdf/wdn_finalreport_dec2009.pdf.
- . 2009c. "Longer-Term Refinancing Operations." Press Release, May 7. https://www.ecb.europa.eu/press/pr/date/2009/html/pr090507_2.en.html.

- . 2011a. “The Implementation of Monetary Policy in the Euro Area.” General documentation on Eurosystem monetary policy instruments and procedures.
- . 2011b. “Statement by the President of the ECB.” Press release, August 7. <https://www.ecb.europa.eu/press/pr/date/2011/html/pr110807.en.html>.
- . 2011c. “ECB Announces Measures to Support Bank Lending and Money Market Activity.” Press release, December 8. https://www.ecb.europa.eu/press/pr/date/2011/html/pr111208_1.en.html.
- . 2012a. “Financial Integration in Europe.” <https://www.ecb.europa.eu/pub/pdf/fie/financialintegrationineurope201204en.pdf?555d50a5346bd5a322c47e8df4f9e223>.
- . 2012b. “Technical Features of Outright Monetary Transactions.” Press release, September 6. https://www.ecb.europa.eu/press/pr/date/2012/html/pr120906_1.en.html.
- . 2013a. “An Assessment of Eurosystem Staff Macroeconomic Projections.” *Monthly Bulletin*, May. <https://www.ecb.europa.eu/pub/pdf/mobu/mb201305en.pdf>.
- . 2013b. “Financial Integration in Europe.” <https://www.ecb.europa.eu/pub/pdf/other/financialintegrationineurope201304en.pdf>.
- . 2013c. “Introductory Statement to the Press Conference (with Q&A).” Press release, July 4. <https://www.ecb.europa.eu/press/pressconf/2013/html/is130704.en.html>.
- . 2014. “Experience with Foreign Currency Liquidity-Providing Central Bank Swaps.” *Monthly Bulletin*, May.
- . 2015a. “Decision of the ECB of 4 March 2015 on a Secondary Markets Public Sector Asset Purchase Programme (ECB/2015/10).” https://www.ecb.europa.eu/ecb/legal/pdf/en_dec_ecb_2015_10_f_sign.pdf.
- . 2015b. “The Transmission of the ECB’s Recent Non-Standard Monetary Policy Measures.” *Economic Bulletin* 7, 32–51. <https://www.ecb.europa.eu/pub/pdf/ecbu/eb201507.en.pdf?3adaa9be8dec4be9a0ed7176ad6030f>.
- . 2015c. “Introductory Statement to the Press Conference (with Q&A).” Press release, January 22. <https://www.ecb.europa.eu/press/pressconf/2015/html/is150122.en.html>.
- . 2015d. “Introductory Statement to the Press Conference (with Q&A).” Press release, December 3. <https://www.ecb.europa.eu/press/pressconf/2015/html/is151203.en.html>.
- . 2015e. “ECB Announces Expanded Asset Purchase Programme.” Press release, January 22. https://www.ecb.europa.eu/press/pr/date/2015/html/pr150122_1.en.html.
- . 2016a. “A Guide to the Eurosystem/ECB Staff Macroeconomic Projection Exercises.” <https://www.ecb.europa.eu/pub/pdf/other/staffprojections/guide201607.en.pdf>.
- . 2016b. “Monetary Policy Decisions.” Press release, March 10. <https://www.ecb.europa.eu/press/pr/date/2016/html/pr160310.en.html>.

- . 2016c. “Introductory Statement to the Press Conference (with Q&A).” Press release, March 10. <https://www.ecb.europa.eu/press/pressconf/2016/html/is160310.en.html>.
- . 2017a. “How Do Professional Forecasters Assess the Risks to Inflation?” *Economic Bulletin* 5. <https://www.ecb.europa.eu/pub/economic-bulletin/html/eb201705.en.html#IDofBox4>.
- . 2017b. “Impact of the ECB’s Non-Standard Measures on Financing Conditions: Taking Stock of Recent Evidence.” *Economic Bulletin* 2. <https://www.ecb.europa.eu/pub/pdf/ecbu/eb201702.en.pdf?5f73fe43bfa3863210291ae55ed6f265>.
- . 2017c. “The Targeted Longer-Term Refinancing Operations: An Overview of the Take-Up and Their Impact on Bank Intermediation.” *Economic Bulletin* 3. <https://www.ecb.europa.eu/pub/pdf/ecbu/eb201703.en.pdf?fbcb755119af8c03bf99092d87d2f7fd>.
- . 2017d. “Introductory Statement.” Press release, June 8. <https://www.ecb.europa.eu/press/pressconf/2017/html/ecb.is170608.en.html>.
- . 2018a. “Euro Area Bank Lending Survey, 2018:Q4.” https://www.ecb.europa.eu/stats/ecb_surveys/bank_lending_survey/html/index.en.html.
- . 2018b. “Structure of the Euro Area Economy.” <https://www.ecb.europa.eu/mopo/eaec/html/index.en.html>.
- . 2018c. “Monetary Policy Decisions.” Press release, June 14. <https://www.ecb.europa.eu/press/pr/date/2018/html/ecb.mp180614.en.html>.
- Ehrhart, Karl-Martin. 2001. “European Central Bank Operations: Experimental Investigation of the Fixed Rate Tender.” *Journal of International Money and Finance* 20, no. 6: 871–93.
- Ehrmann, M., M. Fratzscher, R. Gürkaynak, and E. Swanson. 2011. “Convergence and Anchoring of Yield Curves in the Euro Area.” *Review of Economics and Statistics* 93, no. 1: 350–64.
- Eisenschmidt, Jens, and Frank Smets. 2018. “Negative Interest Rates: Lessons from the Euro Area.” <https://www.bcentral.cl/documents/145129/900210/Eisenschmidt+Smets+%25282017%2529+v2.pdf/da22c99f-b34a-fecf-ce6d-690f6ad3b202>.
- Eisenschmidt, Jens, and Jens Tapking. 2009. “Liquidity Risk Premia in Unsecured Interbank Money Market Rates.” ECB Working Paper 1025. <https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp1025.pdf?68bbf7e8d380ec23151dc9a1f602a1d9>.
- Eser, Fabian, Marta Carmona Amaro, Stefano Iacobelli, and Marc Rubens. 2012. “The Use of the Eurosystem’s Monetary Policy Instruments and Operational Framework since 2009.” ECB Occasional Paper 135. <https://www.ecb.europa.eu/pub/pdf/scpops/ecbocp135.pdf?c18a6c50e4fcefead8aff7a5d82db5f4>.
- Eser, Fabian, and Bernd Schwaab. 2016. “Evaluating the Impact of Unconventional Monetary Policy Measures: Empirical Evidence from the ECB’s Securities Markets Programme.” *Journal of Financial Economics* 119, no. 1: 147–67.

- European Banking Authority. 2011. "Results of the 2011 EU-wide Stress Test under the Adverse Scenario, London, 15 July." <https://eba.europa.eu/risk-analysis-and-data/eu-wide-stress-testing/2011/results>.
- European Commission. 2008. "A European Economic Recovery Plan." Communication from Commission to European Council, Brussels, November 26. http://ec.europa.eu/economy_finance/publications/pages/publication13504_en.pdf.
- . 2015. *Completing Europe's Economic and Monetary Union*. Report by Jean-Claude Juncker in close collaboration with Donald Tusk, Jeroen Dijsselbloem, Mario Draghi, and Martin Schulz. Brussels: European Commission. https://ec.europa.eu/commission/sites/beta-political/files/5-presidents-report_en.pdf.
- . 2017. "Coping with the International Financial Crisis at the National Level in a European Context: Impact and Financial Sector Policy Responses in 2008–2015." European Commission Staff Document SWD (2017) 373. European Commission, Brussels. <https://ec.europa.eu/info/system/files/eucountries-responses-to-financial-crisis.pdf>.
- . 2018. "Spring 2018 Economic Forecast: Expansion to Continue, Amid New Risks." https://ec.europa.eu/info/business-economy-euro/economic-performance-and-forecasts/economic-forecasts/spring-2018-economic-forecast-expansion-continue-amid-new-risks_en.
- European Council. 2012a. "Euro Area Summit Statement, Brussels, 29 June." <https://www.consilium.europa.eu/media/21400/20120629-euro-area-summit-statement-en.pdf>.
- . 2012b. "Towards a Genuine Economic and Monetary Union: Prepared by Herman Van Rompuy, President of the European Council, in Close Collaboration with José Manuel Barroso, President of the European Commission, Jean-Claude Juncker, President of the Eurogroup, and Mario Draghi, President of the European Central Bank, 5 December." <https://www.consilium.europa.eu/media/23818/134069.pdf>.
- . 2012c. "Towards a Genuine Economic and Monetary Union: Report by the President of the European Council Herman Van Rompuy, Brussels, 26 June." <https://www.consilium.europa.eu/media/33785/131201.pdf>.
- European Union. 2012a. "Consolidated Version of the Treaty on European Union." *Official Journal of the European Union*, C326/13, October 26. https://eur-lex.europa.eu/resource.html?uri=cellar:c382f65d-618a-4c72-9135-1e68087499fa.0006.02/DOC_2&format=PDF.
- . 2012b. "Consolidated Version of the Treaty on the Functioning of the European Union." *Official Journal of the European Union*, C326/47, October 26. https://eur-lex.europa.eu/resource.html?uri=cellar:c382f65d-618a-4c72-9135-1e68087499fa.0006.02/DOC_3&format=PDF.
- . 2012c. "Protocol No. 4 on the Statute of the European System of Central Banks and the European Central Bank." *Official Journal of the European Union*, C326/230, October 26. <https://eur-lex.europa.eu/resource>.

- html?uri=cellar:c382f65d-618a-4c72-9135-1e68087499fa.0006.02/DOC_4&format=PDF.
- . 2013. “Council Regulation (EU) No. 1024/2013 of 15 October 2013 Conferring Specific Tasks on the European Central Bank Concerning Policies Relating to the Prudential Supervision of Credit Institutions.” *Official Journal of the European Union*, L287/63, October 29.
- Evanoff, Douglas D., Philipp Hartmann, and George G. Kaufman, eds. 2009. *The First Credit Market Turmoil of the 21st Century*. Hackensack, N.J.: World Scientific Publishing.
- Fahr, Stephan, Roberto Motto, Massimo Rostagno, Frank Smets, and Oreste Tristani. 2013. “A Monetary Policy Strategy in Good and Bad Times: Lessons from the Recent Past.” *Economic Policy* 28, no. 74: 243–88.
- Feld, Lars, Christoph Schmidt, Isabel Schnabel, and Volker Wieland. 2016. “Causes of the Eurozone Crisis: A Nuanced View.” VoxEU article, March 22. <https://voxeu.org/article/causes-eurozone-crisis-nuanced-view>.
- Ferrando, Annalisa, Alexander Popov, and Gregory F. Udell. 2015. “Sovereign Stress, Unconventional Monetary Policy, and SME Access to Finance.” Working Paper 1820. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp1820.en.pdf?1a09169e5aa2d3a6002675b624554421>.
- Fischer, Björn, Michele Lenza, Huw Pill, and Lucrezia Reichlin. 2008. “Money and Monetary Policy: The ECB Experience 1999–2006.” In *Proceedings of the Fourth ECB Central Banking Conference: The Role of Money—Money and Monetary Policy in the Twenty-First Century*, edited by Andreas Beyer and Lucrezia Reichlin. Frankfurt: European Central Bank.
- Fraccaroli, Nicolo, Alessandro Giovannini, and Jean-François Jamet. 2018. “Accounting for Accountability at the ECB.” Vox-EU article. <https://voxeu.org/article/accounting-accountability-evolution-ecb-s-accountability-practices-during-crisis>.
- Francis, Neville R., Laura E. Jackson, and Michael T. Owyang. 2017. “How Has Empirical Monetary Policy Analysis Changed after the Financial Crisis?” Working Paper 2014-19C, Federal Reserve Bank of Saint Louis. <https://s3.amazonaws.com/real.stlouisfed.org/wp/2014/2014-019.pdf>.
- Fuhrer, Jeffrey C., and Giovanni P. Olivei. 2011. “The Estimated Macroeconomic Effects of the Federal Reserve’s Large-Scale Treasury Purchase Program.” Public Policy Brief 11-2, Federal Reserve Bank of Boston. <https://www.bostonfed.org/publications/public-policy-brief/2011/the-estimated-macroeconomic-effects-of-the-federal-reserves-largescale-treasury-purchase-program.aspx>.
- Galati, Gabriele, Steven Poelhekke, and Chen Zhou. 2011. “Did the Crisis Affect Inflation Expectations?” *International Journal of Central Banking* 7, no. 1: 167–207.
- Galí, J. 2003. “Monetary Policy in the Early Years of EMU.” In *EMU and Economic Policy in Europe: Challenges of the Early Years*, edited by M. Buti and A. Sapir. Cheltenham, U.K.: Edward Elgar.
- Garcia-de-Andoain, C., F. Heider, M. Hoerova, and S. Manganelli. 2016. “Lending-of-Last-Resort Is as Lending-of-Last-Resort Does: Central Bank

- Liquidity Provision and Interbank Market Functioning in the Euro Area.” *Journal of Financial Intermediation* 28: 32–47.
- Garcia-Posada, Miguel, and Marcos Marchetti. 2016. “The Bank Lending Channel of Unconventional Monetary Policy: The Impact of the VLTROs on Credit Supply in Spain.” *Economic Modelling* 58: 427–41.
- Gaspar, Vitor, Frank Smets, and David Vestin. 2010. “Is Time Ripe for Price Level Path Stability?” In *Challenges in Central Banking*, edited by Pierre L. Siklos, Martin T. Bohl, and Mark E. Wohar. Cambridge University Press.
- Geraats, Petra. 2002. “Central Bank Transparency.” *Economic Journal* 112: F532–F565.
- Geraats, Petra, Francesco Giavazzi, and Charles Wyplosz. 2008. “Transparency and Governance in the Eurozone.” Vox-EU article, Centre for European Policy Research. <https://voxeu.org/article/transparency-and-governance-eurozone>.
- Gerdesmeier, Dieter, and Barbara Roffia. 2003. “Empirical Estimates of Reaction Functions for the Euro Area.” Working Paper 206. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp206.pdf?7b1765155c69a259582cc3bb5332d66c>.
- Gerlach, Stefan, and John Lewis. 2014. “ECB Reaction Functions and the Crisis of 2008.” *International Journal of Central Banking* 10, no. 1: 137–57.
- Gertler, Mark, and Peter Karadi. 2013. “QE 1 vs. 2 vs. 3 . . . : A Framework for Analyzing Large-Scale Asset Purchases as a Monetary Policy Tool.” *International Journal of Central Banking* 9, no. S1: 5–53.
- Ghysels, Eric, Julien Idier, Simone Manganelli, and Olivier Vergote. 2017. “A High-Frequency Assessment of the ECB Securities Markets Programme.” *Journal of the European Economic Association* 15, no. 1: 218–43.
- Giannone, Domenico, Michele Lenza, Huw Pill, and Lucrezia Reichlin. 2012. “The ECB and the Interbank Market.” *Economic Journal* 122: F467–F486.
- Gorter, Janka, Jan Jacobs, and Jakob de Haan. 2008. “Taylor Rules for the ECB Using Expectations Data.” *Scandinavian Journal of Economics* 110, no. 3: 473–88.
- Gorton, Gary, and Andrew Metrick. 2012. “Securitized Banking and the Run on Repo.” *Journal of Financial Economics* 104, no. 3: 425–51.
- Grishchenko, Olesya, Sarah Mouabbi, and Jean-Paul Renne. 2017. “Measuring Inflation Anchoring and Uncertainty: A U.S. and Euro Area Comparison.” Finance and Economics Discussion Series 2017-102. Washington: Board of Governors of the Federal Reserve System. <https://www.federalreserve.gov/econres/feds/files/2017102pap.pdf>.
- Gürkaynak, Refet S., Andrew Levin, and Eric Swanson. 2010. “Does Inflation Targeting Anchor Long-Run Inflation Expectations? Evidence from the U.S., U.K., and Sweden.” *Journal of the European Economic Association* 8, no. 6: 1208–42.
- Hamilton, James. 2018. “The Effectiveness of Large-Scale Asset Purchases.” Vox-EU article, October 12, Centre for Economic Policy Research. <https://voxeu.org/article/effectiveness-large-scale-asset-purchases>.

- Hammermann, Felix, Kieran Leonard, Stefano Nardelli, and Julian von Landesberger. 2019. "Taking Stock of the Eurosystem's Asset Purchase Programme after the End of Net Asset Purchases." *Economic Bulletin*, March 18. https://www.ecb.europa.eu/pub/economic-bulletin/articles/2019/html/ecb.ebart201902_01~3049319b8d.en.html.
- Hartmann, Philipp. 2010. "How Europe Handled the Financial Crisis," executive address at Sixth Annual FX Week Asia Congress, Singapore, October 12.
- . 2015. "Real-Estate Markets and Macroprudential Policy in Europe." *Journal of Money, Credit and Banking* 47, no. S1: 69–80.
- . 2018. "Policy Considerations on the Prolonged Period of Low Interest Rates: A Euro Area Perspective," panel remarks at 62nd Federal Reserve Bank of Boston Annual Economics Conference, Boston, September 7. <https://www.bostonfed.org/-/media/Images/research-conference-2018/presentations/hartmann-presentation.pdf>.
- Hartmann, P., H. Huang, and D. Schoenmaker, eds. 2018. *The Changing Fortunes of Central Banking*. Cambridge University Press.
- Hartmann, Philipp, Angela Maddaloni, and Simone Manganelli. 2003. "The Euro-Area Financial System: Structure, Integration and Policy Initiatives." *Oxford Review of Economic Policy* 19, no. 1: 180–213.
- Hartmann, Philipp, Michele Manna, and Andrés Manzanares. 2001. "The Micro-structure of the Euro Money Market." *Journal of International Money and Finance* 20, no. 6: 895–48.
- Hartmann, Philipp, and Frank Smets. 2018. "The First Twenty Years of the European Central Bank: Financial Stability." ECB, in progress.
- Heider, F., M. Hoerova, and C. Holthausen. 2015. "Liquidity Hoarding and Interbank Market Rates: The Role of Counterparty Risk." *Journal of Financial Economics* 118, no. 2: 336–54.
- Heider, Florian, Farzad Saidi, and Glenn Schepens. 2018. "Life Below Zero: Bank Lending under Negative Policy Rates." ECB Working Paper 2173. Forthcoming in *Review of Financial Studies*. <https://www.ecb.europa.eu/pub/pdf/scpwp/ ECB.wp2173.en.pdf>.
- High-Level Group on Financial Supervision in the EU. 2009. "De Larosi re Report." Brussels, February 25. http://ec.europa.eu/economy_finance/publications/pages/publication14527_en.pdf.
- Hutchinson, John, and Frank Smets. 2017. "Monetary Policy in Uncertain Times: ECB Monetary Policy since June 2014." *Manchester School* 85, no. S2: e1–e15.
- Inoue, Atsushi, and Barbara Rossi. 2018. "The Effects of Conventional and Unconventional Monetary Policy on Exchange Rates." Economics Working Paper 1638. Barcelona: Universitat Pompeu Fabra. <https://econ-papers.upf.edu/papers/1638.pdf>.
- International Monetary Fund. 2018. *World Economic Outlook, April 2018: Cyclical Upswing, Structural Change*. Washington: International Monetary Fund. <https://www.imf.org/en/Publications/WEO/Issues/2018/03/20/world-economic-outlook-april-2018>.

- Issing, Otmar. 1999. "The Eurosystem: Transparent and Accountable or 'Willem in Euroland'." *Journal of Common Market Studies* 37, no. 3: 503–19.
- . 2003a. "Background Studies for the ECB's Evaluation of Its Monetary Policy Strategy." European Central Bank, Frankfurt. https://www.ecb.europa.eu/pub/pdf/other/monetarypolicystrategyreview_backgrounden.pdf.
- . 2003b. "Monetary and Financial Stability: Is There a Trade-Off?" Speech given at conference on monetary stability, financial stability, and the business cycle, Bank for International Settlements, Basel, March 28–29.
- . 2005. "The Role of Money in the Monetary Policy Strategy of the ECB." Speech given at ECB workshop on what central banks can learn from money and credit aggregates, Eltville, October 28. <https://www.ecb.europa.eu/press/key/date/2005/html/sp051028.en.html>.
- . 2008. *The Birth of the Euro*. Cambridge University Press.
- Issing, Otmar, Vítor Gaspar, Ignazio Angeloni, and Oreste Tristani. 2001. *Monetary Policy in the Euro Area: Strategy and Decision-Making at the European Central Bank*. Cambridge University Press.
- Issing, Otmar, Vítor Gaspar, Oreste Tristani, and David Vestin. 2005. *Imperfect Knowledge and Monetary Policy; The Stone Lectures in Economics*. Cambridge University Press.
- Jackson, Harriet. 2015. "The International Experience with Negative Policy Rates." Staff Discussion Paper 2015-13. Ottawa: Bank of Canada.
- Jasova, Martina, Caterina Mendicino, and Dominik Supera. 2018. "Rollover Risk and Bank Lending Behaviour: Evidence from Unconventional Central Bank Liquidity." Working paper. https://cepr.org/sites/default/files/40021_Supera_roll_over_risk_and_bank_lending_behavior.pdf.
- Joslin, Scott, Kenneth J. Singleton, and Haoxiang Zhu. 2011. "A New Perspective on Gaussian Dynamic Term Structure Models." *Review of Financial Studies* 24, no. 3: 926–70.
- Joyce, Michael, David Miles, Andrew Scott, and Dimitri Vayano. 2012. "Quantitative Easing and Unconventional Monetary Policy: An Introduction." *Economic Journal* 122: F271–F288.
- Kapetanios, George, Haroon Mumatz, Ibrahim Stevens, and Konstantinos Theodoridis. 2012. "Assessing the Economy-wide Effects of Quantitative Easing." *Economic Journal* 122, no. 564: F316–F347.
- Kiley, Michael T., and John M. Roberts. 2017. "Monetary Policy in a Low Interest Rate World." *Brookings Papers on Economic Activity*, Spring: 317–72.
- Kontogeorgos, Georgios, and Kyriacos Lambrias. 2018. "An Analysis of the ECB/Eurosystem Projections." Unpublished paper, European Central Bank.
- Kortela, Tomi. 2016. "A Shadow Rate Model with Time-Varying Lower Bound of Interest Rates." Research Discussion Paper 19/2016. Helsinki: Bank of Finland. https://helda.helsinki.fi/bof/bitstream/handle/123456789/14300/BoF_DP_1619.pdf?sequence=1.
- Krippner, Leo. 2015. *Zero Lower Bound Term Structure Modeling: A Practitioner's Guide*. New York: Palgrave Macmillan.

- Krishnamurthy, Arvind, Stefan Nagel, and Annette Vissing-Jorgensen. 2018. "ECB Policies Involving Government Bond Purchases: Impact and Channels." *Review of Finance* 22, no. 1: 1–44.
- Krugman, Paul R. 2014. "Inflation Targets Reconsidered." In *Proceedings of the ECB Forum on Central Banking Monetary Policy in a Changing Financial Landscape*. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/sintra/ecb.forumcentbank201410.en.pdf>.
- Lane, Philip. 2015. "International Financial Flows and the Eurozone Crisis." Vox-EU article, September 7. <https://voxeu.org/article/international-financial-flows-and-eurozone-crisis>.
- Lemke, Wolfgang, and Andreea Liliانا Vladu. 2017. "Below the Zero Lower Bound: A Shadow-Rate Term Structure Model for the Euro Area." Working Paper 1991. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp1991.en.pdf?29fac2eb9ccb9e00ef5fe6a910c8dc93>.
- Lenza, Michele, Huw Pill, and Lucrezia Reichlin. 2010. "Monetary Policy in Exceptional Times." *Economic Policy* 25, no. 62: 295–339.
- Lyziak, Tomasz, and Maritta Paloviita. 2017. "Anchoring of Inflation Expectations in the Euro Area: Recent Evidence Based on Survey Data." *European Journal of Political Economy* 46: 52–73.
- Mackowiak, B., F. Mongelli, G. Noblet, and F. Smets, eds. 2008. "The Euro at Ten: Lessons and Challenges." Paper presented at Fifth ECB Central Banking Conference, Frankfurt, November 13–14.
- Martin, Philippe, and Thomas Philippon. 2017. "Inspecting the Mechanism: Leverage and the Great Recession in the Eurozone." *American Economic Review* 107, no. 7: 1904–37.
- Martínez Pagés, Jorge, and Antonio Millaruelo. 2016. "The Recent Application of Negative Policy Interest Rates in the Euro Area and in Other Economies: Rationale and Preliminary Evidence on Their Effects." *Economic Bulletin* (Bank of Spain), July–August, 3–13.
- Mouabbi, Sarah, and Jean-Guillaume Sahuc. 2017. "Evaluating the Macroeconomic Effects of the ECB's Unconventional Monetary Policies." Working paper, Centre for Economic Policy Research. https://cepr.org/sites/default/files/40006_MOUABBI%20-%20Evaluating%20the%20Macroeconomic%20Effectcs%20of%20the%20ECBs%20Unconventional%20Monetary%20Policies.pdf.
- Natoli, Filippo, and Laura Sigalotti. 2018. "Tail Co-movement in Inflation Expectations as an Indicator of Anchoring." *International Journal of Central Banking* 14, no. 1: 35–71.
- Nautz, Dieter, and Jörg Oechssler. 2006. "Overbidding in Fixed-Rate Tenders: An Empirical Assessment of Alternative Explanations." *European Economic Review* 50, no. 3: 631–46.
- Neri, S., and S. Siviero. 2019. "The Non-Standard Monetary Policy Measures of the ECB: Motivations, Effectiveness and Risks." Occasional Paper 486. Rome: Banca d'Italia. <http://www.bancaditalia.it/pubblicazioni/qef/2019-0486/index.html?com.dotmarketing.htmlpage.language=1>.
- Nessén, Marianne, and David Vestin. 2005. "Average Inflation Targeting." *Journal of Money, Credit and Banking* 37, no. 5: 837–63.

- Neumann, Manfred J. M. 2010. "Some Observations on the ECB's Monetary Policy." In *The Euro: The First Decade*, edited by Marco Buti, Servaas Deroose, Vítor Gaspar, and João Nogueira Martins. Cambridge University Press.
- Nickell, Stephen. 2006. "Discussion on 'European Monetary Union: The Dark Sides of a Major Success' by Charles Wyplosz." *Economic Policy* 21, no. 46. <https://academic.oup.com/economicpolicy/article/21/46/208/2918717>.
- Organization for Economic Cooperation and Development. 2018. *Economic Outlook*, May. <https://www.oecd.org/eco/outlook/General-assessment-of-the-macro-economic-situation-november-2018-OECD-economic-outlook-chapter.pdf>.
- Orphanides, Athanasios. 2003. "Historical Monetary Policy Analysis and the Taylor Rule." *Journal of Monetary Economics* 50, no. 5: 983–1022.
- . 2006. "The Road to Price Stability." *American Economic Review* 96, no. 2: 178–81.
- Orphanides, Athanasios, and Volker Wieland. 2013. "Complexity and Monetary Policy." *International Journal of Central Banking* 9, issue S1: 167–203.
- Orphanides, Athanasios, and John C. Williams. 2005. "Imperfect Knowledge, Inflation Expectations and Monetary Policy." In *The Inflation Targeting Debate*, edited by B. Bernanke and M. Woodford. University of Chicago Press.
- . 2008. "Learning, Expectations Formation, and the Pitfalls of Optimal Control Monetary Policy." *Journal of Monetary Economics* 55, issue S: S80–S96.
- Paloviita, Maritta, Markus Haavio, Pirkka Jalasjoki, and Juha Kilponen. 2017. "What Does 'Below, but Close to, Two Percent' Mean? Assessing the ECB's Reaction Function with Real Time Data." Research Discussion Paper 29. Helsinki: Bank of Finland. https://helda.helsinki.fi/bof/bitstream/handle/123456789/14932/BoF_DP_1729.pdf?sequence=1&isAllowed=y.
- Paludkiewicz, Karol. 2018. "Unconventional Monetary Policy, Bank Lending, and Security Holdings: The Yield-Induced Portfolio Rebalancing Channel." Discussion Paper 22/2018. Frankfurt: Deutsche Bundesbank. <https://www.econstor.eu/bitstream/10419/180675/1/1027071376.pdf>.
- Papademos, Lucas. 2006. "The Role of Money in the Conduct of Monetary Policy." In *Proceedings of the Fourth ECB Central Banking Conference: The Role of Money—Money and Monetary Policy in the Twenty-First Century*, edited by Andreas Beyer and Lucrezia Reichlin. Frankfurt: European Central Bank.
- Papademos, Lucas D., and Jürgen Stark, eds. 2010. *Enhancing Monetary Analysis*. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/other/enhancingmonetaryanalysis2010en.pdf?13ffb3e72c01de37830173b8a5b8217>.
- Peersman, Gert. 2011. "Macroeconomic Effects of Unconventional Monetary Policy in the Euro Area." Working Paper 1397. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp1397.pdf?99ef78cb1a613f60d8716d0db59b585>.
- Pesaran, M. Hashem, and Ron P. Smith. 2016. "Counterfactual Analysis in Macroeconometrics: An Empirical Investigation into the Effects of Quantitative Easing." *Research in Economics* 70, no. 2: 262–80.

- Pill, Huw, and Lucrezia Reichlin. 2015. "Exceptional Policies for Exceptional Times: The ECB's Response to the Rolling Crises of the Euro Area." In *Handbook of the Economics of European Integration*, edited by Harald Badinger and Volker Nitsch. New York: Routledge.
- Pill, Huw, and Frank Smets. 2013. "Monetary Policy Frameworks after the Great Financial Crisis." In *The Great Recessions*, edited by Jacob Braude, Zvi Eckstein, Stanley Fischer, and Karnit Flug. MIT Press.
- Poole, William. 1970. "Optimal Choice of Monetary Policy Instruments in a Simple Stochastic Macro Model." *Quarterly Journal of Economics* 84, no. 2: 197–216.
- Praet, Peter. 2013. "Forward Guidance and the ECB." VoxEU article, August 6. <https://www.ecb.europa.eu/press/key/date/2013/html/sp130806.en.html>.
- . 2017. "Maintaining Price Stability with Unconventional Monetary Policy Measures." Speech given at MMF Monetary and Financial Policy Conference, London, October 2.
- . 2018. "Delivering on Our Mandate: 20 Years of ECB Monetary Policy." Speech given at an event organized by ECB's representative office in Brussels, July 5. <https://www.ecb.europa.eu/press/key/date/2018/html/ecb.sp180705.en.html>.
- Quint, Dominic, and Oreste Tristani. 2018. "Liquidity Provision as a Monetary Policy Tool: The ECB's Non-Standard Measures after the Financial Crisis." *Journal of International Money and Finance* 80: 15–34.
- Reichlin, Lucrezia. 2014. "Monetary Policy and Banks in the Euro Area: The Tale of Two Crises." *Journal of Macroeconomics* 39: 287–400.
- Reifschneider, David, John C. Williams, Christopher A. Sims, and John B. Taylor. 2000. "Three Lessons for Monetary Policy in a Low-Inflation Era." *Journal of Money, Credit and Banking* 32, no. 4: 936–66.
- Rostagno, Massimo, Ulrich Bindseil, Annette Kamps, Wolfgang Lemke, Tomohiro Sugo, and Thomas Vlassopoulos. 2016. "Breaking Through the Zero Line: The ECB's Negative Interest Rate Policy." Presentation given at Brookings Institution, Washington, June 6.
- Shambaugh, Jay C. 2012. "The Euro's Three Crises." *Brookings Papers on Economic Activity*, Spring: 157–211.
- Smets, Frank. 2003. "Maintaining Price Stability: How Long Is the Medium Term?" *Journal of Monetary Economics* 50, no. 6: 1293–1309.
- . 2010. "Comment on Chapters 6 and 7." In *The Euro: The First Decade*, edited by Marco Buti, Servaas Deroose, Vítor Gaspar, and João Nogueira Martins. Cambridge University Press.
- . 2014. "Financial Stability and Monetary Policy: How Closely Interlinked?" *International Journal of Central Banking* 10, no. 2: 263–300.
- Stolz, Stéphanie Marie, and Michael Wedow. 2010. "Extraordinary Measures in Extraordinary Times: Public Measures in Support of the Financial Sector in the EU and the United States." ECB Occasional Paper 117. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/scpops/ecbocp117.pdf>.
- Surico, Paolo. 2007. "The Monetary Policy of the European Central Bank." *Scandinavian Journal of Economics* 109, no. 1: 115–35.

- Svensson, Lars E. O. 1999a. "Monetary Policy Issues for the Eurosystem." *Carnegie-Rochester Conference Series on Public Policy* 51: 79–136.
- . 1999b. "Price-Level Targeting versus Inflation Targeting: A Free Lunch?" *Journal of Money, Credit and Banking* 31, no. 3: 277–95.
- . 2003. "In the Right Direction, but Not Enough: The Modification of the Monetary Policy Strategy of the ECB." Briefing paper for Committee on Economic and Monetary Affairs of European Parliament. <http://citeseerx.ist.psu.edu/viewdoc/download;jsessionid=7525C6653D727EA72718192ECEFA233D?doi=10.1.1.15.3587&rep=rep1&type=pdf>.
- Swanson, Eric T. 2017. "Measuring the Effects of Federal Reserve Forward Guidance and Asset Purchases on Financial Markets." NBER Working Paper 23311. Cambridge, Mass.: National Bureau of Economic Research. <https://www.nber.org/papers/w23311.pdf>.
- . 2018. "The Federal Reserve Is Not Very Constrained by the Lower Bound on Nominal Interest Rates." NBER Working Paper 25123. Cambridge, Mass.: National Bureau of Economic Research. <https://www.nber.org/papers/w25123.pdf>. Also in this issue of *Brookings Papers on Economic Activity*.
- Szczerbowicz, U. 2015. "The ECB's Unconventional Monetary Policies: Have They Lowered Market Borrowing Costs for Banks and Governments?" *International Journal of Central Banking* 11, no. 4: 91–127.
- Szörfi, Béla, and Máté Tóth. 2018. "Measures of Slack in the Euro Area." *Economic Bulletin* 3. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/ecbu/eb201803.en.pdf?a002c6f56b79f07df072d758fca1025b>.
- Task Force on the Use of Monetary Policy Instruments. 2018. "The Use of the Eurosystem's Monetary Policy Instruments and Its Monetary Policy Implementation Framework Q2 2016–Q4 2017." Occasional Paper 209. Frankfurt: European Central Bank. <https://www.ecb.europa.eu/pub/pdf/scpops/ecb.op209.en.pdf>.
- Taylor, John B. 2007. "Housing and Monetary Policy." NBER Working Paper 13682. Cambridge, Mass.: National Bureau of Economic Research.
- Trebesch, Christoph, and Jeromin Zettelmeyer. 2018. "ECB Interventions in Distressed Sovereign Debt Markets: The Case of Greek Bonds." *IMF Economic Review* 66, no. 2: 287–332.
- Trichet, Jean-Claude. 2003a. "Asset Price Bubbles and Their Implications for Monetary Policy and Financial Stability." In *Asset Price Bubbles: The Implications for Monetary, Regulatory and International Policies*, edited by William C. Hunter, George G. Kaufman, and Michael Pomerleano. MIT Press.
- . 2003b. "The ECB's Monetary Policy Strategy After the Evaluation and Clarification of May 2003." Speech given at Center for Financial Studies, November 20. <https://www.ecb.europa.eu/press/key/date/2003/html/sp031120.en.html>.
- . 2005. "Asset Price Bubbles and Monetary Policy." Mas Lecture in Singapore, June 8. <https://www.bis.org/review/r050614d.pdf>.
- . 2006. "Structural Reforms in Europe." Speech given at forum of Organization for Economic Cooperation and Development, Paris, May 22.

- . 2008. “Some Lessons from the Financial Market Correction.” Speech given at European Banker of the Year 2007 award ceremony, Frankfurt, September 30.
- . 2009. “The ECB’s Enhanced Credit Support.” Speech given at University of Munich, Munich, September 4.
- Tucker, Paul. 2018. *Unelected Power: The Quest for Legitimacy in Central Banking and the Regulatory State*. Princeton University Press.
- Van Bakkum, Sjoerd, Marc Gabarro, and Rustom M. Irani. 2018. “Does a Larger Menu Increase Appetite? Collateral Eligibility and Credit Supply.” *Review of Financial Studies* 31, no. 3: 943–79.
- Van Zandweghe, Willem. 2015. “Monetary Policy Shocks and Aggregate Supply.” *Economic Review* Q3: 31–56.
- Weidmann, Jens. 2015. “Responsibility and Liability in a Monetary Union.” Speech given at conference sponsored by Consiglio Regionale del Veneto, Venice, February 5. <https://www.bis.org/review/r150205f.htm>.
- Wilhelmsen, Björn-Roger, and Andrea Zaghini. 2011. “Monetary Policy Predictability in the Euro Area: An International Comparison.” *Applied Economics* 43, no. 20: 2533–44.
- Woodford, Michael. 2006. “Does a ‘Two-Pillar Phillips Curve’ Justify a Two-Pillar Monetary Policy Strategy?” In *Proceedings of the Fourth ECB Central Banking Conference: The Role of Money—Money and Monetary Policy in the Twenty-First Century*, edited by Andreas Beyer and Lucrezia Reichlin. Frankfurt: European Central Bank.
- Wu, Jing Cynthia, and Fan Dora Xia. 2017. “Time-Varying Lower Bound of Interests in Europe.” Research Paper 17-06. Chicago: Chicago Booth. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2946239.
- Wu, Jing Cynthia, and Ji Zhang. 2017. “A Shadow Rate New Keynesian Model.” NBER Working Paper 22856. Cambridge, Mass.: National Bureau of Economic Research. <https://www.nber.org/papers/w22856.pdf>.
- Wyplosz, Charles. 2001. “Do We Know How Low Should Inflation Be?” In *Proceedings of the First ECB Central Banking Conference: “Why Price Stability?”* edited by Alicia García Herrero, Vítor Gaspar, Lex Hoogduin, Julian Morgan, and Bernhard Winkler. Frankfurt: European Central Bank.
- Yellen, Janet. 2012. “Revolution and Evolution in Central Bank Communications.” Speech given at Haas School of Business, University of California, Berkeley, November 13.
- Zettelmeyer, J., C. Trebesch, and M. Gulati. 2013. “The Greek Debt Restructuring: An Autopsy.” *Economic Policy* 28, no. 75: 513–63.

Comments and Discussion

COMMENT BY

LORENZO BINI SMAGHI My discussion of the very interesting paper by Philipp Hartmann and Frank Smets on the first 20 years of the European Central Bank (ECB) is inevitably influenced by my professional and academic background. First, I was a member of the Executive Board and of the Governing Council of the ECB between June 1, 2005, and December 31, 2011. During that period, I voted in favor of all the decisions that were made by the ECB. Second, I studied monetary theory and policy at the University of Chicago in the early 1980s, and thus have been influenced by Milton Friedman's writings, in particular his 1967 AEA presidential address on the role of monetary policy, in particular when he states that

the first and most important lesson that history teaches about what monetary policy can do . . . is that monetary policy can prevent money itself from being a major source of economic disturbance. This sounds like a negative proposition: avoid major mistakes." (Friedman 1968, 12)

Avoiding making big mistakes is what haunted me during my ECB term. And that is the approach that I would like to take in discussing the paper by Hartmann and Smets.

The biggest mistake that any central bank wants to avoid is to miss its main objective, which is price stability. As Hartmann and Smets confirm in their paper, over the last 20 years, the average rate of inflation in the euro area has been about 1.7 percent, which seems to be within the range of what could be an arithmetic definition of price stability.

I broadly share Hartmann and Smets's conclusions that, overall, the ECB has fulfilled its mandate. It has acquired a high level of credibility as a central bank, in particularly difficult times. Inflation expectations have been firmly anchored. However, precisely because of the high credibility

gained on the ground, the ECB should be slightly more open to assess its performance in response to the various criticisms that have been raised by academics, markets participants, and the public opinion over the last few years. Here, I make a few suggestions concerning the issues that should stimulate further research.

THE DEFINITION OF PRICE STABILITY The ECB has never given a precise numerical definition of price stability. Hartmann and Smets quote Otmar Issing at a 2003 press conference, stating that “a narrow range between roughly 1.7% and 1.9%” should be considered as being consistent with price stability (ECB 2003b). This reminds me of the same sort of calculation that Jean-Claude Trichet was continuously making during his term. As he left the ECB, at the end of October 2011, he was proud to mention—with a certain humor, however—that since the start of the euro, inflation had been on average 1.99 percent, and thus—at least in his view—fully in line with the objective of price stability (Trichet 2011). How bewildered would he now be to learn that—with the benefit of hindsight—1.99 percent was in fact too high, being outside the range of 1.7 to 1.9 percent mentioned by Hartmann and Smets and, in fact, “too close” to 2 percent!

Let us face it, the word “below”—inserted just before 2 percent in the definition of price stability—was added, in my opinion, with a view to emulate the Bundesbank, given that some members of the Governing Council thought at the time that a symmetric target would lead markets to think that the ECB would be excessively tolerant with an inflation rate above 2 percent. And perhaps it was also to reassure the German public that the ECB Governing Council would be as tough as the Bundesbank. In fact, the evidence over the last 20 years shows that the ECB’s performance has been much closer to target than the Bundesbank was in the previous decades, albeit in a different inflationary environment.

The words “close to” were added in 2003, at the time of a review of the monetary policy strategy to avoid the impression that the ECB would tolerate deflation. Not doing like all other central banks—that is, providing a symmetric target of about 2 percent—might have been “prudent” 20 years ago. It is doubtful that it remains appropriate after 20 years of experience. All in all, having a qualitative—rather than a quantitative—definition of price stability has not helped the ECB, and has not even shielded it from criticisms, including those by Otmar Issing himself, who recently stated that an inflation rate of 1 percent was perfectly consistent with the “close to 2 percent” (ECB 2003a, 79), suggesting that he himself had forgotten about his 2003 range. In fact, Hartmann and Smets confirm that the ECB’s reaction function over the last 20 years has been consistent

with a symmetric inflation target. To conclude, a first lesson that could be drawn from the evidence is that the time may be ripe to move to an explicit 2 percent target, which would be not only more credible but also more transparent.

THE LEADS AND LAGS OF MONETARY POLICY Hartmann and Smets's judgment that the ECB did not make big mistakes, having achieved an inflation rate of close to 2 percent, is based on the average performance over 20 years of monetary union. Central banks cannot be held accountable for keeping inflation at target month after month, but over a certain period of time, given that monetary policy operates with long and variable lags. It is not by chance that the words "over the medium term" are an integral part of the ECB's definition of price stability (ECB 2003b, 79). What is thus the appropriate time period for assessing whether inflation has been on target? One year may be too short, but for sure 20 years is too long. The lags with which monetary policy instruments hit their objective range between 18 to 36 months. This is why central banks make forecasts over such a horizon. If this is an appropriate criterion, we may want to test the hypothesis whether the ECB failed to meet its objective between 2013 and 2018. During these six years, as can be seen from the figures in Hartmann and Smets's paper, inflation—both headline and core—is, for the most part, below the range of 1.7 to 1.9 percent. It is thus legitimate to investigate the reasons for such an underperformance, which incidentally is not unique to the ECB. The key question is whether, during this period, monetary policy has been behind the curve—in other words, has been reacting too little too late.

THE FINANCIAL CRISIS AND MONETARY POLICY The financial crisis hit the monetary union after less than 10 years of its young life. The ECB reacted forcefully, but in an environment where it did not always have all the relevant information to fully appreciate the situation or the tools to calibrate its response. Here I point to a few examples, which may deserve greater analysis and a better understanding.

After August 2007. In August 2007, as the money market stopped functioning properly, the ECB intervened by injecting more than €90 billion in one day, accommodating all the demand for liquidity from its counterparties (ECB 2007). In the following months, the money market continued to malfunction, especially at 3-month maturity, which is a key reference rate. The ECB nevertheless kept its tender procedures unchanged, in spite of the growing divergence between market and policy rates. It decided to move to fixed-rates/full-allotment procedures only in October 2008, long after Lehmann Brothers' crash.

This—in my view—might have been a mistake, which derived from a less than complete understanding of the health of the banking system. In that period, the ECB was able to gather information on the euro zone banking industry only indirectly, through the national bank supervisors. This was a large source of inefficiency, because local supervisors had the incentive to underreport the problems of their financial sector.

The July 2008 rate hike. In June 2008, the ECB decided to call for vigilance, which was the catchword for announcing an interest rate rise at its meeting the next month. With the benefit of hindsight, that decision may look like a mistake, and has been widely criticized by observers. The crisis erupted two months later, and the ECB had to rescind its decision, cutting rates in October 2008. Figure 22 in Hartmann and Smets' paper shows that such a decision was not warranted, based on a Taylor rule.

Although no single interest rate decision can constitute a major policy mistake, it is useful to clarify the reasoning underlying this decision. First, the euro zone's headline inflation had been above 3 percent for several months, and inflation expectations were at risk of dis-anchoring. Credit growth was still strong. Conversely, core inflation was still hovering around 2 percent, and the economy was showing signs of slowing down, after a buoyant first quarter. The ECB clearly did not read the signals coming from the real economy, which was decelerating rapidly from the middle of the second quarter. Part of the reason for such a misreading derived from the fact that at the time, the ECB had to rely mainly on national central banks to assess short-term cyclical developments.

The 2008 decision—seen in retrospect—also shows the excessive emphasis that the ECB put on the monetary pillar of its strategy. I will not elaborate further on the two-pillar strategy, an issue extensively discussed by Hartmann and Smets. However, the time may have come to reassess it. The emphasis on monetary indicators, in spite of the lack of stability in the demand for money in the euro zone, may have been a price to be paid at the start of the monetary union, but has become less justified.

The 2011 interest rate hikes. In 2011, the ECB decided to raise rates twice, as announced in March and June. These hikes were reversed after a few months, as the financial crisis deepened. There is a large debate in the literature as to whether these decisions are not to be put in the “big mistake” category, because they may have made the crisis even worse.

My personal judgment is that while the decision announced in March was not a big mistake, the second one might instead have been one. Looking at the data available in the spring of 2011, the euro zone was recovering quite strongly and inflation was moving again, toward 3 percent. Under

these circumstances, a hike of 25 basis points could have been justified. Figure 24 in Hartmann and Smets's paper suggests that a rate hike could have been appropriate even earlier. Other central banks had also raised rates.

At the time of the second hike, the situation had changed substantially, not so much with respect to the real economy but to the risks to financial stability in the euro zone. The restructuring of Greek debt became a clear option at the end of April. Long-term rates started rising gradually but steadily in most peripheral countries. The ECB was opposed to debt restructuring, because of the potential contagion to other countries. It nevertheless made the decision to hike its policy rates, in the expectation that it would not have an impact on the financial situation. It is difficult to assess the extent to which the decision exacerbated the worsening financial conditions. To say the least, it did not help.

SMP versus "whatever it takes." An issue for discussion is why did the ECB wait for more than two years to state that it would do whatever it takes to ensure the stability of the euro and to avoid having a country driven out of the euro against its will. The answer is complex. In 2010, when the ECB started the Securities Market Programme (SMP), the crisis appeared to be circumscribed to three countries; but in 2012, it became systemic. Second, in 2010 the European Stability Mechanism had not yet been established, and the procedure for setting the conditionality for the countries requesting financial support was not yet defined. Third, the institutional framework underlying fiscal discipline had been weakened, especially after the disclosure of Greek budgetary overshooting, thus putting at risk the boundaries between fiscal and monetary policy. The Fiscal Compact, which was adopted in 2012, created the conditions for protecting the ECB from the risk of fiscal dominance.

Overall, the conditions for adopting the Outright Monetary Transactions (OMT) program were not yet mature in 2010. However, the temporary and limited nature of the SMP, which was periodically conveyed to the markets, over time became factors in reducing its effectiveness. One of the OMT's key features is precisely its unlimited nature, which is a fundamental characteristic of a fiat money system, whereby the central bank can create unlimited amounts of central bank money to accommodate demand, and thus stop any panic. This is why the OMT is still untested. If it appeared at any time that there were limits to the OMT, markets would immediately test it. The fact that the SMP was declared to be limited and temporary reduced its effectiveness. Markets interpreted this limit as a sign of the ECB's unwillingness to fully implement the program, and they periodically tested the ECB's resolution.

Negative rates versus quantitative easing. In the spring of 2014, the ECB decided to lower its deposit rates into negative territory. About one year later, it decided to also start quantitative easing (QE). It is fair to ask whether this sequence was right. The decision to cut rates was probably made in the expectation that it would be a sufficiently bold move to allow the ECB to avoid starting QE, which was politically controversial. With the benefit of hindsight, it is legitimate to ask whether and to what extent the ECB underestimated the impact of the financial crisis on the real economy, starting with the recession in 2012–13 and then with the slow pace of the recovery. It also appears that the ECB may have underestimated the extent to which the banking part of the transmission channel of monetary policy was clogged, partly due to the fact that banking union really started only at the end of 2014, when the SSM took full responsibility. It looks like a coincidence that QE started in May 2015, only six months after the start of the banking union.

The argument against QE in Europe was largely based on the assumption that though in the U.S. monetary policy operated mainly through markets, in the euro zone monetary policy operated through the banking system. However, at the zero lower bound, or in negative territory, a fixed-rate/full-allotment tender procedure makes the supply of money entirely demand determined. As Paul Samuelson (1948, 353–54) would remind us, “You can lead a horse to water, but you can’t make him drink.”

The reasons why banks did not drink may not have been fully perceived and understood. To be sure, the Target2 data were providing confusing evidence. Balances increased during the crisis, until July 2012, and then decreased sharply after the “whatever it takes” statement (Draghi 2012). At the time, this was considered as a signal that financial tensions were easing, but it also revealed that the supply of central bank money was remaining stable, as the economy was getting out of the slump, signaling that monetary policy was too restrictive. In fact, the size of the ECB’s balance sheet started rising only when QE was implemented.

These issues should be thoroughly discussed to understand whether indeed, as some may suggest, monetary policy might have reacted too slowly during the crisis and may have maintained an excessively restrictive stance during the recovery. To be sure, these policy decisions were not uncontroversial. However, with the benefit of hindsight, those who thought that monetary policy was being too expansionary and was putting price stability at risk were proved consistently wrong.

OVERSTEPPING THE MANDATE Throughout the global financial crisis, central banks were criticized for having come very close, or even overstepped,

their mandate. The ECB was not immune from this criticism, which came from several sides and different perspectives. The most publicized is the compatibility with the Lisbon Treaty (EU 2007) of the SMP, the OMT and QE—all of which imply the purchase of government bonds. The compatibility of these policies with the ECB's independence and with the prohibition of monetary financing has always been relatively clear, at least from an economic point of view—in particular, because these instruments have been adopted by other independent central banks—and subsequently from a legal point of view.

Other controversial issues have received less attention, but are at least as important for the conduct of monetary policy and the integrity of the monetary union. I only mention three.

The collateral framework. In 2006, the ECB revised its collateral framework to set a minimum standard rating for the assets posted as collateral for monetary policy operations. At that time, the issue was not considered so relevant, because all countries had a rating much above the threshold. The threshold was set in such a way as to make sure that all government bonds could be used as collateral. The decision was not without controversy within the Governing Council. Some raised the issue of arbitrariness and the risk of creating a kink effect that could destabilize financial markets. The ECB is, to my knowledge, the only central bank that may refuse government bonds as collateral and resorts to external rating to set haircuts. This policy produces procyclical effects and may add to financial instability.

Such a policy seems to be based on a priority given to the quality of the balance sheet, and the need to avoid losses to the central bank, at the expense of other priorities that do not concern the ECB directly. However, central banks do not have the maximization of profits as an objective, nor the minimization of losses. The Lisbon Treaty states that, without prejudice to the primary objective, the ECB should support the general economic policies of the Community—as laid down in Article 2 of the treaty, which states that the task of the Community is to “promot[e] throughout the Community a harmonious, balanced and sustainable development of economic activities, a high level of employment and of social protection, equality between men and women, sustainable and non-inflationary growth, a high degree of competitiveness and convergence of economic performance, a high level of protection and improvement in the quality of the environment, the raising of the standard of living and quality of life, and economic and social cohesion and solidarity among Member States” (EC 2006). To sum up, it may be time to revise the ECB's collateral framework, to avoid it being part of the problem rather than the solution.

Emergency Liquidity Assistance. At the start of the monetary union, the Emergency Liquidity Assistance (ELA) policy has been designed for banks that, though solvent, do not have adequate collateral to apply for the ECB's regular monetary policy operations. The ELA policy foresees that the liquidity is provided by the national central bank, with collateral, and thus the risk, posted with that central bank, which are not shared within the Eurosystem. The ECB can only revoke the decision on the basis of a special procedure (ECB 2017). The reason is that the responsibility for assessing whether the bank is solvent was in the hands of the national supervisors. However, with the creation of the Single Supervisory Mechanism at the ECB, the responsibility for declaring a bank solvent has been centralized. It thus appears logical that the risk, and the decision to grant ELA, become centralized.

One specific instance in which the ECB has been strongly criticized is in dealing with the Greek crisis, in particular on the eve of the June 2015 referendum. The ECB limited Greek banks' access to ELA, in a way that might have fueled a run on the banks and caused a loss of confidence. It was obviously difficult for the ECB to consider Greek banks on the same level as other banks a few days before a referendum that was calling Greece's membership in the euro zone into question. Conversely, the ECB's decision had a direct effect on Greece's financial situation, which may not have been fully in line with the mandate of the ECB itself.

Participation in the Troika. Since its inception, the ECB has been part of the Troika—together with the European Commission and the International Monetary Fund—which is in charge of the technical discussions underlying the definition and monitoring of the adjustment program. This role was particularly important with respect to the need to have adequate information about the banking system and making sure that the adjustment program foresaw an adequate capitalization. However, such a role is quite peculiar for a central bank, given that it gets into policies that are not of its competence. There is a risk of getting involved in political discussions, and thus losing degrees of freedom. Now that the banking union has transferred supervisory functions at the ECB, there is much less need for it to participate in the Troika.

CONCLUSION To assess the ECB's performance over the last 20 years on the basis of its primary objective, which is price stability, is necessary but probably not sufficient. The ECB is one of the European Union's institutions, and cannot be immune from the economic, social, and political developments that affect the Union. Although the ECB has demonstrated in a few years that it is an effective and efficient central bank, it should

not fear that its credibility can be undermined by an open discussion of its key decisions over the years. Like other EU institutions, the ECB has been affected by a negative confidence trend, as reflected by the Eurobarometer polls. Although the last Eurobarometer shows that favorable opinions about the euro have gone back above precrisis levels and reached a peak (74 percent), and those against the euro have fallen to a minimum (20 percent), the share of respondents who “trust the ECB” has fallen below those that do not trust it (42 percent against 45 percent; it was 46 percent against 27 percent before the global financial crisis) (EC 2018). Changing these opinions is certainly a challenge for the years to come. Hartmann and Smets’s paper is a good start in this endeavor, but only a start.

With respect to the issues that may need be reassessed, 20 years later and in light of experience, I suggest these: (1) the definition of price stability, symmetric at 2 percent; (2) the further downgrading or elimination of the monetary pillar; (3) centralization of the ELA policy; (4) a review of collateral policy; and (5) an exit from the Troika.

REFERENCES FOR THE BINI SMAGHI COMMENT

- Draghi, Mario. 2012. “Verbatim of the Remarks Made by Mario Draghi.” Speech at Global Investment Conference, London, July 26. <https://www.ecb.europa.eu/press/key/date/2012/html/sp120726.en.html>.
- EC (European Parliament and Council of the European Union). 2006. “Directive 2006/123/EC of the European Parliament and of the Council.” <https://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:32006L0123:en:HTML#d1e39-36-1>.
- . 2018. “Eurobarometer: Support for the Euro Steady at All-Time High Levels November.” https://ec.europa.eu/info/news/eurobarometer-2018-nov-20_en.
- ECB (European Central Bank). 2003a. “The Outcome of the ECB’s Evaluation of Its Monetary Strategy.” *ECB Monthly Bulletin*, June, 79–92. https://www.ecb.europa.eu/pub/pdf/other/pp79_92_mb200306en.pdf?e366490ebb678c0d82b6d62a85e10d44.
- . 2003b. “Press Seminar on the Evaluation of the ECB’s Monetary Policy Strategy.” Press release, May 8. https://www.ecb.europa.eu/press/pressconf/2003/html/is030508_1.en.html.
- . 2007. *Monthly Economic Bulletin*, September.
- . 2017. “Agreement on Emergency Liquidity Assistance.” https://www.ecb.europa.eu/pub/pdf/other/Agreement_on_emergency_liquidity_assistance_20170517.en.pdf.
- EU (European Union). 2007. “Treaty of Lisbon Amending the Treaty on European Union and the Treaty Establishing the European Community, Signed at Lisbon, 13 December 2007.” <https://eur-lex.europa.eu/eli/treaty/lis/sign?eliuri=eli:treaty:lis:sign>.

Friedman, Milton. 1968. "The Role of Monetary Policy." Presidential Address delivered at 80th Annual Meeting of American Economic Association, December 29. Published in the *American Economic Review* 58, no. 1: 1–17.

Samuelson, Paul. 1948. *Economics*. New York: McGraw-Hill.

Trichet, Jean Claude. 2011. "Remarks at the Farewell Event, Frankfurt, October 19." <https://www.ecb.europa.eu/press/key/date/2011/html/sp111019.en.html>.

COMMENT BY

LUCREZIA REICHLIN The paper by Philipp Hartmann and Frank Smets provides a useful narrative of the first 20 years of the European Central Bank (ECB) and an assessment of its performance. Overall, the authors' assessment of the ECB's record is very positive. There are four main conclusions:

1. The ECB has been successful in respecting its price stability mandate throughout its 20 years of history.
2. Its two-pillar strategy and definition of price stability target have served it well.
3. The tools associated with its strategy have evolved over time in a pragmatic way and responded successfully to the challenges of the global financial crisis.
4. Its operational framework revealed itself to be robust to the test of the worst crisis since World War II.

There is a lot to agree with in this assessment, and especially on the broad conclusion that, notwithstanding the global financial crisis and contrary to the expectation of many, the euro has emerged as one of the world's main currencies and the ECB—at least so far—has been a credible custodian of its value.

My own assessment is nevertheless more nuanced. In my view, the main question that should be answered, after 20 years of the life of the euro and 10 years after the global financial crisis, is whether the economic framework on which the European Economic and Monetary Union is based, and the ECB's central role in it as a central bank without a state, is adequate to face periods of particular financial and economic stress. The answer here is not straightforward, and the analysis of the crisis should give elements for reflection on necessary reform.

My remarks are organized in two sections. First, I discuss the ECB's nonstandard policies during the crisis. And second, I comment on interest rate policy during the same period. I base my remarks on my published work on the subject (in particular, Lenza, Pill, and Reichlin 2010; Pill and Reichlin 2014, 2016a, 2016b; and Reichlin 2014, 2018).

NONSTANDARD POLICIES, 2007–14 I analyze four distinct episodes of ECB action. Unlike Hartmann and Smets, I do not follow a chronological order but rather use four case studies to make my points.

The periods 2007–10 and 2011–12: two examples of nonstandard liquidity policies. The first symptoms of the liquidity crisis in the banking sector emerged in the euro zone in August 2007, with tensions in the money markets. The first phase of the crisis can be defined as one of a generalized counterparty risk that generated a significant increase in the demand for liquidity from the central bank by the banks, both fragile and healthy. With the collapse of Lehman Brothers in the United States in 2008, the nature of the crisis changed. A number of banks failed, and the global economy entered its worst recession since the 1930s. The interbank market effectively collapsed.

As I have argued elsewhere (Lenza, Pill, and Reichlin 2010; Pill and Reichlin 2014, 2016a), the response to this liquidity crisis can be considered a success. Hartmann and Smets agree with this view. The ECB acted aggressively and swiftly, thanks to an operating model that was fit to deal with this kind of crisis. As Tommaso Padoa-Schioppa (2004) had predicted in writing years before the crisis, the ECB's operating model was well prepared to act as a "lender of last resort" in cases of a widespread liquidity crisis, because it had both an open market transaction system that was well equipped to deal with systemic crises, and an instrument for providing emergency liquidity assistance to individual institutions. The ECB also had a head start over the leading central banks, including the Federal Reserve in the United States. It had a larger budget (in part because it remunerates bank reserves, a policy only adopted by the Fed after the crisis), and therefore had a greater capacity to absorb liquidity shocks. In addition, it started from a broader definition of eligible collateral in its operations with banks and accepted a broader category of institutions as counterparties in its operations (Pill and Reichlin 2016a). This enabled the ECB to adopt a systemic approach to the crisis right from the start, rather than have recourse to specific rescues. The bank's action at this stage respected the classic Bagehot's rule, according to which the central bank must act as a lender of last resort when counterparty risk blocks the entire system and therefore has an effect on both fragile and robust banks.

As a consequence of the refinancing operations in cooperation with banks with fixed-rate/full-allotment credit operations (in order to meet demand), the ECB's balance sheet increased in size, although the mechanism (and its motivation) was not the same as that implemented at the same time by

other central banks, such as the Fed and the Bank of England. As Huw Pill and I have observed (2016a), the action of the ECB at this stage should be interpreted as aimed at keeping the financial system and its infrastructure working by acting as an intermediary for transactions for which the market had stopped functioning as an intermediary, thus acting as a central counterparty of last resort. These policies need to be seen as complementary to the traditional policies of setting the Main Refinancing Operations interest rate. The motivation was different than that of using balance sheet policies as a substitute for interest rate policy when the latter reaches the zero lower bound. However, as with quantitative easing and credit easing, both the size of the balance sheet and the composition of its assets increased as a result. Quantitative works by Lenza, Pill, and Reichlin (2010) and by Giannone and others (2012) have shown the effectiveness of these policies in supporting lending and economic activity.

However, as the generalized liquidity crisis became a banking crisis and the insolvency of some institutions threatened the stability of the system, this clear distinction between liquidity policy and solvency was blurred. Here is where the limits of the euro area's governance became obvious.

It is interesting to analyze the difference between what happened in the period 2007–9 and in 2011. In late 2011, when Mario Draghi took over from Jean-Claude Trichet as president of the ECB, there was the risk of a new banking crisis. The issue facing the ECB was no longer one of a generalized liquidity drying out, but one of solvency. In this context, without the tools for a comprehensive approach to recapitalization, the ECB found itself as the only institution in the euro zone able to act across the monetary union with the power, if not to resolve the situation, at least to avoid the worst, and thus enable the euro-zone's governments and the European institutions to take the time to devise other solutions.

Against this background, Draghi announced a series of long-term refinancing operations (LTROs) in December 2011 and February 2012 (fixed-rate/full-allotment, 3-year refinancing operations). As for the LTROs adopted by Trichet in 2009, by means of these operations, the ECB became a centralized counterparty in the interbank market, but now for the longer term and therefore with more relevance for financing the banks and not just for managing liquidity. LTROs were also crucial in supporting the public sector at a time of great tensions in the sovereign debt market. In fact, with these measures, the banks were able to borrow funds from the ECB at a much lower rate and reinvest them in government bonds of peripheral countries that yielded much higher rates. In this way the banks not only

made profits but also supported the very market from which foreign investors had fled.

In this situation, the ECB acted as an intermediary for cross-border capital flows in an intra-euro zone market that, given the correlation between bank risk and country risk, was once again segmented by country. This phenomenon, known as the diabolic loop (see Brunnermeier and others 2016), consists of the fact that a country that has difficulties refinancing its debt puts pressure on its banks to purchase national government bonds, while a bank in crisis puts potential pressure on the public finances of its country if it is at risk of failure. The purchasing of their own country's sovereign bonds by banks was made possible by targeted LTROs—that is, loans to banks made over a time frame of up to four years at favorable terms, on the condition that the beneficiary institutions use the funds to provide credit to the real economy—which were introduced in June 2014 and again in March 2016. These targeted LTROs reinforced this correlation between bank risk and sovereign risk, which in turn created heterogeneity between the bank rates to customers, reducing the efficacy of Frankfurt's monetary policy. The ECB provided cheap financing to the banks, and the banks used it to buy sovereign debt (indirectly financing the sovereign) to use as collateral to obtain ECB financing. As a consequence, we saw a substitution in banks' balance sheets—from loans to the private sector to holdings of sovereign bonds.

This episode is an illustration of how the ECB's operations, as they were conceived under its original mandate, can nevertheless lead to it being the conduit for cross-border risk sharing via the portfolio of collateral it comes to hold—with sizable geographical distribution effects.

Bank defaults were avoided or postponed, but the euro zone's economy entered a credit crunch. In this period, there was a far larger fall in the growth of new loans to businesses and households than during the 2008–9 global financial crisis, even after conditioning for the dynamics of industrial production (Reichlin 2014).

There are two lessons from this narrative. The first is about the tension between liquidity policies and the ECB's narrow mandate. Central banks' ability to create liquidity at will means that they are uniquely well placed to resolve liquidity problems in the financial sector. This is the basis for both Bagehot's rule ("Lend freely against good collateral") and Friedman's rule ("Provide central bank liquidity at its marginal social cost—which is zero").

And because liquidity stresses may have solvency concerns at their root, the central bank is bound to monitor the strength of the banking system

overall, as well as the strength of individual institutions. Add to this the fact that central banks have an informational advantage, from their oversight of the payments system, and we can see that it is inevitable that central banks will play a central role in the maintenance of financial stability, whether this is explicitly recognized in their mandates or not.

Bagehot's rule is apparently clear, logical, and consistent with the ECB's narrow mandate, and it was what the ECB applied in the first phase of the crisis. However, in practice the rule is useless because the distinction between illiquidity and insolvency is often impossible to make in real time. In the end, central banks will always act to defend the monetary system, whether it is in their mandate or not, and defending the monetary system will have both monetary and fiscal consequences. The fiscal consequences were clear in the second phase of the crisis.

The second lesson, which is a consequence of the first, is that the governance structure should recognize and anticipate this fact. As Charles Goodhart has argued (1999), the question is not whether or not to act as a lender of last resort, but how best to organize this function so that it preserves the central bank's independence on one hand and ensures its fiscal backing on the other hand. This is a question of institutional design.

The sovereign debt crisis: The Security Market Programme (SMP) and Outright Monetary Transactions (OMTs). I have been describing the ECB's actions in response to a crisis in the banking system. But the ECB was also faced with a crisis in relation to the sovereign states of the euro area.

To understand the effectiveness of the ECB at this juncture, it is useful to compare two programs: the SMP and the OMT. On May 14, 2010, the ECB established the SMP, a program consisting in national central banks buying the government bonds of stressed countries. This program was initially a response to the Greek debt crisis, which gradually developed starting in the autumn of 2009, when the new Greek government first acknowledged the country's poor fiscal situation, to a real funding strike in the early spring of 2010.

This placed the ECB in a bind. On one hand, the ECB was understandably concerned that permitting a default on the sovereign debt of a euro area country threatened that country with financial collapse, given that the banking system held a significant amount of sovereign debt, much of which was used as collateral for ECB operations by this point. Such a financial collapse might then trigger exits from the euro if national authorities were forced to revert to their national currency to sustain payments and provide liquidity.

Moreover, the fear of contagion to other countries was considerable: If Greece were to default and/or exit, then this possibility would be entertained for other peripheral euro area economies, such as Ireland and Portugal. And banks in core countries had significant exposures to Greek sovereign debt.

On the other hand, the ECB was not well equipped on its own to address the solvency problem that threatened Greece. It was subject to institutional constraints that were expressly designed to protect it from pressure to deliver quasi-fiscal support to address solvency problems.

The ECB looked to the euro zone's national governments to provide the necessary fiscal support, but this was challenged on the grounds of the Maastricht Treaty's "no bail out" clause. But by late April 2010, a set of bilateral loans from other euro area countries had been agreed to—a framework that eventually took a stronger institutional form in the European Financial Stability Facility, and ultimately the European Stability Mechanism (ESM), within the context of an adjustment program under the auspices of, and also cofinanced by, the International Monetary Fund.

Yet even this initiative failed to restore market confidence, in part because official loans were to be made senior to private sector holdings. In early May, market tensions in Greece reached fever pitch, cross-border contagion intensified, and the SMP was eventually launched.

Despite this program—on which a total of €223 billion was eventually spent—the effect on sovereign spreads was limited, and the contagion also affected Italy and Spain. Indeed, in August 2011, the SMP was extended to Italy and Spain, but again with not much of an effect on sovereign credit spreads.

The failure of the SMP to calm markets can be attributed to the lack of a solid mandate. We should recall that Axel Weber, the president of the Deutsche Bundesbank, resigned in April 2011 in opposition to the ECB's action. In fact, the ECB itself described it as a limited and temporary program rather than an actual backstop.

This brings me to the second example of the ECB's intervention in the sovereign bond markets: the OMT announcement in July 2012.

In response to the worsening of the sovereign debt crisis, ECB president Draghi declared on July 26, 2012, during a conference in London: "Within our mandate, the ECB is ready to do whatever it takes to preserve the euro. And believe me, it will be enough" (Draghi 2012). He focused his speech on financial fragmentation as the main short-term challenge for restoring the transmission of the ECB's monetary policy.

A few days later, on August 2, the ECB announced outright purchases of sovereign debt in secondary bond markets, and in September it announced

the key parameters of the OMT program. Under the program, the ECB could purchase unlimited amounts of euro zone government bonds with maturities of one to three years, provided that the country whose bonds the ECB would buy met four key conditions. First, it had to receive financial support from the ESM. Second, it had to comply with the reform measures required by the respective ESM program. Third, the OMT program could only be activated if the country had regained complete access to private lending markets. Fourth and finally, the country's government bond yields had to be higher than what could be justified by the fundamental economic data.

As of today, the OMT program has never been used. Yet quantitative evidence in a substantial body of empirical literature, amply acknowledged by Hartmann and Smets, shows that the announcement had a large effect on sovereign yields, much larger than the effect of the actual purchases under the SMP. Equally, the OMT was much more successful in easing the funding conditions for banks in peripheral countries than the LTROs discussed above. This is explained by the fact that those conditions were partly affected by sovereign risks in banks that had had incentives to buy large quantities of domestic sovereign bonds.

So what made the difference? Why was the OMT announcement successful when the SMP was not?

Unlike the SMP, the OMT was conditional on countries entering a "program." This can be seen as a compromise: recognizing, on one hand, that a bad equilibrium resulting from a self-fulfilling crisis is possible; but also recognizing, on the other hand, the moral hazard issue due to the role of underlying solvency problems. In other words, it can be viewed as a mechanism to govern a trade-off between the risk of moral hazard (and therefore price instability) and financial instability via a solution that conditioned policy action to reform.

The fact that the scale of potential bond buying under the OMT was unlimited—and that, by intervening directly in the bond market, the ECB did not make itself a senior claimant—were also important factors contributing to the effectiveness of the OMT program in comparison with the SMP.

However, perhaps more significant than any of these specific aspects was the fact that the institutional context had changed. The ESM had been created, and the banking union had been agreed to. Crucially, the OMT was backed by a political agreement between the major countries; most notably, it was supported by German chancellor Angela Merkel, despite opposition from the Bundesbank.

The conditionality provided greater control over the fiscal dynamics, but the ECB's purchase of sovereign bonds meant taking credit risk onto the balance sheet of the Eurosystem, which implied some degree of fiscal backing. In other words, the OMT's credibility was due to a new bargain with the sovereign fiscal authorities on shared responsibilities involving new institution building.

The role of the lender of last resort for the sovereign in the euro area has been the subject of policy and academic debate. The case for an active ECB role as lender of last resort has been made, for example, by Paul de Grauwe (2012), on the basis of the observation that within a currency union, member states issue bonds in a "foreign" currency—that is, one that they do not themselves control. Hence, these member states cannot give a guarantee equivalent to the one that can be given by a sovereign with its own central bank, and investors may rightly fear that the sovereign will not be able to redeem the bonds when they mature. This means that the market for sovereign bonds of states within a monetary union is prone to liquidity crises and contagion—in much the same way that banking systems were afflicted by such emergencies before central banks stepped in as lenders of last resort.

The extreme behavior of spreads on sovereign bonds—going from around zero up until 2010, then spiking in 2010–12, and then falling again after 2012—is taken as evidence to support this argument. The proponents of this view maintain that even if the probability of default is driven by fundamental solvency issues, the central bank should intervene anyway, because in real time solvency and liquidity problems cannot be distinguished. Giancarlo Corsetti and Luca Dedola (2016) have recently studied this problem using a model with multiple solutions for the interest rate that private investors demand on bonds issued by the fiscal authority. Given that the monetary authority can issue liabilities at a lower interest rate than can a government that is subject to default risk, it can also lower the overall cost of borrowing for the public sector—which makes full repayment via taxation a more likely outcome than default and partial repayment. These researchers thus show that a suboptimal equilibrium can be avoided if the central bank announces its willingness to intervene.

However—and this is the key issue with the simplified version of this argument propounded by De Grauwe (2012)—in some states of the world, default could occur irrespective of whether the central bank made bond purchases. If the state defaults, the monetary authority would then suffer a capital loss and, if its balance sheet is sufficiently impaired, excessive inflation could result. To avoid this inflation scenario, the central bank must

recapitalize, which requires an agreement with the fiscal authorities. It is easy to imagine that if the ECB were to ask national central banks for recapitalization, political questions related to the redistributive effects of monetary policy could lead to paralysis, eventually impairing the ECB's credibility. This problem was originally analyzed by Christopher Sims (1999, 2012), who discusses fiscal backing in the Eurosystem. See also Corsetti and others (2016) for a recent discussion and relevant references.

The comparison between the SMP and the OMT provides a relevant case study showing that the central bank's credibility, and therefore effectiveness, depends on its backing by government. But if the power of the central bank ultimately comes from the backing of the sovereign, there is a problem of institutional design. The OMT in principle provides the fiscal backstop, but the fact that this instrument is in the hands of a central bank rather than democratically elected fiscal authorities could potentially constitute a challenge to the ECB's independence. To design an instrument for the euro area's common fiscal capacity would be more effective and would provide for more accountable governance of the monetary union.

MACRO STABILIZATION: INTEREST RATE POLICY AND THE ECB RULE The paper analyzes the ECB's interest rate policy through the lenses of a policy rule specified by Athanasios Orphanides (2003). According to Hartmann and Smets, this rule captures well the ECB's interest rate setting since 1999, including the crisis years. The implicit inflation target derived by the rule is 1.75 percent, which is very close to the price stability target definition of inflation of "below, but close to, 2 percent in the medium term"—an impressive outcome!

However, it is not clear that this rule was the right one to follow from a normative perspective. Other rules should also have been analyzed—for example, providing measures of the result in terms of inflation and unemployment.

Without such analyses, from a purely descriptive perspective, two episodes are particularly controversial. The first is the interest rate increase of July 2008. At the time, the interest rate increase was motivated by headline inflation (according to the Harmonised Index of Consumer Prices) being well above the target, having reached 3.75 percent. We now know, however, that the euro area had entered a recession in the first quarter of 2008, and of course the financial sector had already given signs of weaknesses on both sides of the Atlantic. The high level of inflation was explained by oil prices. The same was true in 2011, when the ECB

increased interest rates twice. At the time, headline inflation was about 3 percent, while core inflation was well below 2 percent. (For a quantitative assessment of the direct and indirect effects of oil prices on headline inflation, see Reichlin 2018.)

It is interesting to quote the ECB's press statement on April 7, 2011 (ECB 2011):

The adjustment of the current very accommodative monetary policy stance is warranted in the light of upside risks to price stability that we have identified in our economic analysis. . . .

With regard to price developments, euro area annual HICP inflation was 2.6% in March 2011, according to Eurostat's flash estimate, after 2.4% in February. The increase in inflation rates in early 2011 largely reflects higher commodity prices. Pressure stemming from the sharp increases in energy and food prices is also discernible in the earlier stages of the production process. It is of paramount importance that the rise in HICP inflation does not lead to second-round effects in price and wage-setting behaviour and thereby give rise to broad-based inflationary pressures over the medium term. Inflation expectations must remain firmly anchored in line with the Governing Council's aim of maintaining inflation rates below, but close to, 2% over the medium term.

Risks to the medium-term outlook for price developments remain on the upside. They relate, in particular, to higher than assumed increases in energy prices, not least owing to ongoing political tensions in North Africa and the Middle East. More generally, strong economic growth in emerging markets, supported by ample liquidity at the global level, may further fuel commodity price rises. Moreover, increases in indirect taxes and administered prices may be greater than currently assumed, owing to the need for fiscal consolidation in the coming years. Finally, risks also relate to stronger than expected domestic price pressures in the context of the ongoing recovery in activity.

Interestingly, it is recognized that inflation dynamics are explained by commodity prices but, as in July 2008, potential second round effects are emphasized. The first observation is that the ECB has historically given too much weight to headline inflation rather than monitoring measures of underlying inflation, as in other central banks.

Another observation is that the stress on second-round effects was done in a context in which the debt crisis was in full displacement, affecting both banks and sovereigns. Hartmann and Smets comment on these episodes as a possible underestimation of the effect of the credit crunch (for the reasons discussed in the previous section) on the real economy and on underlying inflation.

The question is whether, in this underevaluation, we can identify a problem that again has to do with the narrow interpretation of the mandate, seeing the monetary policy objective and the price stability mandate not

only as separate from the financial stability objective but also as independent. Although it can be argued that, in the spring of 2011, the second recession had not yet started in the euro area, there was ample evidence of a credit crunch, a segmentation of the financial market along national lines, and a substitution of countries' sovereign bonds for loans on the banks' assets that was affecting the real economy.

In the years 2012–14, the interest rate reached the zero bound. At the same time, the balance sheet of the Eurosystem, which had expanded as an endogenous consequence of liquidity operations, began to shrink.

During that period, it can be argued that the ECB was slow to act. Quantitative easing finally found the support of a vast majority of the Governing Council, when it became clear that the inflation target that the ECB is required to meet under the Maastricht Treaty was still not being met, and the European Union was risking entering a period of deflation, as Japan had done in the 1990s. It is of little comfort that the implied inflation target for the first 20 years of the ECB's history, as calculated by the Hartmann-Smets rule, was 1.76 percent.

CONCLUSION Maastricht is the child of the precrisis consensus, which led to the ECB's design—an extreme form of independence, and a constitutional mandate of price stability.

During the global financial crisis, the ECB was confronted with the problem of defending price stability but also defending the stability of the financial system. Because liquidity and solvency concerns cannot be separated in practice, a strict "separation principle" was not always useful for guidance. The ECB, by acting to defend the stability of the financial system, implemented policies with potential fiscal implications and geographical distributional consequences.

In this, it was no different than other central banks. The nature, visibility, and political sensitivity of distributional consequences related to non-standard policies are similar in many countries. Because these policies are likely to also remain in the tool kits of central banks during normal times, new problems of institutional design are likely to emerge. In the euro area, where politics is still largely national and distributional consequences often arise between member states, it is not surprising that these problems are more controversial.

These exceptional policies were seen to be necessary—to stabilize the financial system, and even to save the euro. However, as we have seen, they were less successful when the backing of the fiscal authorities was uncertain. Ultimately, the power of central banks comes from the sovereign, and the ECB case powerfully illustrates this point.

So, if we are to ask whether the ECB did the right thing—whether it did what was necessary in pursuit of its mandated objectives—we must also ask if it needed to overstep the remit given in the Maastricht Treaty. And if so, then what does this imply for the necessary institutional reform?

REFERENCES FOR THE REICHLIN COMMENT

- Brunnermeier, Markus K., Sam Langfield, Marco Pagano, Ricardo Reis, Stijn Van Nieuwerburgh, and Dimitri Vayanos. 2016. “ESBies: Safety in the Tranches.” Working Paper 21. Brussels: European Union, European Systemic Risk Board. <https://www.esrb.europa.eu/pub/pdf/wp/esrbwp21.en.pdf>.
- Corsetti, Giancarlo, and Luca Dedola. 2016. “The Mystery of the Printing Press, Monetary Policy and Self-Fulfilling Debt Crises.” *Journal of the European Economic Association* 14, no. 6: 1329–71.
- Corsetti, Giancarlo, Luca Dedola, Marek Jarocinski, Bartosz Maćkowiak, and Sebastian Schmidt. 2016. “Macroeconomic Stabilization, Monetary-Fiscal Interactions, and Europe’s Monetary Union.” ECB Working Paper 1988. Frankfurt: European Central Bank.
- De Grauwe, Paul. 2012. “The Governance of a Fragile Eurozone.” *Australian Economic Review* 45, no. 3: 255–68.
- Draghi, Mario. 2012. “Verbatim of the Remarks Made by Mario Draghi.” Speech at Global Investment Conference, London, July 26. <https://www.ecb.europa.eu/press/key/date/2012/html/sp120726.en.html>.
- ECB (European Central Bank). 2011. “Introductory Statement to the Press Conference (with Q&A),” by Jean-Claude Trichet and Vítor Constâncio, Frankfurt, April 7. <https://www.ecb.europa.eu/press/pressconf/2011/html/is110407.en.html>.
- Giannone, Domenico, Michele Lenza, Huw Pill, and Lucrezia Reichlin. 2012. “The ECB and the Interbank Market.” *Economic Journal* 122, no. 564: 467–86.
- Goodhart, Charles. 1999. “Myths about the Lender of Last Resort.” *Journal of Finance* 2, no. 3: 339–60.
- Lenza, Michele, Huw Pill, and Lucrezia Reichlin. 2012. “Monetary Policy in Exceptional Times.” *Economic Policy* 25, no. 62: 295–339.
- Orphanides, Athanasios. 2003. “Historical Monetary Policy Analysis and the Taylor Rule.” *Journal of Monetary Economics* 50, no. 3: 983–1022.
- Padoa-Schioppa, Tommaso. 2004. *Regulating Finance: Balancing Freedom and Risk*. Oxford University Press.
- Pill, Huw, and Lucrezia Reichlin. 2016a. “Exceptional Policies for Exceptional Times: The ECB’s Response to the Rolling Crises of the Euro Area.” In *Routledge Handbook of the Economics of European Integration*, edited by H. Badinger and V. Nitsch. New York: Routledge.
- . 2016b. “Non-standard Monetary Policy and Financial Stability: Developing an Appropriate Macrofinancial Policy Mix.” In *Preparing for the Next*

- Financial Crisis: Policies, Tools, and Models*, edited by Esa Jokivuolle and Radu Tunaru. Cambridge University Press.
- Reichlin, Lucrezia. 2014. "Monetary Policy and Banks in the Euro Area: The Tale of Two Crises." *Journal of Macroeconomics* 39, part B: 387–400.
- . 2018. "Comment on 'Slack and Cyclically Sensitive Inflation' by J. Stock and M. Watson." Paper presented at European Central Bank's Annual Forum, "Price and Wage Setting in Advanced Economies, Sintra, June 18–20.
- Sims, Christopher A. 1999. *The Precarious Fiscal Foundations of EMU*. Staff Report 34. Amsterdam: Dutch National Bank.
- . 2012. "Gaps in the Institutional Structure of the Euro Area." *Financial Stability Review* (Banque de France) 16: 217–23.

GENERAL DISCUSSION Athanasios Orphanides began by saying that the European Central Bank (ECB) has made some major mistakes. It is important for this panel to discuss these mistakes and draw lessons from them, he said, so they are not repeated in the future.

During the past several years, the ECB's actions have compromised the safe asset status of sovereign debt in the euro area, he said. This greatly exacerbated the destabilization of the euro area, and he pointed to the increasing Italian spreads over the three months leading up to this conference as the latest example. He highlighted two issues that should be discussed better in the paper by Hartmann and Smets. First, when evaluating the solvency of sovereigns, the ECB has decided to rely exclusively on market interest rates—including unrealistic risk premia that may reflect adverse self-fulfilling equilibria, which has a destabilizing effect. Second, he noted that the ECB is the only central bank that questions whether the government debt of its own member countries is eligible collateral for monetary policy operations, regardless of fundamentals. This happens because of a discretionary decision made before the global financial crisis that delegated the determination of collateral eligibility to rating agencies. He called this an unfortunate decision because it creates destabilizing cliff effects and leads to adverse equilibria. Policymakers know that these policies worked terribly during the crisis, he said, and he wondered whether these mistakes would be corrected.

Jason Furman noted that the paper by Hartmann and Smets treats the neutral interest rate as fixed over the ECB's 20 years. He questioned why the authors made that choice, and if they were to choose differently, wondered if the analysis would show considerably more monetary policy tightening during the examined period relative to the authors' results.

Jón Steinsson found it notable that there was little mention of unemployment rates. The main point he gathered from the paper was that the ECB hit its inflation target, and because the unemployment rate is not part of the ECB's mandate, the authors do not discuss it. However, he noted that there is a view that during the crisis, allowing for a somewhat higher level of inflation would have helped real wages and unemployment rates adjust in Southern Europe. One of the many reasons why this policy was not pursued, he conjectured, is because unemployment is not part of the ECB's mandate. He wondered if this is indicative of a problem with the ECB's mandate itself, and whether a dual mandate like that the U.S. Federal Reserve could work better.

Eric Rosengren asked the authors whether they expect the ECB to hit the zero lower bound frequently in the future, and if that consideration would alter the ECB's policy framework.

Frederic Mishkin said that he found the ECB's framework problematic, particularly regarding the inflation target. Because the language of the mandate indicates that the ECB wants inflation to be slightly lower than 2 percent, it is asymmetric in nature, he said. He thinks that the ECB chose this language because when it was formed, the ECB aimed to inherit the credibility of the Deutsche Bundesbank, and the Bundesbank had an asymmetric inflation target. The unfortunate policy consequence is that the ECB is more concerned about overshoots than undershoots, he said. He thought that this asymmetry was one of the key factors in the ECB's decision to raise interest rates in 2011. He then pointed to the U.S. Federal Reserve, noting how it emphasizes that its target is symmetric. There are even arguments for overshooting for temporary periods. He concluded by acknowledging that the ECB performed well during the initial phases of the global financial crisis, but that the inflation targeting aspect of European monetary policy should be changed.

Robert Gordon referred back to Steinsson's point about the Federal Reserve's dual mandate versus the ECB's inflation mandate. He claimed that the ECB gives a disproportionate amount of attention to inflation movements and inflation expectations, to the exclusion of factors such as unemployment, potential output, actual output, and output gaps—all factors that the Fed considers relevant context for monetary policy. As an evaluation of the ECB's performance over the past 20 years, the paper should have included comparisons on employment rates and actual and potential output growth between the euro area and the United States, he said; but any differences may not be entirely related to monetary policy. He suggested that someone should write a paper comparing actions by the ECB and

the Fed, including the different interest rate sequences and the timing of quantitative easing programs. That paper should also distinguish between fiscal austerity in Southern versus Northern Europe, and document the lack of fiscal coordination in the euro area. He noted that the U.S. did not have to face this problem, despite its own fiscal austerity in 2013 and 2014.

Jay Shambaugh discussed the policy rule—which indicates the central bank’s interest rate response according to economic conditions—that the authors examine in the paper. On one hand, he thought that it was fascinating to see how closely the ECB followed the rule. On the other hand, he wondered if it was desirable for the ECB to follow this particular rule so closely, especially since it closely follows headline inflation. Given that the rule suggests raising rates in 2008 and 2011, he questioned whether this is the right rule for ECB to follow. Further, he wondered if the forecasts incorporated in the rule were biased. For example, if the forecasts systematically underestimate deflation risk, even if the ECB followed the rule, it would systematically prescribe actions that are too tight. Finally, he echoed previous comments that questioned the merits of inflation target asymmetry and, more broadly, whether the ECB’s mandate should be expanded.

Glenn Rudebusch said that the authors’ policy rule choice also puzzled him. Because the rule relates the change in the interest rate to a change in the price level, it is essentially a price level–targeting rule. He also thought that the metric used to assess whether the ECB followed this rule was weak. He suggested that the authors examine the ECB’s response with the Taylor rule, and suspected that it would probably fit just as well as the first-difference rule that the authors used. He noted that it is important to clarify if the rule incorporates an output gap in levels or an output gap in growth terms.

Richard Cooper began by stating that he agrees with the substance of many previous comments. He added that much of this discussion, however, essentially questioned the Maastrich Treaty—something that the ECB cannot change itself. Any revision would need the ratification of all member governments. Although the treaty mandates price stability, Cooper noted that it does not specify an inflation target, nor whether it should be asymmetric or not. Thus, he claimed that the ECB has adopted an asymmetric target around price stability by stating “below, but close to, 2 percent inflation.” Thus, Cooper agreed with the criticisms on the asymmetry point. He noted that he was against the Maastricht Treaty from the beginning, because he believed that it was a poor instrument for achieving monetary union in Europe.

Lorenzo Bini Smaghi addressed two issues: the asymmetry of the 2 percent inflation target, and whether the ECB should adopt a dual mandate. He agreed that the ECB should revisit and discuss the inflation target. Although the ECB was very set on its initial definition of price stability, the ECB should go back and discuss what “close to 2 percent” actually means, he said. However, he does not agree with the comments suggesting that the ECB should adopt a dual mandate like that of the U.S. Federal Reserve. When observing economic variables (like GDP per capita), he said that the euro area and the U.S. have been on similar trends, with the exception of the second recession in the euro area in 2012 and 2013. Indeed, after the second recession, the euro area recovered and fell back in line with the United States. The real difference, he claimed, was the response to the global financial crisis. The crisis was not dealt with well in the euro area, whereas there was a smarter response in the U.S. The underlying issue was financial stability, he said, and the ECB’s mandate did not fully incorporate that.

Further, it is unclear how fiscal policy contributed to financial stability in the euro area, Bini Smaghi noted. He thinks that the ECB is moving to make the euro area’s financial system more resilient by having a fiscal backstop and having a single supervisor. But he urged the audience to remember that long-term economic performance in the euro area has not done worse, on average, than in the U.S. He noted that some European countries have even done better (such as Germany and Finland). It is important to look at the structural issues related to these economies, he concluded.

Philipp Hartmann started out by asking the rhetorical question of whether at the occasion of the European Central Bank’s 20th anniversary, anybody in the room had something positive to say about the ECB or Economic and Monetary Union (EMU). Next, he expressed his impression that a number of the conference participants seem to imply that ECB policymakers should “first, break all the rules,” as suggested by the title of the best-selling book by Marcus Buckingham and Curt Coffman.¹ Even though the true recommendations of those authors for industry leaders had a somewhat different meaning, Hartmann would not find it wise for central bank policymakers to go against the mandates and laws given to them by the democratic political process.

More specifically, he grouped his answers into three main points. First, he addressed the claims that the ECB’s price stability aim was asymmetric.

1. Marcus Buckingham and Curt Coffman, *First, Break All The Rules: What the World’s Greatest Managers Do Differently* (New York: Simon & Schuster, 1999).

He recalled that the paper tested for the asymmetry of the policy rule and rejected that hypothesis. Further, he stressed that the ECB, in its communications, has repeated since the early days of the euro that its aim is symmetric. Given both the scientific analysis and the consistent communication, Hartmann said that the people who still believe that the ECB's inflation aim is asymmetric should provide evidence supporting their perception. Maybe the wording of an inflation rate "below, but close, to 2 percent" in the medium term, which is different from formulations of inflation targets in the academic literature or of other central banks, troubled some observers. If this was the case, Hartmann thought that one could perhaps revisit this wording in the future. All in all, however, the available evidence does not support asymmetry, implying that the issue was not of first order for actual policy, contrary to what some conference participants seem to think.

Second, Hartmann proposed a different narrative for the performance of the euro area's macroeconomy than did those observers who were eager to find ECB "mistakes." Taking the global financial crisis period as an example, he asked what was the main difference between the euro area and the U.S. The main difference was that, due to the European sovereign debt crisis, the euro area had a second deep recession starting in 2011. This recession was mainly caused by the malicious interaction between banking and fiscal instability in (and across) a number of countries, the sovereign-bank nexus. For example, many euro area countries had not addressed their banking problems as swiftly as had the U.S. Moreover, some euro area countries had entered the crisis with high public debt levels—way above the limits prescribed by the fiscal framework for EMU—and some were fiscally weakened by high bank bailout costs. The resulting combined fiscal and banking crisis induced huge obstacles to the ECB for maintaining price stability and forced it to deploy untested and unprecedented unconventional policies in subsequent years. At the time, the ECB was not a banking supervisor and, in general, it is forbidden from financing public debts or government tasks (the prohibition of monetary financing in Article 123 of the Treaty on the Functioning of the European Union), including bank recapitalizations. The underlying sources of the sovereign debt crisis were outside its mandate. At the same time, the paper by Hartmann and Smets provides a scientific analysis of the ECB's interest rate decisions and, through the lens of policy rules, transparently gives a few indications when its monetary policy might have been a bit too loose or too tight during the last 20 years. But overall, this broader perspective

suggests that the first-order reason for the double-dip recession and its depth was not monetary policy but the imperfect handling of banking and fiscal problems. These imperfections included the initial absence of effective euro area institutions and tools for solving the collective action problems of sovereign countries in a monetary union with primarily national prudential and fiscal policies. Such institutions and tools—for example, the European Stability Mechanism or the Banking Union (with Single Supervisory and Resolution Mechanisms)—were only established or put on a credible path with the political agreements that were reached during 2012. Before this time, it is hard to see how the ECB could have run a much more expansionary monetary policy and stayed clear of monetary financing.

Third, Hartmann took up the differences in central bank mandates, laws, and approaches between the euro area and the U.S. He started with the observation that the Federal Reserve is the only Group of Seven central bank with a dual mandate that includes maximum employment as a primary statutory objective. Much like the other central banks, the ECB can pursue employment only without prejudice to the primary price stability objective; the former is clearly subordinated to the latter. Next, he shared his impression that many participants in the conference seem to interpret the role of lender of last resort very broadly. In fact, in the academic literature there is a school of thought that represents the view that the central bank should not only act as lender of last resort for bank liquidity problems but also for fiscal authorities, at least in specific situations. And many of those academics seem to assume that the Federal Reserve would do so when needed. In the EMU, this is not allowed, as reflected in the prohibition of monetary financing. And there are some good reasons for this, one being that it can create a circularity problem between the central bank and the fiscal authority that can contribute to multiple equilibria. Still, Hartmann expressed sympathy with Lorenzo Bini Smaghi's point that the current national approach to providing emergency liquidity assistance to banks could be centralized at the ECB in the future, as ECB president Draghi had also recently hinted in one of the hearings at the European Parliament. In order not to inject money into insolvent banks, however, and therefore not to take over fiscal or bank resolution tasks, this should be accompanied by arrangements ensuring the temporary nature of such operations and adequate fiscal guarantees where solvency cannot be ascertained immediately. Hartmann also agreed with Lucrezia Reichlin that the ECB had a "text-book" lender-of-last-resort reaction to the liquidity problems in the early stages of the crisis.

In response to Jason Furman's and Eric Rosengren's comments, it was clarified that the estimated policy rule does not assume a constant equilibrium real rate (in contrast to the Taylor rule). The paper briefly acknowledges that estimates of the natural rate of interest have been sliding down over time, and that this makes the effective lower bound to policy rates a serious issue.

EMMANUEL FARHI

Harvard University

FRANÇOIS GOURIO

Federal Reserve Bank of Chicago

Accounting for Macro-Finance Trends: Market Power, Intangibles, and Risk Premia

ABSTRACT Real risk-free interest rates have trended down over the past 30 years. Puzzlingly, in light of this decline, (1) the return on private capital has remained stable or even increased, creating an increasing wedge with safe interest rates; (2) stock market valuation ratios have increased only moderately; (3) and investment has been lackluster. We use a simple extension of the neoclassical growth model to diagnose the nexus of forces that jointly accounts for these developments. We find that rising market power, rising unmeasured intangibles, and rising risk premia play a crucial role, over and above the traditional culprits of increasing savings supply and technological growth slowdown.

During the past 30 years, most developed economies have experienced large declines in risk-free interest rates and increases in asset prices such as housing or stock prices, with occasional sudden crashes. At the same time, except for a short period in the 1990s, economic growth, in particular productivity growth, has been rather disappointing, and investment has been lackluster. Earnings growth of corporations has been strong, however, leading in most countries to an increase in the capital share and to stable or slightly rising profitability ratios. Making sense of these trends is a major endeavor for macroeconomists and for financial economists.

Conflict of Interest Disclosure: The authors did not receive financial support from any firm or person for this paper or from any firm or person with a financial or political interest in this paper. They are currently not officers, directors, or board members of any organization with an interest in this paper. No outside party had the right to review this paper before circulation.

Given the complexity of these phenomena, it is tempting to study them in isolation. For instance, a large body of literature has developed that tries to understand the decline in risk-free interest rates. But studying these trends independently may miss confounding factors or implausible implications. For instance, an aging population leads to a higher savings supply, which might well explain the decline in interest rates. However, a higher savings supply should also increase capital accumulation—that is, investment, and hence reduce profitability. Similarly, it should also increase stock prices, as the discount rate falls. Hence, a potential driver that is compelling judged by its ability to explain a single trend may be implausible overall, because it makes it harder to account for the other trends.

Another way to highlight these tensions is to note that the stable profitability of private capital and declining risk-free rate lead to a rising spread, or wedge, between these two rates of return. What gives rise to this spread? A narrative that has recently attracted significant interest is the possibility of rising market power. However, rising risk premia could also account for the wedge. The only way to disentangle these potential causes is to consider additional implications—for instance, everything else being equal, rising market power should imply a lower labor share, and rising risk premia should be reflected in lower prices of risky assets such as stocks.

These simple observations motivate our approach. We believe that a successful structural analysis of the past 30 years should account for these trends *jointly*. A novel feature of our analysis is that we aim to account both for macroeconomic trends and finance trends. The first step of our paper is to document a set of broad macro and finance trends that we believe are of particular interest. We focus on six indicators: economic growth, risk-free interest rates, profitability, the capital share, investment, and valuation ratios (such as the price-dividend or price-earnings ratio).

The paper's second step is to develop an accounting framework to disentangle several potential drivers of these trends. We focus on five narratives that have been put forward to explain some or all of these trends. The first narrative is that the economy experienced a sustained growth decline, owing to lower population growth, investment-specific technical progress, or productivity growth. The second narrative is that the savings supply has increased, perhaps owing to population aging (or to the demand of emerging markets for a store of values). The third narrative involves the rising market power of corporations. The fourth narrative focuses on technological change resulting from the introduction of information technology, which may have favored capital or skilled labor over unskilled labor, or the

rise of hard-to-measure intangible forms of capital. And the fifth narrative, which we emphasize, involves changes in perceived macroeconomic risk, or tolerance of it.

Our approach is simple enough to allow for a relatively clear identification of the impact of these drivers on the facts that we target. Here, our contribution is to propose a simple macroeconomic framework—a modest extension of the neoclassical growth model—that accounts for the “big ratios” familiar to macroeconomists as well as for the “financial ratios” of financial economists. Our model does this in a way that allows for interesting types of feedback between macroeconomic and financial variables. For example, the investment-output ratio is affected by market power and macroeconomic risk, as well as savings supply and technological parameters. At the same time, our framework preserves the standard intuitions and results of macroeconomists and financial economists, and hence is a useful pedagogical device.¹

In our baseline estimation, we abstract from intangibles. Our main empirical result here is that the rising spread between the return on capital is the risk-free rate, which is driven mostly by a confluence of two factors: rising market power *and* rising macroeconomic risk. This rising macroeconomic risk in turn implies that the equity premium, which previous researchers have argued fell in the 1980s and 1990s, may have risen since about 2000. This higher risk is also an important driver of the decline of risk-free rates. We also find little role for technical change. Moreover, we show how previous researchers, who have used models without risk, have attributed too big a role to rising market power. When we incorporate intangibles, we see that a significant increase in their unmeasured component can help explain the rising wedge between the measured marginal product of capital and the risk-free rate. Interestingly, we find that intangible capital reduces the estimated role of market power in our accounting framework, while preserving the role of risk. Overall, our estimates offer a more nuanced understanding of the drivers of investment, profitability, and valuation ratios.

The rest of the paper is organized as follows. Section I discusses the related literature. Section II documents the main trends of interest.

1. Our model, of course, needs to contend with the usual disconnect between macroeconomics and finance—that is, the equity premium puzzle—and hence requires high risk or high risk aversion to generate plausible quantitative implications. Although we do not address the excess volatility puzzle in this paper, the framework can be extended, as done by Gourio (2012), to fit this as well.

Section III presents our model. Section IV explains our empirical methodology and identification. Section V presents the main empirical results. Section VI discusses extensions and robustness. Finally, section VII reviews some outside evidence on the rise in the equity premium, markups, and intangibles. Section VIII concludes.

I. Literature Review

Our paper, given its broad scope, makes contact with many other studies that have separately tried to explain one of the key trends that we document. (In section VII, we discuss in more detail the relation of our results to the recent literature on market power, intangibles, and risk premia.)

First, a large body of literature studies the decline of interest rates on government bonds. James Hamilton and others (2016) provide a long-run perspective, and discuss the connection between growth and interest rates. Łukasz Rachel and Thomas Smith (2017) provide an exhaustive analysis of the role of the many factors that affect interest rates. The role of demographics is studied in detail by Carlos Carvalho, Andrea Ferrero, and Fernanda Nechio (2016); and by Etienne Gagnon, Benjamin Johannsen, and David López-Salido (2016). Marco Del Negro and others (2017) emphasize, as we do, the role of the safety and liquidity premia. Ben Bernanke (2005) and Ricardo Caballero and others (2008) emphasize the role of safe asset supply and demand. Our analysis incorporates all these factors, though in a simple way.

Second, a large body of literature documents and tries to explain the decline of the labor share in developed economies. Michael Elsby, Bart Hobijn, and Ayşegül Şahin (2013) document the facts and discuss various explanations using U.S. data, while Loukas Karabarbounis and Brent Neiman (2014) study international data and argue that the decline is driven by investment-biased technical change. Matthew Rognlie (2015) studies the role of housing. A number of other researchers discuss the impact of technical change for a broader set of facts (Acemoglu and Restrepo, forthcoming; Autor and others 2017; Kehrig and Vincent 2018).

The most closely related papers are by Caballero, Emmanuel Farhi, and Pierre-Olivier Gourinchas (2017); Caballero and Farhi (2018); and by Magali Marx, Benoît Mojon, and François Velde (2018)—as well as the contemporaneous work by Gauti Eggertsson, Jacob Robins, and Ella Wold (2018). Marx and colleagues also find, using a different methodology, that an increase in risk helps explain the rising spread between the marginal product of capital (*MPK*) and the risk-free rate. They do not explicitly

target the evolution of other variables, such as investment or the price-dividend ratio. Conversely, Eggertsson, Robins, and Wold (2018) target some of the same big ratios that we study, but there are differences in methodology and results. Methodologically, our approach uses a simple standard model, which allows a closed-form solution and clear identification. Substantively, we find a more important role for macroeconomic risk, whereas they contend that a rising savings supply and rising market power are the main driving forces.

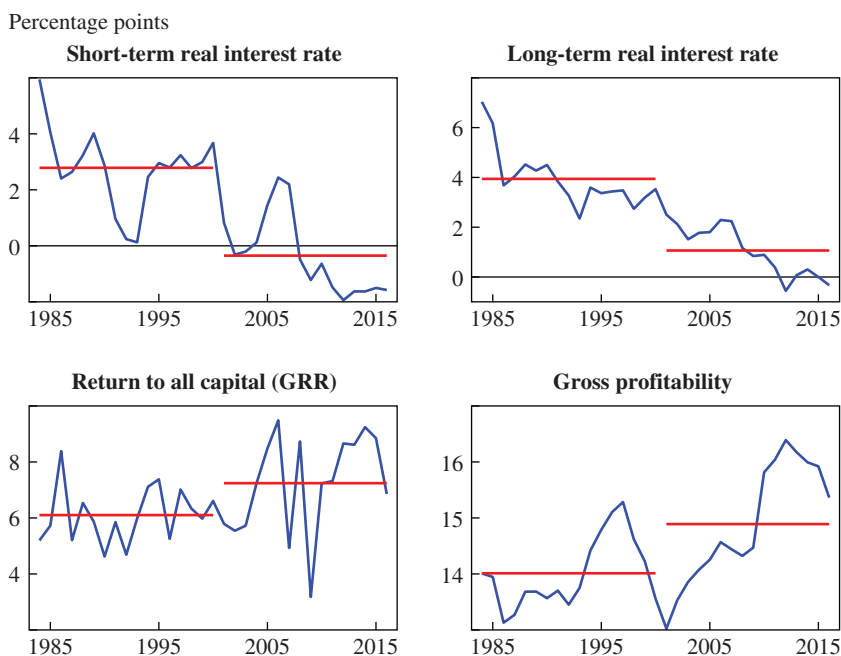
II. Notable Macroeconomic and Finance Trends

This section presents simple evidence on the trends affecting some key macroeconomic and finance moments. We focus on six groups of indicators: interest rates on safe and liquid assets, such as government bonds; measures of the rate of return on private capital; valuation ratios (that is, price-dividend or price-earnings ratio for publicly listed companies); private investment in new capital; the labor share; and growth trends. We first present simple graphical depictions, then add statistical measures.

Our focus is on the United States, but we believe that these facts also hold for other developed economies and hence may reflect worldwide trends.² Like many macroeconomic studies, we mostly consider the post-1984 period, which is associated with low and stable inflation together with relative macroeconomic stability (the “Great Moderation”). We present the changes in the simplest possible way, by breaking our sample equally in the middle, that is, at the millennium. However, we also briefly discuss the longer-range trends and present continuous indicators using moving averages.

One important decision is whether to study the entire private sector or to exclude housing and focus, for instance, on nonfinancial corporations. On one hand, the savings of households include all assets, in particular housing; on the other hand, the housing sector may need to be modeled differently, or we might want to explicitly recognize the heterogeneity of capital goods. In this section, we present indicators that cover both, but our estimation targets cover the entire private sector. For the most part, the trends that we document are apparent both for nonfinancial corporations and in the aggregate.

2. See, for instance, Marx, Mojon, and Velde (2018) for euro area trends.

Figure 1. U.S. Rates of Return, 1984–2016^a

Sources: See the online appendix, section 1, for all data sources.

a. The top left panel displays the difference between the 1-year Treasury bill rate and the median 1-year-ahead Consumer Price Index (CPI) inflation expectations from the Survey of Professional Forecasters (SPF). The top right panel displays the difference between the 10-year Treasury note rate and the median 10-year-ahead CPI inflation expectations from the SPF. The bottom left panel presents the estimate of the pretax return on all capital from Gomme, Ravikumar, and Rupert (2011; GRR). The bottom right panel presents our measure of gross profitability, the ratio of 1 minus the labor share to the capital-output ratio. The horizontal lines represent the mean in the first and second halves of the samples—1984–2000 and 2001–16, respectively.

II.A. Graphical Evidence

We summarize the evolution of the six groups of indicators as six facts.

Fact 1: Real risk-free interest rates have fallen substantially. The top panels of figure 1 present proxies for the 1-year and 10-year real interest rates by subtracting inflation expectations from nominal Treasury yields.³

3. We use median consumer price inflation expectations from the Federal Reserve Bank of Philadelphia's Survey of Professional Forecasters (SPF). Very similar results for the trend are obtained if one uses the mean expectation rather than the median; or the Michigan Survey of Consumers rather than the SPF. For the 1-year rate, one can also replace expectations with ex-post inflation or lagged inflation. For the 10-year rate, one can also use the Treasury Inflation-Protected Securities yield where available (that is, after 1997).

As many researchers have noted before, there has been a strong downward trend in these measures since 1984. The short-term rate exhibits clear cyclical fluctuations, while the long rate has a smoother decline. Table 1 shows that the average 1-year rate falls from almost 2.8 percent in the first half of our sample (1984–2000) to almost –0.3 percent in the second half of our sample (2001–16). The long-term rate similarly falls, from 3.9 percent in the first half to 1.1 percent in the second half.

Fact 2: The profitability of private capital has remained stable or increased slightly. In contrast, there is little evidence that the return on private capital has fallen; if anything, it appears to have increased slightly. Paul Gomme, B. Ravikumar, and Peter Rupert (2011), using data from the National Income and Product Accounts, construct a measure of the aggregate net return on physical capital—roughly, profits over capital. The bottom left panel of figure 1 depicts their series. The rising spread between their measure, which can be thought of as a proxy for the marginal product of capital, and the interest rate on U.S. Treasuries, is an important trend to be explained for macroeconomic and financial economists.

Gomme, Ravikumar, and Rupert (2011) construct their series using detailed data from the National Income and Product Accounts and other sources, but one can construct a simple approximation using the ratio of operating surplus to capital for the nonfinancial corporate sector; table 1 shows that this ratio is also stable, and if anything increases slightly. In our estimation exercise, we focus on gross profitability, and, to ensure consistency between our measures, we construct it simply as the ratio of the profit-output ratio that we use (that is, 1 minus the labor share) to the capital-output ratio. For this measure, which is depicted in the bottom right panel of figure 1, the overall level is higher, in part because it is gross rather than net; but the trend is similar to the measure used by Gomme and colleagues.

Fact 3: Valuation ratios are stable or have increased moderately. The top two panels of figure 2 present measures of valuation ratios for the U.S. stock market. The top left panel shows the ratio of price to dividends from the Center for Research in Security Prices, while the top right panel shows the price-operating earnings ratio for the Standard & Poor's 500 Index (S&P 500).⁴ The latter is essentially trendless, while

4. We focus on operating earnings that exclude exceptional items such as write-offs and hence are less volatile. In particular, total earnings were negative in 2008:Q4 because banks marked down the values of their assets substantially.

Table 1. Macroeconomic and Financial Trends^a

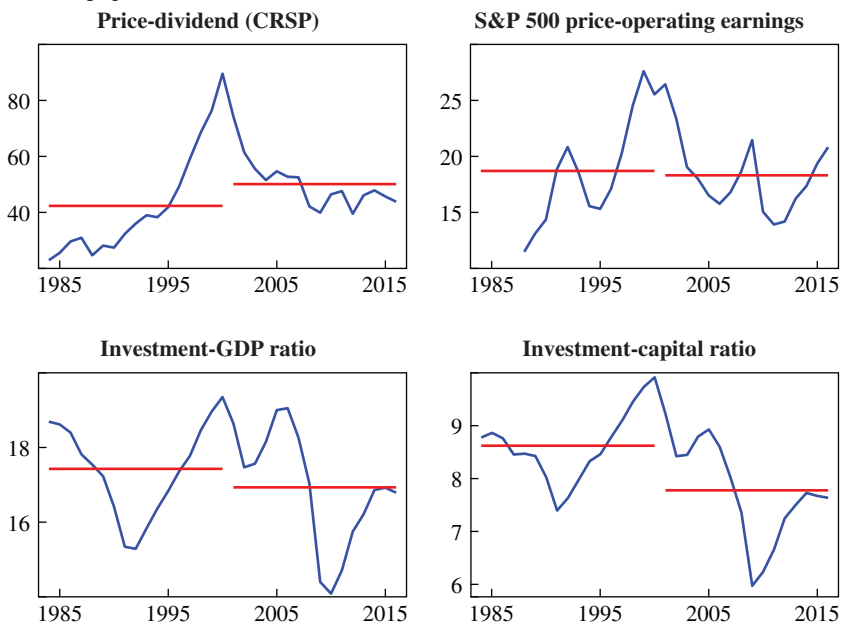
Group	Variable	Averages				Trend	
		1984–2000	SE	2001–16	SE	Difference	SE
Real interest rate	One-year maturity*	2.79	.45	–.35	.59	–3.14	.75
	Ten-year maturity	3.94	.41	1.06	.46	–2.88	.69
Return on capital	AA rate	4.69	.48	1.09	.57	–3.6	.8
	Ten-year adjustment for term premium	1.52	.26	–.09	.35	–1.61	.4
Valuation ratios	GRR: all, pretax	6.1	.2	7.24	.45	1.14	.45
	GRR: business, pretax	8.59	.32	10.46	.62	1.87	.62
Investment	Nonfinancial corporations, GOS/NRK	7.59	.34	7.87	.36	.27	.51
	Gross profitability* (see text)	14.01	.26	14.89	.49	.88	.6
Capital-output	Price-to-dividend ratio* CRSP	42.34	8.56	50.11	3.4	7.78	8.39
	Price-operating earnings, S&P 500	18.7	2	18.31	1.09	–.39	1.75
Labor share	Price-smoothed earnings, Shiller	22.07	4.41	24.36	1.25	2.29	4.5
	Investment share of GDP	17.43	.53	16.93	.65	–.5	.76
Growth	Nonresidential investment share of GDP	12.94	.40	12.79	.18	–.15	.43
	Investment capital ratio: all*	8.1	.25	7.23	.35	–.88	.38
Labor share	Investment capital ratio: nonresidential	10.95	.39	10.2	.24	–.76	.4
	Fixed assets	2.13	.03	2.28	.03	.15	.04
Capital-output	Real index (BLS)	1.06	.02	1.18	.01	.13	.02
	Nonfarm business (BLS) gross	62.07	.31	58.56	1.01	–3.51	1.11
Growth	Nonfinancial corporations, gross*	70.11	.34	66.01	1.21	–4.1	1.29
	Output per worker	1.80	.22	1.22	.23	–.58	.29
Labor share	Total factor productivity*	1.10	.31	.76	.32	–.34	.36
	Population*	1.17	.08	1.1	.06	–.07	.08
Growth	Price of investment: all*	–1.77	.15	–1.13	.34	.64	.26
	Price of investment: nonresidential	–2.38	.19	–1.75	.29	.63	.25
Capital-output	Price of investment: equipment	–3.62	.60	–3.27	.53	.34	.72
	Price of investment: intellectual property products	–1.71	.30	–2.15	.36	–.44	.52
Labor share	Employment–population ratio*	62.34	.58	60.84	0.94	–1.51	1.06
						–.07	.06

Sources: See the online appendix, section 1, for all data sources.

a. This table reports, for each variable, the mean in the 1984–2000 sample, in the 2001–16 sample, their difference, and the coefficient on a linear time trend, all with standard errors. * = a moment targeted in our estimation exercise. SE = standard error; CRSP = Center for Research in Security Prices; GRR = Gomme, Ravikumar, and Rupert (2011); GOS = gross operating surplus; NRK = nonresidential capital. Variables' construction is detailed in section 1 of the online appendix (the online appendixes for this and all other papers in this volume may be found at the *Brookings Papers* web page, www.brookings.edu/bpea, under “Past BPEA Editions”).

Figure 2. U.S. Investment and Valuation Ratios, 1984–2016^a

Percentage points



Sources: See the online appendix, section 1, for all data sources.

a. The top left panel displays the price-dividend ratio from the Center for Research in Security Prices (CRSP). The top right panel shows the ratio of price to operating earnings for the S&P 500. The bottom left panel shows the ratio of nominal investment spending to nominal GDP. The bottom right panel shows the ratio of nominal investment to capital (at current cost). The horizontal lines represent the mean in the first and second halves of the samples—1984–2000 and 2001–16, respectively.

the former exhibits a large boom and bust in about 2000, before settling down to a higher value. Another commonly used valuation ratio is the price-smoothed earnings ratio of Shiller (the Cyclically Adjusted Price-Earnings Ratio), which divides the S&P 500 price by a 10-year moving average of real earnings, and is reported in table 1. Though all these ratios are quite volatile, overall, they exhibit only a moderate increase from the first period to the second period. Our analysis emphasizes that this limited increase is puzzling, given the large decline of the risk-free rate (fact 1).

Fact 4: The share of investment in output or in capital has fallen slightly. The bottom two panels of figure 2 depict the behavior of investment. As several researchers have noted recently (Lewis and Eberly 2016; Gutiérrez

and Philippon 2017), investment has been relatively lackluster over the past decade or more; but the magnitude of this decline is quite different depending on exactly how one measures it. Because the price of investment goods falls relative to the price of consumption goods, it is simpler to focus on the expenditure share of GDP (the bottom left panel of figure 2) or the ratio of nominal investment to capital (evaluated at current cost; the bottom right panel). Both ratios ought to be stationary in standard models, and they appear nearly trendless over long samples. Investment spending exhibits a strong cyclical pattern, increasing faster than GDP during expansions and falling faster than GDP during recessions; but overall, both ratios appear to exhibit small to moderate declines across our two subsamples. Table 1 also reports the ratios for the nonresidential sector (that is, business fixed investment), which behaves very similarly, indicating that our results are not driven by housing. Note that business fixed investment includes equipment, structures, and intellectual property products. The table also reports two measures of the evolution of the capital-output ratio: first, the ratio of capital at current cost to GDP; and second, the ratio of a real index of capital services (from the Bureau of Labor Statistics, BLS) to real output (which we normalize to 1 in 1984).⁵ Both ratios exhibit an increase of about 0.15 and 0.13, respectively.⁶

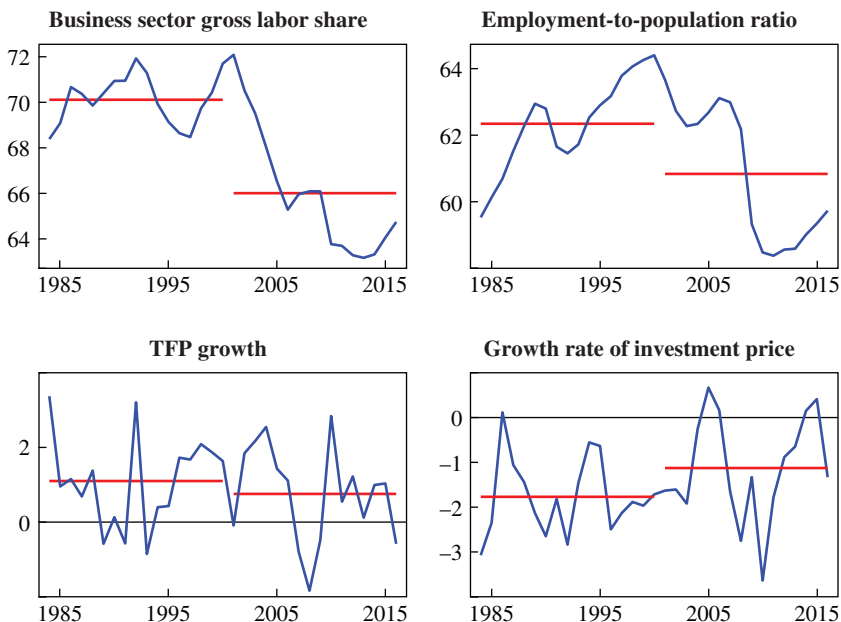
Fact 5: Total factor productivity and investment-specific growth have slowed down, and the employment-to-population ratio has fallen. There has been much public discussion that overall GDP growth has declined over the past couple of decades. This decline is in part attributable to a decline in the employment-to-population ratio, largely due to demographic factors (Aaronson and others 2015), shown in the top right panel of figure 3. However, the decline between the two samples in output per worker growth is still large, from about 1.8 to 1.2 percent a year, according to table 1. This decline is largely driven by lower total factor productivity (TFP) growth and lower investment-specific technical progress. Table 1 shows that the growth rate of John Fernald's (2015) TFP measure goes from 1.1 percent a year to less than 0.8 percent a year, while the growth rate of the relative

5. This index aggregates underlying capital goods using rental prices, which is the correct measure for an aggregate production function. In contrast, capital at current cost is a nominal value that sums purchase prices.

6. Over the long term, these ratios behave differently. The BLS index has exhibited an upward trend since the mid-1970s due to the decline in the price of investment goods, but this trend has slowed down recently. In contrast, the current cost capital-output ratio is nearly trendless.

Figure 3. U.S. Macroeconomic Trends, 1984–2016^a

Percentage points



Sources: See the online appendix, section 1, for all data sources.

a. The top left panel shows the gross labor share for the nonfinancial corporate sector, measured as the ratio of nonfinancial business labor compensation to gross nonfinancial business value added. The top right panel is the employment-to-population ratio. The bottom left panel shows the growth rate of total factor productivity (TFP). The bottom right panel is the growth rate of the relative price of investment goods and consumption goods. The horizontal lines represent the mean in the first and second halves of the samples—1984–2000 and 2001–16, respectively.

price of investment goods to nondurable and service consumption goes from about –1.8 percent to –1.1 percent a year. These series are depicted in the bottom panels of figure 3.

Fact 6: The labor share has fallen. Finally, the top left panel of figure 3 presents a measure of the gross labor share for the nonfinancial corporate sector; table 1 also includes a measure that covers the entire U.S. economy. As has been noted by many researchers (Karabarbounis and Neiman 2014; Elsby, Hobijn, and Şahin 2013; Rognlie 2015), the labor share exhibits a decline, especially after 2000 in the United States.

Of course, all these facts are somewhat difficult to ascertain graphically, given the short-term samples and the noise in some series. This leads us to evaluate the statistical significance of these changes.

II.B. Statistical Evaluation

To summarize the trends in these series in a more formal way, table 1 reports several statistics for the series presented in figures 1 through 3 as well as for alternative series that capture the same concepts. The first through fourth columns of table 1 report the means in the first and second subsamples, which are depicted in figures 1 through 3 as horizontal lines, together with standard errors. The fifth column of table 1 reports the difference between the means in the second and first samples, and the sixth column is the associated standard error. The seventh column is the regression coefficient of the variable of interest on a linear time trend, and the eighth column is the associated standard error. (All standard errors are calculated using the Newey-West method with five annual lags.)

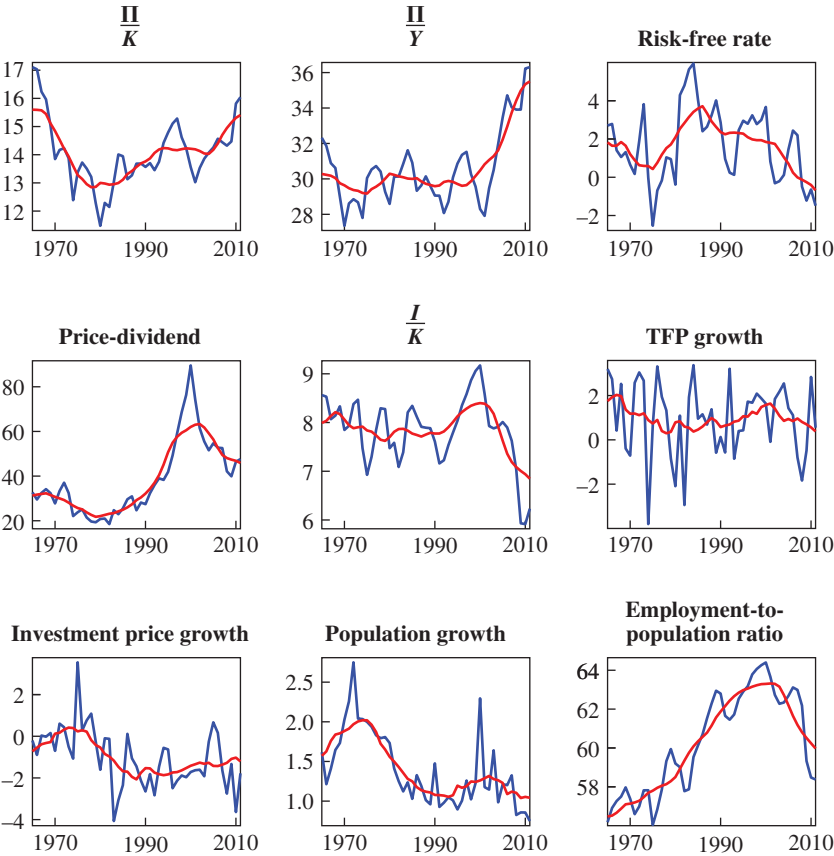
Given the persistence of the series and the relatively short sample, statistical significance should be assessed cautiously. With this caveat, table 1 shows that for some indicators, there is little evidence of a break between the samples, while for others, there is clear evidence of a break. Specifically, interest rates, the labor share, and the investment-capital ratios are markedly lower in the second sample. Conversely, valuation ratios and the return on capital appear fairly stable. Growth measures, such as TFP growth, are substantively smaller in the second sample, but the change is not necessarily statistically significant.

II.C. Longer Historical Trends

Figure 4 presents the evolution of nine of the moments we described above, but over a longer sample, since 1950. (These nine moments will be our estimation targets below.) For clarity, we add an 11-year centered moving average to each series, so we depict the evolution from 1955 to 2011. One motivation for studying a longer sample is that real interest rates were also low in the 1970s and to some extent the 1960s, and hence one question is whether the abnormal period is the early 1980s, when real interest rates were high. The figure shows, however, that the similarities between the 1960s or 1970s and the 2000s are limited to a few variables. It is true that profitability was high in the 1960s, but the price-dividend ratio was lower, and the labor share and the investment-capital ratio were relatively high, in contrast to the more recent period. Overall, neither the 1960s nor the 1970s are similar in all respects to the 2000s. Moreover, a serious consideration of the role of inflation is warranted to study the 1970s and early

Figure 4. Macroeconomic and Financial Trends, 1965–2011^a

Percentage points



Sources: See the online appendix, section 1, for all data sources.

a. This figure presents the nine series used in our estimation exercise over the 1965–2011 sample, together with an 11-year centered moving average.

1980s, as inflation likely affected many of the macroeconomic aggregates depicted here. This is why, for now, we focus on the post-1984 sample. However, below we present some results starting in 1950 to illustrate what our approach implies for these earlier periods.

III. The Model

This section introduces a simple model to account for the macroeconomic and finance moments. Our framework adds macroeconomic risk and monopolistic competition to the standard neoclassical growth model. Given our focus on medium-run issues, we abstract from nominal rigidities and adjustment costs.

III.A. The Model

We consider a standard dynastic model with inelastic labor supply. To highlight the role of risk, we use Epstein-Zin preferences:

$$(1) \quad V_t = \left((1 - \beta) L_t c_{pc,t}^{1-\sigma} + \beta E_t (V_{t+1}^{1-\theta})^{\frac{1-\sigma}{1-\theta}} \right)^{\frac{1}{1-\sigma}}$$

where V_t is utility, L_t is population size (which is exogenous and deterministic), $c_{pc,t}$ is per capita consumption at time t , σ is the inverse of the intertemporal elasticity of substitution of consumption (IES), and θ is the coefficient of relative risk aversion. We assume that labor supply is exogenous and equal to $N_t = \bar{N}L_t$, where \bar{N} is a parameter that captures the employment-population ratio.

Final output is produced using a constant return to scale from differentiated inputs,

$$Y_t = \left(\int_0^1 y_{i,t}^{\frac{\varepsilon-1}{\varepsilon}} di \right)^{\frac{\varepsilon}{\varepsilon-1}}$$

where $\varepsilon > 1$ is the elasticity of substitution. These intermediate goods are produced using a Cobb-Douglas production function,

$$y_{i,t} = Z_t k_{i,t}^\alpha (S_t n_{i,t})^{1-\alpha}$$

where $k_{i,t}$ and $n_{i,t}$ are capital and labor in firm i at time t , Z_t is an exogenous deterministic productivity trend, and S_t is a stochastic productivity process, which we assume to be a martingale:

$$(2) \quad S_{t+1} = S_t e^{\chi_{t+1}}$$

where χ_{t+1} is independent and identically distributed (*iid*).

Capital is accumulated using a standard investment technology, but is subject to an aggregate “capital quality” shock ψ_{t+1} , which we also assume to be *iid*:

$$k_{i,t+1} = ((1 - \delta)k_{i,t} + Q_i x_{i,t}) e^{\psi_{t+1}}.$$

Here Q_i is an exogenous deterministic trend reflecting investment-specific technical progress, as given by Jeremy Greenwood, Zvi Hercowitz, and Per Krusell (1997). The relative price of investment and consumption goods is $\frac{1}{Q_i}$.

Capital and labor can be reallocated frictionlessly across firms at the beginning of each period after the shocks X and ψ have been realized. Given the constant-return-to-scale technology, firms then face a constant (common) marginal cost. It is easy to see that the economy aggregates to a production function (see the online appendix, section 2, for details):⁷

$$(3) \quad Y_t = Z_t K_t^\alpha (S_t N_t)^{1-\alpha}$$

and that markups distort the firms’ first-order conditions, leading to

$$(4) \quad \frac{(1 - \alpha)Y_t}{N_t} = \mu w_t$$

$$(5) \quad \alpha \frac{Y_t}{K_t} = \mu R_t$$

where $\mu = \frac{\varepsilon}{\varepsilon - 1} > 1$ is the gross markup, w_t is the real wage, and R_t is the rental rate of capital.

Moreover, the law of motion for capital accumulation also aggregates,

$$(6) \quad K_{t+1} = ((1 - \delta)K_t + Q_t X_t) e^{\psi_{t+1}}.$$

The choice of investment is determined by the (common) marginal product of capital, leading to the Euler equation:

$$(7) \quad E_t [M_{t+1} R_{t+1}^K] = 1$$

7. The online appendixes for this and all other papers in this volume may be found at the *Brookings Papers* web page, www.brookings.edu/bpea, under “Past BPEA Editions.”

where M_{t+1} is the real stochastic discount factor and R_{t+1}^K is the return on capital, which is given by

$$(8) \quad R_{t+1}^K = \left(\frac{\alpha Y_{t+1}}{\mu K_{t+1}} + \frac{1 - \delta}{Q_{t+1}} \right) Q_t e^{\Psi_{t+1}}.$$

This expression is a standard user cost formula, which incorporates the rental rate of capital of equation 5 but also depreciation, the price of investment goods, and the capital quality shock. Given the preferences assumed in equation 1, the stochastic discount factor is

$$(9) \quad M_{t+1} = \beta \left(\frac{c_{pc,t+1}}{c_{pc,t}} \right)^{-\sigma} \left(\frac{V_{pc,t+1}}{E_t (V_{pc,t+1})^{\frac{1}{1-\theta}}} \right)^{\sigma-\theta}$$

where $V_{pc,t}$ is the utility normalized by population, $V_{pc,t} = \frac{V_t}{1/L_t^{-\sigma}}$.

The resource constraint reads

$$(10) \quad C_t + X_t = Y_t$$

where $C_t = L_t c_{pc,t}$ is total consumption, and X_t are investment expenses measured in consumption good units.

The equilibrium of this economy is $\{c_{pc,t}, C_t, X_t, K_t, Y_t, R_{t+1}^K, M_{t+1}, V_{pc,t}, V_t\}$, which solves the system of equations 1 through 10, given the exogenous processes $\{L_t, Z_t, Q_t, S_t, \chi_{t+1}, \Psi_{t+1}\}$. As is well known, in general such a model admits no closed-form solution. Many researchers build their intuition by studying either the nonstochastic steady state or numerical approximations. This makes it somewhat difficult to explain the role that macroeconomic risk plays. We show, in contrast, that for an interesting special case, our model can be solved easily for a “risky balanced growth path.”

III.B. Risky Balanced Growth

We make two simplifying assumptions. First, to obtain a balanced growth path, we make the usual assumption that the exogenous trends (population, L_t ; TFP, Z_t ; and investment-specific technical progress, Q_t) all grow at possibly different constant rates, so that $\frac{L_{t+1}}{L_t} = 1 + g_L$, $\frac{Z_{t+1}}{Z_t} = 1 + g_Z$,

$\frac{Q_{t+1}}{Q_t} = 1 + g_Q$ for all $t \geq 0$. Second, we assume that the productivity shock and capital quality shock are equal:

$$\chi_{t+1} = \psi_{t+1}.$$

In this case, it is straightforward to verify that the equilibrium has the following structure:

$$X_t = T_t S_t x^*$$

$$Y_t = T_t S_t y^*$$

and similarly for C_t , while for capital and utility, we have $K_t = T_t S_t Q_t k^*$ and $V_t = L_t^{\frac{\sigma}{1-\sigma}} T_t S_t v^*$. Here, the lowercase, starred values denote constants; S_t is the stochastic trend defined in equation 2 corresponding to the accumulation of past productivity / capital quality shocks X_t ; and T_t is a deterministic trend, defined as

$$T_t = L_t Z_t^{\frac{1}{1-\alpha}} Q_t^{\frac{\alpha}{1-\alpha}}$$

whose growth rate is denoted g_T and satisfies the usual condition:

$$(11) \quad 1 + g_T = (1 + g_L)(1 + g_Z)^{\frac{1}{1-\alpha}}(1 + g_Q)^{\frac{1}{1-\alpha}}$$

where α is the Cobb-Douglas parameter, g_Q is the rate of growth of investment-specific technical progress, g_L is population growth, and g_Z is productivity growth. The trend growth rate of output per capita is


$$1 + g_{PC} = \frac{1 + g_T}{1 + g_L}.$$

Finally, the stochastic discount factor is

$$(12) \quad M_{t+1} = \beta(1 + g_{PC})^{-\sigma} e^{-\theta \chi_{t+1}} E(e^{(1-\theta)\chi_{t+1}})^{\frac{\theta-\sigma}{1-\theta}}$$

where θ is risk aversion and σ is the inverse of the IES. We can then easily calculate all objects of interest in the model, including x^* , y^* , as we show in the next section and in section 1 of the online appendix.

Figure 5 presents an example of the time series produced by the model. The equilibrium corresponds to a “balanced growth path,” but one where

Figure 5. An Example of the Time Series Produced by the Model^a

Source: Authors' calculations.

a. The figure presents an example of the time series generated by the model—in the top panel, output, consumption, and investment (in log); in the bottom panel, return on capital and the risk-free rate. In this example, the economy is affected by two realizations of χ shocks, at $t = 4$ and $t = 57$.

macroeconomic risk still affects decisions and realizations. Specifically, the realization of the macroeconomic shock χ_{t+1} affects the stochastic trend S_{t+1} and hence X_{t+1} , Y_{t+1} , and so on, while the effect of risk, conversely, is reflected in the constants x^* , y^* . The bottom line is that the “big ratios”—such as $\frac{X_t}{Y_t}$, $\frac{\Pi_t}{Y_t}$, $\frac{\Pi_t}{K_t/Q_t}$, and the like—are constant, as in the standard Kaldor calculations, but now incorporate risk; we discuss these ratios in the next section.⁸ This result holds regardless of the probability distribution of χ_{t+1} .

8. Of course, the economy can also exhibit transitional dynamics if its initial capital is too low or too high, before it reaches the “risky balanced growth” path.

The treatment of deterministic trends is completely standard. What is less standard is that in our model, a common stochastic trend affects all variables equally, which generates great tractability. In the standard real business cycle model, there are no capital quality shocks—that is, $\psi_{t+1} = 0$, and a (permanent) productivity shock χ_{t+1} leads to a transition as the economy adjusts its capital stock to the newly desired level, before eventually reaching the new steady state. By assuming $\chi_{t+1} = \psi_{t+1}$, this transition period is eliminated because the capital stock “miraculously” adjusts by the correct amount. This simplifies the solution of the model because agents’ expectations of future paths are now easy to calculate.⁹ The capital quality shock is also important if the economy is to generate a significant equity premium, for it makes the return on capital volatile rather than bounded below by $1 - \delta$.

III.C. Model Implications

This subsection presents model implications for the “big ratios” and other key moments of interest along the risky balanced growth path. We present the Euler equation, which leads to a standard user cost calculation, and then discuss valuation ratios and rates of return.

It is useful to define the composite parameter

$$\beta^* = E_t(M_{t+1}e^{\chi_{t+1}})$$

which equals

$$(13) \quad \beta^* = \beta(1 + g_{PC})^{-\sigma} \times E(e^{(1-\theta)\chi_{t+1}})^{\frac{1-\sigma}{1-\theta}}$$

and its rate of return version $r^* = \frac{1}{\beta^*} - 1 \simeq -\log \beta^*$, which satisfies

$$(14) \quad r^* \simeq \rho + \sigma g_{PC} + \sigma \frac{1 - \frac{1}{\theta}}{1 - \theta} \log E(e^{(1-\theta)\chi_{t+1}})$$

9. Because we do not study the actual responses to χ_{t+1} shocks, there is little loss in this simplification; what is key for us is that agents regard the future as uncertain, and that bad realizations of χ_{t+1} will have reasonable consequences (for example, a low return on capital), which lead agents ex ante to adjust their choices, such as for investment. This argument—formulated by Gabaix (2011) and Gourio (2012)—can be applied to larger models; for instance, for New Keynesian models with disaster risk, see Gourio, Kashyap, and Sim (2018); and Isoré and Szczerbowicz (2017).

where $\rho = \frac{1}{\beta} - 1 \simeq -\log \beta$.¹⁰ The parameter r^* will turn out to equal in equilibrium the expected return on capital, and to be a “sufficient statistic” to solve for the “big ratios”—that is, we do not need to know ρ (that is, β) θ , σ , or the distribution of χ , but only r^* .

CAPITAL ACCUMULATION To solve the model, we use the Euler equation 7, which along the risky balanced growth path reads

$$(15) \quad \frac{1}{\beta^*} = \left(\frac{\alpha}{\mu} Q^* \left(\frac{k^*}{N} \right)^{\alpha-1} \frac{1}{1+g_Q} + \frac{1-\delta}{1+g_Q} \right)$$

where Q^* is the level of investment technical progress Q_t , that is, $Q_t = Q^*(1+g_Q)^t$; so $\frac{1}{Q^*}$ affects the level of the relative price of investment and consumption. This equation pins down k^* and the capital-labor ratio, and it generalizes the familiar condition of the neoclassical growth model to incorporate risk, through β^* . We can rewrite this as the equality of the user cost of capital and marginal revenue:

$$(16) \quad \frac{1}{Q^*} (r^* + \delta + g_Q) \approx \frac{\alpha}{\mu} \left(\frac{k^*}{N^*} \right)^{\alpha-1}.$$

Equation 16 directly shows how higher market power or a higher required *risky* return lowers the desired capital-labor ratio.

To calculate the other big ratios, first note that K_t/Q_t is the capital stock, evaluated at current cost. The capital-output ratio is obtained from equation 16 as

$$(17) \quad \frac{K_t/Q_t}{Y_t} \approx \frac{\alpha}{\mu} \frac{1}{r^* + \delta + g_Q}$$

10. Here and thereafter, the \simeq sign reflects the first-order approximation $\log(1+x) \simeq x \simeq \frac{1}{1-x}$.

and the investment-capital ratio is

$$(18) \quad \frac{X_t}{K_t/Q_t} \approx g_\varrho + g_\tau + \delta$$

which reflects the familiar balanced growth relation. Last, the investment-output ratio is obtained by combining equations 17 and 18:

$$(19) \quad \frac{X_t}{Y_t} \approx \frac{\alpha}{\mu} \frac{g_\tau + \delta + g_\varrho}{r^* + \delta + g_\varrho}.$$

INCOME DISTRIBUTION The labor share in gross value added is, using equation 4,

$$(20) \quad s_L = \frac{w_t N_t}{Y_t} = \frac{1 - \alpha}{\mu}$$

and hence the measured capital share is

$$s_K = 1 - s_L = \frac{\mu + \alpha - 1}{\mu}.$$

This capital share can be decomposed into a pure profit share, which rewards capital owners for monopoly rents, and a true capital remuneration share, corresponding to rental payments to capital, that is, $s_K = s_\pi + s_C$, with

$$(21) \quad s_\pi = \frac{\mu - 1}{\mu}$$

and

$$(22) \quad s_C = \frac{\alpha}{\mu}.$$

VALUATION RATIOS The firm value is the present discounted value of the dividends $D_t = \Pi_t - X_t$. In equilibrium, this value equals the value of installed capital plus monopolistic rents. Formally, the ex-dividend firm value P_t satisfies the standard recursion

$$P_t = E_t(M_{t+1}(P_{t+1} + D_{t+1})).$$

Given that the equilibrium is *iid*, the price-dividend ratio is constant, and satisfies the familiar Gordon growth formula:

$$(23) \quad \frac{P^*}{D^*} = \frac{\beta^* (1 + g_T)}{1 - \beta^* (1 + g_T)} \approx \frac{1 + g_T}{r^* - g_T}.$$

Tobin's Q is defined as

$$(24) \quad \frac{P_t}{K_t/Q_t} \approx (1 + g_T) \left(1 + \frac{\mu - 1}{\alpha} \frac{r^* + \delta + g_Q}{r^* - g_T} \right).$$

Because we do not incorporate adjustment costs, Tobin's Q equals (approximately) 1 when there is no market power—that is, $\mu = 1$.¹¹ But if there is some market power, the value of Tobin's Q depends on several parameters, which affect (1) the size of the economy and hence the rents, and (2) the discount rate applied to all future rents.

RATES OF RETURN We now compare three benchmark rates of return in this economy: the risk-free rate, the return on equity, and the profitability of capital, which is often used in macroeconomics as a proxy for the marginal product of capital. The gross risk-free rate (which can be priced, even though it is not traded in equilibrium) is

$$RF = \frac{1}{E(M_{t+1})} = \frac{E(e^{(1-\theta)\chi_{t+1}})}{\beta^* E(e^{-\theta\chi_{t+1}})}.$$

which we can rewrite as the net risk-free rate, that is, $r^f = RF - 1$:

$$(25) \quad r_f \approx r^* + \log E(e^{(1-\theta)\chi_{t+1}}) - \log E(e^{-\theta\chi_{t+1}})$$

The average profitability of capital can be inferred—as by Gomme, Ravikumar, and Rupert (2011) and Casey Mulligan (2002)—as the ratio of

11. Tobin's Q is usually defined as $\frac{P_t}{K_{t+1}/Q_{t+1}}$, but with capital quality shocks K_{t+1} unknown at time t , leading us to adopt this definition, which creates the $1 + g_T$ wedge. One could also define Tobin's Q as $\frac{P_t}{E_t K_{t+1}/Q_{t+1}}$, which eliminates the wedge provided that $E_t e^{\chi_{t+1}} = 1$, an assumption that we maintain through most of the paper.

(measured) profits to the stock of capital. We denote it *MPK* because it is often used as a proxy for the marginal product of capital, though this holds only under constant return to scale and perfect competition. This *MPK* can be calculated either gross or net of depreciation. For instance, in gross terms, we have

$$(26) \quad MPK = \frac{\Pi_t}{K_t/Q_t} = \frac{\mu + \alpha - 1}{\alpha} (r^* + \delta + g_Q).$$

Conceptually, this *MPK* exceeds the risk-free rate for three reasons: first, it is gross of both physical and economic depreciation; second, it incorporates profit rents; and third, it is risky. We can decompose the spread between the *MPK* and the risk-free rate to reflect these three components:

$$(27) \quad MPK - r_f = \delta + g_Q + \frac{\mu - 1}{\alpha} (r^* + \delta + g_Q) + r^* - r_f.$$

A main goal of our empirical analysis is to evaluate the importance of these different components.

The expected equity return is defined as

$$E(R_{t+1}) = E\left(\frac{P_{t+1} + D_{t+1}}{P_t}\right)$$

and it is easy to show using equation 23 that

$$(28) \quad E(R_{t+1}) = \frac{1}{\beta^*} E(e^{z_{t+1}}).$$

In the case where $E(e^{z_{t+1}}) = 1$, which we use in our applications, the gross expected return on equity is exactly $\frac{1}{\beta^*}$, and the net return is r^* .

The same expected return also applies the return on physical capital

$R_{t+1}^K = \left(\frac{aY_{t+1}}{\mu K_{t+1}} + \frac{1 - \delta}{Q_{t+1}}\right) Q_t e^{z_{t+1}}$ defined in equation 8. Conceptually, the firm

value here stems from capital and rents, but it turns out that both components have equal risk exposure and hence equal expected returns.

Finally, the equity risk premium (*ERP*) is obtained by combining equations 25 and 28:

$$ERP = \frac{E(R_{t+1})}{R_{f,t+1}} = \frac{E(e^{-\theta\chi_{t+1}})E(e^{\chi_{t+1}})}{E(e^{(1-\theta)\chi_{t+1}})}.$$

III.D. Comparative Statistics

We now use the expressions developed in the previous subsection to illustrate key comparative statics of the risky balanced growth path. These statics are useful for understanding the identification of our model. Most of the parameters have the usual effects; we focus on parameters that are typically absent from the neoclassical growth model, or parameters that play an important role in our empirical results.

THE EFFECT OF RISK The effect of higher risk on macroeconomic variables is mediated through β^* . The cleanest thought experiment is to consider a shift in the distribution of the shock χ in the sense of second-order stochastic dominance, so that χ becomes more risky. Such a shift reduces $E(e^{(1-\theta)\chi})^{\frac{1}{1-\theta}}$, and hence leads to a lower β^* if and only if $\sigma < 1$, that is, the IES is greater than unity. A lower β^* in turn leads to a lower capital-output ratio, a lower investment-output ratio, and a higher profit-capital ratio, according to equations 17, 19, and 26, respectively. The logic is that risk deters investment in this case, leading to less capital accumulation. This reduction in the supply of capital increases *MPK*, given a stable demand for capital. Moreover, as is well known in the macroeconomic and finance literature, and as shown by equation 23, higher risk decreases the *PD* ratio if the IES is greater than unity. Conversely, if the IES is lower than unity, higher risk leads to a lower expected return, and hence to higher capital accumulation and a higher price-dividend ratio. In the knife-edge case of a unit IES, corresponding to log preferences, risk does not affect the required return on capital r^* and hence does not affect capital accumulation. In all cases, risk has no effect on the labor share or long-term growth (though higher risk has a level effect on capital and GDP, that is, k^* and y^*). The equity risk premium $r^* - r_f$ is increasing in risk, regardless of the IES. The spread between the *MPK* and the risk-free rate is hence increasing in risk, at least if μ is small enough so that the middle term of equation 7 does not dominate the third term.

We have not specified the distribution of the shock χ ; but for some particular distributions, one can obtain exact formulas. For instance, if χ is

normal with variance σ_χ^2 and mean $\mu_\chi = -\frac{\sigma_\chi^2}{2}$, so that an increase in σ_χ is a pure increase in risk, we have, denoting $\hat{\beta} = \beta(1 + g_{PC})^{-\sigma}$,

$$\log \beta^* = \log \hat{\beta} - (1 - \sigma)\theta \frac{\sigma_\chi^2}{2}$$

$$\log RF = -\log \hat{\beta} - (1 + \sigma)\theta \frac{\sigma_\chi^2}{2}$$

$$\log ERP = \theta \sigma_\chi^2$$

These formulas capture the usual effect of risk aversion and the quantity of risk on the *ERP* and the risk-free rate, but are now valid in a production economy, and furthermore β^* links macroeconomic risk to macroeconomic variables such as the capital-output ratio, as discussed above. We provide more discussion in section 2 of the online appendix for different assumptions about the distribution of χ .

THE EFFECT OF SAVINGS SUPPLY In our model, the effects of a change in the discount factor β are the same as a change in risk, because both are mediated through β^* . The one exception is the risk-free rate, which is affected directly by β^* but also directly by risk measures, for example, risk aversion θ or the quantity of risk χ . In the case where the IES is greater than unity, higher β has the same implications as lower risk. Hence, higher savings supply leads to higher capital accumulation, a higher investment-output ratio and a lower marginal product of capital, and a higher price-dividend ratio, while the risk-free rate falls. The spread between the *MPK* and the risk-free rate, shown in equation 27, is little affected by β : β only affects the quantity of rents through r^* , while the equity risk premium $r^* - r_f$ is independent of β .

THE EFFECT OF MARKET POWER One potentially important factor that has been invoked to explain the trends we document is market power. In our model, an increase in μ has no effect on long-term growth, the risk-free rate, or the price-dividend ratio; but it has a significant effect on other variables. Higher markups reduce both the labor share and the “true capital share,” s_c , but increase the pure profit share, s_π . According to equations 19 and 17, higher market power also reduces investment-output and capital-output ratios, as firms have less incentive to build capacity. The spread between the *MPK* and the risk-free rate is increasing in market power (equation 27). Finally, higher market power reduces the level of GDP by reducing capital accumulation.

THE EFFECT OF TREND GROWTH Trend growth, g_T —which can be traced back to productivity growth, population growth, or investment-specific technical growth—affects β^* but also independently affects the ratios of interest. Higher growth generally increases the investment-capital and investment-output ratios and increases the risk-free rate and valuation ratios, while the effect on profitability ratios depends on the exact source of growth.

IV. The Accounting Framework

This section describes our empirical approach and discusses identification.

IV.A. Methodology

We use a simple method of moment estimation. In the interest of clarity and simplicity, we perform an exactly identified estimation with nine parameters and nine moments. In a first exercise, we estimate the model separately over our two samples: 1984–2000 and 2001–16. We then discuss which parameters drive variation in each moment. In a second exercise, we estimate the model over 11-year rolling windows, starting with 1950–61 and ending with 2006–16. In all cases, we fit the model’s risky balanced growth path to the model’s moments. In doing so, we abstract from business cycle shocks, in line with our focus on longer frequencies.¹²

The moments we target are motivated by the observations in the introduction and section I:

(M1) the gross profitability, $\frac{\Pi}{K}$,¹³

(M2) the gross capital share, $\frac{\Pi}{Y}$;

(M3) the investment-capital ratio, $\frac{I}{K}$;

(M4) the risk-free rate, RF ;

12. This exercise involves some “schizophrenia,” because our model assumes that parameters are constant, even though they are estimated to change over time; and when parameters change, the model would exhibit some transitional dynamics, which we abstract from for now; see section VI. Further, the agents inside our model do not understand that parameters might change, let alone anticipate some of these changes.

13. From here on, we denote measured average profitability $\frac{\Pi}{K}$ and the investment rate $\frac{I}{K}$ —that is, we omit Q ; and we denote investment with X .

(M5) the price-dividend ratio, PD ;
 (M6–M8) the growth rates of population, TFP, and investment prices;
 and

(M9) the employment-population ratio.

As we show here, these moments lead to a clear identification of our nine parameters, which are:

(P1) the discount factor, β ;

(P2) risk, modeled as the probability of an economic crisis or “disaster,” p ;

(P3) the markup, μ ;

(P4) the depreciation rate of capital, δ ;

(P5) the Cobb-Douglas parameter, α ;

(P6–P8) the growth rates of TFP, g_z , investment-specific progress, g_Q , and population, g_L ; and

(P9) the labor supply parameter, \bar{N} .

The choice of moments is motivated, of course, by the questions of interest—explaining the joint evolution of interest rates, profitability, investment, valuation, and trend growth—but also by the clarity with which these moments map into estimated parameters. For instance, because we target $\frac{\Pi}{K}$, $\frac{\Pi}{Y}$, and $\frac{I}{K}$ (and because we have taken care to construct these moments in a consistent manner), the model will mechanically match the evolution of the investment-output ratio $\frac{I}{Y}$ or the capital-output ratio $\frac{K}{Y}$. Hence, we could have taken $\frac{I}{Y}$ as a targeted moment, which would have led to the exact same estimates and implications, but the identification is clearer with $\frac{I}{K}$. Beyond this, some changes in identification strategy are possible, however; for instance, one could target the price-earnings ratio instead, or GDP growth per worker; these yield quite similar results.

We also note that the parameters can be mapped into the narratives often put forth when discussing the trends, at least at a high level; in particular, changes in longevity map into a change in the discount factor β ; more generally, changes in savings supply can be captured as changes in β ; changes in the competitive environment are captured by a change in μ ; changes in technology should be reflected in α, δ , or the growth rates of the technological factors g_z and g_Q ; and so on. However, it is also possible that some economic factors affect all our parameters at the same time.

There are three parameters that we do not estimate; we discuss why, and how this affects our results in the next section on identification. The three parameters are the IES $\frac{1}{\sigma}$, the coefficient of risk aversion θ , and the size of macroeconomic shocks b . Specifically, we assume that χ_{t+1} follows a “disaster risk” three-point distribution, that is,

$$\chi_{t+1} = 0 \text{ with probability } 1 - 2p$$

$$\chi_{t+1} = \log(1 - b) \text{ with probability } p$$

$$\chi_{t+1} = \log(1 + b_H) \text{ with probability } p$$

where b_H is chosen so that $E(e^{\chi_{t+1}}) = 1$. We estimate p but fix b (and hence b_H).

IV.B. Identification

In this subsection, we provide a heuristic discussion of identification, and make two main points. First, the identification is nearly recursive, so that it is easy to see which moments affect which parameters. Second, and consequently, the identification of some parameters does not depend on all the data moments.¹⁴

The identification is easily seen to be nearly recursive. First, some parameters are obtained directly as their counterparts are assumed to be observed: population growth, investment price growth (the opposite of g_Q), and the employment-population ratio. The growth rate g_z is next chosen to match measured TFP.¹⁵ One hence obtains g_T , the trend growth rate of GDP,

14. Section 3 of the online appendix includes the matrix of sensitivity of parameters to moments, as suggested by Andrews, Gentzkow, and Shapiro (2017).

15. This step is, however, not completely straightforward, which is why we only say that the identification is nearly recursive. TFP in the data is measured using the revenue-based labor share, which in the model is $s_L = \frac{1 - \alpha}{\mu}$ rather than the cost-based labor share, which in the model is $1 - \alpha$. As a result, the TFP that an economist would measure in our model is

$$g_T - s_L g_N - (1 - s_L) g_K = \left(\frac{s_L}{1 - \alpha} g_z + \left(\frac{s_L \alpha}{1 - \alpha} - (1 - s_L) \right) g_Q \right)$$

and hence is not equal to g_z because $s_L \neq 1 - \alpha$. In particular, matching TFP requires knowing the value of α , which is why it is not fully recursive. This turns out to have relatively small effects in our empirical work.

given by equation 11. The depreciation rate δ is then chosen to match $\frac{I}{K}$ according to the familiar balanced growth relation (equation 18):

$$\frac{I}{K} \simeq \delta + g_\varrho + g_\tau.$$

The model then uses the Gordon growth formula (equation 23) to infer the expected return on risky assets, r^* , given the observed price-dividend ratio:

$$\frac{P^*}{D^*} \simeq \frac{1 + g_\tau}{r^* - g_\tau}.$$

Importantly, to infer r^* , we do not need data on the risk-free rate, or assumptions about the value of β , risk aversion θ , or the distribution of χ .

The next step is to identify the parameters α and μ to match the profit share of output and the ratio of profits to capital, using equations 20 and 27, that is:

$$s_L = \frac{1 - \alpha}{\mu}$$

and

$$MPK = \frac{\mu + \alpha - 1}{\alpha} (r^* + \delta + g_\varrho)$$

where s_L and $MPK = \frac{\Pi}{K}$ are the observables and α and μ the unknowns.

The solution is, denoting by $uc = r^* + \delta + g_\varrho$ the frictionless user cost of capital, to set

$$\mu = \frac{MPK}{s_L MPK + (1 - s_L) uc}$$

and

$$\alpha = \frac{uc(1 - s_L)}{s_L MPK + (1 - s_L) uc}.$$

Intuitively, the first equation infers market power (here, the Lerner index) from the discrepancy between MPK and the frictionless user cost of capital uc . The parameter α is then obtained to fit the observed labor share. A key remark is that our identification of α and μ does not require data on the risk-free rate or making any assumption about risk aversion θ or the distribution of χ —we simply use the sufficient statistic r^* , which has been previously identified.

Economically, our approach boils down to using the traditional Gordon growth formula—which holds in our standard neoclassical framework—to deduce the required return on capital from the price-dividend ratio and the growth rate, and hence to construct a user cost of capital $r^* + \delta + g_Q$ that incorporates risk.¹⁶

At this point, we can also bring in data on the risk-free rate to infer the equity premium $r^* - r_f$. Here again, note that the behavior of the equity premium is therefore inferred without making assumptions about risk aversion θ or the distribution of χ . However, to understand what drives the risk-free rate, one needs to separately infer β , risk aversion θ , and the quantity of risk χ . Doing so requires extra assumptions about these variables and about the IES (which is not identified in our model, given that growth rates are *iid*), as can be seen from equation 14:

$$r^* \simeq \rho + \sigma g_{PC} + \sigma \frac{1 - \frac{1}{\theta}}{1 - \theta} \log E(e^{(1-\theta)\chi_{t+1}}).$$

We present our baseline result with an IES of 2, a rare disaster distribution¹⁷ for χ with a shock of 15 percent ($e^b = 0.85$), a probability p that we estimate, and a risk aversion coefficient of 12. As should be clear by now, none of these choices affects our inferences about α , μ , or the equity premium. Concretely, given these additional assumptions, we can solve for the quantity of risk p that satisfies

$$r^* - r_f = \log E(e^{-\theta\chi_{t+1}}) - \log E(e^{(1-\theta)\chi_{t+1}})$$

16. Our procedure is closely related to the approach of Barkai (2016), the main difference being the way we incorporate risk. Barkai (2016) simply uses a Treasury rate or corporate bond yield to construct the user cost.

17. The asset pricing disaster literature—Barro (2006), Gabaix (2012), Gourio (2012), and Wachter (2013)—often models disasters as much larger shocks; here, the 15 percent decline we assume is roughly in line with the U.S. experience after 2008 (for example, the level of GDP as of 2016 is about 15 percent below what would have been predicted based on a log-linear trend in 2007).

Table 2. Estimated Parameters: Baseline Model^a

<i>Parameter name</i>	<i>Symbol</i>	<i>Estimate</i>		
		<i>1984–2000</i>	<i>2001–16</i>	<i>Difference</i>
Discount factor	β	0.961	0.972	0.012
Markup	μ	1.079	1.146	0.067
Disaster probability	p	0.034	0.065	0.031
Depreciation	δ	2.778	3.243	0.465
Cobb-Douglas	α	0.244	0.243	–0.000
Population growth	g_N	1.171	1.101	–0.069
Total factor productivity growth	g_Z	1.298	1.012	–0.286
Investment in technical growth	g_Q	1.769	1.127	–0.643
Labor supply	\bar{N}	0.623	0.608	–0.015

Source: Authors' calculations.

a. This table reports the estimated parameters in our baseline model for each of the two subsamples, 1984–2000 and 2001–16, and the change between subsamples.

and we can then use the equation above for r^* to deduce ρ —that is, β . In section VI, we present the results when the IES is instead assumed to be 0.5, and we also discuss results when we choose other distributions for χ , or if we instead fix the amount of risk and estimate the risk aversion coefficient θ .

V. Empirical Results

We first compare the two subsamples, and then we contrast the results with more standard macroeconomic approaches that do not entertain a role for risk. Finally, we present results over rolling windows in a long sample.

V.A. A Comparison of Two Subsamples

Table 2 shows the estimated parameters for each subsample and the change of parameters between subsamples. Overall, our results substantiate many of the narratives that have been advanced and that we mention in the introduction. The discount factor β rises by about 1.2 points, reflecting higher savings supply. Market power increases significantly, by about 6.7 points. Technical progress slows down, and the labor supply falls (relative to population). The model also estimates a significant increase in macroeconomic risk (the probability of a crisis), which goes from about 3.4 percent to 6.5 percent a year. We will return to the interpretation of this result below. Conversely, there is only moderate technological change: Depreciation increases, reflecting the growing importance of high-depreciation capital such as computers, but the Cobb-Douglas parameter remains fairly stable. This stability of the production function is an

interesting result. Overall, the model gives some weight to four of the most popular explanations (β , μ , p , g s). But exactly how much does each story explain?

Table 3 provides one answer. By construction, the model fits perfectly all nine moments in each subsample using the nine parameters. We can decompose how much of the change in each moment between the two subsamples is accounted for by each parameter. Because our model is non-linear, this is not a completely straightforward task; in particular, when changing a parameter from a first subsample value to a second subsample value, the question is at which value to evaluate the other parameters (for example, the first or second subsample value). If the model were linear, or the changes in parameters were small, this would not matter; but such is not the case here, in particular for the price-dividend ratio. In this table, we simply report the average over all possible orders of changing parameters, as we move from the first to the second subsamples.¹⁸

Overall, we see that the decline in the risk-free rate of about 3.1 percent (314 basis points) is explained mostly by two factors: higher perceived risk p , and higher savings supply β , with lower growth playing only a moderate role.¹⁹ Why does the model not attribute all the change in the risk-free

18. Formally, let $\theta^a = (\theta_1^a, \dots, \theta_K^a)$ and $\theta^b = (\theta_1^b, \dots, \theta_K^b)$ denote the parameter vectors in subsamples a and b respectively, and consider a model moment that is a function of the parameters: $m = f(\theta)$. Consider a permutation $\sigma: [1, K] \rightarrow [1, K]$ that describes an order in which we change parameters from their initial to final value; we first change $\theta_{\sigma(1)}$, then $\theta_{\sigma(2)}$, and so on. Then calculate the change implied when we change parameter $l \in [1, K]$ along this order, that is,

$$\Delta_l(\sigma) = f(\theta_{z_2}^b; \theta_{-z_2}^a) - f(\theta_{z_1}^b; \theta_{-z_1}^a)$$

where $z_2 = \sigma(1:\sigma^{-1}(l))$ are the parameters that have been switched already from initial to final values, and $z_1 = \sigma(1:\sigma^{-1}(l) - 1)$ the ones which are not switched yet. The change in m due to parameter $l \in [1, K]$ is defined as

$$\Delta_l = \frac{1}{N_\sigma} \sum_{\sigma} \Delta_l(\sigma)$$

where the sum ranges over all possible permutations. By construction, $\sum_{l=1}^K \Delta_l = f(\theta^b) - f(\theta^a)$

accounts exactly for the model implied change in the moment, which, because the model fits the targeted moments perfectly, and also accounts exactly for the change in the data: $f(\theta^b) - f(\theta^a) = m^b - m^a$. In the online appendix, we also report the upper and lower bounds when we consider all possible combinations of other parameters. This provides a way to bound the importance of each factor.

19. This conclusion does depend somewhat on our assumed IES, as we discuss in detail below.

Table 3. Contributions of Parameters to Changes in Targeted Moments^a

<i>Parameter</i>	<i>Targeted moment</i>			<i>Contribution of each parameter to change in moment</i>							
	<i>1984–2000</i>	<i>2001–16</i>	<i>Difference</i>	β	μ	p	δ	α	g_N	g_Z	g_O
Gross profitability	14.01	14.89	0.88	-1.88	2.76	0.76	0.68	0.00	-0.00	-0.29	-1.15
Capital share	29.89	33.99	4.10	0.00	4.13	0.00	0.00	-0.03	0.00	0.00	0.00
Risk-free rate	2.79	-0.35	-3.14	-1.22	0.00	-1.62	0.00	-0.00	-0.00	-0.19	-0.10
Price-dividend ratio	42.34	50.11	7.78	30.67	0.00	-13.19	0.00	-0.02	-1.86	-5.07	-2.76
Investment-capital ratio	8.10	7.23	-0.88	0.00	0.00	0.00	0.47	-0.00	-0.07	-0.39	-0.88
Growth of TFP	1.10	0.76	-0.34	0.00	-0.14	0.00	0.00	-0.00	-0.00	-0.26	0.06
Growth of investment price	-1.77	-1.13	0.64	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.64
Growth in population	1.17	1.10	-0.07	0.00	0.00	0.00	0.00	0.00	-0.07	0.00	0.00
Employment-population ratio	62.34	60.84	-1.51	0.00	0.00	0.00	0.00	0.00	0.00	0.00	-1.51

Source: Authors' calculations.

a. TFP = total factor productivity. This table reports the targeted moments in each of the two subsamples, 1984–2000 and 2001–16, as well as the change between samples, and the contribution of each parameter to each change in moment, so that the 3rd column equals the sum of the 4th through 12th columns. See the text for details.

rate to savings supply? Simply because it would make it impossible to match other moments, in particular the *PD* ratio. Even as it is, if only the change in savings supply β were at work, the *PD* ratio would increase by over 30 points. The model attributes offsetting changes to risk and growth, explaining in this way that the *PD* ratio increased only moderately over this period, despite the lower interest rates.

Similarly, profitability would decrease by almost 2 percentage points if the change in β was the only one at work—all rates of return ought to fall if the supply of savings increases. The model reconciles the stable profitability with the data by inferring higher markups and higher risk. Overall, we see how the model needs multiple forces to account for the lack of changes observed in some ratios. The higher capital share is attributed entirely to higher markups, as capital-biased technical change appears to play little role.

We can now use these model estimates to explain the evolution of some other moments; these are reported in table 4. First, as we discussed in section III (equation 27), the spread between the measured *MPK* and the risk-free rate can be decomposed in three components:

$$MPK - r_f = \delta + g_Q + \frac{\mu - 1}{\alpha}(r^* + g_Q + \delta) + r^* - r_f$$

where the three components are depreciation ($\delta + g_Q$), rents, and risk ($r^* - r_f$). We can calculate this decomposition in the model using the estimated parameters. The table reveals that depreciation changed little overall—faster physical depreciation is offset by slower economic depreciation—but the rents and risk components both rise by about 2 percentage points. (An alternative way to decompose the change in spread is to read, in the first row, the decomposition of the change in spread due to each parameter change; this yields a similar answer, as the increases in μ and in p account for the bulk of the increase in the spread.)

We also report the model implied equity return and equity premium. Though not a direct target, we estimate a sizable equity premium, of nearly 5 percent a year in the recent sample. (This premium assumes no leverage; see section VI for a discussion of leverage.) More interestingly, the premium has increased by about 2 percentage points since 2000. In total, expected equity returns have fallen by almost 1 percentage point because the decline in the risk-free rate is larger than this increase in the equity premium.

Regarding valuation ratios, we have already emphasized the moderate increase of the price-dividend ratio due to offsetting factors. Table 4 also

Table 4. Contributions of Parameters to Other Moments^a

Parameter	Model-implied moments			Contribution of each parameter								
	1984–2000	2001–16	Difference	β	μ	p	δ	α	g_N	g_Z	g_Q	\bar{N}
A. MPK–RF spread												
Total spread	11.22	15.24	4.02	–0.66	2.76	2.39	0.68	0.00	–0.00	–0.10	–1.05	–0.00
Depreciation	4.55	4.37	–0.18	0.00	0.00	0.00	0.47	0.00	0.00	0.00	–0.64	0.00
Market power	3.39	5.55	2.17	–0.59	2.73	0.24	0.21	0.00	–0.00	–0.09	–0.35	0.00
Risk premium	3.15	5.23	2.08	–0.05	0.00	2.14	0.00	–0.00	0.00	–0.01	–0.00	0.00
B. Rate of return												
Equity return	5.85	4.90	–0.96	–1.22	0.00	0.56	0.00	–0.00	–0.00	–0.19	–0.10	0.00
Equity premium	3.07	5.25	2.18	0.00	0.00	2.18	0.00	0.00	0.00	0.00	0.00	0.00
Risk-free rate	2.79	–0.35	–3.14	–1.22	0.00	–1.62	0.00	–0.00	–0.00	–0.19	–0.10	0.00
C. Valuation ratios												
Price-dividend	42.34	50.11	7.78	30.67	0.00	–13.19	0.00	–0.02	–1.86	–5.07	–2.76	0.00
Price-earnings	17.85	25.79	7.94	10.16	5.08	–4.57	–0.35	0.00	–0.59	–1.47	–0.34	–0.00
Tobin's Q	2.50	3.84	1.34	1.05	1.34	–0.48	0.11	0.00	–0.08	–0.28	–0.31	–0.00
D. Income shares												
Share of labor	70.11	66.01	–4.10	0.00	–4.13	0.00	0.00	0.03	0.00	0.00	0.00	0.00
Share of capital	22.59	21.24	–1.35	0.00	–1.33	0.00	0.00	–0.03	0.00	0.00	0.00	0.00
Share of profit	7.30	12.76	5.46	0.00	5.46	0.00	0.00	0.00	0.00	0.00	0.00	0.00
E. Macroeconomy												
K/Y	2.13	2.28	0.15	0.29	–0.13	–0.12	–0.11	–0.00	0.00	0.04	0.18	–0.00
I/Y	17.28	16.50	–0.78	2.20	–1.03	–0.90	0.23	–0.02	–0.16	–0.52	–0.59	0.00
Trend Y (% change)	—	—	–0.30	4.18	–1.95	–1.70	–1.52	–0.07	0.00	0.65	2.56	–2.45
Trend I (% change)	—	—	–4.95	17.18	–8.02	–6.98	–0.12	–0.20	–0.94	–2.45	–0.96	–2.45

Source: Authors’ calculations.

a. MPK = marginal product of capital; RF = risk-free rate. This table reports some moments of interest calculated in the model using the estimated parameter values for each of the two subsamples, 1984–2000 and 2001–16, as well as the change between samples, and the contribution of each parameter to each moment change.

shows the analysis of the price-earnings ratio and Tobin's Q . The latter increases significantly, from about 2.5 to 3.8 between the two samples, reflecting both the increase in market power and the effect of the change in discount rates at which these rents are discounted.

The model also speaks to the income distribution between labor, capital, and rents. The approach taken here is that we accurately observe the payments to labor in the data, and cannot easily split the remainder between capital and profits. In the model, we can study the decomposition and how it changes between the two subsamples. The decline of about 4 points in the labor share is accompanied by an even larger increase in the profit share, of about 5 points, so that the capital share actually declines slightly.

Finally, we can use the model to see the effect of these changes on macroeconomic variables—for instance, the capital-output and investment-output ratios. On one hand, a higher savings supply pushes investment up, leading to more capital accumulation. For instance, the change in β would push the investment-output ratio up by over 2 percentage points, while in the data it fell. On the other hand, rising market power and rising risk push investment down. Our model hence accounts for the coexistence of low investment and low interest rates. Note also that higher depreciation also requires more investment along the balanced growth path, while lower growth implies less investment. The model hence produces a fairly nuanced decomposition for the evolution of this ratio.

We can also ask what is the effect of each parameter on the level of GDP or investment.²⁰ For instance, higher market power discourages capital accumulation and reduces output. It is easy to show that the elasticity of GDP to markups in this model is $-\frac{\alpha}{1-\alpha}$, or about -0.32 for our estimate of α . Given the fact that estimated markups rise by 6.2 percent ($= 6.7/1.079$), the effect on GDP is about -0.32×6.2 , or about -2 percentage points (-1.95 percent in table 4). Here, too, there are several counteracting factors, however, which imply that the overall level effect on GDP is small (about -0.30 percent). In particular, a higher savings supply and lower economic depreciation lead to higher capital accumulation, while higher risk leads to lower capital accumulation. Investment is more negatively affected by the changes, with a level effect of about -5 percentage points, owing largely

20. By level of GDP we mean y^* , that is, the level of GDP once the proper deterministic and stochastic trends have been removed. We abstract from the growth effects—for example, a higher g_z or g_Q has the mechanical effect of steepening the overall path of GDP.

to markups and risk, but also to lower growth and a lower employment-population ratio.

V.B. A Comparison with Macroeconomic Approaches

It is interesting to compare our results with alternative procedures followed by macroeconomists. Indeed, our empirical exercise is essentially the calibration of the “steady state” of a very-bare-bones dynamic stochastic general equilibrium (DSGE) model. Any DSGE model writer faces the same issues as we do to fit these key moments.

Indeed, real business cycle modelers are aware of a trade-off between fitting the capital-output ratio and the risk-free interest rate. Because these models also target the labor share, the discrepancy precisely reflects the gap between the *MPK* (the profit-capital ratio) and the risk-free interest rate. Often, modelers reject short-term Treasury interest rates as measures of the rate of return on capital, noting that these securities have special safety and liquidity attributes, which are not explicitly modeled.²¹ Mechanically, these models consider that the observed risk-free rate equals the model risk-free rate times an unobserved convenience yield e^{ξ} . This yields an additional parameter ξ to estimate. At the same time, these models have traditionally abstracted from aggregate market power, setting $\mu = 1$,²² and from risk, so that $p = 0$, and have not explicitly targeted the price-dividend ratio. The assumptions lead to a well-defined exactly identified exercise with eight moments (our baseline, minus the price-dividend ratio) and eight parameters (our baseline, plus the liquidity wedge ξ , less market power μ and risk p), which is an alternative to our approach. The last three columns of table 5 present the results from this exercise, which we call the “macro-without-markups” approach.

This approach leads to a much higher value of α and “explains” the decline of the labor share by an increase of α . The decline of the Treasury rate, and the growing gap between the *MPK* and this rate, are fully accounted for by a very large, and growing, liquidity premium, which equals about $-\xi = 6.1$ percent in the first sample and about 10.2 percent in the second sample. We find both the level and change in this wedge to be implausible.

21. See, for instance, Campbell and others (2017) for a presentation of the Federal Reserve Bank of Chicago’s DSGE model, which, based on Fisher (2015), introduces a liquidity wedge that accounts for the discrepancy between the rate of return of capital and the risk-free rate.

22. New Keynesian models are an important exception, but market power is often set on an a priori basis in these studies (for example, a markup of 15 percent), and profits are offset in a steady state by fixed costs.

Table 5. Parameter Estimates: Baseline versus Macroeconomic Approaches^a

Variable	Baseline approach			Macro-with-markups			Macro-without-markups		
	1984–2000	2001–16	Difference	1984–2000	2001–16	Difference	1984–2000	2001–16	Difference
β	0.961	0.972	0.012	0.984	1.012	0.028	0.925	0.913	-0.012
μ	1.079	1.146	0.067	1.165	1.330	0.166	1	1	0
p	0.034	0.065	0.031	0	0	0	0	0	0
δ	2.778	3.243	0.465	2.778	3.243	0.465	2.778	3.243	0.465
α	0.244	0.243	-0.000	0.183	0.122	-0.061	0.299	0.340	0.041
g_p	1.171	1.101	-0.069	1.171	1.101	-0.069	1.171	1.101	-0.069
g_z	1.298	1.012	-0.286	1.544	1.358	-0.187	1.074	0.738	-0.335
g_Q	1.769	1.127	-0.643	1.769	1.127	-0.643	1.769	1.127	-0.643
\bar{N}	62.344	60.838	-1.507	62.344	60.838	-1.507	62.344	60.838	-1.507
ξ	0	0	0	0	0	0	-0.061	-0.102	-0.041

Source: Authors' calculations.

a. This table reports the estimated parameters in each of the two subsamples, 1984–2000 and 2001–16, in our baseline model; in the macro model with markups; and in the macro model without markups.

An alternative approach is to abstract from this liquidity but to allow for markup, while still omitting the *PD* ratio from the list of targets and risk from the potential parameters. This is also a well-posed exercise with eight moments and eight parameters, which we call the “macro-with-markups” approach. In this case, the spread between the *MPK* and the risk-free rate must reflect depreciation or rents. Intuitively, this approach assumes that the risk-free rate can be used to infer the cost of capital, and hence rents are deduced as a residual. The approach is conceptually quite similar to that taken by Simcha Barkai (2016), though we present it in a slightly more structural framework. The results are shown in the middle two columns of table 5. There are a number of differences between these results and our baseline results. First, the level of markups is much higher, and the increase in markups is much stronger (about 16.6 points instead of 6.7 points). Second, the increase in markups is so large that the model requires a sharp decline in α (from about 0.18 to 0.12) to keep the labor share from falling too much. This estimate suggests that technical progress has been biased toward labor over the past 30 years—a somewhat implausible conclusion. Conversely, this model also implies that β rose significantly. Below, we discuss further differences for a longer sample.

Table 6 presents the implications of these different “calibrations.” Notably, our approach offers a balanced view where increases in markups and risk premia jointly explain the rising spread, while the macroeconomic model without markups accounts for all of it with an unmodeled liquidity premium and the macro model with markups accounts for all of it with rising market power. As a result, the macro model with markups implies a sharp decline in the level of GDP, by about 8 percentage points. Moreover, the share of income going to capital falls, while the share of profits surges. Conversely, the macro model without markups predicts an increase in the level of GDP relative to trend—the liquidity premium does not discourage capital accumulation in that model as much as markups or risk premia do in the other versions of the model.

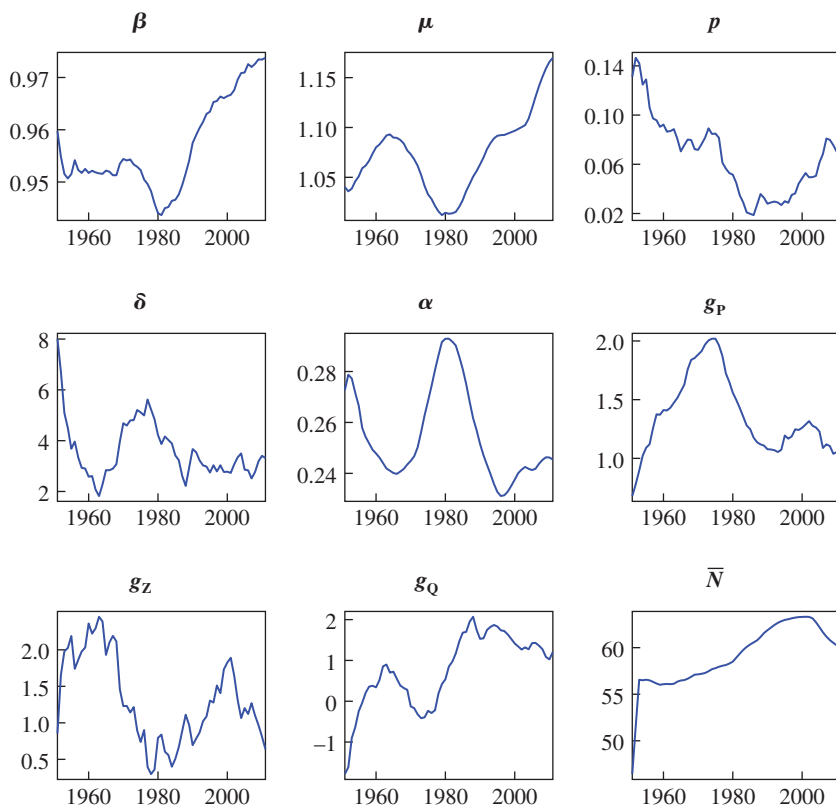
Another interesting implication is that Tobin’s *Q*, which increase significantly in our baseline, in a way that is broadly consistent with the data, is actually undefined in the macro-with-markups approach, because the low discount rates make the firm value infinite. In this sense, that model cannot match the evolution of valuation ratios, given its target of interest rates. Furthermore, the macro-without-markups approach implies decreasing valuation ratios, which are at odds with the data, owing to the very large, and rising, liquidity premium. These results provide indirect support for our baseline model.

Table 6. The Model's Implications: Baseline versus Macroeconomic Approaches^a

Parameter	Baseline approach			Macro-with-markups			Macro-without-markups		
	1984–2000	2001–16	Difference	1984–2000	2001–16	Difference	1984–2000	2001–16	Difference
<i>A. MPK–RF spread</i>									
Total spread	11.22	15.24	4.02	11.22	15.24	4.02	11.22	15.24	4.02
Depreciation	4.55	4.37	–0.18	4.55	4.37	–0.18	4.55	4.37	–0.18
Market power	3.39	5.55	2.17	6.58	10.89	4.30	0.00	0.00	0.00
Risk premium	3.15	5.23	2.08	0.00	0.00	0.00	0.00	0.00	0.00
Liquidity premium	0.00	0.00	0.00	0.00	0.00	0.00	6.51	10.75	4.24
<i>B. Rates of return</i>									
Equity return	5.85	4.90	–0.96	2.79	–0.35	–3.14	9.30	10.40	1.10
Equity premium	3.07	5.25	2.18	0.00	0.00	0.00	0.00	0.00	0.00
Risk-free rate	2.79	–0.35	–3.14	2.79	–0.35	–3.14	2.79	–0.35	–3.14
<i>C. Valuation ratios</i>									
Price-dividend	42.34	50.11	7.78	NA	NA	NA	17.82	13.57	–4.25
Price-earnings	17.85	25.79	7.94	NA	NA	NA	7.52	6.98	–0.53
Tobin's Q	2.50	3.84	1.34	NA	NA	NA	1.05	1.04	–0.01
<i>D. Income shares</i>									
Share of labor	70.11	66.01	–4.10	70.11	66.01	–4.10	70.11	66.01	–4.10
Share of capital	22.59	21.24	–1.35	15.75	9.17	–6.58	29.89	33.99	4.10
Share of profit	7.30	12.76	5.46	14.14	24.82	10.69	0.00	0.00	0.00
<i>E. Macroeconomy</i>									
K/Y	2.13	2.28	0.15	2.13	2.28	0.15	2.13	2.28	0.15
I/Y	17.28	16.50	–0.78	17.28	16.50	–0.78	17.28	16.50	–0.78
Detrend Y (% change)	—	—	–0.30	—	—	–8.00	—	—	7.77
Detrend I (% change)	—	—	–4.95	—	—	–12.65	—	—	3.12

Source: Authors' calculations.

a. This table reports some moments of interest calculated in the baseline model, in the macro model with markups, and in the macro model without markups, using the estimated parameter values for each of the two subsamples 1984–2000 and 2001–16, as well as the change between samples.

Figure 6. Estimated Parameters over Rolling windows, 1955–2011^a

Source: Authors' calculations.

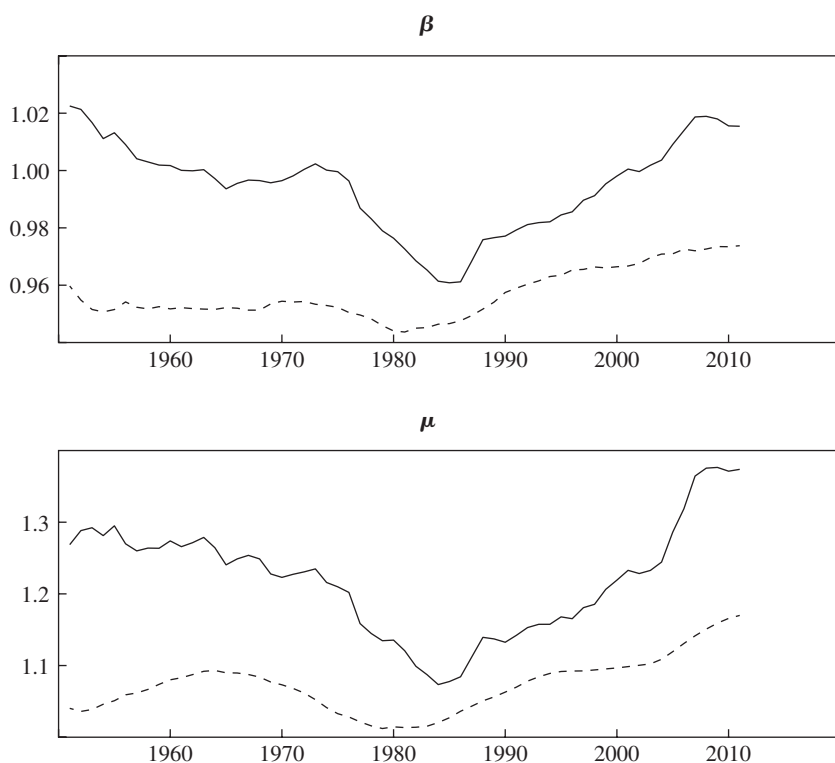
a. This figure plots the estimated parameters for each year. The target moments are the local moving averages over the 11 surrounding years.

V.C. Rolling Estimation

An alternative approach to fitting the model is to estimate it using rolling windows rather than two subsamples. In this spirit, figure 6 presents the estimated parameters when we estimate the model each year using an 11-year centered moving average to calculate the targeted moments. (That is, we target the smooth lines shown in section II, in figure 4.) We start our analysis in 1950 to avoid World War II.²³ As noted above, this calculation assumes that agents are myopic, in the sense that they believe that

23. We thank Matthew Rognlie for proposing (and executing) this exercise in his discussion at the National Bureau of Economic Research's Summer Institute.

Figure 7. Estimated Parameters β , μ : Baseline versus “Macro-with-Markups” Approach, 1951–2011^a




Source: Authors' calculations.

a. The figure plots the estimated β and μ over rolling windows for the baseline model (dashed line) and for the macro approach with markups (solid line).

the currently observed targeted moments will be constant forever, and it abstracts from transitional dynamics.

We find a U shape in the parameter β (savings supply) and in macro-economic risk p . Hence, our results suggest that risk premia declined in the 1970s and in the early to middle 1980s, before rising. Markups also have a U shape but also an initial increase in the 1950s and 1960s. The capital parameter α has an increase in the late 1970s, which is later reversed. Figure 7 compares the evolution of our parameters β , μ with the parameters estimated using the macro-with-markups approach. Our estimated parameters are significantly more stable over time—the U shape is much weaker. We find this interesting because accounting for stock market valuation

Figure 8. Decomposition of Income, 1951–2011^a

Source: Authors' calculations.


a. This figure presents the model-implied distribution of income, using the parameters estimated in each year using the rolling window estimation. The labor share is 1 minus the sum of capital and rents.

ratios might be expected to lead to more unstable parameters—but we find the opposite.

We can then use these rolling estimates to study the income distribution, the return spread $MPK - RF$, and their drivers. Figure 8 presents the share of pure profits, the true capital share, and the sum of the two for each year. By construction, the total equals 1 minus the labor share, and matches the data exactly.

The figure shows that the share of pure profits is estimated to have risen in the 1960s, then fallen in the 1970s and risen since 1980. Inversely, the capital share fell, then rose and fell. This picture reflects the puzzling pattern of a U shape in profits and an inverse U shape in α emphasized by Karabarbounis and Neiman (2019). However, we find it interesting that the U shape is significantly less strong with our estimation strategy than if one follows the macro-with-markups strategy. Karabarbounis and Neiman (2019) note that the strong negative correlation between the interest rate and the capital share, and the strong positive correlation between the interest rate and the profit share, are suggestive of measurement problems in the cost of capital. Figure 9 shows the capital share and

Figure 9. The Capital Share and the Pure Profit Share: Baseline versus “Macro-with-Markups” Approach, 1951–2011^a




Source: Authors' calculations.

a. This figure presents the distribution of income, using the parameters estimated at each point in time, for both the “macro-with-markups” and macrofinance (baseline) estimations. The top panel shows the true capital share, and the bottom panel shows the profit share; the dashed lines correspond to the macro estimation, and the solid lines to the macrofinance (baseline) estimation.

the pure profit share implied by the two estimations. There is clearly less volatility for the macroeconomic and finance estimates.

Figure 10 presents the $MPK-RF$ spread and its three subcomponents: economic and physical depreciation, rents, and risk. The spread falls in the 1970s before rising in the 1980s. The depreciation component moves, if anything, in the opposite direction from the spread, and hence does not help explain its movements. Rents are estimated to fall and then rise, and so is risk. The empirical success here is that the risk premium—which is

Figure 10. Decomposition of the Spread $MPK-RF$, 1951–2011^a

Source: Authors' calculations.

a. This figure presents the model-implied spread between the average product of capital and the risk-free rate, and the three components that explain this wedge—depreciation, rents, and risk—using the parameters estimated for each year using the rolling windows moments.

estimated without looking at the MPK , but rather by single-mindedly observing the PD ratio and growth rates—helps explain some of this variation.


Figure 11 again compares these results with those obtained with the more standard macroeconomic estimation. Both estimation approaches infer the same depreciation component. The macro approach attributes none of the spread to risk by construction, and hence infers a large and highly volatile rent (or profit) component. Finally, figure 12 depicts the implied risk-free rate, expected equity return, and equity risk premium. The risk-free rate exactly matches our data target, by construction. The equity premium mimics the evolution of p depicted in figure 6.

VI. Extensions and Robustness

This section presents some extensions of our baseline framework. We first discuss the interpretation of rising risk premia and alternative approaches to modeling them. We next analyze how financial leverage, the IES, alternative interest rates that adjust for liquidity or term premia, and capital

Figure 11. Rents and Risk Premium Components of the Spread Between the Marginal Product of Capital and the Risk-Free Rate, 1951–2021^a

Percent



Source: Authors' calculations.


a. This figure presents the three components of the model-implied spread between the marginal product of capital and the risk-free rate, for both the baseline (macrofinance) calibration and the macroeconomic calibration. The left panel shows the rent (profit) component and the right panel shows the risk component.

mismeasurement affect our results. Finally, we present an example to evaluate the importance of transitional dynamics.

VI.A. Interpretation of Rising Risk Premia

Our baseline results are obtained using a parameterization of χ as a rare “disaster” corresponding to a permanent decline of 15 percent in the level of GDP. Our estimates suggest that the risk of such a large shock was low in the 1990s but rose gradually in the 2000s and 2010s. Part of this increase may be attributed to a recognition after 2008 that financial crises are recurrent events that affect even developed economies.²⁴ But part of this increase occurs before the financial crisis. One interpretation is that this increase

24. Kozlowski, Veldkamp, and Venkateswaran (2018) offer a quantitative theory along these lines.

Figure 12. Risk-Free Rate, Expected Equity Return, and Equity Risk Premium, 1951–2011^a

Source: Authors' calculations.

a. By construction, the risk-free rate matches the data.

corresponds to a higher perception of risk starting in the late 1990s and early 2000s, owing to the combination of the Asian financial crisis, the Long-Term Capital Management crisis, and the 2001 crash in the United States. We must acknowledge, however, that it is not straightforward to relate our estimate of the probability of a “disaster” to data on beliefs or other asset prices.²⁵ This leads us to study alternative risk modeling in this section. For instance, the aging of developed economies, or the desire of emerging markets to accumulate safe reserves, might be interpreted in a reduced form as higher effective risk aversion. Alternatively, one may interpret the time-varying risk premium as reflecting time-varying pessimism—that is, a “behavioral” interpretation.

As explained in subsection IV.B, the precise specification of the risk model is theoretically irrelevant for some conclusions, such as the value of markups μ or the Cobb-Douglas parameter α , or the estimated equity premium, ERP . We now illustrate that even for the objects where this specification is potentially relevant, it may not be quantitatively first-order.

25. The issues also arise when studying the 1960s and 1970s, where our model says the risk of disaster was larger. The 1970s were a volatile decade, so it is perhaps not surprising that the perceived tail risk was high.

Table 7. Parameter Estimates for Different Risk Assumptions^a

<i>Assumption</i>		β	<i>Risk</i>	b	θ	σ
Baseline	1984–2000	0.961	0.034	0.163	12	0.5
	2001–16	0.972	0.065	0.163	12	0.5
Baseline with drift	1984–2000	0.960	0.038	0.163	12	0.5
	2001–16	0.971	0.071	0.163	12	0.5
Baseline with no offset	1984–2000	0.962	0.034	0.163	12	0.5
	2001–16	0.974	0.066	0.163	12	0.5
Lognormal	1984–2000	0.962	0.050	0.163	12	0.5
	2001–16	0.974	0.065	0.163	12	0.5
Time-varying disaster size	1984–2000	0.960	0.020	0.192	12	0.5
	2001–16	0.970	0.020	0.229	12	0.5
Time-varying risk aversion	1984–2000	0.960	0.020	0.163	15.316	0.5
	2001–16	0.970	0.020	0.163	19.560	0.5
$IES = 1$	1984–2000	0.966	0.034	0.163	12	1
	2001–16	0.970	0.065	0.163	12	1
$IES = 0.5$	1984–2000	0.976	0.034	0.163	12	2
	2001–16	0.965	0.065	0.163	12	2

Source: Authors' calculations.

a. IES = intertemporal elasticity of substitution. This table reports the estimated parameters in each of the two subsamples 1984–2000 and 2001–16 in the baseline model and in some variants: disaster risk with certain small offsets rather than rare windfalls; disaster risk without offset; lognormal risk; time-varying risk aversion; time-varying disaster size; IES = 1; and IES = 0.5.

Table 7 presents estimates of parameters in the first and second samples under different assumptions. The table's first row presents the baseline model. The second and third rows present alternative disaster models where, rather than a “bonanza” to offset the disaster risk, we introduce a small positive drift (the second row) or simply do not offset the disaster (the third row). The results are nearly identical. The fourth row considers a log-normal process for χ rather than a rare disaster. That model requires a large, and rising, standard deviation σ_χ of the lognormal shock to account for the data; but as we will see, it behaves quite similarly overall. The fifth and sixth rows display estimates when the disaster size b (respectively, risk aversion θ), rather than the disaster probability, is allowed to vary. Unsurprisingly, these models require rising disaster size or risk aversion to account for the data.²⁶ But all these models generate the same perfect fit of the data moments. Finally, the seventh and eighth rows present estimates of the baseline model when the IES is set to unity or 0.5 rather than 2; we discuss these below.

Table 8 presents the “causal” decomposition along the lines of tables 3 and 4; that is, they show the effect of the changes in β , the risk parameter

26. The estimated rising risk aversion could reflect wealth reallocation between agents of different risk aversion as studied, for instance, by Barro and others (2016) and Hall (2017).

Table 8. Contribution of the Parameters to Changes in Financial Variables for Different Risk Assumptions^a

Assumption	Risk-free rate			Price-dividend			Tobin's Q		
	β	Risk	Others	β	Risk	Others	β	Risk	Others
Baseline	-1.22	-1.62	-0.29	30.67	-13.19	-9.70	1.05	-0.48	0.77
Baseline with drift	-1.17	-1.67	-0.29	28.86	-11.62	-9.47	0.99	-0.42	0.77
Baseline with no offset	-1.32	-1.53	-0.29	34.55	-16.52	-10.25	1.17	-0.60	0.77
Lognormal	-1.26	-1.59	-0.29	32.07	-14.40	-9.89	1.09	-0.52	0.77
Time-varying disaster size	-1.05	-1.80	-0.29	24.82	-8.03	-9.01	0.85	-0.29	0.78
Time-varying risk aversion	-1.04	-1.81	-0.29	24.45	-7.70	-8.97	0.84	-0.28	0.78
$IES = 1$	-0.43	-2.12	-0.59	9.30	0.00	-1.52	0.33	0.00	1.01
$IES = 0.5$	1.14	-3.11	-1.17	-35.66	27.75	15.68	-1.32	0.95	1.71

Source: Authors' calculations.

a. IES = intertemporal elasticity of substitution. This table reports for each variant of the baseline model, the decomposition of the risk-free rate, the price-dividend ratio, and Tobin's Q , into the changes driven by (1) the discount factor, (2) the risk parameter, and (3) all the other parameters.

used in the variant (p , θ , b , or σ_χ), or the other parameters (all grouped together for simplicity) on some model moments. We know already that the implications for α , μ , and so on are unchanged; so we focus here on three key financial variables: the risk-free rate, the price-dividend ratio, and Tobin's Q . The table shows that across a range of specifications, the decline of the risk-free rate is driven in significant parts by β and by the risk parameter—the probability of disaster, or the risk aversion or disaster size, regardless of the exact specification. Similarly, the increase in the price-dividend ratio and in Tobin's Q is the result of offsetting effects of the decline of β , the increase of the risk factor, and the decline of growth factors (“others”). Hence, our results are insensitive to the exact way risk is modeled.

VI.B. Leverage

Our model calculations assume an all-equity-financed firm. In reality, corporations are leveraged, which in particular may affect the price-dividend ratio, which we use as an input in our estimation strategy. In this subsection, we propose a simple approach to bound the effect of leverage. To take this into account, we assume a Modigliani-Miller world where corporate leverage has no effect on real quantities, and only affects prices and dividends. We assume that corporate debt is fully risk-free. We then adjust the price-dividend ratio of the model given an exogenous leverage decision, which we take directly from the data.²⁷ We then reestimate the model and obtain the results shown in the third set of columns in tables 9 and 10.²⁸

Qualitatively, the findings are quite similar to those of the model without leverage: β , μ , and p all go up, and are important contributors to the observed changes in the risk-free rate, profitability, and the price-dividend ratio. However, the role of risk is somewhat smaller than in our baseline version. The logic is clear from the Gordon formula: With leverage, the change in r^* required to account for the change in valuation ratio is smaller. (Going in the other direction, however, is that in

27. Specifically, we use S&P 500 data and define *leverage* as short-term debt plus long-term debt less cash, divided by market value of equity; see the online appendix.

28. As an alternative approach, one can adjust the r^* from the model directly to account for leverage, noting that the r^* identified by the model from the PD ratio is actually $(1 + \omega)r^* - \omega r'$ where ω is the observed debt-equity ratio. This approach yields nearly identical results to the one where we adjust the PD ratio directly.

Table 9. Parameter Estimates for the Two Subsamples: Robustness^a

Parameter	IES = 0.5				Leverage				AA rate as risk-free rate				10-year Treasury adjusted for term premium as risk-free rate			
	1984–2000	2001–16	Difference		1984–2000	2001–16	Difference		1984–2000	2001–16	Difference		1984–2000	2001–16	Difference	
β	0.976	0.965	-0.011		0.968	0.982	0.014		0.957	0.969	0.012		0.963	0.972	0.008	
μ	1.079	1.146	0.067		1.106	1.191	0.084		1.079	1.146	0.067		1.079	1.146	0.067	
p	0.034	0.065	0.031		0.021	0.044	0.023		0.012	0.043	0.031		0.052	0.061	0.009	
δ	2.778	3.243	0.465		2.778	3.243	0.465		2.778	3.243	0.465		2.778	3.243	0.465	
α	0.244	0.243	-0.000		0.224	0.214	-0.010		0.244	0.243	-0.000		0.244	0.243	-0.000	
g_P	1.171	1.101	-0.069		1.171	1.101	-0.069		1.171	1.101	-0.069		1.171	1.101	-0.069	
g_Z	1.298	1.012	-0.286		1.378	1.096	-0.282		1.298	1.012	-0.286		1.298	1.012	-0.286	
g_θ	1.769	1.127	-0.643		1.769	1.127	-0.643		1.769	1.127	-0.643		1.769	1.127	-0.643	
\bar{N}	62.344	60.838	-1.507		62.344	60.838	-1.507		62.344	60.838	-1.507		62.344	60.838	-1.507	

Source: Authors' calculations.

a. IES = intertemporal elasticity of substitution. This table reports the estimated variable values of each of the two subsamples, 1984–2000 and 2001–16, in the baseline model; in the model with IES = 0.5; in the model with financial leverage; and in the model estimated with a different interest rate target (AA).

Table 10. Model Implications: Robustness^a

<i>Moment</i>	<i>IES = 0.5</i>			<i>Leverage</i>		
	<i>1984–2000</i>	<i>2001–16</i>	<i>Difference</i>	<i>1984–2000</i>	<i>2001–16</i>	<i>Difference</i>
<i>A. MPK–RF spread</i>						
Total spread	11.22	15.24	4.02	11.22	15.24	4.02
Depreciation	4.55	4.37	–0.18	4.55	4.37	–0.18
Market power	3.39	5.55	2.17	4.47	6.99	2.52
Risk premium	3.15	5.23	2.08	2.08	3.81	1.73
<i>B. Rates of return</i>						
Equity return	5.85	4.90	–0.96	5.77	4.84	–0.93
Equity premium	3.07	5.25	2.18	2.99	5.19	2.20
Risk-free rate	2.79	–0.35	–3.14	2.79	–0.35	–3.14
<i>C. Valuation ratios</i>						
Price-dividend	42.34	50.11	7.78	NA	NA	NA
Price-earnings	17.85	25.79	7.94	NA	NA	NA
Tobin's <i>Q</i>	2.50	3.84	1.34	NA	NA	NA
<i>D. Income shares</i>						
Share of labor	70.11	66.01	–4.10	70.11	66.01	–4.10
Share of capital	22.59	21.24	–1.35	20.26	17.96	–2.30
Share of profit	7.30	12.76	5.46	9.62	16.03	6.40
<i>E. Macroeconomy</i>						
<i>K/Y</i>	2.13	2.28	0.15	2.13	2.28	0.15
<i>I/Y</i>	17.28	16.50	–0.78	17.28	16.50	–0.78
Detrend <i>Y</i> (% change)	—	—	–0.30	—	—	–1.88
Detrend <i>I</i> (% change)	—	—	–4.95	—	—	–6.52

Source: Authors' calculations.

a. IES = intertemporal elasticity of substitution; *MPK* = marginal product of capital; *RF* = risk-free rate; *TP* = term premium. This table reports some moments of interest calculated in the baseline model, in the model with IES = 0.5, in the model with financial leverage, and in the model estimated with a different interest rate target (AA), using the estimated parameter values for each of the two subsamples, 1984–2000 and 2001–2016, as well as the change between samples.

our data, aggregate leverage declines from the first sample to the second one.) In particular for the spread decomposition *MPK–RF* in table 10, the share of the spread due to risk is smaller (about 2.1 and 3.8 percentage points in the first and second samples, respectively). However, the share of the increase in the spread due to risk remains substantial. Moreover, in terms of the implied equity premium, the increase is actually similar, because leverage now amplifies the variation in r^* . These results are conservative, because we have assumed that corporate debt pays the same return as the risk-free asset; in reality, corporate debt yields

AA rate as risk-free rate			10-year Treasury adjusted for term premium as risk-free rate		
1984–2000	2001–16	Difference	1984–2000	2001–16	Difference
9.32	13.80	4.48	12.49	14.98	2.49
4.55	4.37	–0.18	4.55	4.37	–0.18
3.39	5.55	2.17	3.39	5.55	2.17
1.25	3.79	2.54	4.42	4.97	0.55
5.88	4.84	–1.05	5.87	4.88	–0.99
1.19	3.75	2.56	4.35	4.97	0.62
4.69	1.09	–3.60	1.52	–0.09	–1.61
42.34	50.11	7.78	42.34	50.11	7.78
17.85	25.79	7.94	17.85	25.79	7.94
2.50	3.84	1.34	2.50	3.84	1.34
70.11	66.01	–4.10	70.11	66.01	–4.10
22.59	21.24	–1.35	22.59	21.24	–1.35
7.30	12.76	5.46	7.30	12.76	5.46
2.13	2.28	0.15	2.13	2.28	0.15
17.28	16.50	–0.78	17.28	16.50	–0.78
—	—	–0.30	—	—	–0.30
—	—	–4.95	—	—	–4.95

are higher than Treasury securities yields, which would reduce the adjustment to the *PD* ratio.

VI.C. The Intertemporal Elasticity of Substitution

We have assumed an IES equal to 2 in our baseline estimation. The IES cannot be identified, given that the model generates *iid* growth rates for all macroeconomic variables. As noted above, the assumed value for the IES does not affect estimates of α , μ , r^* , or the equity premium. This can be verified in tables 9 and 10, where we present parameter estimates

for an elasticity equal to 0.5. Our conclusions that risk and market power increased are hence completely unaffected by this assumption. However, changing the IES does affect the counterfactual decompositions studied above; for instance, the effect of an increase in risk on capital accumulation depends on the assumed IES.

Table 8 presents decompositions for three financial variables, and section 3 of the online appendix provides the decompositions of all variables. With a low IES, the effect of the decline of growth in accounting for the decline of the risk-free rate is larger. The model hence does not require an increase in β —rather, β falls. The change of the risk-free rate due to uncertainty is now larger. In this sense, a lower IES gives a larger role for risk. The low IES implies very different decompositions of the changes in the *PD* ratio. As emphasized by Ravi Bansal and Amir Yaron (2004), with a low IES, higher risk and lower growth both raise the *PD* ratio because of their strong effect on the risk-free rate.

VI.D. Liquidity and Term Premia

As a risk-free rate proxy in the data, we use the 1-year Treasury rate (minus lagged core inflation). One concern is that our model abstracts from the liquidity premium, which makes this rate especially low. To gauge the role of the liquidity premium, we instead use as a risk-free rate proxy the rate on AA corporate bonds, minus the SPF median Consumer Price Index inflation over the next 10 years. This is a rate for securities that do not possess the same unique liquidity attributes as a U.S. Treasury security. We then repeat our estimation. Tables 9 and 10 show the results. Given the identification provided by the model, changing the risk-free rate does not affect α , μ , or r^* . However, the different risk-free rate target will affect the value of β and the amount of risk identified by the model, and their respective changes. Indeed, we see that both the estimated β and the estimated p are lower than in our baseline model; but crucially, our model still estimates that β and p increased significantly between the two samples. Our conclusion about the relative importance of risk and markups is also not affected by this change in target, suggesting that liquidity considerations do not play a very large role in these trends.

A related concern is that long-term rates reflect term premia that may be driven by an inflation or real rate premium which is not present in the model. We hence consider as a target for the risk-free rate the 10-year Treasury constant maturity rate, less SPF-expected inflation, less the term

premium estimate made by Tobias Adrian, Richard Crump, and Emanuel Moench (2013), which they obtained from a statistical term structure model. Because the term premia estimate declines strongly during this period, the decline in this measure of the risk-free rate is only about 1.5 points rather than over 3 points. The resulting estimates imply a smaller increase in macroeconomic risk. Moreover, the spread $MPK-RF$ is also increasing by a smaller amount, and the contribution of risk premia is smaller there as well. We view these results as somewhat less plausible because the decline of the term premium implied by this model is very large—we are unaware of macroeconomic models that can rationalize this. Also, to the extent that the decline of the term premium is related to macroeconomic risk, it may not be sound to adjust for it.

VI.E. Capital Mismeasurement

One natural explanation for the rising spread $MPK-RF$ is that K is mismeasured, and in particular is underestimated by the U.S. Bureau of Economic Analysis (BEA) analysts, who traditionally focus on tangible assets. To get a sense of how much mismeasurement of capital matters, we present a simple approach in this subsection. In section 4 of the online appendix, we then estimate a more detailed model of intangible accumulation. We are interested in two questions: First, can a plausible amount of mismeasurement explain the rising spread? Second, is this mismeasurement also consistent with the other observed features of the data?

In this section, we simply assume that the BEA measures only a fraction, λ , of total investment. When $\lambda = 1$, there is no mismeasurement, corresponding to our baseline model. When $\lambda < 1$, however, this mismeasurement of investment affects our targeted moments, and hence possibly our parameter estimates. We denote with a superscript m the measured values of the model variables.²⁹ Measured investment is $x^m = \lambda x$, and hence along the balanced growth path $k^m = \lambda k$. Moreover, GDP and the profit share are now underestimated, because the unmeasured investment $(1 - \lambda)x$ is treated as an intermediate input by BEA accountants. As a

29. We do the algebra for detrended variables, but one can obviously also apply the same adjustments to the level variables.

result, measured GDP is $y^m = y - (1 - \lambda)x$. Measured profits equal measured GDP, less labor compensation, or $\pi^m = \pi - (1 - \lambda)x$. The profit share is hence underestimated as

$$\frac{\pi^m}{y^m} = \frac{\pi - (1 - \lambda)x}{y - (1 - \lambda)x} < \frac{\pi}{y}.$$

However, dividends are correctly measured because the unmeasured investment reduces both profits and investment: $d = \pi - x = \pi^m - x^m$. Hence, the asset price is unaffected by measurement error (even if investors do not observe intangible investment).

It is easy to extend our formula 27 for the spread:

$$(29) \quad MPK - r_f = \delta + g_o + \frac{\mu - 1}{\alpha} (r^* + \delta + g_o) + r^* - r_f + \frac{1 - \lambda}{\lambda} \frac{d}{k}$$

and we see that mismeasurement ($\lambda < 1$) now adds an additional component to the measured spread, which is consistent with basic intuition.

How important is this mismeasurement wedge? First, note that the measured ratio $\frac{d}{\lambda k} = \frac{d}{k^m}$ can be calculated as the difference between profitability and the investment rate, and hence equals about 6 percent in the first sample and 7.5 percent in the second sample. Hence, with $\lambda = 0.8$, or a 20 percent undermeasurement, the wedge is about 1.2–1.5 points, which is significant. Our focus, however, is on the *increase* in the spread. To explain this increase requires a rising mismeasurement. Though there is wide agreement that intangibles play a critical role in modern economies, it is not as clear if mismeasurement has increased over the past few decades. Suppose however, that one wanted to generate an increase in the spread by 2 percentage points (or about half the increase in the spread observed during our sample, and about the same as what is explained by risk premia or markups according to our baseline results), the model requires λ to go, for instance, from 1 (perfect measurement) to $\lambda = 0.73$, a 27 percent underestimation of investment. This rising mismeasurement would reduce measured GDP by about 4.4 percent, and the profit share by about 4 percentage points.³⁰ One

30. This calculation is based on the formulas of the previous page, $y^m = y - (1 - \lambda)x$ and $\pi^m = \pi - (1 - \lambda)x$, assuming a measured investment-output ratio equal to 0.17, as in our data.

tension, hence, is that rising intangibles lead to a measured labor share going up rather than down, as in the data.

To evaluate more precisely how this mismeasurement affects our results, we estimate three versions of our baseline model corresponding to different assumptions about mismeasurement. In the first version, mismeasurement is constant at 10 percent in both samples ($\lambda = 0.9$). In the second version, mismeasurement starts at 10 percent in the first subsample and then rises to 20 percent in the second subsample. In the third version, mismeasurement starts at 10 percent and then rises to 30 percent. These numbers are largely illustrative; note, however, that the share in capital of measured “intangibles”—that is, intellectual property products—is about 6 percent recently.³¹ We are hence assuming that the unmeasured stock of intangible capital is significantly larger than the current measured stock, and has been rising significantly over the past 15 years.

Table 11 reports the parameter estimates, and table 12 reports the implied moments corresponding to different scenarios. There are a few interesting results. First, all parameters are completely unaffected, except for μ and α . In particular, the increase in β and in risk are not affected by these assumptions. Second, when mismeasurement is constant at 10 percent, the model has similar implications to our baseline model (the level of α is higher and the level of μ lower, but the changes between two subsamples are nearly identical). Third, the estimated increase in markup is smaller when there is an increase in mismeasurement. For instance, with a mismeasurement rising to 30 percent of capital, the markup rises by only about 4.1 points instead of 6.6 points when mismeasurement is constant and 6.7 points in the baseline model. This is intuitively consistent with the simple formula 29: With more mismeasurement, there is less of a gap between the *MPK* and the risk-free rate to explain. The other implication is that the estimated α rises. This is because the labor share rises with mismeasurement; to offset this, the model needs an increase in capital-biased technical change—that is, α .

Overall, in our most generous calibration, the rising mismeasurement explains about a 1.65-point increase in the wedge, the markup now only 0.47 point, and the risk premium 2.08 points. Of course, the magnitude of the mismeasurement is difficult to ascertain. But it is interesting that

31. This number is obtained by dividing line 7 by line 3 in table 1.1 of the Fixed Asset Accounts of the United States.

Table 11. Parameter Estimates for the Two Subsamples: Capital Mismeasurement^a

Variable	Baseline			Constant bias: 10 percent			Rising bias: 10–20 percent			Rising bias: 10–30 percent		
	1984–2000	2001–16	Difference	1984–2000	2001–16	Difference	1984–2000	2001–16	Difference	1984–2000	2001–16	Difference
β	0.961	0.972	0.012	0.961	0.972	0.012	0.961	0.972	0.012	0.961	0.972	0.012
μ	1.079	1.146	0.067	1.070	1.136	0.066	1.070	1.125	0.055	1.070	1.111	0.041
p	0.034	0.065	0.031	0.034	0.065	0.031	0.034	0.065	0.031	0.034	0.065	0.031
δ	2.778	3.243	0.465	2.778	3.243	0.465	2.778	3.243	0.465	2.778	3.243	0.465
α	0.244	0.243	–0.000	0.264	0.263	–0.000	0.264	0.287	0.023	0.264	0.315	0.051
g_p	1.171	1.101	–0.069	1.171	1.101	–0.069	1.171	1.101	–0.069	1.171	1.101	–0.069
g_z	1.298	1.012	–0.286	1.217	0.956	–0.262	1.217	0.889	–0.328	1.217	0.809	–0.408
g_o	1.769	1.127	–0.643	1.769	1.127	–0.643	1.769	1.127	–0.643	1.769	1.127	–0.643
N	62.344	60.838	–1.507	62.344	60.838	–1.507	62.344	60.838	–1.507	62.344	60.838	–1.507

Source: Authors' calculations.

a. This table reports the estimated parameters in each of the two subsamples, 1984–2000 and 2001–16, in the baseline model and in the model with mismeasured capital, for different values of the mismeasurement parameters, using the estimated parameter values for each of the two subsamples, as well as the change between samples.

incorporating realistic mismeasurement would reduce further the implied markup, while leaving the role of risk unaffected.

VI.F. Transitional Dynamics

Our calculations so far assume that the economy remains along its “risky balanced growth path.” However, if the model parameters such as the discount factor or markup change, the economy will experience a transition before it reaches its new balanced growth path. This transition may affect our estimation results.

To evaluate the importance of this bias, we estimated the model, taking into account the transitional dynamics. Specifically, we make the following assumptions. We use the baseline version of the model and assume that the economy starts in 1992 in balanced growth with the parameters that we estimate over the first sample.³² We then assume that the nine parameters change linearly over 24 years (to end in 2016), from the value we estimated in the first sample to a final value that we will estimate (and that may not be our estimate for the second sample).

We then calculate the transitional dynamics for this economy using a standard shooting method. A key issue is agents’ expectations. With perfect foresight, the model cannot fit the data, because agents see the lower interest rates coming, which leads to a boom in the price-dividend ratio. (Furthermore, the long-term interest rate would fall significantly more than the short-term rate, unlike what we see in the data.) We hence assume myopic expectations: In each period, agents observe the new values of the parameters, and they assume (incorrectly, at least for the first 24 years) that these parameters will remain constant forever.³³

We then numerically find the final parameters such that, when calculating the transition, this procedure yields an average time series for our targets (over the period 2001–16) that matches what we measured in

32. We use 1992 to take into account that these parameters are estimated over the period 1984–2000.

33. Agents consequently make investment choices that would, eventually, lead to converge to a new steady state corresponding to today’s parameter values. However, the next period, new parameter values (unexpectedly) arrive, leading to new choices and a revised transition path. This process continues until the parameters are indeed constant, and the economy then converges to its final steady state.

Table 12. Implications: Baseline versus Capital Mismeasurement^a

<i>Moment</i>	<i>Baseline</i>			<i>Constant bias: 10 percent</i>		
	<i>1984–2000</i>	<i>2001–16</i>	<i>Difference</i>	<i>1984–2000</i>	<i>2001–16</i>	<i>Difference</i>
<i>A. MPK–RF spread</i>						
Total spread	11.22	15.24	4.02	11.22	15.24	4.02
Depreciation	4.55	4.37	–0.18	4.55	4.37	–0.18
Market power	3.39	5.55	2.17	2.80	4.79	1.99
Risk premium	3.15	5.23	2.08	3.15	5.23	2.08
Mismeasurement	0.13	0.09	–0.05	0.72	0.85	0.13
<i>B. Rates of return</i>						
Equity return	5.85	4.90	–0.96	5.85	4.90	–0.96
Equity premium	3.07	5.25	2.18	3.07	5.25	2.18
Risk-free rate	2.79	–0.35	–3.14	2.79	–0.35	–3.14
<i>C. Valuation ratios</i>						
Price-dividend	42.34	50.11	7.78	42.34	50.11	7.78
Price-earnings	17.85	25.79	7.94	17.85	25.79	7.94
Tobin's <i>Q</i>	2.50	3.84	1.34	2.50	3.84	1.34
<i>D. Income shares</i>						
Share of labor	70.11	66.01	–4.10	68.79	64.82	–3.97
Share of capital	22.59	21.24	–1.35	24.63	23.17	–1.46
Share of profit	7.30	12.76	5.46	6.58	12.01	5.43
<i>E. Macroeconomy</i>						
<i>K/Y</i>	2.13	2.28	0.15	2.13	2.28	0.15
<i>I/Y</i>	17.28	16.50	–0.78	17.28	16.50	–0.78
Detrend <i>Y</i> (% change)	—	—	–0.30	—	—	0.05
Detrend <i>I</i> (% change)	—	—	–4.95	—	—	–4.60

Source: Authors' calculations.

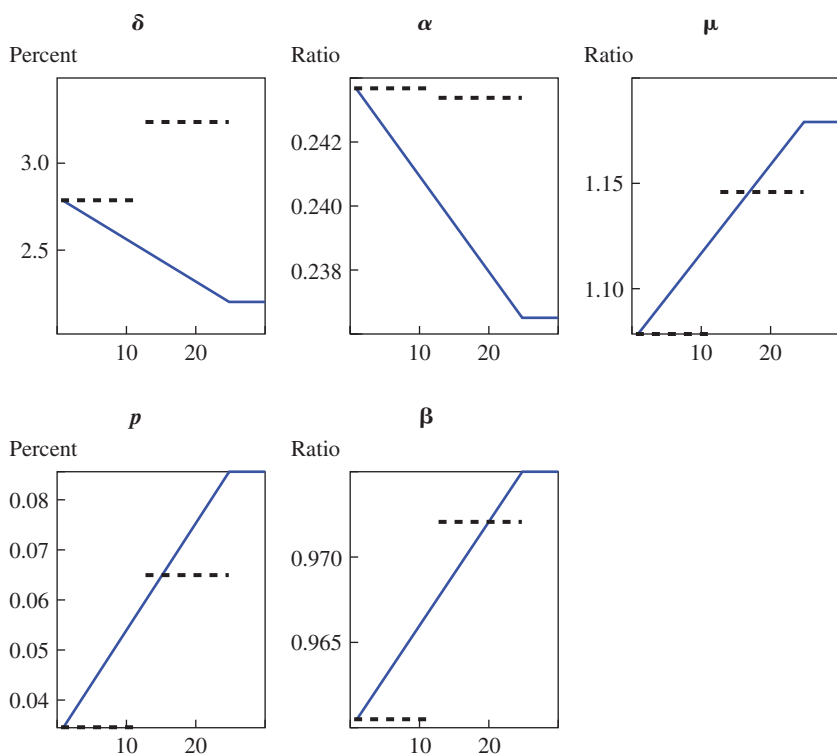
a. *MPK* = marginal product of capital; *RF* = risk-free rate. This table reports some moments of interest calculated in the baseline model and in the model with mismeasured capital, for different values of the mismeasurement parameters, using the estimated parameter values for each of the two subsamples, 1984–2000 and 2001–16, as well as the change between samples.

the data. Figure 13 presents the path obtained for parameter values, and figure 14 shows the path for the moments targeted (we abstract here from parameters that map directly into moments). The dashed lines in these tables represent the parameters and moments from the baseline estimation for the two samples. Table 13 presents the numerical counterpart to these graphs.

As can be eyeballed in figure 14, the model moments, averaged over periods 13–25 (corresponding to the second sample), match reasonably well the targeted moments for the second sample (the darker line). The more

<i>Rising bias: 10–20 percent</i>			<i>Rising bias: 10–30 percent</i>		
<i>1984–2000</i>	<i>2001–16</i>	<i>Difference</i>	<i>1984–2000</i>	<i>2001–16</i>	<i>Difference</i>
11.22	15.24	4.02	11.22	15.24	4.02
4.55	4.37	−0.18	4.55	4.37	−0.18
2.80	4.03	1.23	2.80	3.27	0.47
3.15	5.23	2.08	3.15	5.23	2.08
0.72	1.61	0.89	0.72	2.37	1.65
5.85	4.90	−0.96	5.85	4.90	−0.96
3.07	5.25	2.18	3.07	5.25	2.18
2.79	−0.35	−3.14	2.79	−0.35	−3.14
42.34	50.11	7.78	42.34	50.11	7.78
17.85	25.79	7.94	17.85	25.79	7.94
2.50	3.84	1.34	2.50	3.84	1.34
68.79	63.39	−5.40	68.79	61.65	−7.14
24.63	25.49	0.87	24.63	28.33	3.71
6.58	11.11	4.53	6.58	10.02	3.44
2.13	2.28	0.15	2.13	2.28	0.15
17.28	16.50	−0.78	17.28	16.50	−0.78
—	—	5.74	—	—	13.60
—	—	1.10	—	—	8.95

surprising result is in figure 13, where we see that the parameter values estimated in this way are quite similar to these obtained in the simple baseline model, which assumes balanced growth. To see this, note that the full line, averaged over periods 13–25, is economically quite similar to the darker line (results from the baseline model). The one exception is δ , which now falls slightly instead of rising. Table 13 shows the same result: Comparing the third and fourth columns, the estimated parameters are quite similar, except for δ . We view these results as suggesting that, at least in the myopic case, perhaps not much is lost by focusing on the risky

Figure 13. Estimated Path for the Parameters^a

Source: Authors' calculations.

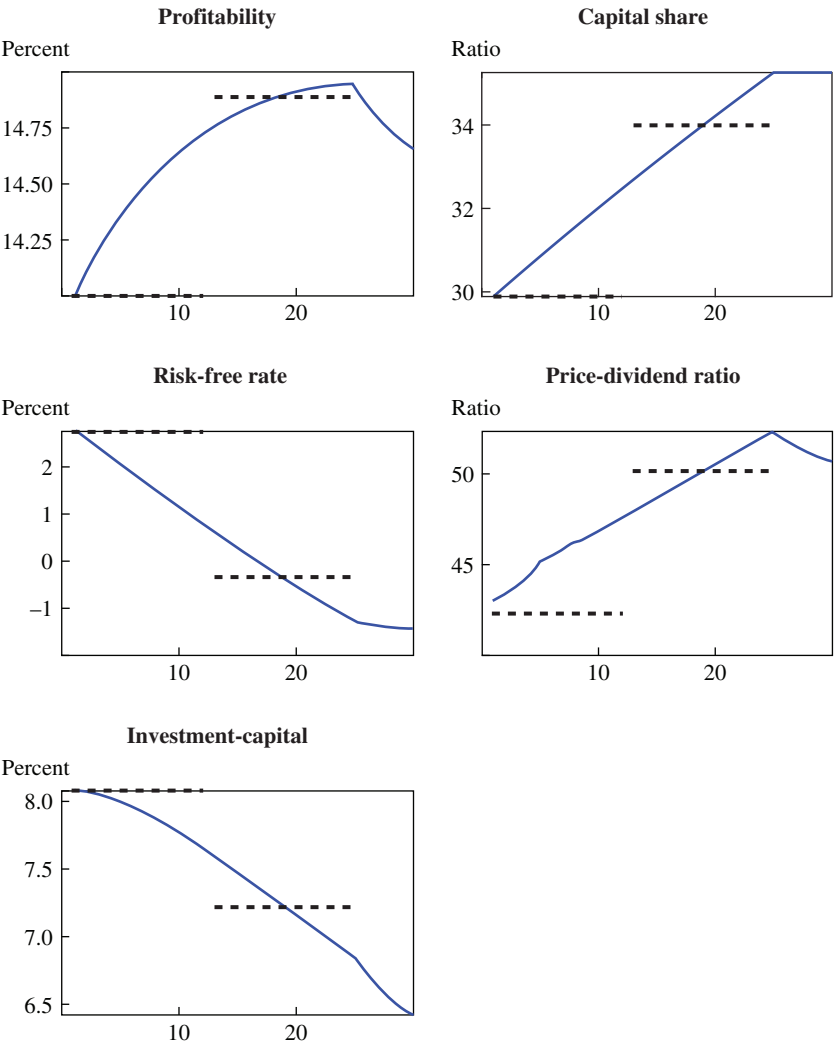
a. This figure plots the estimated path for the parameters using the transitional dynamics method. The dashed lines denote the values estimated in the baseline approach in the first and second samples.

balanced growth path. This conclusion might not hold true for all models, however—in particular, with intangibles, if there is significant accumulation during the transition.

VII. Other Evidence on Market Power, Risk Premia, and Intangibles

Our empirical results show that rising risk premia and rising market power appear to be two of the significant drivers of some of the macroeconomic and finance trends on which we focus, and intangibles have a

Figure 14. Fitted Path for the Targeted Moments^a



Source: Authors' calculations.
a. This figure plots the estimated path for the target moments using the transitional dynamics method. The dashed lines denote the values targeted in the baseline approach.

Table 13. Average Parameter Estimates and Moments: Transitional Dynamics^a

Variable	1984–2000		2001–16	
	Baseline	Transition	Baseline	Transition
Moments				
Π	14.012	14.426	14.890	14.890
$\frac{K}{Y}$				
Π	29.887	31.194	33.992	33.991
$\frac{RF}{PD}$	2.787	1.785	–0.350	–0.350
$\frac{I}{K}$	42.336	45.451	50.115	50.115
	8.103	7.932	7.227	7.227
Estimated parameters				
α	0.244	0.242	0.243	0.238
μ	1.079	1.102	1.146	1.154
β	0.961	0.964	0.972	0.971
p	0.034	0.046	0.065	0.073
δ	2.778	2.642	3.243	2.334

Source: Authors' calculations.

a. This table reports the average value of the targeted moments and the average values of the estimated parameters over the first and second samples using the transitional dynamics method. The final parameter values are chosen such that the average values of the moments match the targeted moments in the second sample. See the text for details.

potential contribution as well. In this section, we step outside the model and present independent evidence for these two phenomena. We also discuss related estimates presented by other researchers, which tend to support our conclusions.

VII.A. Empirical Estimates of the Equity Risk Premium

We first present reduced-form estimates of the equity premium. Estimating the equity premium is notoriously difficult, even retrospectively. Using realized excess equity returns is essentially pointless over short-term samples, because returns are noisy, and because an increase in the risk premium may lead, by itself, to lower realized returns.³⁴ But methods that use standard forecasting return regressions have also been found to be very

34. For instance, suppose a researcher has a sample of 16 years (as we do) and that the excess equity return has a mean of 8 percent with a volatility of 16 percent. The 95 percent confidence interval for the mean excess equity return is [0%, 16%]. It is clearly impossible to detect a change of the equity premium of even several percentage points based solely on realized returns.

unstable; Ivo Welch and Amit Goyal (2008) argue that none of them outperforms the simple mean out of sample. Here, we follow a few approaches that have been shown to be somewhat more successful empirically.

Our first approach is simply to use the static Gordon growth formula, which states that the price-dividend ratio is the inverse of the difference between the return on the asset and the dividend growth rate:

$$\frac{P}{D} = \frac{1}{R - G}$$

where R is the expected equity return, which can be decomposed into $R = RF + EP$, with RF risk-free and EP the equity premium, and G the growth rate of dividends. This approach can be used at any point in time, given the observed PD and RF and given an assumption about G . The main difference with our structural estimation above is that here we use data on dividends.

Our second approach builds on the research of Eugene Fama and Kenneth French (2002), who argue that, if the dividend yield or earnings yield is stationary, as each one ought to be, one can advantageously estimate the mean of $\frac{P_{t+1}}{P_t}$ by $\frac{D_{t+1}}{D_t}$ by $\frac{E_{t+1}}{E_t}$ (which are less volatile). As a result, they suggest estimating

$$ERP = E\left(\frac{D_{t+1}}{P_t}\right) + E\left(\frac{D_{t+1}}{D_t}\right) - E(RF),$$

which amounts to the Gordon growth formula, or replacing dividend growth with earnings growth,

$$ERP = E\left(\frac{D_{t+1}}{P_t}\right) + E\left(\frac{E_{t+1}}{E_t}\right) - E(RF).$$

This approach is best thought of as applying to a long-sample average.

Our third approach follows that of John Campbell (2008) and Campbell and Samuel Thompson (2008), who show how combining the current dividend yield and the return on book equity can be used to create a real-time estimate of the equity premium:

$$ERP = \frac{D}{E} \frac{E}{P} + \left(1 - \frac{D}{E}\right) ROE$$

and where they suggest smoothing the payout ratio $\frac{D}{E}$, earnings-price ratio $\frac{E}{P}$, and the return on book equity ROE to reduce the effect of influential but transitory observations.

These formulas can be applied either using arithmetic averages or using geometric averages. We report both in table 14, though we like Campbell and Thompson's recommendation to use the geometric averages. We then incorporate an adjustment of half the variance of stock returns to produce an estimate of the arithmetic equity premium.

The key observation from table 14 is that, though the estimates of the equity premium are clearly different across models and methods, most calculations suggest that the *ERP* increased from the first sample to the second sample. Specifically, all nine estimates are positive, ranging from about 1.8 percent to 7.2 percent. This reflects the fact that valuation ratios increased moderately, while earnings or dividend growth increased more significantly, and the risk-free rate fell. (For this exercise, we take the risk-free rate to be the 10-year Treasury yield minus SPF inflation expectations over the next 10 years.)

Figure 15 graphically presents estimates of the equity risk premium for each of the three approaches, obtained over centered 11-year rolling windows. We smooth the estimates using a 3-year moving average. Here, too, the exact numbers vary quite a bit across models, but all models suggest some increase over the past 15 years or so. (A particular difficulty is how one deals with the very low corporate earnings in 2008 or 2009, which affect the Fama-French Earnings Model significantly, leading to the extreme arithmetic implication in the middle panel.)

VII.B. Other Measures of Changes in Risk Premia

We now discuss other evidence on the changes in the risk premium. Fernando Duarte and Carlo Rosa (2015) provide an exhaustive survey of the different methods that can be used to estimate the equity premium in real time. They distinguish between different methods based on variants of the Gordon Growth Model, on predictive regressions, and on cross-sectional regressions. Overall, the conclusion is that the equity premium has risen, in line with our findings.³⁵ Campbell and Thompson (2008)


35. An earlier body of literature documented a decline of the equity premium during the 1980s and 1990s (Blanchard 1993; Jagannathan and McGrattan 2000; Heaton and Lucas 1999; Lettau and Ludvigson 2007), which is not inconsistent with our results.

Table 14. Reduced-Form Estimates of the Equity Risk Premium^a

<i>Variable</i>	<i>Arithmetic average</i>			<i>Geometric average</i>		
	<i>1984–2000</i>	<i>2001–16</i>	<i>Change</i>	<i>1984–2000</i>	<i>2001–16</i>	<i>Change</i>
Real dividend growth	2.03	4.7	2.67	2.04	8.21	6.17
Real earnings growth	6.22	16.97	10.75	10.25	12.06	1.81
Return on book equity	10.94	10.0	–0.94	10.7	9.4	–1.3
Dividend yield D/P	2.78	1.92	–0.86	—	—	—
Payout ratio, D/E	0.49	0.47	–0.02	—	—	—
Earnings-price ratio, E/P	5.74	4.69	–1.05	—	—	—
ERP, Gordon Growth Model	0.87	5.56	4.69	1.91	9.16	7.25
ERP, Fama-French Earnings Model	2.43	4.78	2.35	4.61	8.66	4.05
ERP, Campbell-Thompson	1.47	4.11	2.64	1.84	3.65	1.81
ERP, Gordon, with variance adjustment	—	—	—	2.43	8.26	5.83
ERP, Fama-French, with variance adjustment	—	—	—	4.81	10.3	5.49
ERP, Campbell-Thompson, with variance adjustment	—	—	—	2.31	5.56	3.25

Source: Authors' calculations.

a. ERP = equity risk premium. This table reports the estimated equity premium as well as the mean of some variables used to construct our estimates, for the samples 1984–2000 and 2001–16. See the text for details.

Figure 15. Reduced-Form Estimates of the Equity Risk Premium, 1989–2011^a

Source: Authors' calculations.

a. This figure depicts some reduced-form estimates of the equity risk premium. The left panel shows the Gordon growth model; the middle panel shows the Fama-French earnings model; and the right panel shows estimates from the Campbell-Thompson method. The dotted line = arithmetic average; the long-dashed line = geometric; and the short-dashed line = geometric + variance adjustment.

propose a method to estimate the equity premium in real time. Their estimate also shows a small increase after 2000. Using a very different methodology, based on a maximum-likelihood estimation of a structural model, Efsthios Avdis and Jessica Wachter (2017) reach a fairly similar conclusion. Another important contribution is Ian Martin (2017), who uses an ingenious argument to provide, under a relatively weak condition, a lower bound on the equity premium based on option data. His lower bound has a very high correlation with the Chicago Board Options Exchange's Volatility Index (VIX). The estimate is very elevated during the global financial crisis, and remains at a higher level after the crisis. However, his lower bound is quite low in the mid-2000s. If the lower bound has a constant bias with the mean, then this series does not behave like the other estimates we discussed above. However, it is possible that the bias between the lower bound he finds and the true expected equity premium is time-varying.

Table 15 presents evidence on the evolution of some other measures of risk: the Gilchrist–Zakrajšek (2012) spread, the standard BAA and AAA

Table 15. Other Measures of Risk Premia^a

<i>Measure</i>	<i>Mean</i>			<i>Difference</i>		
	<i>1984–2000</i>	<i>2001–16</i>	<i>2001–16, excluding</i>			
	<i>(1)</i>	<i>(2)</i>	<i>global financial crisis</i>	<i>(2) – (1)</i>	<i>(3) – (1)</i>	<i>SE</i>
			<i>(3)</i>			
Spread, Gilchrist–Zakrajšek	1.5	2.54	2.31	1.04	0.81	0.16
Spread, BAA–10 year	1.94	2.74	2.61	0.80	0.67	0.15
Spread, AAA–10 year	1.01	1.64	1.61	0.63	0.60	0.12
VIX	18.92	20.22	18.62	1.3	–0.3	1.98
Realized volatility	13.36	17.43	15.34	4.07	1.98	1.62

Sources: See the online appendix, section 1, for all data sources.

a. SE = standard error; VIX = Chicago Board Options Exchange Volatility Index. This table reports the mean of various credit spreads and volatility measures for the samples 1984–2000; 2001–16; and 2001–16, excluding the June 2007–June 2009 period of the global financial crisis. The table also reports the difference between these means and an SE (calculated using the Newey–West method, with 12 monthly lags).

spreads, the VIX, and stock-market-realized volatility (calculated using daily data). The table reports the mean in the two samples, as well as the mean in the second sample excluding the period of the global financial crisis. We see that all these credit spreads have increased between the two samples, and this conclusion is true even excluding this period. Realized volatility is also somewhat higher. The VIX exhibits little trend (but is only available starting in 1996). These results are consistent with Del Negro and others (2017), who show that the premia for safe and liquid assets increased over time.³⁶

VII.C. Independent Evidence on Rising Markups

A number of recent contributions, using different methods, have found that average markups have been increasing. For example, Barkai (2016) uses aggregate data and implements a user cost approach à la Robert Hall and Dale Jorgenson (1967) to decompose the nonlabor share into a true capital share and a profit share. The true capital share is computed by multiplying the capital-output ratio by the user cost of capital. The profit share is a residual. The aggregate markup can be directly inferred from the profit share. Because his measure of user cost does not incorporate a meaningful risk premium, Barkai finds that the evolutions of the user cost track those of the interest rate, so the user cost declined substantially over the period 1984–2014. This implies a large decrease in the capital share, a large increase in the profit share, and a large increase in the aggregate markup of about 20 percent, roughly in line with our macroeconomic estimation.

Jan De Loecker and Jan Eeckhout (2016) use firm-level data and estimate firm-level markups using a production function approach that recovers markups as the ratio of the elasticity of production to a flexible input share of that input in revenues, where the former is computed by estimating the production function. The aggregate markup, computed as a harmonic sales-weighted average of firm-level markups, increases by about 25 percent. James Traina (2018) criticizes the measure of costs used by De Loecker and Eeckhout. Using a broader measure, he finds that the increase in average markups is much smaller. Germán Gutiérrez and Thomas Philippon (2017) also use firm-level data, but they estimate firm-level markups using a user cost approach allowing for sizable and variable risk premia. They also find a sizable increase in aggregate markups

36. One caveat is that the underlying riskiness of the firms issuing corporate bonds may have changed over time, even within credit ratings.

of about 10 percent over the period 1984–2014, somewhat above our baseline results.

VII.D. Rising Intangible Capital

There is a growing body of literature that recognizes the importance of intangible capital in the U.S. economy. Carol Corrado, Charles Hulten, and Daniel Sichel (2005, 2009) and Leonard Nakamura (2010) present estimates of the size of intangible capital. Anmol Bhandari and Ellen McGrattan (2017) also contribute to this measurement. Dongya Koh, Raül Santaaulàlia-Llopis, and Yu Zheng (2015) argue that rising intangibles help explain the evolution of the labor share. Nicolas Crouzet and Janice Eberly (2018) argue that growing intangibles help explain both the rising market power and lower capital investment. Andrea Caggese and Ander Perez (2018) show how growing intangibles may help account for some of the same macroeconomic trends on which we focus in this paper.

VIII. Conclusion

We provide a simple accounting framework that allows decomposing the changes observed over the past 30 years in some key macroeconomic and finance trends into “semistructural” parameters using a fairly clear identification. We say “semistructural” because, allowing these parameters to vary over time flexibly suggests they are not microfounded and invariant to policy. Yet we find the results useful because deeper explanations need to be consistent with the changes of parameters implied by our approach.

We find that about half the increase in the spread between the return on private capital and the risk-free rate is due to rising market power, and half is due to rising risk premia. Technical change plays little role. Higher savings supply and higher risk premia are the prime proximate contributors to the decline in the risk-free rate. Rising market power helps explain the evolution of the capital share, profitability, and capital accumulation, but its contribution is substantially overstated if the model is estimated using a macroeconomic approach that abstracts from risk. Finally, taking into account intangibles reduces further the estimated increase in market power.

One limitation of our approach is that we treat the parameter changes as independent causal factors, but they might actually be driven by common causes; for instance, higher market power might reduce innovation and hence productivity growth, but we treat these as independent. Our analysis

also does not incorporate some factors that could help explain the evolution of some of the big ratios that we study. In particular, we abstract from taxes and from agency issues (for example, external finance or corporate governance frictions) or market incompleteness, that could also give rise to wedges that might vary over time. Our study of transitional dynamics is only scratching at the vast possibilities. Finally, it would be interesting to study these issues taking into account the specific open economy considerations or at least to study these same facts for a variety of countries.

ACKNOWLEDGMENTS The views expressed here are those of the authors and do not necessarily represent those of the Federal Reserve Bank of Chicago or the Federal Reserve System. We thank participants for their comments during presentations at the Fall 2018 Brookings Panel on Economic Activity, the European Central Bank's 2018 research conference, the Society for Economics Dynamics' 2017 annual meeting, the 2017 conference of the Society for the Advancement of Economic Theory, the Centre for Economic Policy Research's 2018 conference on asset prices and the macroeconomy, the National Bureau of Economic Research's 2018 conference on capital markets and the economy, the 2018 financial research conference of the Federal Reserve Bank of San Francisco, and presentations at the Federal Reserve Bank of Chicago and Federal Reserve Bank of Minneapolis, and at the Hoover Institution's conference in honor of John Cochrane. For their comments, we especially thank Stefania d'Amico; Marco Bassetto; Robert Barsky; Gadi Barlevy; Jeffrey Campbell; John Cochrane; Nicolas Crouzet; Lars Hansen; Monika Piazzesi; Martin Schneider; Sam Schulhofer-Wohl; Venky Venkateswaran; François Velde; our discussants, Riccardo Colacito, Martin Lettau, Matthew Rognlie, and Jaume Ventura; our Brookings discussants, Mark Gertler and Dimitris Papanikolaou; and the *Brookings Papers* editors, Janice Eberly and James Stock.

References

- Aaronson, Daniel, Luojia Hu, Arian Seifoddini, and Daniel G. Sullivan. 2014. "Declining Labor Force Participation and Its Implications for Unemployment and Employment Growth." *Economic Perspectives—Federal Reserve Bank of Chicago* 38, 4th quarter.
- Acemoglu, Daron, and Pascual Restrepo. Forthcoming. "The Race Between Man and Machine: The Implications of Technology for Growth, Factor Shares and Employment." *American Economic Review*.
- Adrian, Tobias, Richard K. Crump, and Emanuel Moench. 2013. "Pricing the Term Structure with Linear Regressions." *Journal of Financial Economics* 110, no. 1: 110–38.
- Andrews, Isaiah, Matthew Gentzkow, and Jesse Shapiro. 2017. "Measuring the Sensitivity of Parameter Estimates to Estimation Moments." *Quarterly Journal of Economics* 132, no. 4.
- Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. 2017. "The Fall of the Labor Share and the Rise of Superstar Firms." NBER Working Paper 23396. Cambridge, Mass.: National Bureau of Economic Research. <https://www.nber.org/papers/w23396.pdf>.
- Avdis, Efsthios, and Jessica A. Wachter. 2017. "Maximum Likelihood Estimation of the Equity Premium." *Journal of Financial Economics* 125, no. 3: 589–609.
- Bansal, Ravi, and Amir Yaron. 2004. "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles." *Journal of Finance* 59, no. 4: 1481–1509.
- Barkai, Simcha. 2016. "Declining Labor and Capital Shares." Job Market Paper. <http://home.uchicago.edu/~barkai/doc/BarkaiDecliningLaborCapital.pdf>.
- Barro, Robert. 2006. "Rare Disasters and Asset Markets in the Twentieth Century." *Quarterly Journal of Economics* 121, no. 3: 823–66.
- Barro, Robert, Jesús Fernández-Villaverde, Oren Levintal, and Andrew Mollerus. 2016. "Safe Assets." NBER Working Paper 20652. Cambridge, Mass.: National Bureau of Economic Research.
- Bernanke, Ben. "The Global Saving Glut and the U.S. Current Account Deficit." Sandridge Lecture, delivered to Virginia Association of Economists. <https://www.federalreserve.gov/boarddocs/speeches/2005/200503102/>.
- Bhandari, Anmol, and Ellen R. McGrattan. 2018. "Sweat Equity in U.S. Private Business." NBER Working Paper 24520. Cambridge, Mass.: National Bureau of Economic Research. <https://www.nber.org/papers/w24520>.
- Blanchard, Olivier. 1993. "Movements in the Equity Premium." *Brookings Papers on Economic Activity* 24, no. 2: 75–138.
- Caballero, Ricardo J., and Emmanuel Farhi. 2018. "The Safety Trap." *Review of Economic Studies* 85, no. 1.
- Caballero, Ricardo J., Emmanuel Farhi, and Pierre-Olivier Gourinchas. 2008. "An Equilibrium Model of 'Global Imbalances' and Low Interest Rates." *American Economic Review* 98, no. 1: 358–93.

- . 2017. “Rents, Technical Change, and Risk Premia Accounting for Secular Trends in Interest Rates, Returns on Capital, Earning Yields, and Factor Shares.” *American Economic Review* 107, no. 5: 614–20.
- Caggese, Andrea, and Ander Perez. 2018. “Capital Misallocation and Secular Stagnation.” Pompeu Fabra University working paper.
- Campbell, Jeffrey R., Jonas D. M. Fisher, Alejandro Justiniano, and Leo Melosi. 2017. “Forward Guidance and Macroeconomic Outcomes Since the Financial Crisis.” In *NBER Macroeconomics Annual 2016, Volume 31*, edited by Martin Eichenbaum and Jonathan A. Parker. University of Chicago Press.
- Campbell, John. 2008. “Estimating the Equity Premium.” *Canadian Journal of Economics* 41, no. 1: 1–21.
- Campbell, John Y., and Samuel B. Thompson. 2008. “Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?” *Review of Financial Studies* 21, no. 4: 1509–31.
- Carvalho, Carlos, Andrea Ferrero, and Fernanda Nechio. 2016. “Demographics and Real Interest Rates: Inspecting the Mechanism.” *European Economic Review* 88: 208–26.
- Corrado, Carol, Charles Hulten, and Daniel Sichel. 2005. “Measuring Capital and Technology: An Expanded Framework.” In *Measuring Capital in the New Economy*, edited by Carol Corrado, John Haltiwanger, and Dan Sichel. University of Chicago Press.
- . 2009. “Intangible Capital and U.S. Economic Growth.” *Review of Income and Wealth* 55, no. 3: 661–85.
- Crouzet, Nicolas, and Janice Eberly. 2018. “Understanding Weak Capital Investment: The Role of Market Concentration and Intangibles.” In *Economic Policy Symposium Proceedings: Changing Market Structures and Implications for Monetary Policy*. Jackson Hole, Wyo.: Federal Reserve Bank of Kansas City.
- De Loecker, Jan, and Jan Eeckhout. 2017. “The Rise of Market Power and the Macroeconomic Implications.” NBER Working Paper 23687. Cambridge, Mass.: National Bureau of Economic Research. <https://www.nber.org/papers/w23687>.
- Del Negro, Marco, Domenico Giannone, Marc Giannoni, and Andrea Tambalotti. 2017. “Safety, Liquidity, and the Natural Rate of Interest.” *Brookings Papers on Economic Activity*, Spring: 235–94.
- Duarte, Fernando, and Carlo Rosa. 2015. “The Equity Risk Premium: A Review of Models.” Staff Report 714, Federal Reserve Bank of New York. https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr714.pdf.
- Eggertsson, Gauti B., Jacob A. Robbins, and Ella Getz Wold. “Kaldor and Piketty’s Facts: The Rise of Monopoly Power in the United States.” NBER Working Paper 24287. Cambridge, Mass.: National Bureau of Economic Research. <https://www.nber.org/papers/w24287>.
- Elsby, Michael W. L., Bart Hobijn, and Ayşegül Şahin. 2013. “The Decline of the U.S. Labor Share.” *Brookings Papers on Economic Activity*, Fall: 1–63.

- Fama, Eugene F., and Kenneth R. French. 2002. "The Equity Premium." *Journal of Finance* 57, no. 2: 637–59.
- Fernald, John. 2015. "Productivity and Potential Output Before, During, and After the Great Recession." NBER Working Paper 20248. Cambridge, Mass.: National Bureau of Economic Research.
- Fisher, Jonas D. M. 2015. "On the Structural Interpretation of the Smets–Wouters 'Risk Premium' Shock." *Journal of Money, Credit and Banking* 47, nos. 2–3: 511–16.
- Gabaix, Xavier. 2011. "Disasterization: A Simple Way to Fix the Asset Pricing Properties of Macroeconomic Models." *American Economic Review* 101, no. 3: 406–9.
- . 2012. "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance." *Quarterly Journal of Economics*, May, 645–700.
- Gagnon, Etienne, Benjamin K. Johannsen, and David López-Salido. 2016. "Understanding the New Normal: The Role of Demographics." Working Paper 2016–080. Washington: Board of Governors of the Federal Reserve System.
- Gilchrist, Simon, and Egon Zakrajšek. 2012. "Credit Spreads and Business Cycle Fluctuations." *American Economic Review* 102, no. 4: 1692–1720.
- Gomme, Paul, B. Ravikumar, and Peter Rupert. 2011. "The Return to Capital and the Business Cycle." *Review of Economic Dynamics* 14, no. 2: 262–78.
- Gourio, François. 2012. "Disaster Risk and Business Cycles." *American Economic Review* 102, no. 6: 2734–66.
- . 2013. "Credit Risk and Disaster Risk." *American Economic Journal: Macroeconomics* 5, no. 3: 1–34.
- Gourio, François, Anil K. Kashyap, and Jae W. Sim. 2018. "The Trade-Offs in Leaning Against the Wind." *IMF Economic Review* 66, no. 1: 70–115.
- Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell. 1997. "Long-Run Implications of Investment-Specific Technological Change." *American Economic Review* 87, no. 3: 342–62. <http://www.jstor.org/stable/2951349>.
- Gutiérrez, Germán, and Thomas Philippon. 2017. "Investmentless Growth: An Empirical Investigation." *Brookings Papers on Economic Activity*, Fall: 89–190.
- Hall, Robert. 2017. "The Role of the Growth of Risk-Averse Wealth in the Decline of the Safe Real Interest Rate." NBER Working Paper 22196. Cambridge, Mass.: National Bureau of Economic Research.
- Hall, Robert, and Dale W. Jorgenson. 1967. "Tax Policy and Investment Behavior." *American Economic Review* 57, no. 3: 391–414.
- Hamilton, James D., Ethan S. Harris, Jan Hatzius, and Kenneth D. West. 2015. "The Equilibrium Real Funds Rate: Past, Present and Future." *IMF Economic Review* 64, no. 4: 660–707.
- Heaton, John, and Deborah Lucas. 2000. "Stock Prices and Fundamentals." In *NBER Macroeconomics Annual 1999*, vol. 14, edited by Ben S. Bernanke and Julio J. Rotemberg. University of Chicago Press.

- Isoré, Marlène, and Urszula Szczerbowicz. 2017. "Disaster Risk and Preference Shifts in a New Keynesian Model." *Journal of Economics Dynamics and Control* 79: 97–125.
- Jagannathan, Ravi, Ellen R. McGrattan, and Anna Scherbina. 2000. "The Declining U.S. Equity Premium." *Quarterly Review of the Federal Reserve Bank of Minneapolis*, Fall, 3–19.
- Karabarbounis, Loukas, and Brent Neiman. 2014. "The Global Decline of the Labor Share." *Quarterly Journal of Economics* 129, no. 1: 61–103.
- . 2019. "Accounting for Factorless Income." In *NBER Macroeconomics Annual 2019, Volume 33*, edited by Martin Eichenbaum and Jonathan Parker. University of Chicago Press.
- Kehrig, Matthias, and Nicolas Vincent. 2018. "The Micro-Level Anatomy of the Aggregate Labor Share Decline." NBER Working Paper 25275. Cambridge, Mass.: National Bureau of Economic Research. <https://www.nber.org/papers/w25275.pdf>.
- Koh, Dongya, Raúl Santaella-Llopis, and Yu Zheng. 2015. "Labor Share Decline and Intellectual Property Products Capital." Working Paper ECO 2015/05. Florence: European University Institute. http://cadmus.eui.eu/bitstream/handle/1814/36837/ECO_2015_05.pdf?sequence=1.
- Kozłowski, Julian, Laura Veldkamp, and Venky Venkateswaran. 2015. "The Tail That Wags the Economy: Belief-Driven Business Cycles and Persistent Stagnation." Working Paper 15-25. New York: New York University Law and Economics. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2687558.
- . 2018. "The Tail That Keeps the Riskless Rate Low." In *NBER Macroeconomics Annual 2018, Volume 33*, edited by Martin Eichenbaum and Jonathan A. Parker. University of Chicago Press. <https://www.nber.org/chapters/c14073>.
- Lettau, Martin, Sydney Ludvigson, and Jessica Wachter. 2008. "The Declining Equity Premium: What Role Does Macroeconomic Risk Play?" *Review of Financial Studies* 21, no. 4: 1653–87.
- Lewis, Alexander, and Janice Eberly. 2018. "Investment Hollowing Out." *IMF Economic Review* 66, no. 1: 5–30.
- Martin, Ian. 2017. "What Is the Expected Return on the Market?" *Quarterly Journal of Economics* 132, no. 1: 367–433.
- Marx, Magali, Benoît Mojon, and François R. Velde. 2018. "Why Have Interest Rates Fallen Far Below the Return on Capital?" Working Paper 2018-01, Federal Reserve Bank of Chicago.
- Mulligan, Casey B. 2002. "Capital, Interest and Aggregate Intertemporal Substitution." NBER Working Paper 9373. Cambridge, Mass.: National Bureau of Economic Research. <https://www.nber.org/papers/w9373.pdf>.
- Nakamura, Leonard I. 2010. "Intangible Assets and National Income Accounting." *Review of Income and Wealth* 56, no. 1: S135–S155.
- Rachel, Łukasz, and Thomas Smith. 2017. "Are Low Real Interest Rates Here to Stay?" *International Journal of Central Banking* 13, no. 3: 1–42.

- Rietz, Thomas. 1988. "The Equity Premium: A Solution." *Journal of Monetary Economics* 22, no. 1: 117–31.
- Rognlie, Matthew. 2015. "Deciphering the Fall and Rise in the Net Capital Share: Accumulation or Scarcity?" *Brookings Papers on Economic Activity*, Spring: 1–54.
- Traina, James. 2018. "Is Aggregate Market Power Increasing? Production Trends Using Financial Statements." Working Paper 17. Chicago: Chicago Booth Stigler Center for the Study of the Economy and the State. <https://research.chicagobooth.edu/-/media/research/stigler/pdfs/workingpapers/17isaggregatemarketpowerincreasing.pdf>.
- Wachter, Jessica. 2013. "Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?" *Journal of Finance* 68: 987–1035.
- Welch, Ivo, and Amit Goyal. 2008. "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction." *Review of Financial Studies* 21, no. 4: 1455–1508.

Comments and Discussion

COMMENT BY

MARK GERTLER This very nice paper by Emmanuel Farhi and François Gourio certainly clarified many key issues in the literature for me. The most useful way I can use my space here is to describe what I think are the paper's key arguments. Then I offer a few suggestions.

The paper's goal is to account for a variety of macroeconomic trends over the past several decades. Farhi and Gourio describe nine trends. But I think these three facts are central to their analysis:

1. Declining real interest rates,
2. A rising capital income share, and
3. A slightly increasing average return to capital.

The first two facts are widely known, and each is the subject of a large independent body of literature. The third fact is well known by insiders in the area. Consistent with a very recent wave of literature, the authors note that the macroeconomic trends are interdependent phenomena and thus need to be studied within a unified framework. The distinctive methodological aspect of their approach is to integrate finance explicitly within their macroeconomic model. By including finance, they mean allowing for a role for risk and risk premia.

How Farhi and Gourio account for various phenomena ranges from less to more controversial. Chief among the less controversial results is the way they account for fact 1: the declining real interest rate. Here, they find that a combination of an increased propensity to save (a rising discounted factor) and increased demand for safe assets (due to increasing risk) does the job. These findings are consistent with the range of explanations in the literature (Bernanke 2005; Caballero, Farhi, and Gourinchas 2008; Del Negro and others 2017).

On the more controversial side is the way the authors account for facts 2 and 3: the rising capital income share, and the stable average return to capital. Within their baseline model, they allow for a tug of war between technology, market power, and risk. (In extensions of the baseline model, they consider other factors, such as intangible capital.) What makes the analysis somewhat controversial is the gold rush of recent literature that emphasizes rising market power and how this phenomenon can account for a variety of important phenomena, including the increasing capital share. The authors push back a bit on this euphoria by emphasizing the role of increasing risk premia. Their key message is that allowing for increasing risk premia dampens significantly (though does not eliminate) the measured increase in market power.

ACCOUNTING FOR TRENDS: THE IDENTIFICATION PROBLEM To understand the problem of disentangling the relative importance of technology, market power, and risk, it is first useful to examine the expression for the capital income share. Let W be the average wage, N total employment, R the rental rate to capital, K the capital stock, and Π monopoly profits. Then we can express the capital income share S^K as

$$(1) \quad S^K = 1 - \frac{WL}{Y} = \frac{RK + \Pi}{Y}.$$

The key point to note is that capital income is the sum of the rental income to capital RK and monopoly profits Π . Accordingly, one can categorize theories of the rising capital income share into whether they yield increasing rental income or increasing monopoly power. For example, the early literature emphasized capital-biased technological change, which involved a reallocation of rents from labor to capital. Intangible capital provides another way to account for rising rental income. Stories based on rising market power appeal to increasing markups to explain increasing profits.

The challenge in sorting out these different theories is that the division of total capital income between rents and profits is not directly observed, as Loukas Karabarbounis and Brent Neiman (2018) emphasize. A very nice paper by Simcha Barkai (2016) attempts to solve this problem directly by measuring capital rental income and then using this measure along with the total measure of capital income to impute profits. One of the problems is that the capital rental rate is not directly observed. Barkai effectively assumes that the rental rate equals the risk-free rate plus a fixed equity

premium. As a result, the measured rental rate declines with the risk-free rate. The net effect is that the measured composition of capital income shifts in favor of monopoly profits. For this reason, he finds that a large increase in the markup is required to explain the increasing capital share.

Where the authors step in is to argue that the equity premium may have increased, implying that the rental income to capital may have not fallen nearly as much as Barkai suggests, and, conversely, that monopoly profits may not have increased as much. It is largely for this reason that the authors find a much smaller increase in markups.

THE FARHI–GOURIO FRAMEWORK The model the authors develop to analyze trends is elegantly simple. It is a variant of a standard neoclassical growth model, modified to include monopoly power and risk. The way they include market power is to allow for monopolistically competitive final goods producers. These producers use intermediate goods as an input to make a differentiated final product. Intermediate goods producers, in turn, make output Y using capital K and labor N , according to this Cobb–Douglas production function:

$$(2) \quad Y_t = Z_t K_t^\alpha (S_t N_t)^{1-\alpha}$$

and where Z_t and S_t reflect productivity disturbances. To include risk, the authors add a time-varying disaster probability. Finally, they restrict the shocks to the economy to ensure that the economy is always on a balanced growth path, absent any changes in parameters. Doing so makes the model appropriate for analyzing trends.

There are three key parameters of interest:

1. $\alpha \equiv$ output elasticity of capital
2. $\mu \equiv$ gross markup
3. $\chi \equiv$ equity premium

Each parameter reflects one of the factors driving the macroeconomic trends. The output elasticity of capital α , which comes from the production function, reflects technology. We refer to a rise in α as capital-biased technical change, given that the marginal product of capital rises, everything else being equal. The gross markup μ measures market power (and is a function of the elasticity of substitution between the differentiated final output goods). Finally, the equity premium χ captures risk. Note that the primitive model parameter is the disaster probability p . However, given χ , one can use the model equations to back out p .

Over a given sample, three moment conditions pin down the parameter vector (μ, α, χ) . Let r^f denote the riskless rate, g trend growth, P the

price of stocks, and D dividends. Then the three moments conditions are given by

1. Capital income share

$$(3) \quad S^K = \frac{\alpha}{\mu} + \frac{\mu - 1}{\mu}$$

2. Average return to capital

$$(4) \quad \frac{RK + \Pi}{K} = \left(1 + \frac{\mu - 1}{\alpha}\right)(\chi + r^f)$$

3. Gordon growth formula

$$\frac{P}{D} = \frac{1}{\chi + r^f - g}$$

→

$$(5) \quad \frac{D}{P} + g = \chi + r^f$$

where r^f and g are given by data, as are the three target variables S^K , $\frac{RK + \Pi}{K}$, and $\frac{P}{D}$.¹

It is useful to give the intuition underlying each of the moment conditions. The capital income share depends on two terms: The first is the rental income share, which is increasing in α . The second is monopoly profits, which is increasing in μ . The average return to capital is a multiple of the expected return to capital, which is the sum of the risk premium and the risk-free rate, $\chi + r^f$. In the absence of market power ($\mu = 1$), the average return simply equals the expected equity return. With market power, there is an extra term that reflects monopoly profits.

Observe that conditional on the trend equity premium χ , conditions 3 and 4 determine the technology and market power parameters, α and μ . To solve for χ , the authors use the familiar Gordon growth formula, which relates the price-dividend ratio along a balanced growth path to the inverse of the expected equity return net of the steady state growth rate of output.

1. For simplicity, I am abstracting from the effects of depreciation and investment-specific technical change, which do not appear to affect the results significantly.

From rearranging the Gordon formula, one can express the trend expected return to equity as the sum of the price-dividend ratio and the steady state growth rate.

IMPLEMENTATION AND RESULTS The authors first compute averages of the three target variables over each of the two subsamples: 1984–2000 versus 2001–16. They find that across subsamples:

1. S^K increases
2. $\frac{RK + \Pi}{K}$ increases slightly
3. $\frac{D}{P} + g$ decreases slightly

They next compute model parameters over each subsample. The key findings are that across subsamples:

1. The gross markup μ increases 700 basis points
2. Technology as measured by α is unchanged
3. The equity premium χ increases 200 basis points (from 300 to 500)

I have several observations about the findings: First, the estimate of the markup increase is well below that of similar studies using aggregate data. It is about half the number estimated by Gauti Eggertsson, Jacob Robbins, and Ella Getz Wold (2018), and a third of what Barkai (2016) finds. Second, it is interesting that technology is not a factor in the declining labor share, given the widespread view that there has been significant capital-biased technological change. (Perhaps this kind of technological change mainly affects the distribution of income between skilled and unskilled labor.) Finally, the estimate of the increase in the risk premium is not without controversy, given the absence of clear indicators of increased risk since the Great Recession. I return to this issue shortly.

What is the intuition for the authors' findings? First, because the Gordon measure of the expected return to equity, $\frac{D}{P} + g$, falls by much less than the risk-free rate, r^f , the equity premium χ increases as required by equation 5. Second, the increase in χ offsets much of the effect of decline in r^f on the expected return to equity. As a result, the increase in the markup μ required to account for the uptick in the average return to capital is smaller than would be the case otherwise, as equation 4 suggests. Finally, the resulting rise in μ is sufficient to account for the rise in the labor share without any change in α , as plugging the number into equation 3 will confirm.

We now get to perhaps the central message of the paper. If we were to ignore the increase in the risk premium, the model would predict a much


larger increase in the markup. Intuitively, a much larger rise in μ would be required to account for the slight increase in the average return to capital (given the sharp decline in r^f). There is a significant corollary implication of failing to account for the rising risk premium: The overestimate of the markup leads to an underestimate of the technology parameter α . What this implies is that failing to account for the increasing risk premium leads to estimating a decline in α , suggesting that recent technical change has been labor biased, which clearly goes against conventional wisdom.

A FEW ISSUES WORTH FURTHER INVESTIGATION The paper's overall message is sensible and reasonably persuasive. It is likely that the cost of capital has not fallen nearly as much as the risk-free rate. Not taking this into account is likely to substantially overstate the increase in markup. Along these lines, it is important to take account of the role of risk in measuring the cost of capital.

Several issues, however, merit further investigation. The first involves the measure of the required expected return to capital. Over each subsample, the authors use the Gordon formula to compute the expected return to capital as the sum of the average dividend-price ratio and the average growth rate. By using subsample averages, the calculation masks a high degree of variability of the dividend-price ratio. In addition, the average growth rate may be a poor indicator of future growth expectations, especially toward the end of each subsample.

Accordingly, in my figure 1, I use annual data to compute a "real-time" Gordon measure of expected return to equity. For each year, I calculate the expected return to equity as the sum of the dividend-price ratio and the expected long-run average growth rate of output. To measure the latter I use the median 10-year average growth rate from the Survey of Professional Forecasters. As with the standard Gordon formula, two assumptions underlie the calculations: (1) the required return to equity at any time t is expected to be constant (think of it as evolving as a random walk); and (2), dividends are cointegrated with output, so expected output growth is also a measure of expected dividend growth. Think of this real-time Gordon measure as providing a benchmark estimate of the expected return to equity. To the extent that the two assumptions are violated, the expected return will differ from this benchmark.

The dashed line in my figure 1 is the dividend-price ratio, while the dotted line is the measure of the expected return given by the sum of the dividend-price ratio and the expected long-run growth rate. Because the survey data only go back to 1992, we use the 1992 forecast to measure expected output growth in the earlier years. Throughout the early

Figure 1. Real-Time “Gordon” Expected Return on Equity, 1985–2015^a

Sources: Center for Research in Security Prices; Survey of Professional Forecasters; author's calculations.

a. The expected return on equity (the upper, dashed-and-dotted line) is defined as the dividend yield (the lower, long-dashed line) plus the expected long-term growth rate: $\frac{D}{P} + \frac{E_t(g_{10})D}{P} + E_t(g_{10})$. The dividend yield is computed by the Center for Research in Security Prices. The expected 10-year growth rate is from the Survey of Professional Forecasters, and is extrapolated backward from 1992 (the dashed section of the upper, dashed-and-dotted line).


subsample, there is a downward trend in the required return to equity, which accelerates due to the stock market boom in the later 1990s (which reduces the dividend-price ratio). The stock market correction in the early 2000s reverses this downward trend. The net effect is that though the measured expected return in the second subsample is lower than in the first one, the difference is not dramatic, consistent with the authors' argument.

In particular, the decline in the measured expected return to equity is much less over the sample than is the drop in the expected 1-year Treasury yield, as my figure 2 shows. To the extent that we can take as an estimate of the equity premium the gap between the Gordon measure of the expected return to equity and the expected 1-year Treasury yield, then it is clear from the figure that the equity premium has widened nontrivially over the sample, as the authors suggest.

But two concerns arise. First, to calculate the equity premium using the Gordon approach, investors must expect the current 1-year yield to persist.²

2. Otherwise, for example, a high dividend-price ratio could reflect an expected increase in future interest rates as opposed to a high equity premium.

Figure 2. Real-Time “Gordon” Equity Premium, 1985–2015^a




Sources: Adrian, Crump, and Moench (2013); Survey of Professional Forecasters; author's calculations.

a. The equity premium is the expected return on equity $\frac{D}{P} + E_t(g_{10})$ (dotted line) minus the expected real 1-year Treasury yield $E_t(rf - \pi)$ (dashed line). The expected equity return is computed by the Center for Research in Security Prices and the Survey of Professional Forecasters. The 1-year nominal Treasury yield is from Adrian, Crump, and Moench (2013). The expected 1-year Consumer Price Index inflation rate is from the Survey of Professional Forecasters.

Not only is there a downward trend in the real rate over the sample; there are also clear cyclical patterns: Relative to trend, the short-term real rate increases in expansions and decreases in recessions. An open question is how much investors perceive the low real rates after the Great Recession as reflecting a trend versus a cycle. As I discuss below, this matters for the calculation of the benchmark equity premium using the Gordon formula. The second issue involves identifying where the increase in risk in the system may be that could account for the increasing risk premium.

I address the two issues in reverse. First, where is the risk? The puzzle is that some traditional indicators of risk, such as the Chicago Board Options Exchange's Volatility Index market indicator, are down. I think the most natural source of greater risk is the perceived increase in risk to the banking system. Within the authors' model, the relevant risk is that of a disaster, which would lead to an exogenous decline in real activity. In practice, at the core of most economic disasters are banking crises. My figure 3, which is adapted from a paper by Darrell Duffie (2019), plots the average credit default swap (CDS) rate for banks from 2004 to 2018. The CDS rate

Figure 3. Where Is the Risk? Bank Credit Default Swaps, 2005–17^a


Sources: Bloomberg; Duffie (2019).

a. Average 5-year credit default swap rates (in basis points) of the five major U.S. dealer banks: Bank of America Merrill Lynch, Citigroup, Goldman Sachs, JPMorgan Chase, and Morgan Stanley.

increases from below 50 basis points before the Great Recession to a peak of 250 basis points at the recession's height. Importantly, the rate fluctuates between 150 and 250 basis points through 2013. It eventually declines a bit, but remains elevated relative to its pre–Great Recession value by a factor of roughly three (about 100 basis points, versus roughly 30 pre–Great Recession). Accordingly, the CDS data suggest that market perceptions of the probability of a banking crisis are elevated relative to the pre–Great Recession period. As Duffie notes, the experience of the recent crisis has led market participants to attach a higher probability to a future crisis than might otherwise have been the case. Also relevant are new restrictions on the extent to which the government can protect banks and bank creditors. The elevated perception of bank risk could account for the authors' observation that credit spreads are high after relative to before the Great Recession. It similarly could be a factor accounting for an increase in the equity risk premium.

Finally, given the real-time Gordon measure of the return to equity, I address the issue of which real rate to use to calculate the equity premium. Because the Gordon measure is effectively a trend measure of the return to equity at each point in time, the real rate with which to compare this return should similarly be a trend measure. A natural candidate for the latter is the

Figure 4. Short-Term Rates versus Risk-Neutral Long-Term Rates, 1985–2015^a




Sources: Adrian, Crump, and Moench (2013); Survey of Professional Forecasters; author's calculations.

a. The expected real 1-year Treasury yield $E_t(rf - \pi)$ (dotted line) is the nominal 1-year yield minus the expected 1-year Consumer Price Index inflation rate. The expected real risk-neutral longer-term Treasury yield $E_t(rn^{10} - \pi)$ (dashed-and-dotted line) is the nominal 10-year yield minus the term premium (per Adrian, Crump, and Moench) and the expected Consumer Price Index inflation rate. Yields are from Adrian, Crump, and Moench (2013). Inflation expectations are from the Survey of Professional Forecasters.

10-year government bond rate adjusted to eliminate the term premium. After eliminating the term premium, the 10-year bond rate reveals the market expectation of the average long-term real rate. Accordingly, the dashed-and-dotted line in my figure 4 plots the long-term real rate, measured as the nominal 10-year government bond rate adjusted to eliminate the term premium as measured by Michael Abrahams and others (2016), minus the 10-year forecast of inflation from the Survey of Professional Forecasters. Though the 10-year rate exhibits a secular decline similar to the 1-year rate (the dotted line), it is not as steep. In addition, not surprisingly, the cyclical deviations from trend are smaller than for the 1-year rate. An important consequence is that the long-run rate is below the short-run rate at the beginning of the sample, a period when monetary policy was still tight. Conversely, it is significantly above the short term rate at the end of the sample, a period of easy monetary policy.

As my figure 5 shows, if we use the 10-year real interest rate to compute the trend equity premium, we get a different perspective on the behavior of relative returns. The trend equity premium looks reasonably stable over

Figure 5. Equity Return versus Real Short- and Long-Term Yields, 1985–2015^a

Sources: Adrian, Crump, and Moench (2013); Center for Research in Security Prices; Survey of Professional Forecasters; author's calculations.

a. The expected return on equity $\frac{D}{P} + E_t(g_{10})$ (solid line) is the dividend yield plus the expected long-term growth rate. The expected real 1-year Treasury yield $E_t(rf - \pi)$ (dotted line) is the nominal 1-year yield minus the expected 1-year Consumer Price Index inflation rate. The expected real risk-neutral longer-term Treasury yield $E_t(rn^{10} - \pi)$ (dashed-and-dotted line) is the nominal 10-year yield minus the ACM term premium and the expected Consumer Price Index inflation rate. The dividend yield is computed by the Center for Research in Security Prices. Treasury yields are from Adrian, Crump, and Moench (2013). Expectations are from the Survey of Professional Forecasters.

the sample, except for a decrease over the period of the stock market boom in the late 1990s that is reversed over the next few years. It is important to emphasize, however, that the authors' estimates of the markup and technology parameters remain valid, as does their argument that previous studies have likely overestimated the increase in markups. What matters for the estimation of these parameters is the estimate of the return on equity and not how this return is divided between the risk premium and the risk-free rate. My only point here is that if one is going to use the Gordon formula to back out an equity premium, it matters which real rate is used, and it may make more sense to use the 10-year rate adjusted for the term premium.

CONCLUDING REMARKS This paper makes a compelling case that in analyzing macroeconomic trends, it is important to think carefully about measuring the cost of capital. By doing so, further, one is likely to obtain much lower estimates of the rise in markups than the previous literature has suggested.

REFERENCES FOR THE GERTLER COMMENT

- Abrahams, Michael, Tobias Adrian, Richard K. Crump, Emmanuel Moench, and Rui Yu. 2016. "Decomposing Real and Nominal Yield Curves." *Journal of Monetary Economics* 84: 182–200.
- Adrian, Tobias, Richard K. Crump, and Emanuel Moench. 2013. "Pricing the Term Structure with Linear Regressions." *Journal of Financial Economics* 110, no. 1: 110–38.
- Barkai, Simcha. 2016. "Declining Labor and Capital Shares." Unpublished paper, London Business School.
- Bernanke, Ben. 2005. "The Global Saving Glut and the U.S. Current Account Deficit." Speech at Federal Reserve Board, March 10.
- Caballero, Ricardo J., Emmanuel Farhi, and Pierre-Olivier Gourinchas. 2008. "An Equilibrium Model of 'Global Imbalances' and Low Interest Rates." *American Economic Review* 98: 358–93.
- Del Negro, Marco, Domenico Giannone, Marc Giannoni, and Andrea Tambalotti. 2017. "Safety, Liquidity, and the Natural Rate of Interest." *Brookings Papers on Economic Activity*, Spring: 235–94.
- Duffie, Darrell. 2019. "Prone to Fail: The Pre-Crisis Financial System." *Journal of Economic Perspectives* 33, no. 1: 81–106.
- Eggertsson, Gauti B., Jacob A. Robbins, and Ella Getz Wold. 2018. "Kaldor and Piketty's Facts: The Rise of Monopoly Power in the United States." NBER Working Paper 24287. Cambridge, Mass.: National Bureau of Economic Research.
- Karabarbounis, Loukas, and Brent Neiman. 2018. "Accounting for Factorless Income." In *NBER Macroeconomics Annual 2018*. University of Chicago Press.

COMMENT BY

DIMITRIS PAPANIKOLAOU This paper by Emmanuel Farhi and François Gourio illustrates how taking into account financial market data helps explain some recent stylized features of the data: the decline in the labor share of output; the decline in interest rates; the increase in the average product of capital in excess of the riskless rate; and the relatively low levels of corporate investment as a share of output. Previous explanations have relied on a combination of a rise in the importance of intangibles and/or an increase in firms' market power (Barkai 2017; De Loecker and Eeckhout 2017; Gutierrez and Philippon 2017). But in this paper, Farhi and Gourio show that stable equity valuation ratios and declining risk-free rates strongly suggest that the equity premium has increased in recent decades. A structural macroeconomic model attributes a considerable role to an increase in risk in accounting for these recent trends—and a much more modest role for an increase in market power. Interestingly, allowing for the

presence of intangibles—here, mismeasured capital—weakens the case for rising markups, but not for rising risk premia. Given that the main contribution of the paper is to provide new evidence on the rising equity premium, my comment mostly focuses on this aspect of the paper.


Farhi and Gourio have written an important paper that illustrates how asset markets can be a useful source of information on macroeconomic models. Overall, I am highly sympathetic to the authors' goal, and I find their main argument broadly convincing. That said, there needs to be some scope for clarifying the limitations of their approach: the equity premium is essentially unobservable, and can only be inferred from the data based on additional assumptions. Hence, the authors' argument would be greatly strengthened if they were to empirically link the imputed equity premium with observable measures of risk. Absent this link, the imputed increase in the equity premium can only be rationalized as an increase in risk aversion—and because shifts in preference parameters are unobservable, they are ultimately unsatisfying as explanations of economic phenomena.¹

The novel part of the paper infers the equity premium from equity valuations. To understand the authors' identification strategy, consider the familiar Gordon growth formula. It can be rewritten as

$$(1) \quad \frac{D}{P} - r_f = E[R_m^e] - E[g].$$

The Gordon growth formula links two observable quantities on the left side (the dividend-price ratio and the real risk-free rate) to two unobservable quantities on the right side: the expected excess return on equity $E[R_m^e]$ and the expected growth rate of dividends $E[g]$. The two panels of my figure 1 plot the dividend-price ratio and two measures of the real risk-free rate: the yield on a 10-year Treasury Inflation-Protected Security, and the difference between the 10-year yield of the Constant Maturity Rate series of the Federal Reserve Bank of Saint Louis and 10-year inflation expectations from the Survey of Professional Forecasters (SPF). Examining these two panels brings the main point of the paper into sharp focus: We see that, in terms of levels, stock valuation ratios are at the same level as in 2003,

1. That said, risk aversion in these models is often a metaphor that can be a stand-in for other types of frictions. Specifically, models with financial constraints often imply that economic agents exhibit risk-averse behavior, even if their underlying utility is linear (He and Krishnamurthy 2013). Thus, an alternative route would be to link the imputed equity premium with measures of the health of financial intermediaries.

Figure 1. Interest Rates and Valuation Ratios, 1985–2015^a

Sources: Center for Research in Security Prices; Federal Reserve Bank of Saint Louis, Constant Maturity Rate series; Survey of Professional Forecasters.

a. The left panel plots the price-dividend ratio from the Center for Research in Security Prices. The right panel plots estimates of the real interest rate: the solid line plots the difference between the 10-year nominal rate (yield on Constant Maturity Rate series bonds) and the expected inflation over the next 10 years from the Survey of Professional Forecasters; the dashed line plots the yield of 10-year Treasury Inflation-Protected Securities.


b. Constant Maturity Rate series, expected inflation on Treasury Inflation-Indexed Securities.

even though the real rate of interest declined from about 4 percent in 2003 to less than 1 percent in 2016.

These patterns are consistent with equation 1, as long as either expected dividend growth rates have declined or the equity risk premium has risen. The authors equate dividend to output growth, and assume that expected growth is equal to average realized growth in each period. Because average realized growth was about 30 basis points lower in 2001–16 than in 1984–2000, they conclude that the difference needs to be accounted for by an increase in the equity premium. But is it always reasonable to equate expectations with average realizations? If we were to estimate the expected return on equity based on the average realized return of stocks in excess of bonds in each period, we would have arrived at the opposite conclusion: During the 1984–2000 period, stocks outperformed bonds by 10.5 percent compared with 7.3 percent in 2001–17.² Now, there are some very good reasons why estimating the equity premium based on average realizations is fraught with pitfalls; not only are realized stock returns quite noisy, but they are also inversely related to changes in expectations for future returns. Nevertheless, perhaps we should not completely discard this information.

2. Estimates based on data from Kenneth French's website, <http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/>.

Figure 2. Imputed Dividend Growth Rate, Assuming Constant Equity Premium, 1990–2016^a




Sources: Center for Research in Security Prices; Federal Reserve Bank of Saint Louis, Constant Maturity Rate series; Survey of Professional Forecasters; author's calculations.

a. This figure plots the imputed expected growth rate of dividends $E[g]$, given equation 1, the Center for Research in Security Prices' price-dividend ratio, and the real risk-free rate—measured as the difference between the 10-year nominal rate (the yield on Constant Maturity Rate bonds) and expected inflation over the next 10 years from the Survey of Professional Forecasters.

For the sake of argument, suppose that we were to assume a constant equity premium and back out the expected growth rate $E[g]$ from equation 1, together with the realizations of D/P and r_f .³ I plot the resulting series in my figure 2. We see that the data would imply a secular decline in expected growth rates after 2000. Is the resulting expectations series reasonable? Without additional work, it is rather difficult to ascertain whether that is the case. One possibility would be to extend the estimation exercise to allow households' prior beliefs about future productivity to vary from average realizations. One could then infer the extent to which these differences in beliefs could account for additional features of the data—for instance, the decline in corporate investment.

Data on expectations of future economic growth and asset returns could shed some light on these issues. I use expectations of future output growth

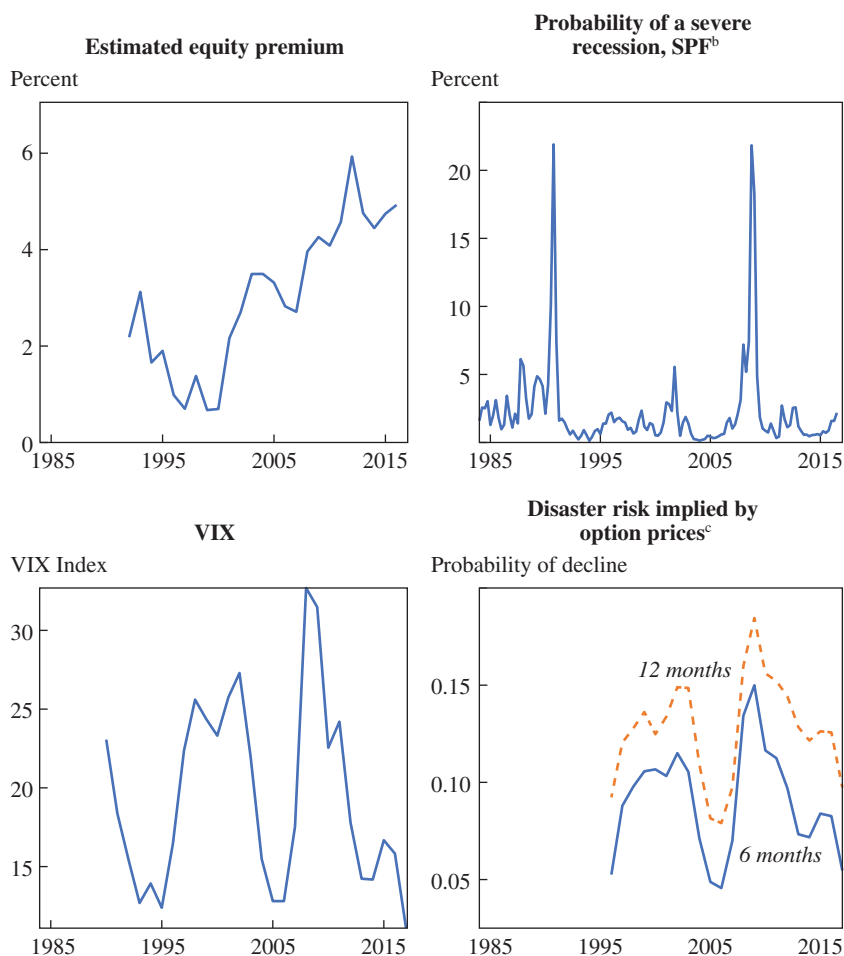
3. One could object to this exercise on the grounds that the price-dividend ratio does not appear to forecast future dividend growth very well (Campbell and Shiller 1988). However, recent work by Van Binsbergen and Koijen (2010) shows that, using a different empirical methodology, dividend growth may be predictable.

Figure 3. Expectations of Future Growth and Excess Return on Equity, 1990–2015

Sources: Survey of Professional Forecasters (SPF); Duke University Quarterly Survey of Chief Financial Officers; author's calculations.

over the next 10 years from the SPF. To measure expectations of future excess returns on equity, I use data from Duke University's quarterly survey of chief financial officers (CFOs) (Duke 2019). In the survey, CFOs are asked what they think the average excess return of the Standard & Poor's 500 will be over the next 10 years. Given that these CFOs are responsible for the capital budgeting decision of the largest firms in the economy, their beliefs about risk premia are likely consequential.

I plot these two series of expected growth and stock market returns, respectively, in the two panels of my figure 3. Examining the left panel, we see that survey expectations of future growth display a qualitatively similar pattern as the imputed growth rate in my figure 2, but the magnitudes are off by a considerable amount. Of course, we should keep in mind that the resulting series are not directly comparable—we are ignoring leverage, taxes, and all other distinctions between cash dividends and output. In the right panel, I plot the equity premium implied by the CFO survey data. The series starts in 2000, hence it is not possible to make comparisons with the pre-2000 period. But we can compare the resulting series with the rolling estimate of the equity premium in Farhi and Gourio's figure 6—or the top left panel of my figure 4. We see that the survey-based measure of the equity premium declines between 2000 and 2006, but then exhibits a secular increase in the 2007–16 period. Naturally, we can quibble on what exactly these surveys measure—hopes about future market performance versus required rates of return. But the point remains that inferring required rates of return from equity valuations is not straightforward.

Figure 4. Estimates of Disaster Risk, 1985–2015^a

Sources: Survey of Professional Forecasters (SPF); Federal Reserve Bank of Saint Louis, Constant Maturity Rate series; Chicago Board Options Exchange's Volatility Index (VIX); Martin (2017); author's calculations.

a. This figure presents estimates of macroeconomic risk from several sources. The top left panel plots a point-in-time version of the equity premium based on SPF forecasts on inflation and growth over the next 10 years, and the nominal yield on 10-year Treasury bonds (using the Constant Maturity Rate series). The top right panel plots the forecasted probability (from the SPF) of a decline in output in at least three out of the next four quarters. The bottom left panel plots the VIX. And the bottom right panel plots the perceived likelihood of a 15 percent decline in the stock market, from the perspective of a log investor who is fully invested in the market portfolio, from Martin (2017); the solid line uses options of 6-month maturity, and the dashed line uses options of 12-month maturity.

b. GDP decline in three out of the next four quarters.

c. For maturities of 6 months and 12 months; see note a.

To strengthen the main point of Farhi and Gourio's paper, it would be useful to connect the imputed increase in the equity risk premium to observed measures of risk and uncertainty. In the paper, risk is modeled as a (small) possibility of a (large) disaster—that is, destruction of 15 percent of the capital stock. Hence, examining empirical measures of disaster risk is a useful place to start. Naturally, this is easier said than done. Part of the difficulty lies with the fact that rare disasters are, by definition, rare. In the postwar sample, there has been not a single event when the capital stock declined by 15 percent, but given the low estimated probabilities of disaster (3–6 percent), such lucky stretches are not implausible. It is therefore extremely difficult for an econometrician to estimate a time-varying likelihood of a rare disaster from data on real outcomes. However, we have access to additional sources of data: macroeconomic surveys and—consistent with the spirit of the paper—data from financial markets.

I consider three empirical measures of disaster risk. First, I use data from the SPF; I construct the average probability, across survey participants, of a severe recession, which I define as a decline in real output in at least 3 quarters over the next year. Second, I use the Chicago Board Options Exchange's Volatility Index (VIX); this variable, often referred to as "the fear index" in the popular press, is the implied volatility of the Standard & Poor's 500 stock market index that is consistent with traded options on the index. The VIX is an amalgam of the perceived risk in investing in the stock market and the degree of risk aversion of a representative investor. If one is willing to make additional assumptions, one can recover investors' beliefs about the risk of rare disasters from option prices. Ian Martin (2017) derives the perceived probability of a 15 percent drop in the underlying index over the next year, from the perspective of an investor who is 100 percent invested in the stock market and has log utility preferences. I use these implied probabilities, based on 6-month and 12-month equity options, as my third measure of disaster risk.

My figure 4 compares these three estimates of disaster risk to the estimates implied by the paper. Specifically, the top left panel of figure 4 plots a point-in-time version of the equity premium in the paper that uses equation 1 above, along with point-in-time estimates of the real risk-free rate and expected (output) growth using the yield on 10-year Treasury securities and forecasts of inflation and output from the SPF. We see a significant upward trend in the equity premium after 2000. In contrast, as we see in the top right panel of figure 4, survey estimates of disaster risk provide rather weak support for a low-frequency increase in perceived macroeconomic risk. Survey estimates of risk spike during recessions, but

there are no differences in the average probability between the 1984–2000 and 2000–2015 subsamples. Using different definitions of a “severe recession” yields similar results.

Prices of financial options are reliably available only after the mid-1990s, so we cannot reliably compare the pre-2000 to the post-2000 period. However, we can examine whether they imply a secular increase in disaster risk relative to 2000. The bottom left panel of my figure 4 plots the time series of the VIX. The VIX spiked considerably in the late 1990s and during the Great Recession. Though the average level is somewhat higher during the 2001–15 period relative to 1990–2000, the difference is not statistically significant—probably because the VIX itself is quite volatile. The bottom right panel plots the option-implied estimates of disaster risk, using the methodology of Martin (2017). We see that the resulting series resembles the VIX, and again reveals no evidence of a secular increase in disaster risk after 2000.

In sum, we see that data from macroeconomic surveys and financial markets indicate a transitory increase in the likelihood of a rare disaster during the financial crisis. However, there is no evidence for a secular increase in disaster probabilities after 2000. Here, however, it is helpful to step a bit outside the exact structure of the model; rare disasters are a convenient device to model risk that delivers a realistic equity premium, but they are not the only possibility. A credible alternative is that macroeconomic risk takes the form of uncertainty about long-term economic growth—that is, “long-run risk,” as described by Ravi Bansal and Amir Yaron (2004).

Is it possible that perceived uncertainty about long-run growth rates has increased over the last few decades? Perhaps it has; but unfortunately, obtaining direct evidence for small but persistent sources of fluctuations in output is as challenging as obtaining evidence for the changing likelihood of rare disasters. One possibility is to estimate such risk using a structural model—in a way that is similar to what is done by Farhi and Gourio in this paper. Along these lines, Frank Schorfheide, Dongho Song, and Amir Yaron (2018) estimate a structural model in which consumption and dividends are modeled in reduced form. Importantly, there is uncertainty about the long-run mean of consumption growth, and the level of uncertainty varies over time in a persistent fashion. Schorfheide, Song, and Yaron (2018) estimate this time-varying volatility using a particle filter (a non-linear version of the Kalman filter) that uses asset returns, and the growth rates of consumption and dividends. In sum, Schorfheide, Song, and Yaron (2018) and Farhi and Gourio both rely on asset return data, but their methodologies are quite different.


Figure 5. Estimates of Long-Run Risk versus the Equity Premium, 1985–2015^a

Sources: Schorfheide, Song, and Yaron (2018); Federal Reserve Bank of Saint Louis, Constant Maturity Rate series; Survey of Professional Forecasters (SPF); author's calculations.

a. The solid line in this figure plots the filtered volatility of the long-run risk component from Schorfheide, Song, and Yaron (2018). The dashed line plots a point-in-time estimate of the equity risk premium constructed using SPF forecasts on inflation and growth over the next 10 years, and the nominal yield on 10-year Treasury bonds (using the Constant Maturity Rate series).

In my figure 5, I compare Schorfheide, Song, and Yaron's (2018) estimate of long-run uncertainty (the solid line) with the imputed point-in-time estimate of the equity premium implied by the Gordon growth formula. Interestingly, even though the two papers use different data and methodologies, they display similar behavior. That is, both methodologies imply a secular increase in macroeconomic risk after 2000. Though this correlation is comforting, it still does not fully settle the matter—what aspects of the data identify an increase in uncertainty here is not fully transparent.

However, once we move beyond the notion that disaster risk is the primary determinant of risk premia, we can expand the sources of data that can be used to directly measure risk. Fiscal and monetary policy likely have a measurable impact on economic quantities. Yet another possibility is that perceptions of political risk have shifted since 2000. To explore this idea further, I use the political uncertainty index of Scott Baker, Nicholas Bloom, and Steven Davis (2016). Specifically, Baker and colleagues construct an estimate of the degree of uncertainty about economic policy,

Figure 6. Economic Policy Uncertainty, 1985–2015^a

Sources: Baker, Bloom, and Davis (2016); Federal Reserve Bank of Saint Louis, Constant Maturity Rate series; Survey of Professional Forecasters (SPF); author's calculations.

a. The solid line in this figure plots the Economic Policy Uncertainty Index of Baker, Bloom, and Davis (2016). The dashed line plots a point-in-time estimate of the equity risk premium constructed using SPF forecasts for inflation and growth over the next 10 years, and the nominal yield on 10-year Treasury bonds (using the Constant Maturity Rate series).

based on an analysis of news articles. Their index captures uncertainty not only about which policies will be implemented but also on their economic impact—about half the articles discuss uncertainty about the economic effect of past, current, or future policy actions.

I plot Baker and colleagues' index in my figure 6. We see an increase in the average level of economic policy uncertainty in the 2001–15 period relative to 1984–2000. Some of this increase can be attributed to the financial crisis and uncertainty about the short- and long-run outcomes of the economic policies that were undertaken to remedy its effects. But their index is also high in the few years after 2000, partly due to the September, 11, 2001, terrorist attacks; the collapse of the tech “bubble”; and the second Gulf War—all of which could have plausibly increased the level of uncertainty about future economic growth. Interestingly, the policy uncertainty series exhibits behavior that is similar to the implied equity risk premium.

In brief, I think the main point of Farhi and Gourio's paper is most likely correct. Financial market data seem to indicate an increase in risk premia

after 2000. In any reasonable macroeconomic model, an increase in risk will lead to lower investment in risky projects; a higher capital share; lower interest rates; and a higher average return on capital. I find these forces equally plausible explanations as an increase in market power. My only reservation is that it is not immediately obvious how exactly the economy became riskier after 2000. Perhaps increased political uncertainty—and polarization—played a role. To lend further credibility to the argument that risk premia played an important role for recent trends, I think more work on measurement is needed.

More broadly, I believe that the economic interpretation of these accounting decompositions has been underexplored. In the context of a model, these decompositions quantify the extent to which certain shifts in the data can be accounted for by changes in parameters. But the interpretations of these parameter shifts are not obvious, and the same economic forces may account for all these changes. For instance, brand value is a form of intangible capital that gives firms some measure of market power. Thus, a rise in market power could be driven by an increased importance of intangibles (Crouzet and Eberly 2018a, 2018b). Similarly, one could argue that intangible capital is more fragile than physical capital; it is perhaps easier to argue that 15 percent of the value of a brand is lost than, say, a 15 percent destruction of machines. As the composition of the economy shifts between tangibles and intangibles, so will risk in the economy change endogenously. Understanding the fundamental causes driving these changes is worthwhile.

REFERENCES FOR THE PAPANIKOLAOU COMMENT

- Baker, Scott R., Nicholas Bloom, and Steven J. Davis. 2016. “Measuring Economic Policy Uncertainty.” *Quarterly Journal of Economics* 131, no. 4: 1593–1636.
- Bansal, Ravi, and Amir Yaron. 2004. “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles.” *Journal of Finance* 59, no. 4: 1481–1509.
- Barkai, Simcha. 2017. “Declining Labor and Capital Shares.” Working paper, University of Chicago.
- Campbell, John Y., and Robert J. Shiller. 1988. “Stock Prices, Earnings, and Expected Dividends.” *Journal of Finance* 43, no. 3: 661–76.
- Crouzet, Nicholas, and Janice Eberly. 2018a. “Intangibles, Investment, and Efficiency.” *American Economic Association: Papers and Proceedings* 108: 426–31.
- . 2018b. “Understanding Weak Capital Investment: The Role of Market Concentration and Intangibles.” Technical report for Jackson Hole Symposium,

- Federal Reserve Bank of Kansas City. <https://www.kansascityfed.org/~media/files/publicat/sympos/2018/papersandhandouts/824180816crouzeteberlyhandout.pdf?la=en>.
- De Loecker, Jan, and Jan Eeckhout. 2017. "The Rise of Market Power and the Macroeconomic Implications." NBER Working Paper 23687. Cambridge, Mass.: National Bureau of Economic Research.
- Duke University. 2019. "Duke CFO Global Business Outlook." <https://www.cfosurvey.org/>.
- Gutierrez, German, and Thomas Philippon. 2017. "Investment-less Growth: An Empirical Investigation." *Brookings Papers on Economic Activity*, Fall: 89–169.
- He, Zhiguo, and Arvind Krishnamurthy. 2013. "Intermediary Asset Pricing." *American Economic Review* 103, no. 2: 732–70.
- Martin, Ian. 2017. "What Is the Expected Return on the Market?" *Quarterly Journal of Economics* 132, no. 1: 367–433.
- Schorfheide, Frank, Dongho Song, and Amir Yaron. 2018. "Identifying Long-Run Risks: A Bayesian Mixed-Frequency Approach." *Econometrica* 86, no. 2: 617–54.
- Van Binsbergen, Jules H., and Ralph S. J. Koijen. 2010. "Predictive Regressions: A Present-Value Approach." *Journal of Finance* 65, no. 4: 1439–71.

GENERAL DISCUSSION James Stock began by noting that it would be useful to get a better sense of what the authors' macroeconomic risk variable reflects, because its historical time series behavior does not necessarily square with what are conventionally thought of as risky periods.

Robert Hall commented that he has found evidence for growth in average market power in some of his own recent research. But there is a low correlation between growth in market power and growth in concentration. He explained that the two phenomena can coexist in terms of oligopoly theory.¹ His research finds there has been a considerable rise in both rents and Tobin's Q —a finding that can be reconciled with little growth in market power if intangible assets have become more important to firms. He recommended a Jackson Hole paper by Janice Eberly and Nicolas Crouzet that corroborates the importance of intangibles, and cited research by James

1. Robert Hall, "New Evidence on the Markup of Prices over Marginal Costs and the Role of Mega-Firms in the U.S. Economy," NBER Working Paper 24574 (Cambridge, Mass.: National Bureau of Economic Research, 2018).

Traina that provides a strong critique of the evidence that market power has grown significantly since the 1980s.²

Hall said he was surprised that none of the presenters discussed the Campbell-Shiller method of measuring the equity premium, and that the hypothesis that there has been a persistent increase in the equity premium would not be supported by what he regards as the mainstream finance literature.³

Steven Davis remarked that the paper's dividend-price ratio, a key input into its analysis, mirrors the time series history of influxes of newly listed firms in the 1980s and 1990s, and that this may present a challenge for their calculation of the ratio. Research by Eugene Fama and Kenneth French shows that the flow of newly listed firms in the United States represented a large share of public firms in the 1980s and 1990s.⁴ Later research by Davis and his colleagues calculated that firms first listed in the 1980s and 1990s accounted for more than 40 percent of all employment at publicly listed firms as of 2000.⁵ Thus, Davis concluded, there may be a significant role for selection in the evolution of the paper's measured dividend-price ratio, because firms that were first listed in the 1980s and 1990s were likely to have high prices and low dividends. Moreover, this trend reversed after the dot-com bubble burst in the early 2000s. He suggested that the authors recalculate the dividend-price ratio using microeconomic data to construct an index of changes in the ratio based on firms that are listed in consecutive years.

Olivier Blanchard noted that the authors ought to be careful in distinguishing between markups and rents, given that monopolistic competition

2. James Traina, "Is Aggregate Market Power Increasing? Production Trends Using Financial Statements," Stigler Center New Working Paper 17, 2018; Nicolas Crouzet and Janice Eberly, "Understanding Weak Capital Investment: The Role of Market Concentration and Intangibles," technical report for Jackson Hole Symposium, Federal Reserve Bank of Kansas City (<https://www.kansascityfed.org/~media/files/publicat/sympos/2018/papersandhandouts/824180816crouzeteberlyhandout.pdf?la=en>).

3. John Campbell and Robert Shiller, "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies* 1, no. 3 (1988): 195–228.

4. Eugene Fama and Kenneth French, "New Lists: Fundamentals and Survival Rates," *Journal of Financial Economics* 73, no. 2 (2004): 229–69.

5. Steven Davis, John Haltiwanger, Ron Jarmin, and Javier Miranda, "Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms," in *NBER Macroeconomics Annual 2006*, edited by Kenneth Rogoff and Daron Acemoglu (Cambridge, Mass.: National Bureau of Economic Research, 2016).

with free entry leads to markups, which cover fixed costs of entry, but not to rents. As a result, some markets could have seen large increases in markups but small increases in rents.

Blanchard observed that Tobin's Q has increased substantially for non-financial firms in the United States, and that this could either be the result of measurement methods or increasing rents. In contrast to Robert Hall's view, he argued that mismeasurement of capital due to an increase in intangibles investment would need to be implausibly large to explain the increase in Tobin's Q , and thus that increasing rents must make up a large portion of the increase.

Eric Swanson noted that an increase in the savings supply is a key explanatory factor in the authors' analysis, but that this increase in savings is modeled as coming from a change in the domestic discount factor rather than as a capital inflow from abroad. Thus, the authors are studying a "domestic savings glut" rather than a "global savings glut," and the effects of the latter in an open economy can be different in important ways (such as the effect on domestic consumption growth). Swanson also observed that many of the trends the authors describe were present in Europe over the same period, and he suggested that the authors fit their model using European data as a second set of observations to check the robustness of their findings.

Jason Furman remarked that much of the literature on changes in the capital share assume it is a description of technology and nothing more. He noted the importance, thus, of the authors finding a significant role for markups in explaining changes in the capital share. He suggested that the authors consider exploiting variation in concentration across industries to test whether their findings about markups hold across industries.

Janice Eberly responded to the comments by Hall, Blanchard, and Furman, noting that her research with Nicolas Crouzet found not only a role for intangibles and investment but also that they appear to be co-related to both markups and productivity growth. She explained that intangibles should be treated as having different properties from physical capital, and that their properties may vary across industries. In the health industry, for example, intangibles appear to be closely related to markups but not to productivity; in contrast they appear to be correlated with productivity growth in the retail sector.

John Haltiwanger observed that measures of risk in fixed-income markets were declining both before and after the financial crisis, and he asked the authors to comment on why returns in debt markets could have been so low while they were rising in equity markets.

Mark Gertler responded that Baa- and A-rated bond yields have remained elevated since the financial crisis relative to their levels before the crisis. Haltiwanger responded again, noting that high-yield bonds in particular have low yields relative to precrisis levels, and that these provide a closer measure of fixed-income risk.

François Gourio began by thanking the commenters and participants for their observations. He noted that many commented on what has driven macroeconomic risk perception to increase alongside the equity premium. He pointed out that the paper tries to provide some evidence on this question by looking at other measures of risk, such as realized volatility and credit spreads. Another possible set of explanations focuses on changes in risk preferences. For example, he described how aging populations may have higher risk aversion and a larger demand for safe assets. Also, some countries appear to have larger preferences for safe assets, and these may be driving estimates of risk premia.

Responding to comments about estimating the equity premium, Gourio noted that Campbell proposes a method that differs from the Campbell-Shiller approximation. In section VII of their paper, Farhi and Gourio provide an alternative estimate of risk premia according to this method, and they find that it appears to increase after 2000, consistent with their own estimates.⁶ He acknowledged that estimating the premium involves some uncertainty, and he suggested that further research could explore the differences between estimation methods.

Gourio acknowledged that modeling one closed economy (that of the United States) is a potential limitation of the paper. However, he argued that one could conceivably treat the model as applying to the global economy, given that many trends observed in the United States are consistent with those observed globally.

Gourio agreed with comments that many of the parameters in the model are reduced-form, to some extent, and that they may be driven by another factor not included in the model, or they may be jointly driven by one common underlying factor. However, the contribution of the paper is to recover these reduced-form parameters, and to decompose their relative importance within the model. Deeper analyses that try to explain what drives these changes in parameters are of course warranted, but they will need to be consistent with the authors' reduced-form findings.

6. John Campbell, "Estimating the Equity Premium," *Canadian Journal of Economics* 41, no. 1 (2008): 1–21.

Regarding Tobin's Q , Gourio said that it is important to note that, though it is equal to 1 regardless of the risk premium if there are no rents, it is actually quite sensitive to the risk premium (and to other parameters) if there are rents, because the risk premium affects the discounting of future rents. As a result, he said, the model is consistent with an increase in Tobin's Q .

Gourio concluded by agreeing with comments about distinguishing markups from rents, considering cross-industry evidence, and taking into account firm selection when estimating the dividend-price ratio.

The Real Effects of Disrupted Credit: Evidence from the Global Financial Crisis

ABSTRACT Economists both failed to predict the global financial crisis and underestimated its consequences for the broader economy. Focusing on the second of these failures, this paper makes two contributions. First, I review research since the crisis on the role of credit factors in the decisions of households, firms, and financial intermediaries and in macroeconomic modeling. This research provides broad support for the view that credit market developments deserve greater attention from macroeconomists, not only for analyzing the economic effects of financial crises but in the study of ordinary business cycles as well. Second, I provide new evidence on the channels by which the recent financial crisis depressed economic activity in the United States. Although the deterioration of household balance sheets and the associated deleveraging likely exacerbated the initial economic downturn and the slowness of the recovery, I find that the unusual severity of the Great Recession was due primarily to the panic in funding and securitization markets, which disrupted the supply of credit. This finding helps to justify the government's extraordinary efforts to stem the panic in order to avoid greater damage to the real economy.

The horrific financial crisis of a decade ago, and the deep recession that followed it, exposed two distinct failures of forecasting by economists and economic policymakers. First, although many economists (Greenspan 2005; Rajan 2005; Shiller 2007) worried about low risk premiums, misaligned incentives for risk-taking, high house prices, and other

Conflict of Interest Disclosure: Ben S. Bernanke is a Distinguished Fellow in residence with the Economic Studies Program at the Brookings Institution, as well as a senior adviser to the Pacific Investment Management Company LLC and Citadel. The author did not receive financial support from any firm or person for this paper or from any firm or person with a financial or political interest in this paper. With the exception of the aforementioned, he is currently not an officer, director, or board member of any organization with an interest in this paper. No outside party had the right to review this paper before circulation.

excesses in the run-up to the crisis, the full nature and dimensions of the crisis—including its complex ramifications across markets, institutions, and countries—were not anticipated by the profession. Second, even as the severity of the financial crisis became evident, economists and policy-makers significantly underestimated its ultimate impact on the real economy, as measured by indicators like GDP growth, consumption, investment, and employment.

Do these failures imply that we need to remake economics, particularly macroeconomics, from the ground up, as has been suggested in some quarters? Of course, it is essential that we understand what went wrong. However, I think the failure to anticipate the crisis itself and the underestimation of the crisis's real effects have somewhat different implications for economics as a field. As I argued in a speech some years ago (Bernanke 2010), the occurrence of a massive, and largely unanticipated, financial crisis might best be understood as a failure of economic engineering and economic management, rather than of economic science. I meant by that that our fundamental understanding of financial panics—which, after all, have occurred periodically around the world for hundreds of years—was not significantly changed by recent events. (Indeed, the policy response to the crisis was importantly informed by the writings of 19th-century authors, notably Walter Bagehot.) Rather, we learned from the crisis that our financial regulatory system and private sector risk management techniques had not kept up with changes in our complex, opaque, and globally integrated financial markets; and, in particular, that we had not adequately identified or understood the risk that a classic financial panic could arise in a historically novel institutional setting. The unexpected collapse of a bridge should lead us to try to improve bridge design and inspection, rather than to rethink basic physics. By the same token, the response to our failure to predict or prevent the crisis should be to improve regulatory and risk management systems—economic engineering—rather than to seek to reconstruct economics at a deep level.

However, the second shortcoming, the failure to adequately anticipate the economic *consequences* of the crisis, seems to me to have somewhat different, and more fundamental, implications for macroeconomics. To be sure, historical and international experience strongly suggested that long and deep recessions often follow severe financial crises (Reinhart and Rogoff 2009). As a crisis-era policymaker, I was inclined by this evidence—as well as by my own academic research on the Great Depression (Bernanke 1983) and on the role of credit market frictions in macroeconomics (Bernanke and Gertler 1995)—toward the view that the crisis

posed serious risks to the broader economy. However, this general concern was not buttressed by much in the way of usable quantitative analyses. For example, as Donald Kohn and Brian Sack (2018) note in their recent study of crisis-era monetary policy, and as I discuss further below, Federal Reserve forecasts significantly underpredicted the rise in unemployment in 2009, even in scenarios designed to reflect extreme financial stress. This is not an indictment of the Fed staff, who well understood that they were in uncharted territory; indeed, almost all forecasters at the time made similar errors. Unlike the failure to anticipate the crisis, the underestimation of the impact of the crisis on the broader economy seems to me to implicate basic macroeconomics and requires some significant rethinking of standard models.

Motivated by this observation, the focus of this paper is the relationship between credit market disruptions and real economic outcomes. I have two somewhat related but ultimately distinct objectives. The first is to provide an overview of postcrisis research on the role of credit factors in economic behavior and economic analysis. There has indeed been an outpouring of such research. Much of the recent work has been at the microeconomic level, documenting the importance of credit and balance sheet factors for the decisions of households, firms, and financial institutions. The experience of the crisis has generated substantial impetus for this line of work, not just as motivation but also by providing what amounts to a natural experiment, allowing researchers to study the effects of a major credit shock on the behavior of economic agents. Moreover, as I discuss, the new empirical research at the microeconomic level has been complemented by innovative macro modeling, which has begun to provide the tools we need to assess the quantitative impact of disruptions to credit markets. Based on this brief review, I argue that the case for including credit factors in mainstream macroeconomic analysis has become quite strong, not only for understanding extreme episodes like the recent global crisis but possibly for the analysis and for forecasting of more ordinary fluctuations as well.

The second objective of the paper is to provide new evidence on the specific channels by which the recent crisis depressed economic activity in the United States. Why was the Great Recession so deep? (My focus here is on the severity of the initial downturn rather than the slowness of the recovery, although credit factors probably exacerbated the latter along with the former.) Broadly, various authors have suggested two channels of effect, each of which emphasizes a different aspect of credit market disruptions. David Aikman and others (2018) describe these two sources of

damage from the crisis as (1) fragilities in the financial system, including excessive risk-taking and reliance on “flighty” wholesale funding, which resulted in a financial panic and a credit crunch; and (2) a surge in household borrowing, of which the reversal, in combination with the collapse of housing prices, resulted in sharp deleveraging and depressed household spending.

In the former, “financial fragility” narrative, mortgage-related losses triggered a large-scale panic, including runs by wholesale funders and fire sales of credit-related assets, particularly securitized credit (Brunnermeier 2009; Bernanke 2012). The problems were particularly severe at broker-dealers and other nonbank credit providers, which had increased both their market shares and their leverage in the years leading up to the crisis. Like the classic financial panics of the 19th and early 20th centuries, the recent panic—in wholesale funding markets, rather than in retail bank deposits—resulted in a scramble for liquidity and a devastating credit crunch. In this narrative, the dominant problems were on the supply side of the credit market; and the implied policy imperative was to end the panic and stabilize the financial system as quickly as possible, to restore more normal credit provision.

The alternative, “household leverage” narrative focuses on the buildup of household debt, especially mortgage debt, during the housing boom of the early 2000s. This buildup reflected beliefs (on the part of both borrowers and lenders) that rapid increases in house prices would continue, which in turn promoted a loosening of credit standards, speculative home purchases (“flipping”), and the extraction of home equity through second mortgages. Given the large increase in leverage, the decline in house prices beginning in 2006 sharply reduced household wealth and put many homeowners into financial distress, leading to precipitate declines in consumer spending (Mian and Sufi 2010). Relative to the financial fragility narrative, this approach emphasizes the decline in the effective demand for credit, rather than the effective supply. From a policy perspective, this narrative does not deny the necessity of restoring calm in financial markets, but it places relatively greater importance on policies aimed at stabilizing housing markets, modifying troubled mortgages, and helping consumers (Mian and Sufi 2014a). To be sure, the two narratives are complementary, not mutually exclusive. For example, household leverage and mortgage delinquencies affected the financial health of lenders, increasing the risk of panic; while restrictions on the supply of credit lowered house prices and employment and ultimately affected household finances as well. But the two narratives do have somewhat different implications both for policy

and for macroeconomic analysis, so assessing their relative importance is worthwhile.

Some recent work has compared the macroeconomic effects of the two channels in the crisis, finding a significant role for each (Gertler and Gilchrist 2018; Aikman and others 2018). In the second part of the paper, I present some new evidence on this issue, comparing the real effects of the financial panic to those arising from deteriorating balance sheets, including household balance sheets. I proceed in two steps. First, I apply factor analysis to daily financial data to identify stages of the financial crisis, beginning with the loss of investor confidence in subprime mortgages, followed by the broad-based run on short-term funding, the panic in securitization markets, and the declining solvency of the banking system. Each of these stages involved disruptions to the operation of credit markets, and so should have had real consequences, as suggested by the research I review in the first portion of the paper. In the second step, I compare the ability of the estimated factors (which are orthogonal by construction) to forecast monthly macroeconomic indicators over the period 2006 through 2012. I find that the factors most strongly associated with the financial panic—the run on short-term funding and the panic in securitization markets—are also by far the best predictors of adverse economic changes in a range of macroeconomic indicators, and that ending the panic is likewise associated with relative economic improvement. The macroeconomic forecasting ability of factors associated with housing and mortgage quality is much more modest. As I discuss, these results do not rule out important effects through each of the identified channels, including channels linked to household balance sheets, but they do highlight the central role of the panic in setting off the Great Recession.

I draw several conclusions. For macroeconomists, recent experience and research highlight the need for greater attention to credit-related factors in modeling and forecasting the economy. Standard models used by central banks and other policymakers include basic financial prices—such as interest rates, stock prices, and exchange rates—but do not easily accommodate financial stresses of the sort seen in 2007–09, including the evident disruption of credit markets. Plausibly, this omission explains why standard approaches seriously underestimated the economic impact of the crisis. Moreover, if variations in the efficiency of credit markets were important determinants of economic performance during the Great Recession, they may deserve greater attention in the analysis of “garden-variety” business cycles as well.

For policymakers, a better understanding of why financial stresses are economically costly could help inform efforts to prevent and respond to crises. In particular, the policy response to the financial crisis of 2007–09 focused heavily on ending the financial panic and protecting the banking system, and it included some highly unpopular measures, including the bailouts of financial institutions with taxpayer funds. The rationale that policymakers gave for their apparent favoritism to the financial industry—despite its culpability in many of the problems that gave rise to the crisis in the first place—was that stabilizing Wall Street was necessary to prevent an even more devastating blow to Main Street. The results of this paper support this rationale. More generally, the results support reforms that improve the resilience of the financial system to future bouts of instability, and that increase the capacity of policymakers to respond effectively to panics, even if such reforms involve some costs in terms of credit extension or growth.

Although some of the empirical studies I discuss bear on the international transmission of the crisis, the focus of this paper is on the experience of the United States. Extending the analysis to other countries and considering aspects of the crisis more prominent outside the U.S., such as sovereign debt problems, are important directions for future research.

1. Credit Markets and the External Finance Premium

The first objective of this paper is to review recent research on the real effects of credit market disruptions and to discuss some implications for macroeconomics. As background, I begin with some simple theory. The key concept to be developed is the existence of an external finance premium (EFP), which may vary over time and depends on the financial health of both borrowers and lenders.

The starting point is the familiar observation that the process of credit extension is rife with problems of asymmetric information between borrowers and lenders. Potential lenders are only imperfectly informed about the characteristics of borrowers, including their skills and trustworthiness; nor can they easily observe borrowers' investment opportunities or effort levels. Asymmetric information in the borrower–lender relationship implies that the extension of credit involves costs above the cost of funding, including the costs of screening and monitoring by the lender and the dead-weight losses arising from adverse selection or principal–agent problems. Moreover, even a fully informed lender may face costs of transmitting and verifying its information about borrowers to third parties, forcing the

lender to bear liquidity risk and idiosyncratic return risk. These various costs contribute to the existence of a transaction-specific EFP, the difference between the all-in cost of borrowing and the return to safe, liquid assets like Treasury securities.

In much of economics (for example, in corporate finance), the assumption of asymmetric information and theoretical frameworks (principal–agent models, incomplete contracting) based on this assumption are central to the analysis of credit relationships. Mainstream macroeconomic analyses have paid less attention to these ideas. Certainly, to be relevant to macroeconomics, the EFPs associated with diverse transactions must have an aggregate or common component that is quantitatively significant, varies over time, and is linked to broad economic conditions. I use the term *credit factors* to refer to economic variables that affect the aggregate component of the EFP, in contrast to broader financial factors, such as the levels of equity prices and interest rates.

What affects the EFP? The EFP depends, *inter alia*, on the financial health (broadly defined) of both potential borrowers and financial intermediaries.

1.A. Borrowers

On the borrowers' side, the key intuition is that problems of asymmetric information are less severe when potential borrowers have skin in the game—that is, when they have sufficient net worth, equity, or collateral at risk to align their incentives with the goals of lenders and to reduce lenders' exposure to losses. For example, a large down payment by a homebuyer not only protects the lender from price declines; it also reduces the lender's need to investigate the borrower's income prospects in detail and incentivizes the borrower to maintain the home properly. Thus, a borrower who can make a substantial down payment can expect easier access to credit and terms that are more favorable. Likewise, an entrepreneur able to contribute substantial equity to his or her startup is more likely to obtain outside financing and will face fewer intrusions on her business decisionmaking by lenders.

In a macroeconomic setting, aggregate descriptors of the average financial health of borrowers (net worth, collateral, leverage) are state variables that, at least in principle, can affect the economy-wide component of the EFP and, consequently, macroeconomic dynamics. In the *financial accelerator* model of Bernanke and Mark Gertler (1989), endogenous deterioration of the net worth of borrowers in an economic downturn, and improvements in an upturn, make the aggregate EFP countercyclical. The endogenous variation in the EFP in turn increases the responsiveness of the economy

to exogenous shocks. Nobuhiro Kiyotaki and John Moore (1997) and John Geanakoplos (2010) describe related mechanisms.

1.B. Lenders

The EFP can also be affected by the financial health of lenders. Financial intermediaries (“banks”) are institutions that specialize in reducing the costs of making loans. Bank employees acquire both general lending skills and specific knowledge about particular industries, firms, communities, or individual borrowers. Complementarities in the provision of financial services—for example, a bank has more information about a potential borrower who also holds a checking account with the bank—further reduce the costs of lending. Banking organizations, by holding many illiquid loans, may also achieve greater diversification of lending risks.

Although banks serve to reduce the net cost of lending, banks are themselves borrowers as well, in that they must raise funds from the ultimate savers in order to make loans. Consequently, the financial health of banks also matters for the EFP. For example, if banks suffer loan losses in an economic downturn, the depletion of capital will reduce their ability to attract funding, on the margin. Weakened banks will become choosier in their lending, raising the aggregate EFP and reinforcing the financial accelerator mechanism. (Loss of bank capital will not deter government-insured depositors, but it may lead the deposit insurance agency, acting on behalf of at-risk taxpayers, to insist on tighter lending standards.) Michael Woodford (2010) discusses, in the context of a simple macro model, how reductions in bank capital and thus the effective supply of intermediary services can depress the economy. Similarly, because liquid assets facilitate lending and risk-taking, increased cost or reduced availability of funding (due to tighter monetary policy, for example) also reduces the supply of bank credit. This is a variant of the so-called bank-lending channel of monetary policy (see Drechsler, Savov, and Schnabl 2018).¹

1.C. Panics

The simple balance sheet perspective is also useful for understanding the real effects of financial panics—that is, systemwide runs on banks or

1. Early work on the bank lending channel includes that of Kashyap, Stein, and Wilcox (1993) and Van den Heuvel (2002). Gertler and Karadi (2011) interpret unconventional monetary policies, like quantitative easing, as a means by which the central bank can partially offset the decline in commercial banks’ lending capacity in a downturn.

other credit intermediaries. Generally, panics may arise in situations when longer-term, illiquid assets are financed by very short-term liabilities, for example, bank loans financed by demand deposits. A large body of literature has examined why such financing patterns persist and why panics sometimes erupt. In the classic work by Douglas Diamond and Philip Dybvig (1983), these arrangements allow society to marshal the necessary resources for long-term investment while simultaneously allowing individual savers to insure against unexpected needs for liquidity. The benefits of this setup must be weighed against the possibility of Pareto-inferior, self-fulfilling (“sunspot”) panics. In contrast, Charles Calomiris and Charles Kahn (1991) see short-term financing as a mechanism for lenders to use to discipline borrowers. In their framework, a run or panic is simply investors exercising their prerogative of withdrawing funding from borrowers in whom they have lost confidence.

An approach that seems particularly useful for understanding the recent financial crisis, and that fits nicely with the idea of a variable EFP, comes from Gary Gorton and coauthors (Gorton and Pennacchi 1990; Dang, Gorton, and Holmstrom 2015, 2018). In the Gorton setup, intermediaries meet a substantial part of their financing needs by issuing “information-insensitive” liabilities, that is, liabilities structured in a way that makes their value constant over almost all states of the world. Besides demand deposits, examples of information-insensitive liabilities in modern finance include short-term, overcollateralized loans (for example, many repo agreements), asset-backed commercial paper (ABCP), shares in low-risk money market mutual funds, and the most senior tranches of securities constructed from diverse underlying credits.

From the perspective of ultimate investors, the advantage of information-insensitive liabilities is that they can be held without incurring the costs of evaluating the individual credits that back these claims—a task at which most investors are at a comparative disadvantage—and without concern about principal–agent problems, adverse selection, and other costs that often arise in lender–borrower relationships. Moreover, information-insensitive liabilities will tend to be liquid, because potential buyers likewise do not have to incur high costs of evaluating them or worry about adverse selection among sellers. Consequently, investors who face unpredictable needs for liquidity (as in the Diamond–Dybvig setup) will benefit from holding such claims. Investor risk and transaction costs are reduced further when the information-insensitive liabilities have short maturities, because, rather than selling the assets when liquidity is needed, investors can simply stop rolling over their claims as they mature. From the issuer’s point of view,

the benefit of information-insensitive liabilities is their lower required yield and their attractiveness to broad classes of investors. Much of the financial innovation of the precrisis period reflected issuer efforts to create information-insensitive liabilities from risky underlying assets.²

Panics emerge in this setup when, as the result of unexpected events or news, investors begin to worry that the intermediary liabilities are not money-good, that is, those liabilities are no longer information-insensitive. Investors continuing to hold these claims face the unattractive alternatives of either making independent evaluations of the underlying credits—which they are not well equipped to do—or bearing the costs of uncertainty, illiquidity, and adverse selection. If the claims are contractually short term in nature, many investors will decide not to roll them over, resulting in a panic.

Panics raise the aggregate EFP because they can result in a violent disintermediation, which overturns the normally efficient division of labor in credit extension. In normal times, banks and other intermediaries make loans, manage existing credits, and hold most of the credit risk on their balance sheets. In a panic, intermediaries lose their funding, and as a result (assuming the funding cannot be replaced), they must dispose of existing loans and stop making new ones. The resulting fire sales of existing loans depress prices to the point where they can be voluntarily held by the subset of savers who are most able to evaluate and manage these assets, or who have the greatest tolerance for illiquidity (Shleifer and Vishny 2010). Because these asset holders are not specialists at making and monitoring loans, and because they are satiated with risky credits in the disintermediated equilibrium, the cost of new credit—the EFP—spikes during a panic (Gertler and Kiyotaki 2015). Increases in the EFP can help to explain the adverse macroeconomic effects of financial crises (Bernanke 1983; Reinhart and Rogoff 2009).³

2. Hanson and Sunderam (2013) provide a model of this process, arguing that, because of informational externalities, information-insensitive securities are overissued in good times. Caballero, Farhi, and Gourinchas (2017) discuss the global “shortage” of safe assets, which motivates financial engineers to create such assets. Sunderam (2015) discusses the creation of safe assets through shadow banking. Relatedly, Peek and Rosengren (2016) discuss the evolution of financial markets in recent decades, pointing out that many of the changes increased the dependence of the system on “runnable” wholesale funding.

3. A secondary effect of the sharp increases in risk aversion and liquidity preference is that normal relationships among asset prices break down as arbitrage capital declines (Krishnamurthy 2010).

Panic-type phenomena occurred in a variety of contexts in the recent financial crisis.⁴ The most intense pressures were felt in the so-called shadow banking system, which experienced runs on ABCP (Covitz, Liang, and Suarez 2009; Kacperczyk and Schnabl 2010; Schroth, Suarez, and Taylor 2014); structured investment vehicles and other conduits (Gorton 2008); securities lending (Keane 2013); and money market funds (McCabe 2010). Of particular concern were funding pressures in the critical market for repurchase agreements (repos), which are used heavily by broker-dealers and others to finance credit holdings. The repo market is dichotomized into two major components: triparty repo, intermediated by two large clearing banks; and the bilateral market, involving direct borrowing and lending among broker-dealers and other participants. The triparty market experienced less overt panic during the crisis, except, crucially, when borrowers like Bear Stearns and Lehman Brothers were close to the brink of failure (Copeland, Martin, and Walker 2010).⁵ The bilateral market, in contrast, appears to have suffered runs on multiple dimensions, including not only refusals to roll over loans but also a narrowing of the types of collateral accepted, increases in the amount of collateral required (haircuts), and reductions in the maturities of loans. Overall, the sharp contraction in funding in the shadow-banking sector forced a painful disintermediation, which in turn depressed prices and raised yields on virtually all forms of private credit, not just troubled mortgages (Longstaff 2010; Scott 2016).

Although the most severe disintermediation occurred at broker-dealers and other shadow banks, commercial banks also faced pressures, including from uninsured depositors (Rose 2015), in wholesale funding and interbank loan markets (Afonso, Kovner, and Schoar 2011), and from borrowers taking down precommitted credit lines in order to hoard liquidity (Ivashina and Scharfstein 2009). Banks were also (explicit or implicit) backstop liquidity providers for structured investment vehicles, ABCP programs, and other conduits, and were consequently forced to replace much of

4. Bao, David, and Han (2015) provide comprehensive time series of “runnable” liabilities. They calculate that, during the financial crisis, runnable liabilities fell from about 80 percent of nominal GDP to about 60 percent.

5. Concerns also arose in the triparty market that the intermediating banks would refuse to accept the credit risk during the daily period when repo funding is rolled over. The failure of one or both of the banks to accept this exposure would have been equivalent to a massive run on repo borrowers.


their funding as it ran out (Arteta and others 2013). Viral Acharya and Nada Mora (2015) find that liquidity was a significant issue for banks from the beginning of the crisis until after the collapse of Lehman Brothers, when government capital became available. However, commercial banks generally had more stable funding sources than broker-dealers—including insured deposits, advances from Federal Home Loan Banks (Gissler and Narajabad 2017, part 1), and access to the Fed’s discount window. Consequently, as the crisis wore on, banks were able to take advantage of fire sale prices to increase holdings of some forms of credit (He, Khang, and Krishnamurthy 2010).

1.D. Measures of the EFP

The simple analysis thus far makes two basic predictions about the aggregate EFP: that it should be countercyclical, rising in downturns when the balance sheets of lenders and borrowers deteriorate; and that it should rise sharply during periods of financial instability. To evaluate these predictions, we need measures of the EFP. Of course, although in macro modeling we may speak of “the” EFP (as we often speak of “the” interest rate), in practice the EFP is heterogeneous, depending not only on the balance sheets of individual prospective borrowers and lenders but also on borrower type (household versus firm) and other characteristics that bear on the costs of lending, like firm size.

With these caveats in mind, figure 1 shows two related measures of borrowing costs for nonfinancial corporations developed by Simon Gilchrist and Egon Zakrajšek (2012a), following earlier work by Andrew Levin, Fabio Natalucci, and Zakrajšek (2004). The series in figure 1 labeled *GZ spread* is essentially the difference between the yield on nonfinancial corporate bonds and comparable-maturity Treasury obligations, constructed from data on individual issues to match durations and to adjust for call options and other features. The second series, labeled *EBP* for the excess bond premium, subtracts from the *GZ* credit spread a measure of issue-specific default probabilities, based on the “distance to default” methodology of Robert Merton (1974). Gilchrist and Zakrajšek (2012a) interpret the EBP as a measure of investor appetite for corporate debt, holding constant estimated default risk. They find that both measures are highly predictive of real economic activity but that, interestingly, the bulk of the predictive power lies in the excess bond premium rather than in the default probability. We will use the EBP in later analysis. For now, I note that both indicators are generally countercyclical (shaded bars in the figure show the National Bureau of Economic Research’s recession dates), and both spike during the 2008

Figure 1. Two Measures of the External Finance Premium for Nonfinancial Corporations, 1973–2017^a



Sources: Gilchrist and Zakrajšek (2012a); updated data from Favara and others (2016).

a. Shaded bars indicate the National Bureau of Economic Research's recession dates.

crisis, consistent with the theory. The cyclical nature of these measures also appears to have increased over time, consistent with the general perception that financial factors have played a larger role in business cycles since the 1980s.

The Gilchrist-Zakrajšek measures, derived from observed yields, reflect the “price” of credit for certain classes of borrowers. Students of credit markets have long noted that, consistent with the complex agency and monitoring problems that affect lender–borrower relationships, loans often involve many nonprice elements, including limits on loan size, covenants, call provisions, and so on. In principle, the shadow value of nonprice terms should be included in the EFP. Studies suggest that these nonprice terms move in the same way as more directly observable spreads, and, moreover, that nonprice terms have predictive power for economic activity. For example, using bank-level responses to the Federal Reserve’s Loan Officer Opinion Survey, William Bassett and others (2014) constructed an indicator of changes in lending standards, adjusted for factors affecting loan demand, and found that their indicator forecasts lending and output. Carlo Altavilla, Matthieu Darracq Paries, and Giulio Nicoletti (2015) found similar results for the euro area.

1.E. Credit Factors in Precrisis Mainstream Macroeconomics

Before the financial crisis, mainstream macro models (including models used by central banks for forecasting and policy analysis) did not include much role for credit factors, of the type described in the previous section. Notably, the FRB/US model of the U.S. economy, the Fed's workhorse model, provided little guidance to the staff on how to think about the likely economic effects of the crisis, despite having (relative to the models most used in academic work) an extensive financial sector. The staff supplemented FRB/US with various ad hoc adjustments, based on historical case studies, anecdotes, and judgment. However, the staff and the Federal Open Market Committee (FOMC) still systematically underpredicted the economic impact of the crisis, as mentioned above.

For example, as noted by Kohn and Sack (2018), in August 2008, a year into the crisis, the Fed staff predicted (in the FOMC briefing document known as the Greenbook) that unemployment would peak at under 6 percent. In reality, the unemployment rate would rise to nearly 10 percent. This underprediction partly reflected excessive optimism about the evolution of financial conditions. However, an alternative Greenbook forecast scenario that hypothesized "severe financial stress," and that assumed in particular that house prices would fall further than they ultimately did, saw unemployment remaining below 7 percent. Moreover, even in October 2008, well after the collapse of Lehman Brothers and the rescue of AIG, the staff saw unemployment peaking at about 7.25 percent.⁶

What accounts for this important blind spot—which, I emphasize again, was shared by all major forecasters? Although the basic theoretical framework outlined above existed before the crisis, in the view of many economists the benefits of incorporating credit factors into macro models did not exceed the costs. Most macroeconomic modeling focused on explaining the behavior of the postwar U.S. economy, a period that until 2007 had been without a major financial crisis.⁷ From a modeling perspective, adding credit factors required allowing heterogeneity among agents (including savers, borrowers, and intermediaries), which added technical complexity.

6. Kohn and Sack (2018) also report an exercise, conducted by Bob Tetlow of the Federal Reserve Board, which calculates what the forecast of the FRB/US model would have been if the staff had had perfect foresight about the financial variables included in the model. Even with this information, according to this exercise, FRB/US would have significantly underpredicted the magnitude and speed of the rise in the unemployment rate.

7. Del Negro, Hasegawa, and Schorfheide (2016) show formally that a dynamic stochastic general equilibrium (DSGE) model that incorporates financial frictions produces better forecasts in periods of financial distress but underperforms in samples without such periods.

Arguments from parsimony and computational simplicity thus worked against the addition of credit factors to the standard model.

Deficiencies in the received credit literature also played a role. The financial accelerator literature, which incorporated credit factors into otherwise standard macro models, showed that such factors could improve the fit of models to data (Bernanke, Gertler, and Gilchrist 1999). However, this literature, like other new Keynesian modeling of the time, focused on the dynamics of normal business cycles rather than on financial crises and their effects.

Another barrier to the incorporation of credit factors was that the use of microeconomic data to measure credit effects, an essential element in building quantitative macro models, was bedeviled by identification problems. Credit-focused theories posit relationships between measures of financial health—like net worth, leverage, or collateral values—and aspects of economic behavior, such as borrowing, consuming, or investing. However, measures of financial health are generally themselves endogenous, complicating identification. For example, theory suggests that, all else being equal, a firm with more internal funds available should face a lower EFP and thus be willing to invest more. In practice, however, a finding that internal cash flow and investment are correlated across firms (Fazzari, Hubbard, and Petersen 1988) is subject to the potential critique that causality may flow in both directions. In particular, although higher cash flows may promote investment, it is likely also true that firms endowed with better investment opportunities will tend to enjoy higher profits and stronger cash flows, even if no credit market frictions are present.

However, the recent crisis has significantly changed economists' views on the importance of credit factors. The Great Recession was the worst downturn since the Great Depression of the 1930s, and its severity seems impossible to explain except as the result of credit market dysfunction, broadly construed (Stock and Watson 2012). Explanation of recent events thus requires incorporation of credit factors into otherwise standard models, and there has been much activity in this area. Studies at the microeconomic level have also proliferated, as economists have tried to better understand the links between credit factors and aspects of household, firm, and bank behavior. An interesting side effect of the crisis is that it helped solve the perennial identification problem, by creating what is in effect a natural experiment. Because the crisis was plausibly an exogenous event for most economic units, differences in behavior that correlate with initial financial health provide better-identified estimates of the effects of credit market shocks.

In the next section, I briefly review this postcrisis literature. Collectively, the research provides substantial support for the view that factors affecting the costs of credit extension have an important independent influence on credit flows and, crucially, on the economic choices of households and businesses as well.

II. Recent Research on Credit Factors and Real Economic Activity


This section first reviews new microeconomic evidence on the role of credit factors, then turns to postcrisis research in macroeconomic modeling that includes such factors.

II.A. Microeconomic Evidence: Households

The run-up to the crisis showed a significant expansion in household debt, especially mortgage debt. As aspiring homeowners pressed to get into the hot housing market, weakening lending standards gave more households access to mortgages, and existing homeowners borrowed against built-up home equity. Figure 2 shows the ratio of mortgage debt service to income and the Fannie Mae single-family mortgage delinquency rate for the period 2002–12. Evident in the figure is both the buildup in debt service burdens before the crisis and the financial stress placed on households by the reversal of the housing boom in 2006 and thereafter.

In a frictionless world, with no credit constraints, declining house prices would have only small effects on consumer spending, because households would be able to borrow and save as needed to smooth over time the effects of wealth changes. Moreover, the negative impact of a house price decline on wealth should, in principle, be largely offset by a corresponding decline in the user cost associated with living in the house. In short, with no credit constraints, the marginal propensity to consume (MPC) out of housing wealth should be small.

However, when households face an EFP that in turn depends on the states of their balance sheets, declines in housing wealth can have much larger effects on spending, for two related reasons. First, declining housing wealth depletes the pool of net worth that the household could draw upon to smooth spending if needed; and, second, declines in net worth and the collateral value of the home raise the effective cost of credit (the EFP) for the homeowner. Note that the effects of rising and falling house prices on consumption may be asymmetric. Starting from a level of home equity at

Figure 2. Household Debt Service and Delinquencies, 2002–12^a

Sources: Haver Analytics; Fannie Mae; Federal Reserve Board, Z.1 Financial Accounts of the United States.

a. Mortgage debt service is measured relative to disposable personal income. The delinquency rate refers to the share of conventional single-family home mortgages that are 90+ days past due or in foreclosure.


which credit constraints do not bind very tightly, the MPC out of additional housing wealth is likely to be small, while declines in housing wealth that cause the constraints to bind can reduce consumption significantly. This asymmetry helps explain why the positive effects of the housing boom on consumption appear to have been outweighed by the negative effects of the housing bust (Guerrieri and Iacoviello 2017).

The period since the crisis has seen a great deal of new research on the links between household balance sheets and household spending. Atif Mian and Amir Sufi, with their coauthors, have been especially prolific on this topic. For example, using county-level and zip-code-level data, Mian, Kamalesh Rao, and Sufi (2013) confirmed the basic predictions of the theory that MPCs out of housing wealth are much higher than can be explained in standard life cycle frameworks, and that these MPCs are relatively higher for poorer, more-leveraged households. Consistent with a link between home equity and credit access, they also found that areas with larger declines in house prices saw, on average, relatively larger deteriorations in credit scores and credit limits, along with greater declines in the likelihood of mortgage refinancing.

Mian and Sufi have emphasized the role of weakening household balance sheets in triggering the Great Recession. For example, they showed that, in counties where housing booms were accompanied by large increases in household leverage from 2002 to 2006, durables consumption declined relatively more sharply beginning in the second half of 2006 (Mian and Sufi 2010). Similarly, Mian and Sufi (2014b) found that, in a cross section of U.S. counties, deterioration in household balance sheets was an important correlate of declining employment in the recession period 2007–9. Much of this work treats the housing boom and bust as given, focusing on the economic consequences. However, in their most recent research, Mian and Sufi (2018a) also explore the credit market sources of the boom, finding that zip codes that were most exposed to the 2003 acceleration of the private-label mortgage securitization market saw a sudden subsequent increase in mortgage originations and house prices, followed by sharp housing price collapses.

Other researchers have also explored the links between households' balance sheets and their spending decisions. Notably, while Mian and Sufi have mostly used data aggregated over geographic units, a study by Scott Baker (2018) employed data on millions of individual households, matched with employers. He considered household income changes associated with shocks to their employers, which are therefore arguably exogenous to the households. He found that the consumption of highly indebted households is meaningfully more sensitive to income, and that these differences are almost entirely driven by borrowing and liquidity constraints. He estimated that consumption in the 2007–9 recession dropped by 20 percent more than it would have if household balance sheets' positions had been comparable to those in the 1980s. Also consistent with the Mian-Sufi findings, Aditya Aladangady (2014) reported that homeowners with high debt service ratios have significantly higher MPCs out of housing wealth. Greg Kaplan, Kurt Mitman, and Giovanni Violante (2016) also found a high MPC out of housing wealth, although—in contrast to Mian and Sufi and other authors—they did not find an independent role for leverage. Claudia Sahm, Matthew Shapiro, and Joel Slemrod (2015) found that the condition of a household's balance sheet was a key determinant of its spending and saving behavior in response to a change in fiscal policy.

As has been known for some time, household balance sheets influence entrepreneurial activity, as many small business startups are financed from personal resources, including borrowing against home equity. Consistent with this “collateral channel,” Manuel Adelino, Antoinette Schoar, and

Figure 3. Corporate Debt Service and Delinquency, 2002–12^a

Sources: Haver Analytics; Call Report, Bank for International Settlements.

a. The debt service of nonfinancial corporations is measured relative to pretax profits. The delinquency rate is the share of commercial and industrial loans at commercial banks that are 30+ days past due.

Felipe Severino (2015) found that, in the period leading up to the crisis, small business starts and small firm employment growth were highest in areas with rising house prices and leverage. They did not find the same relative increase in employment in large firms, which presumably do not rely on household collateral for financing.

II.B. Microeconomic Evidence: Nonfinancial Firms

The balance sheets of nonfinancial firms did not deteriorate as dramatically as those of households in the periods before and during the recession, but nonfinancial firms certainly did experience increased stress. Figure 3 shows corporate debt service and delinquencies during the period around the crisis. Corporate balance sheets improved in the period after the 2001 recession. However, starting in about 2006, nonfinancial corporate debt service began to rise, to be followed by a spike in delinquencies in commercial and industrial loans after the recession began.

Similar to studies of households, cross-sectional studies of nonfinancial firms during the crisis era have provided an opportunity to observe how differing balance sheet conditions affected the responses of those firms to the downturn. Analogous to the responses of households to changes

in wealth or income, firms with initially weaker balance sheets (higher leverage, less internal cash, less usable collateral) would be expected to react more sensitively—for example, in terms of hiring and investment—to changes in revenue or demand. Likewise, smaller or younger firms, which typically require more lender screening and monitoring per dollar of lending, should be more sensitive to deteriorating financial conditions.

Postcrisis research has generally confirmed these predictions. For example, Xavier Giroud and Holger Mueller (2017) found that, during the Great Recession, highly leveraged firms cut employment significantly more than other firms did, in response to a given decline in local consumer demand. They concluded that firms' balance sheets were an essential part of the link between final demand and employment. Similarly, Ran Duchin, Oguzhan Ozbas, and Berk Sensoy (2010) found that the crisis affected investment the most in companies with low cash reserves or high net short-term debt. In a novel application of the theory, Gilchrist and others (2017) considered the effects of firms' balance sheets on their pricing behavior, finding that firms with limited internal liquidity and high operating leverage raised rather than lowered their prices in the face of the 2008 contraction. Interpreting price cuts as investments in maintaining customer relationships, the paper found that financially stressed firms were relatively less able to make such investments.

An interesting aspect of the recent literature on nonfinancial firms is the variety of identification strategies that researchers have applied. For example—following precrisis work by Giovanni Dell'Ariccia, Enrica Detragiache, and Raghuram Rajan (2005)—quite a few studies have compared firms in industries that are normally more dependent on external finance with firms in industries that are normally more self-sufficient for credit. Studies that use this approach (among others) find that firms in industries more dependent on external finance also reacted more sharply to the crisis include, among others, the aforementioned Duchin, Ozbas, and Sensoy (2010); Luc Laeven and Fabian Valencia (2013); and Samuel Haltenhof, Seung Jung Lee, and Viktors Stebunovs (2014). In another approach to identification, Thomas Chaney, David Sraer, and David Thesmar (2012) used local variations in real estate prices as a proxy for the change in the value of collateral of firms owning real estate, finding a strong association of new capital investment at the firm level with changes in collateral values. Following yet another identification strategy, in a sample of firms with long-term debt, Heitor Almeida and others (2009) found that firms with large portions of long-term debt maturing right at the time of the crisis reduced investment by considerably more than


otherwise similar firms whose debt was not scheduled to mature. However, in a contrarian study, Kathleen Kahle and René Stulze (2013) found that firms relatively more dependent on bank-provided credit did not decrease capital expenditures more than otherwise similar firms in the early stages of the crisis.

Researchers studying firm behavior have also made use of survey data. For example, based on a survey of 1,050 chief financial officers around the world, Murillo Campello, John Graham, and Campbell Harvey (2010) reported that firms describing themselves as credit-constrained during the crisis planned relatively deeper cuts in employment and capital spending, including bypassing otherwise attractive opportunities and canceling or postponing planned investments.

Small firms are likely to be more sensitive to reductions in credit supply, and the research confirms that this sector was hit hard during the crisis. For example, using firm-level data, Michael Siemer (2014) found that, during the 2007–9 recession, financial constraints substantially reduced employment in small firms relative to large ones, controlling for aggregate demand and other factors. Other studies documenting the impact of restricted credit on the entry, growth, and survival of smaller firms include Traci Mach and John Wolken (2012); Arthur Kennickell, Myron Kwast, and Jonathan Pogach (2015); and Burcu Duygan-Bump, Alexey Lekov, and Judit Montoriol-Garriga (2015). Brian Chen, Samuel Hanson, and Jeremy Stein (2017) found that the largest U.S. banks pulled back sharply and differentially from small business lending in 2008–10, as they grappled with the stresses of the crisis.

II.C. Microeconomic Evidence: Banks and Nonbank Lenders

As discussed above, the theory suggests that the balance sheets of financial intermediaries should also affect the EFP and the flow of credit. The postcrisis research generally confirms this prediction, finding in particular that cross-sectional differences among lenders in initial capital, funding sources, and exposure to mortgage-related losses affected their willingness or ability to make nonmortgage loans. Although some borrowers were able to shift to other sources of credit, including trade credit, the available evidence suggests that many could not, or had to pay much higher rates. Consequently, shocks to the financial health of lenders had consequences for the real economy, including for consumption, investment, and employment. Figure 4 shows capital and nonperforming loans at U.S. commercial banks in the period around the crisis. Despite capital raises, the ratio of bank Tier 1 common equity capital to loans dropped

Figure 4. Capital and Nonperforming Loans at Commercial Banks, 2002–12^a

Sources: Haver Analytics; Federal Reserve Bank of New York.

a. Nonperforming loans are defined as the share of loans with payments that are 90+ days past due.

precipitously in 2007 and 2008 as delinquencies rose. Gertler and Gilchrist (2018, fig. 3) document the rapid deleveraging of investment banks during the crisis.

Once again, for many studies, the shock of the crisis provided a natural experiment that helped to sharpen identification. For example, for a variety of reasons, banks differed in their exposures to mortgage losses arising from the housing and subprime busts. Absent balance sheet effects, there is no evident reason that these differential exposures should have affected the willingness of individual banks to make *nonmortgage* loans. However, many studies have found that there was a linkage between mortgage exposures and nonmortgage lending, presumably because mortgage-related losses depleted bank capital. For example, controlling for firm-specific factors, João Santos (2011) found that firms borrowing from banks that suffered larger subprime losses paid higher spreads and received smaller loans than those borrowing from other banks. Lu Zhang, Arzu Uluc, and Dirk Bezemer (2017) obtained similar results for the United Kingdom, finding that British banks that were more exposed to residential mortgages before the crisis reduced their nonmortgage lending by relatively more during and after the crisis. Jose Berrospide, Lamont Black, and William Keaton (2016) found that, all else being equal, banks serving a number

of metropolitan areas reduced their local mortgage lending in response to mortgage losses in other markets.

Earlier in this paper, I cited evidence that the effects of balance sheet conditions on household spending are not symmetric, with balance sheet deterioration having a larger effect than improvements. Analogous effects appear to occur for banks. For example, Mark Carlson, Hui Shan, and Missaka Warusawitharana (2013), using matched samples of banks and controlling for a variety of factors, found that the effect of changes in bank capital on lending is nonlinear—modest when capital is at high levels, but large when capital is low, as predicted by the theory.

Researchers have linked banks' willingness to lend to their sources of liquidity, as well as to their levels of capital. Notably, quite a few studies report that banks able to fund through retail deposits, rather than wholesale funding, cut their lending by relatively less during the crisis (Ivashina and Scharfstein 2009; Cornett and others 2011; Dagher and Kazimov 2015; Irani and Meisenzahl 2014).

Changes in loan supply by individual banks would not matter much if borrowers could easily compensate, for example, by switching to other lenders or other sources of credit, such as trade credit. As noted, however, this does not seem to have been the case in most instances. In a nice study, Gabriel Chodorow-Reich (2014) used the dispersion in lender health following the Lehman Brothers crisis as a source of exogenous variation in credit availability to borrowers. Using data on 2,000 nonfinancial firms with precrisis banking relationships, he found that firms with weaker lenders borrowed less, paid higher rates when they borrowed, and reduced employment more than other firms. The strongest employment effects were at small and medium-sized firms. Other studies making the explicit linkages among bank health, credit extension, and real economic activity include those by Martin Goetz and Juan Gozzi (2010); Antonio Falato and Nellie Liang (2016); John Kandrak (2014); and Laura Alfaro, Manuel Garcia-Santana, and Enrique Moral-Benito (2018). Tobias Adrian, Paolo Colla, and Hyun Song Shin (2012) found that some large nonfinancial firms were able to make up part of the reduction in bank lending through bond issuance, but only by paying high rates. Those authors argue that the impact of the credit crisis on real activity came through the associated spike in risk premiums rather than a contraction in the total quantity of credit. However, that finding is consistent with an approach centered on the EFP, which, as figure 1 suggests, rose sharply during the crisis.

In the United States, nonbank lenders are important credit providers, and many nonbanks were severely affected by the crisis. A number of

interesting studies have identified links between nonbank lending and economic activity. For example, using a data set linking every U.S. car sale to an associated supplier of auto credit, Efraim Benmelech, Ralf Meisenzahl, and Rodney Ramcharan (2017) drew an empirical connection between the collapse of the ABCP market and automobile sales. The collapse of the ABCP market hit the financing capacity of nonbank auto lenders, like captive leasing companies, particularly hard. These authors found that counties in which nonbank lenders had traditionally been dominant suffered deeper declines in car sales than other counties. In another interesting analysis, Ramcharan, Skander van den Heuvel, and Stephane Verani (2016) used the unique tiered structure of national credit unions to identify credit supply effects. Losses in the asset-backed securities (ABS) market at top-tier institutions imposed costs on local credit unions, in ways plausibly uncorrelated with local market conditions. However, these authors found that credit unions suffering such losses contracted their extensions of consumer credit to local customers by more than credit unions without such losses.

II.D. Microeconomic Evidence: Cross-Border Banking

Cross-border effects, whereby financial stresses in one country affect credit supply and economic activity in another, are a potentially important channel of international transmission of crises. Documenting such effects also provides another tool for identifying the links between bank balance sheets, lending, and economic outcomes.

Joe Peek and Eric Rosengren (2000), in a classic paper, were among the first to use cross-border linkages to identify balance sheet effects. They used the facts that (1) Japanese banks were active lenders in the United States during the 1990s and that (2) the Japanese banking crisis of that decade could reasonably be viewed as exogenous to economic developments in the U.S. to construct a natural experiment. Using the variation in the lending shares of Japanese banks across various U.S. commercial real estate markets, they showed that loan supply shocks emanating from Japan had real effects on economic activity in the United States.

In a similar vein, for the recent crisis, the evidence suggests that banks experiencing losses abroad, or that were dependent on foreign sources of funding that came under pressure, reduced their domestic lending by more than other banks. For example, Manju Puri, Jörg Rocholl, and Sascha Steffen (2011) examined the domestic retail lending of German savings banks during the years 2006–8, comparing savings banks with substantial indirect exposures to U.S. subprime mortgages with savings banks without such exposures. They found that the exposed banks rejected substantially

more loan applications than banks not so affected. Also for Germany, Kilian Huber (2018) studied the effects of domestic lending cuts by Commerzbank, a large bank that suffered significant losses in its international trading book. He found that cuts to Commerzbank's lending in Germany were not offset by other sources of credit. Rather, they resulted in persistent adverse effects on output, employment, and productivity in firms and regions where the bank had a relatively larger market share before the crisis.

Studies with analogous findings exist for many other countries, including the United Kingdom (Aiyar 2011, 2012); Italy (Albertazzi and Marchetti 2010); Portugal (Iyer and others 2014); and Denmark (Jensen and Johannesen 2017). In a multicountry study, Ralph De Haas and Neeltje Van Horen (2012) analyzed cross-border syndicated lending by 75 banks to 59 countries after the collapse of Lehman Brothers, finding that banks that had to write down subprime assets or refinance large amounts of long-term debt reacted by curtailing their lending abroad. Not all cross-border studies look at the effects of events in the United States on foreign economies: For example, Ricardo Correa, Horacio Saprizo, and Andrei Zlate (2013) found that the European sovereign debt crisis affected the United States, as U.S. branches of euro area banks, hit by liquidity strains, reduced lending to U.S. firms by more than did the U.S. branches of foreign banks headquartered outside Europe. Shin (2011) emphasizes the role of global banks in transmitting changes in financial conditions internationally.

II.E. The Great Depression

Interestingly, the recent crisis appears also to have inspired new research on another worldwide financial and banking crisis, the Great Depression of the 1930s. My research on the Depression discussed the real effects of the deterioration of both bank and borrower balance sheets (Bernanke 1983). I also drew on international comparisons for evidence (Bernanke and James 1991; Bernanke 1994). However, my empirical work on the period relied heavily on aggregate time series, making it subject to the usual concerns about endogeneity and identification. Remarkably, recent research has developed new microeconomic, cross-sectional databases for the 1930s, allowing for something closer to the natural experiment approach.

For example, using newly collected data on large industrial firms, Benmelech, Carola Frydman, and Dimitris Papanikolaou (2017) exploited preexisting variation in the need to raise external funds at a time when bond markets were frozen and banks were failing. They found a large, negative effect of financing frictions on employment at large firms. Building on earlier work by Calomiris and Joseph Mason (2003), who

found that bank distress in the 1930s reduced loan supply and economic activity in the regions where the banks operated, Kris James Mitchener and Gary Richardson (2016) examined the effects of correspondent relationships that played an important role in interwar banking. They found that a bank's financial distress reduced credit available not only to the bank's own customers but also to the customers of their (regionally dispersed) correspondents, who had to accommodate sharp increases in the demand for liquidity. Other, related papers using cross-sectional data to study the effects of bank distress during the Depression include those by Carlson and Jonathan Rose (2015), Ramcharan and Rajan (2014), and Jon Cohen, Kinda Cheryl Hachem, and Richardson (2017). In general, this literature supports the view that disruptions in banking and credit markets help to explain the depth, duration, and international incidence of the Depression.

II.F. Credit Factors in Quantitative Macroeconomic Models

Microeconomic studies provide evidence that household, firm, and bank behavior are affected by balance sheet conditions and asymmetric information about creditworthiness. However, such studies are inherently partial equilibrium in nature. It is possible that balance sheet effects, though important in the cross section, “wash out” in aggregate time series (Jones, Midrigan, and Philippon 2018). For example, it could be that, for the economy as a whole, reduced investment or hiring by financially constrained firms is offset by greater activity at less-constrained firms. Assessing the importance of credit factors for macroeconomic outcomes inevitably requires the incorporation of such factors into quantitative, general equilibrium models of the economy.

As noted above, before the crisis, a modest body of literature incorporated credit factors into otherwise standard models, generally finding that doing so could improve the fit of the models to the data (Carlstrom and Fuerst 1998; Bernanke, Gertler, and Gilchrist 1999). However, these papers did not argue that credit factors were a dominant source of variation in output and employment. More important, the earlier models did not capture the phenomenon of the occasional large, discontinuous crisis, or other nonlinear effects.

Work since the crisis has made substantial progress in accommodating credit factors in dynamic macro models. This research supports two separate, though related, substantive conclusions. The first of these is that credit factors are essential for understanding the Great Recession specifically. In the words of Lawrence Christiano, Martin Eichenbaum, and Mathias Trabandt (2014, 110), “The vast bulk of movements in aggregate

real economic activity during the Great Recession were due to [in their terminology] financial frictions interacting with the zero lower bound [on short-term interest rates].” Many other papers have reported similar conclusions. The finding that the Great Recession was in large part the result of financial and credit market dysfunction is of course not really a surprise at this point; but it is nevertheless important to confirm that quantitatively realistic economic effects of credit shocks can be rationalized in what are otherwise largely standard models.

This observation, together with the conclusion of James Stock and Mark Watson (2012) that the Great Recession differed from other postwar business cycles in magnitude but not in kind, leads to the second conclusion: that credit factors may play a more important role than previously thought even in “garden-variety” business cycles. Complementary, model-based analyses finding central roles for credit shocks in both the Great Recession and in business cycles generally include (in a very partial listing) those by Charles Nolan and Christoph Thoenissen (2009); Robert Hall (2010, 2011); Urban Jermann and Vincenzo Quadrini (2012); Gilchrist and Zakrajšek (2012b); Matteo Iacoviello (2014); and Marco Del Negro and others (2017). In related research, Mian and Sufi (2018b) have recently argued that periodic, excessive expansions in the supply of credit to households are a major source of business cycles globally, not just the U.S. Great Recession. Cristina Arellano, Yan Bai, and Patrick Kehoe (2016) show that credit market frictions can help models match cross-sectional aspects of the macro data (such as the dispersion of investment and hiring across firms) as well as time-series aspects. In a stylized macro model, Gauti Eggertsson and Paul Krugman (2012) discuss the interaction of household leverage and the zero lower bound on interest rates. Philippe Bacchetta and Eric van Wincoop (2016) use a two-country model to study the transmission of the panic between economies.

The paper by Bernanke, Gertler, and Gilchrist (1999), and other papers of that genre, studied log-linear approximations around steady states, which facilitated the analysis of credit factors in normal cyclical dynamics but ruled out large, discontinuous shifts in economic activity. As discussed earlier in this paper, financial panics are inherently discontinuous (for example, the economy shifts from one equilibrium to a quite different one), and the empirical work to be presented later in this paper will rely on these discontinuities for identification. Recent modeling has shown how to reproduce this important feature of the data. Notably, Gertler and Kiyotaki (2015) and Gertler, Kiyotaki, and Andrea Prestipino (2017) incorporate banking panics in quantitative macro models, finding that panics can produce severe,

highly nonlinear contractions in economic activity. The mechanism of this effect, as discussed above, is the sharp disintermediation and rise in the EFP associated with a panic. Markus Brunnermeier and Yuliy Sannikov (2014) analyze a theoretical model in which financial frictions create highly nonlinear contractions in economic activity and lead to occasional crisis episodes. Nonlinear outcomes also emerge from the models of Zhiguo He and Arvind Krishnamurthy (2013) and Frédéric Boissay, Fabrice Collard, and Frank Smets (2016). Recent work has also made progress in modeling housing booms and busts in a general equilibrium context (see, for example, Favilukis, Ludvigson, and Van Nieuwerburgh 2010).


In sum, there has been substantial recent progress in the development of quantitative macro models incorporating credit factors, including the potentially large and nonlinear effects of financial crises. This literature represents an important step toward remedying the weaknesses of empirical modeling and forecasting that became evident during the crisis.

III. New Evidence on the Effects of the Financial Crisis on the Real Economy

Research since the financial crisis suggests that credit factors matter. However, credit was disrupted in a number of ways during the crisis, including through the two broad mechanisms described in the introduction: (1) the loss of investor confidence in financial institutions and securitized credit, which triggered a financial panic that choked off credit supply; and (2) the weakening of household balance sheets, which resulted in deleveraging and the constriction of household spending. This section provides new evidence on the links between the financial crisis and the Great Recession and, in particular, on the relative importance of these two channels. The empirical strategy is to use financial data to identify points of discontinuity in the evolution of the crisis, and then to evaluate the extent to which these shifts predict movements in a standard set of macroeconomic variables.

The analysis here is loosely motivated by figures presented by Gorton and Andrew Metrick (2012); see especially their figures 8 and 9. Similar to their figures, this paper's figure 5 uses four representative (daily) financial data series to illustrate informally the principal stages of the crisis. The four series shown in figure 5 are:

—ABX BBB spread (2006:Q1) is a market-traded index of the value of BBB-rated, 2006-vintage subprime mortgage-backed securities. It is a proxy for investor views of housing and mortgage markets.

Figure 5. Stages of the Financial Crisis, 2006–12

Sources: Bloomberg; IHS Markit; Haver Analytics.

—LIBOR–OIS spread is the interest rate on one-month London Inter-bank Offered Rate loans (LIBOR) less an indicator of expected safe rates (overnight indexed swaps, or OIS). This variable is an indicator of stress in the interbank lending market and, more generally, in wholesale funding.

—The spread on ABS backed by credit card receivables (Bloomberg/Barclays index) shows the yield (relative to Treasuries) on securities backed by an important class of nonmortgage credit. This spread measures investors' willingness to hold nonmortgage credit, especially in the form of securitizations.

—The credit default swap (CDS) spread of a large bank (Bank of America) reflects the perceived risk of default on that bank's bonds, and is thus a measure of the banking system's solvency.

By means of these four representative financial variables, figure 5 illustrates the stages of the financial crisis. Stage 1, captured here by the ABX index of subprime mortgage values, is the deflation of the housing bubble and the growing concerns about the mortgage market. That variable takes an index value near 100 through 2006, showing that through that year,

investors remained sanguine about the prospects for subprime mortgages. As reflected in the ABX indicator, that confidence began to wane in early 2007 and ratcheted downward thereafter. Worsening conditions in mortgage markets corresponded to a deterioration of household balance sheets and, ultimately, also in the balance sheets of mortgage lenders.

Stage 2 of the crisis, indicated by the LIBOR–OIS spread shown in figure 5, was the inception of liquidity pressures on financial institutions that began in the summer of 2007. As Gorton–Metrick point out, the initial loss of investor confidence in the mortgage market (ABX) was not mirrored by any investor concerns about lenders or securitization markets. However, after BNP Paribas announced in August 2007 that it was no longer able to value the subprime mortgages in its sponsored funds, wholesale funding markets came under pressure, beginning with ABCP conduits and other off-balance-sheet vehicles. Funding pressures, as proxied by LIBOR–OIS, continued to build through the second half of 2007 and in 2008, spiking after Lehman Brothers’ failure and AIG’s rescue in September 2008. Funding pressures eased by the end of 2008, presumably reflecting the active policy response, and declined further after the bank stress test results were announced in the spring of 2009.

Stage 3 of the crisis, according to this taxonomy, corresponds to the sharp rise in the ABS spreads on nonmortgage credit (specifically, in figure 5, on credit card receivables) that occurred after the sale of Bear Stearns in March 2008 and, especially, after the collapse of Lehman Brothers and the rescue of AIG. Gorton and Metrick (2012) describe this episode as the “run on repo,” in which repo lenders (particularly in the bilateral repo market) stopped lending against private credit securitizations, except at very short terms and with very large haircuts. The pullback from securitized credit was, I think, somewhat broader than Gorton–Metrick suggest, in that it reflected runs by almost all forms of wholesale funding, not just repo, as well as dumping of credit-backed securities by some investors and also by dealers and other intermediaries. A spike in risk aversion also exacerbated the pullback. In any case, a particularly critical aspect of stage 3, indicative of panic and contagion, was that investors had begun to flee from non-mortgage-related assets as well as mortgage-related ones, despite the fact that nonmortgage credit quality never deteriorated to the extent that most lower-rated mortgages did. As discussed above, the panic led to disintermediation and fire sales, driving up yields on existing credits, as is evident from the behavior of the ABS spread shown in figure 5. These stresses also moderated at about the end of 2008 but continued well into the next year.

The combination of mortgage losses, funding problems, and mark-downs of nonmortgage credit took its toll on the banking system, although government interventions ranging from capital injections to debt guarantees shored up banks as well. Stage 4 of the crisis, capital losses at banks and other lenders, is represented in figure 5 by the CDS spread for the Bank of America. As this variable shows, banks' health worsened steadily through early 2009 (higher values imply a higher risk of default), improved following the stress tests of that spring, but then worsened again at about the time of the credit downgrade of the U.S. government in 2011 and with continuing pressures in Europe.

As suggested by this four-stage simple theory, each stage of the crisis potentially affected real economic activity. In stage 1, falling house prices and rising mortgage payments relative to income pressured household balance sheets and consumer spending, as documented by Mian and Sufi and others. Stage 2 showed the first signs of the panic, as wholesale funders pulled back from lenders, including off-balance-sheet vehicles and conduits. Tighter funding conditions would have been reflected in restrictions on credit supply. Stage 3 was the most violent stage of the panic, as investors refused to fund even nonmortgage securitizations, driving up the yield on nonmortgage credit. As noted, the expansion of the panic to include nonmortgage credit along with mortgages was arguably a turning point of the crisis, with broad ramifications for both firm and household borrowers. Finally, in stage 4, the commercial banking system weakened further, perhaps adding to the constraints on the supply of credit. Powerful feedback effects operated throughout—for example, among the solvency of mortgage lenders, the supply of mortgage credit, household balance sheets, and house prices, with each affecting the others. There were also strong feedback effects between financial and economic developments, as financial disruptions slowed the economy, which in turn worsened financial and credit conditions.

Figure 5 is only illustrative, of course—a vehicle for laying out a narrative of the crisis. (As I have noted, I am focusing here on the United States; additional stages of the crisis could be identified as problems continued and spread in Europe and emerging market economies.) I have two reasons for presenting this figure in detail.

First, as we will see, the four variables shown in figure 5 are not idiosyncratic but instead are stand-ins for larger groupings of financial variables. That is, the narrative I have summarized shows up in a much larger set of financial indicators than the four seemingly arbitrary choices shown above.

Second, figure 5 shows clearly the sharp discontinuities and non-linearities that characterized the crisis. These discontinuities are the basis for the identification strategy of this section. Although there is little doubt, for example, that mortgage problems (stage 1) were an important ultimate source of the subsequent stages of the crisis, the precise size and timing of the subsequent stages depended on many contingencies, ranging from the capital and mortgage exposures of particular firms to the psychology of market participants. These discontinuities should allow us to identify the effects of the various stages of the crisis on the real economy. Put another way, we can ask what would have happened in the real economy if the housing/mortgage crisis had occurred, say, but for some unrelated reason the panic in nonmortgage securitization markets had been avoided. This identification should shed light on the mechanisms by which the crisis affected the economy and help in evaluating policy responses.

III.A. Identifying Stages of the Crisis: Methodology and Data

The methodology employed in the rest of this paper is factor analysis, a data reduction technique that can be used to represent n time series variables as linear combinations of k underlying, orthogonal factors plus idiosyncratic noise, with k much smaller than n . Motivated by figure 5, I applied factor analysis to a set of financial variables, observed daily over the period 2006–12. Because the period of financial distress is relatively short, the hope is that daily data will allow greater insight into the sources of covariation among the indicators and to identify the stages of the crisis with greater precision. Financial variables are used because they are available at high frequency and because they are likely to quickly embody new information about the outlook for financial markets and the economy. I consider 75 series, grouped in four broad categories of roughly equal size. The categories and groupings reflect the narrative of the stages of the crisis given above. Qualitative descriptions of the included variables are as follows (for a more detailed listing of data and sources, see the data appendix):

—*Housing and mortgages* (17 series): Indexes of securitized mortgage values (ABX); ABS spreads for securities backed by home equity loans; homebuilder stock prices; real estate investment trust stock prices; subprime lender stock prices (all stock prices are relative to the Standard & Poor's 500 index).

—*Short-term funding* (15 series): LIBOR–OIS spreads of various maturities; TED (difference between 3-month LIBOR and treasury yields)

spreads; ABCP spreads; financial commercial paper spreads; repo spreads (yields on general collateral financing mortgage-backed securities and agency securities over Treasury repo).

—*Nonmortgage credit* (22 series): ABS spreads (credit cards, auto loans, student loans); ABS indexes (consumer loans); corporate bond spread indexes; A2P2 (lower-rated) commercial paper rates, relative to OIS.

—*Bank solvency* (21 series): For the largest U.S. commercial and investment banks, CDS spreads and stock prices (relative to the Standard & Poor's 500).

To interpret these data, I performed two exercises. First, I applied factor analysis to the full sample of 75 variables, an exercise I refer to as *full-sample factor analysis*. This analysis, which makes no prior distinctions among the four groups of financial variables, shows that at least three orthogonal factors are required to adequately describe the data, with a borderline case for including a fourth factor (see below for further discussion).

Second, I applied factor analysis to each of the four groups of variables separately, extracting a single factor from each group. I call this procedure *subsample factor analysis*. I found that one estimated factor per group seemed adequate, with a single factor typically explaining about 70 percent of the sum of squared residuals in each subsample. Unlike the full-sample factors, the subsample factors reflect my prior groupings of the 75 variables into descriptive categories.

As a general matter, for the purposes of summarizing and, potentially, interpreting these data, both the full-sample and subsample factor analyses have advantages. The full-sample analysis uses and describes all the data simultaneously, without imposing prior categories; and, because the estimated full-sample factors are orthogonal by construction, decomposing economic forecasts into components attributable to each factor is straightforward. Conversely, without further assumptions, the economic interpretations of the full-sample factors may not be clear. In contrast, the factors estimated in individual subsamples have more obvious economic interpretations, by construction. The factor extracted from the group of mortgage and housing variables, for example, is naturally viewed as a summary measure of housing developments, as reflected in financial markets. However, the subsample analyses would generally be expected to have their own shortcomings. In particular, the factors estimated separately in the subsamples are not guaranteed to be mutually orthogonal, making more difficult the attribution of forecasting power or causality to one factor versus another.

Importantly, however, for reasons that are discussed below, in the present application the sets of factors extracted by the two methods turn out to be quite similar. Figure 6 graphically compares the four factors estimated jointly from the full sample with those estimated separately in the subsamples. In the figure, factor 1 is the estimated factor explaining the greatest share of the variance of the 75 variables; factor 2 explains the greatest share of the remaining variance after controlling for factor 1; factor 3 explains the most variance after controlling for factors 1 and 2; and so on.⁸ The factors estimated independently from the four subsamples are designated in the figure as the “housing,” “nonmortgage credit,” “funding,” and “bank solvency” factors.

The comparison shown in figure 6 between the estimated full-sample and subsample factors is striking. The first factor estimated from the full sample (factor 1) lines up nearly perfectly with the factor estimated from the housing subgroup (figure 6, upper left-hand panel). Likewise, the second estimated factor from the full sample (factor 2) looks very similar to the factor estimated from only the financial variables related to nonmortgage credit, and the third full-sample factor (factor 3) lies nearly on top of the factor estimated from short-term funding variables only. The fourth full-sample factor, which as noted above explains a relatively small amount of the variance of the full set of data, is evidently correlated with the factor estimated from the bank solvency variables (as can be seen in the lower right panel of figure 6), but the overall relationship is weaker.

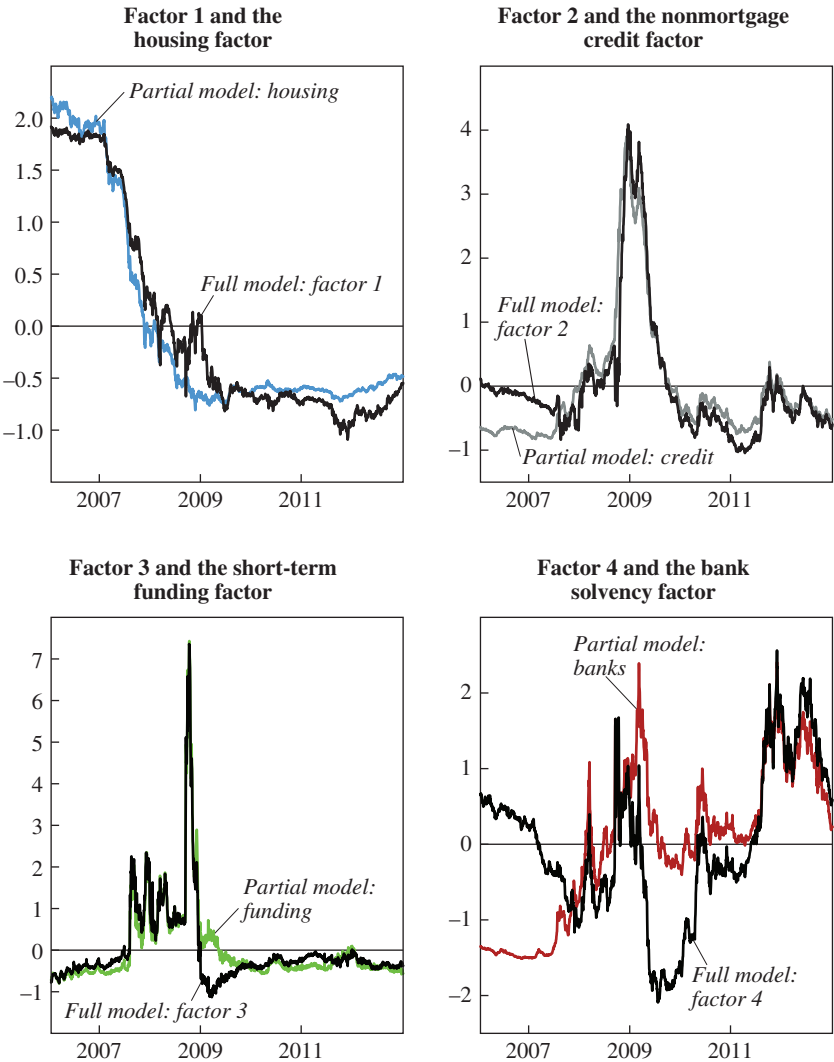
The correlations of the full-sample and subsample factors, shown in table 1, confirm the visual impressions of figure 6. The correlations of factors 1, 2, and 3 with the housing, credit, and funding factors are 0.97, 0.95, and 0.92, respectively, despite the noisiness of the daily data. But the correlation of factor 4 with the bank solvency factor is only 0.40. Interestingly, however, the bank solvency factor has a correlation with factor 1 of -0.86 . Economically, interpreting factor 1 as the housing factor suggests that deterioration in the housing and mortgage markets is an important driver of investor assessments of bank solvency over this period.

What accounts for the close correlation of the full-sample factors, estimated in an unconstrained way from all 75 variables, and the subsample factors, each estimated from about one-fourth of the variables? To answer this question, first, note that estimated factors in general are identified only

8. The shares of variance explained by factors 1 through 4 are 0.32, 0.26, 0.19, and 0.08, respectively.

Figure 6. Estimated Factors: Full Sample versus Subsamples, 2006–12^a

Standard deviation from the mean



Source: Author's calculations.

a. The panels compare estimated factors from the full sample and from the subsamples.

Table 1. Correlations of Full-Sample Factors and Subsample Factors

<i>Factor</i>	<i>Factor 1</i>	<i>Factor 2</i>	<i>Factor 3</i>	<i>Factor 4</i>
Housing	0.97	-0.13	-0.12	0.09
Funding	-0.07	0.30	0.95	-0.02
Credit	-0.26	0.92	0.30	-0.03
Banks	-0.86	0.30	0.01	0.40

Source: Author's calculations.

up to an orthogonal rotation, as any linear combinations of the estimated factors that preserve their orthogonality will explain precisely the same fraction of the variability of the data. To pick a normalization, in our full-sample estimation I applied a standard procedure called a varimax rotation. By design, this procedure tends to favor normalizations in which some variables have very high loadings on a given factor and near-zero loadings on the other factors.⁹ In effect, the varimax procedure tends to associate estimated factors with groups of observed variables that are highly correlated within the group but have relatively low correlations with variables outside the group.


I suggested above that the four variables shown in figure 5 were representative of a broader set of data. The factor analysis confirms this claim. The full-sample factor analysis sorts the larger data set into three, or possibly four, groups of variables, with relatively high intragroup correlations and lower intergroup correlations. Comparing the full-sample and subsample factors in turn suggests that these groups are economically interpretable and correspond to our description of the stages of the financial crisis. In particular, figure 7, which shows the estimated full-sample factors, looks qualitatively very similar to figure 5, which described the stages of the crisis in terms of a few, apparently arbitrarily chosen, variables. In short, the story told using a few chosen variables in figure 5 can also be told by considering the common factors in larger groups of financial variables.

Further motivation for equating the estimated factors with stages of the financial crisis is given by figure 8, which shows the squared factor scores for the full-sample factors. Loosely, the figure shows the average variability of the financial data and the share of this variability accounted

9. More specifically, this procedure chooses the particular orthogonal combination of factors that maximizes the sum of the variances of the squared correlations between the explained variables and the estimated factors.

Figure 7. Estimated Factors from the Full Sample, 2006–12^a

Standard deviation from the mean




Source: Author's calculations.

a. Data show full-sample estimated factors computed from 75 standardized variables over the period 2006–12.

for by each factor over the 2006–12 period. The periods during which each factor is dominant correspond closely to the stages of the crisis discussed above. For example, factor 1, which from now on we identify with housing and mortgages, is the dominant source of variability from the beginning of the sample through mid-2007, while factor 3, which corresponds to short-term funding stresses, becomes important after the BNP Paribas announcement, spiking after the Lehman Brothers failure and the AIG rescue. Factor 2 (nonmortgage credit) is the dominant factor beginning shortly after Lehman Brothers/AIG into early 2009, and factor 4 (bank solvency) lags the other stages. Based on our economic interpretations of the estimated factors, we use them in the next stage of the analysis, where we examine how well they forecast aspects of real activity.

Figure 8. Squared Factor Scores, 2006–12^a

Standard deviation from the mean, squared



Source: Author's calculations.

a. Data show full-sample estimated factors computed from 75 standardized variables over the period 2006–12.

Before turning to these results, there is one further issue of interpretation to discuss. Factor 2, the second-most-important estimated factor in the full data set, is associated with the deterioration of nonmortgage credit—as reflected, for example, in wider spreads for securities backed by nonmortgage assets. However, even within a framework that emphasizes credit frictions and asymmetric information, there are at least two alternative economic interpretations of this factor. First, the weakening of the economy, and the associated deterioration of household and nonfinancial firm balance sheets, clearly worsened the creditworthiness of consumers and firms; in principle, this deterioration in borrowers' financial health could account for the blowout in nonmortgage spreads. A second possibility is that the rise in nonmortgage spreads primarily reflected a change in investor

behavior, as investors lost confidence in all forms of private (and especially securitized) credit. In this interpretation, the panicky pullback from mortgage-related and securitized credit (including the Gorton-Metrick “run on repo”) and the subsequent fire sales led to sharply depressed prices and also to higher spreads on nonmortgage credit. In short, in principle, the movements in factor 2 could reflect developments on either the demand side of credit markets (borrower financial health) or the supply side (lender health and investor confidence).

Although these two interpretations of factor 2 are not mutually exclusive, the evidence favors the second, investor-led explanation. First, aggregate balance sheets evolve relatively slowly, which seems inconsistent with the sharp deterioration in the nonmortgage credit factor after the Lehman Brothers failure, and (given the slow pace of deleveraging and financial recovery) looks especially inconsistent with the sharp improvement in this factor that began just a few months later. Additional evidence on this point is given by figure 9, which shows factors estimated separately for the household and nonfinancial corporate components of the nonmortgage credit subsample. As the figure shows, the two estimated factors lie almost on top of each other, indicating the virtually identical behavior of spreads on these two categories of credit. The correlation of the two series in daily data is 0.96. Because household and corporate balance sheets certainly evolved differently during the crisis (compare figures 2 and 3 above), the high correlation strongly suggests a common determinant, which I take to be the general run on credit products by panicked investors and the subsequent fire sales. Consistent with this assessment, Francis Longstaff (2010) finds strong evidence of contagion from subprime mortgages to other markets, and Alberto Manconi, Massimo Massa, and Ayako Yasuda (2012) find contagion from toxic securities to corporate bonds arising from changes in investor demands for liquidity.


III.B. How Do the Stages of the Crisis Forecast the Economy?

We turn now to a key question: To what extent do the factors, estimated strictly from financial variables and intended to reflect the stages of the financial crisis, predict aspects of real economic activity?

To answer this question, I began with a list of monthly economic indicators, and I aggregated the daily financial factors to monthly averages. (See the appendix for details and sources of economic data. Here, “GDP” is a monthly measure of real output constructed by Macroeconomic Advisers. All other series are from official sources.) For each economic

Figure 9. Estimated Factors for Household and Corporate Credit Variables, 2006–12^a

Standard deviation from the mean



Source: Author's calculations.

a. Data show the first factors estimated from consumer and corporate nonmortgage credit separately.

indicator, I estimated a prediction equation over the 2006–12 sample. Prediction equations, estimated by ordinary least squares, include a constant, two monthly lags of the predicted indicator, and the current value and two monthly lags of each of the factors sequentially.¹⁰

Table 2 shows the statistical significance of each full-sample factor in the respective prediction equations, compared with the simple AR2 baseline. As the table shows, factor 2 (which we identify with nonmortgage credit) and factor 3 (short-term funding) are statistically significant at the 5 percent

10. The results are qualitatively similar when multiple factors are included in the same prediction equation. Note that the factors are orthogonal by construction in daily data, but for sampling reasons are not precisely orthogonal when aggregated to monthly series.

Table 2. *F* Statistics for Exclusion of Each Factor from Prediction Equations^a

<i>Forecasted variable</i>	<i>Factor 1: housing</i>	<i>Factor 2: credit</i>	<i>Factor 3: funding</i>	<i>Factor 4: banks</i>
GDP	0.12	4.95***	3.26**	0.42
Industrial production	0.04	7.82***	4.76***	1.60
Employment, excluding construction	1.66	8.75***	2.52*	0.29
Unemployment	1.16	11.35***	2.56*	1.09
Real PCE	0.41	4.20***	3.69**	0.77
Real PCE—durables	0.18	3.12**	3.67**	0.46
Retail sales	0.15	11.02***	4.54***	2.79**
Housing starts	1.34	1.69	0.96	1.74
Capital goods orders	0.40	9.45***	2.89**	3.07**
ISM Manufacturing Index	0.62	23.97***	13.09***	1.42
Core PCE inflation	0.99	1.9	0.83	0.44
Degrees of freedom	(3;76)	(3;76)	(3;76)	(3;76)


Source: Author's calculations; see the appendix.

a. *F* statistics test the exclusion of each factor sequentially from a prediction equation that also contains two monthly lags of the forecasted variable. Statistical significance is indicated as ***, **, and * for $p < 0.01$, $p < 0.05$, and $p < 0.1$. PCE = personal consumption expenditures.

or 1 percent level for most variables—basically, everything except housing starts and core inflation. By this metric, factor 1 (housing) and factor 4 (bank solvency) do much worse. Factor 1 is not significant at the 10 percent level in any prediction equation, and factor 4 is predictive (at the 5 percent level) only for retail sales and capital goods orders.

Table 2 reports the *statistical* significance of the factors as predictors. To assess *economic* significance, I used the estimated prediction equations to create simulated forecasts of each macro variable for the sample period 2006–12. The simulated forecasts are dynamic; that is, I simulate each prediction equation forward from the beginning of the sample, applying the autoregressive coefficients dynamically to simulated rather than actual values of the macro variables. Note that, in order to assess the importance of each factor in isolation, the dynamic forecasts use one factor at a time, implicitly assuming that other factors are zero.

Figure 10 shows graphically the results of the dynamic simulation exercise for one macroeconomic variable, industrial production. This variable is selected because the results are typical. In both panels of the figure, the black line shows the actual, historical path of (the growth rate of) industrial production for the period 2006–12. The other lines in the figure show the dynamically forecast path of industrial production based on each full-sample factor, taken one at a time. As the top panel of figure 10 shows, dynamic forecasts based on factor 1 (housing) and factor 4 (banks) do not

Figure 10. Dynamic Simulations: Industrial Production, 2006–12^a

Source: Author's calculations.

a. Data show dynamic simulations of a model regressing industrial production on two lags of itself and on each factor and two lags of each factor. Forecasts are dynamic, in that the lagged values are predicted rather than realized. Dependent variables are in year-over-year percentage changes.

Table 3. Correlation of Actual Values of Forecasted Variables with Simulated Values^a

<i>Forecasted variable</i>	<i>Factor 1: housing</i>	<i>Factor 2: credit</i>	<i>Factor 3: funding</i>	<i>Factor 4: banks</i>
GDP	0.30	0.88	0.80	0.12
Industrial production	0.41	0.90	0.86	0.22
Employment, excluding construction	0.61	0.93	0.69	0.44
Unemployment	0.88	0.99	0.93	0.91
Real PCE	0.63	0.87	0.86	0.21
Real PCE (durables)	0.45	0.83	0.86	-0.23
Retail sales	0.47	0.95	0.79	0.35
Housing starts	0.75	0.54	0.60	0.45
Capital goods orders	0.46	0.88	0.69	0.42
ISM Manufacturing Index	0.51	0.93	0.87	0.02
Core PCE inflation	0.66	0.74	0.62	0.54

Source: Author's calculations; see the appendix.

a. PCE = personal consumption expenditures.

capture much of the variation in industrial production. This result is not surprising, given the low statistical significance for these factors seen in table 2. In contrast, the bottom panel shows the better fit of the forecasts conditional on factor 2 (nonmortgage credit) and factor 3 (funding). In particular, both factors capture much of the decline in activity in the second half of 2008 and the recovery in mid-2009. The funding factor captures a bit less of this decline than the nonmortgage credit factor but also leads the downturn by a bit more. Again, these qualitative results are typical for these simulations. Table 3 shows the correlations of the forecasted macro variables with the dynamic simulations of these variables. The highest correlations are with the credit factor or the funding factor for every macro variable except housing starts, which is most correlated with the housing factor.

Rather than show analogous figures for each of the macro variables, I next consider a somewhat different comparison. At some level, all the major elements of the crisis were driven by the housing boom and bust and the associated mortgage lending. However, as discussed in the introduction to this paper, the housing and mortgage bust affected the economy through at least two broad channels. First, as in the “financial fragility” narrative of the introduction, actual and potential mortgage losses, together with vulnerabilities such as high leverage and dependence on short-term funding, collapsed investor confidence not only in mortgages but in a much broader set of securities. The loss of investor confidence led to indiscriminate runs, disintermediation, and fire sales that sharply reduced the prices and

increased the yields on most forms of private credit, not just residential mortgages. In the present analysis, this “panic” channel can be represented by the combination of factor 3 (which reflects stresses in markets for whole-sale funding) and factor 2 (which captures the broader run on securitized credit, especially nonmortgage credit).

Second, even in the absence of a panic, the housing and mortgage bust would have affected the economy by damaging sectoral balance sheets. The damage to household balance sheets was particularly severe—this is the “household leverage” narrative of the introduction—and presumably constrained consumer spending. In addition, even in the absence of a panic, mortgage losses would have reduced the capital of banks and other lenders and thus limited the supply of credit. In the analysis presented here, the “nonpanic” effects of developments in housing and mortgage markets are represented by full-sample factor 1, and additional developments regarding the solvency of banks are captured by factor 4. In the horse races below, we combine the predictive power of factors 1 and 4 and refer to them in tandem as the “balance sheet channel”; that is, together they reflect developments in the balance sheets of both households and banks. However, the inclusion of factor 4 makes only a modest difference, and the results reported below are not much changed if only factor 1 is included.

To compare the economic importance of these two channels, we look at the predictive power for our list of economic indicators of the “panic factors” (factors 2 and 3) versus the “balance sheet factors” (factors 1 and 4). Again, we estimate prediction equations for each monthly economic indicator. Each equation includes two lags of the predicted variable, plus the current value and two lags of (1) both panic factors or (2) both balance sheet factors. Table 4 shows the resulting F statistics for the joint inclusion of the factors against an AR2 baseline.

Not surprisingly, given the earlier results, the predictive power of the two panic factors greatly exceeds that of the two balance sheet factors. Exclusion of the panic factors from the prediction equations is rejected at the 1 percent level for all the economic indicators, except for housing starts and core inflation. The balance sheet factors are significant at the 5 percent level only for capital goods orders.

Figure 11 shows the results of running dynamic simulations for representative economic indicators, conditional on the estimated values of, separately, the balance sheet factors and the panic factors. Each panel of this figure shows, for the 2006–12 sample period, the actual path of the economic indicator in question, compared with the simulated values.

Table 4. *F* Statistics for Exclusion of Pairs of Factors in Prediction Equations^a

<i>Forecasted variable</i>	<i>Panic factors (factors 2 and 3)</i>	<i>Balance sheet factors (factors 1 and 4)</i>
GDP	3.58***	0.25
Industrial production	5.82***	0.86
Employment, excluding construction	4.65***	1.16
Unemployment	7.67***	1.69
Real PCE	4.31***	0.88
Real PCE—durables	5.34***	0.34
Retail sales	9.19***	1.77
Housing starts	1.37	1.50
Capital goods orders	5.34***	1.86**
ISM Manufacturing Index	17.53***	0.99
Core PCE inflation	1.13	0.99
Degrees of freedom	(6;73)	(6;73)

Source: Author's calculations; see the appendix.

a. *F* statistics are relative to an AR2 baseline. Statistical significance is shown as *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$. PCE = personal consumption expenditures.

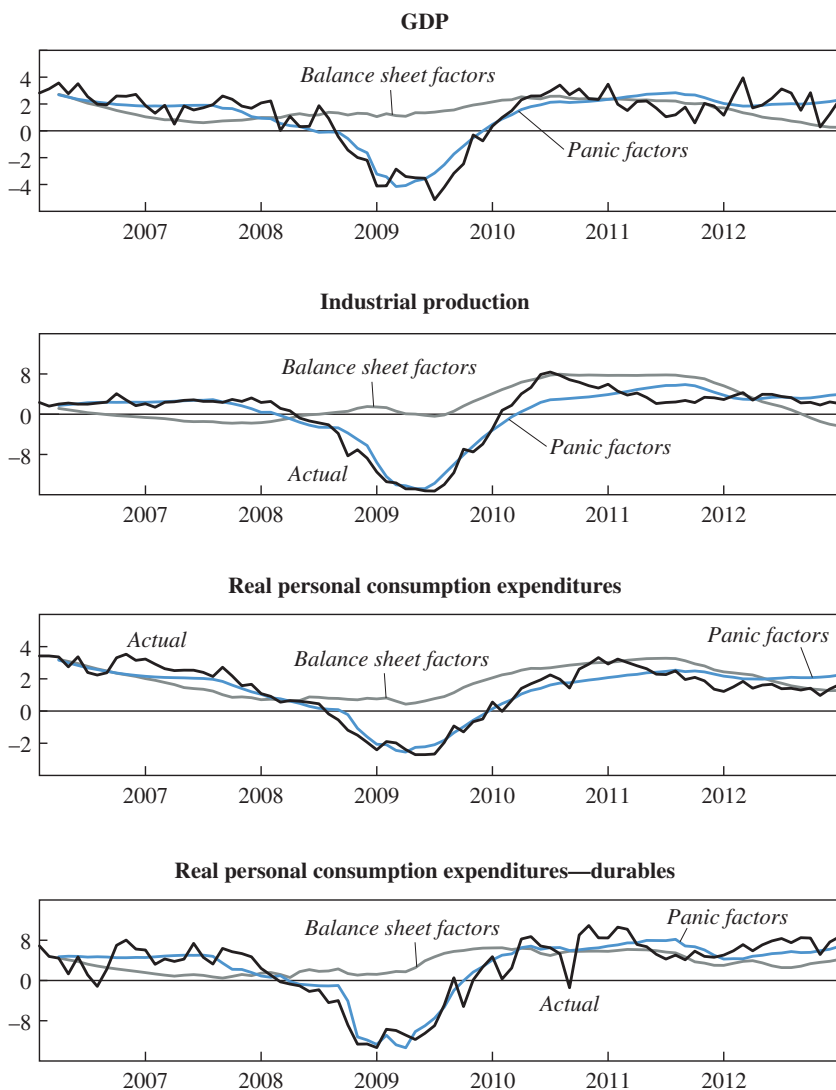
Consistent with table 4, the comparisons are quite one-sided. For housing starts (in the bottom right panel of the figure), the balance sheet variables provide a better fit in the first part of the sample, but not after late 2008. For all the other variables shown, along with those omitted for lack of space, the panic factors provide uniformly better (and quite close) fits.

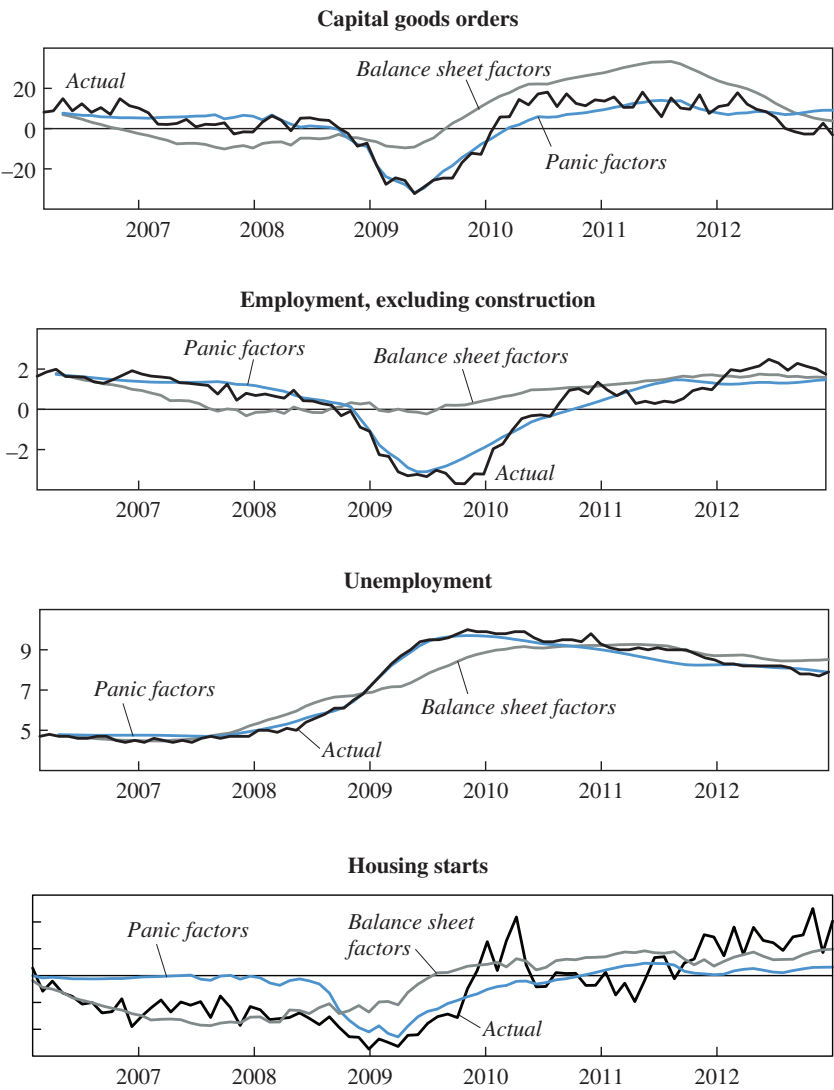
The *F* statistics shown in tables 2 and 4, and the dynamic simulations shown in figure 11, are the main results of this part of the paper. I interpret these results (and the robustness checks discussed below) primarily as an affirmation of the role of the panic in explaining the severity of the economic downturn in late 2008 and early 2009. In intuitive terms, we see that financial markets showed large, discontinuous breaks at certain points during the sample period; these breaks were closely associated with variables indicative of panic in funding and securitization markets; and these shifts in turn are strongly predictive of a range of macroeconomic variables. The finding of the centrality of the panic helps to explain why the recession, which looked moderate in its early stages, became so deep.

Importantly, although balance sheet factors do not forecast economic developments well in my setup, I do not think we should conclude that these channels of transmission were unimportant, even putting aside the point that the housing boom and bust helped to trigger the panic in the first place. First, the full-sample factor analysis finds that the factor most closely identified with housing (factor 1) explains the largest share of the variation of the financial variables considered over the 2006–12 sample period; and,

Figure 11. Dynamic Simulations: Panic Factors and Balance Sheet Factors, 2006–12^a

Percent





Source: Author's calculations.
a. Panic factors include factor 2: nonmortgage credit, and factor 3: funding. Balance sheet factors include factor 1: housing, and factor 4: bank solvency. Simulations are as shown in figure 10.

in particular, that the housing factor dominates this variability during the first part of the sample (figure 8). Evidently, market participants viewed developments in housing and mortgages as having significant economic consequences, even during the period before they became concerned about broader financial instability. Second, as already discussed, diverse empirical studies have found significant links between household leverage and employment, including those by Mian and Sufi (2010, 2014b); Jan Hatzius (2008); Haltenhof, Lee, and Stebunovs (2014); and Mikael Juselius and Mathias Drehmann (2015). Beyond work based on the U.S. experience, several studies have used international and historical data to draw connections between household leverage buildups and subsequent recession (Jordà, Schularick, and Taylor 2016; Mian and Sufi 2018). With all this (and other) evidence taken into account, a plausible conclusion is that the deterioration of household balance sheets exacerbated the early declines in consumer spending, particularly on consumer durables, and proved a drag on the pace of recovery, while the panic explains the acute phase of the economic downturn. Likewise, I would not conclude from the poor predictive performance of factor 4 that the balance sheets of banks (outside their effects on the probabilities of panic) were not economically important, for very much the same reasons. It may be that both household and bank balance sheets evolve too slowly and comparatively smoothly for their effects to be picked up in the type of analysis presented in this paper.

III.C. Two Robustness Checks

I briefly report next on a couple of robustness checks of this paper's key finding, that the panic phase of the crisis was central to explaining the damage that the crisis wrought on the real economy. First, the results presented above use the factors estimated from the full sample of 75 financial variables. Alternatively, we can use the factors estimated separately on each of the four subsamples to represent the stages of the crisis. Because the subsample factors (unlike the full-sample factors, by construction) are not orthogonal, we orthogonalize them in this order: housing, funding, nonmortgage credit, and bank solvency. This ordering is consistent with the hypothesized sequencing of the crisis (see the discussion of figure 5). In particular, by ordering first the factor estimated in the housing subsample, this procedure attributes co-movements of the housing variables and other variables entirely to the housing factor. This procedure will likely lead us to understate the economic effects of the panic, because it excludes the possibility that the panic itself was the cause of some of the deterioration in the housing market.

Table 5. Correlations of Full-Sample Factors with Orthogonalized Subsample Factors

<i>Factor</i>	<i>Factor 1: housing</i>	<i>Factor 2: credit</i>	<i>Factor 3: funding</i>	<i>Factor 4: banks</i>
Housing	0.97	-0.13	-0.12	0.09
Funding (orth.)	0.14	0.28	0.94	0.01
Credit (orth.)	0.09	0.95	-0.28	0.01
Banks (orth.)	-0.13	0.02	-0.04	0.91

Source: Author's calculations.

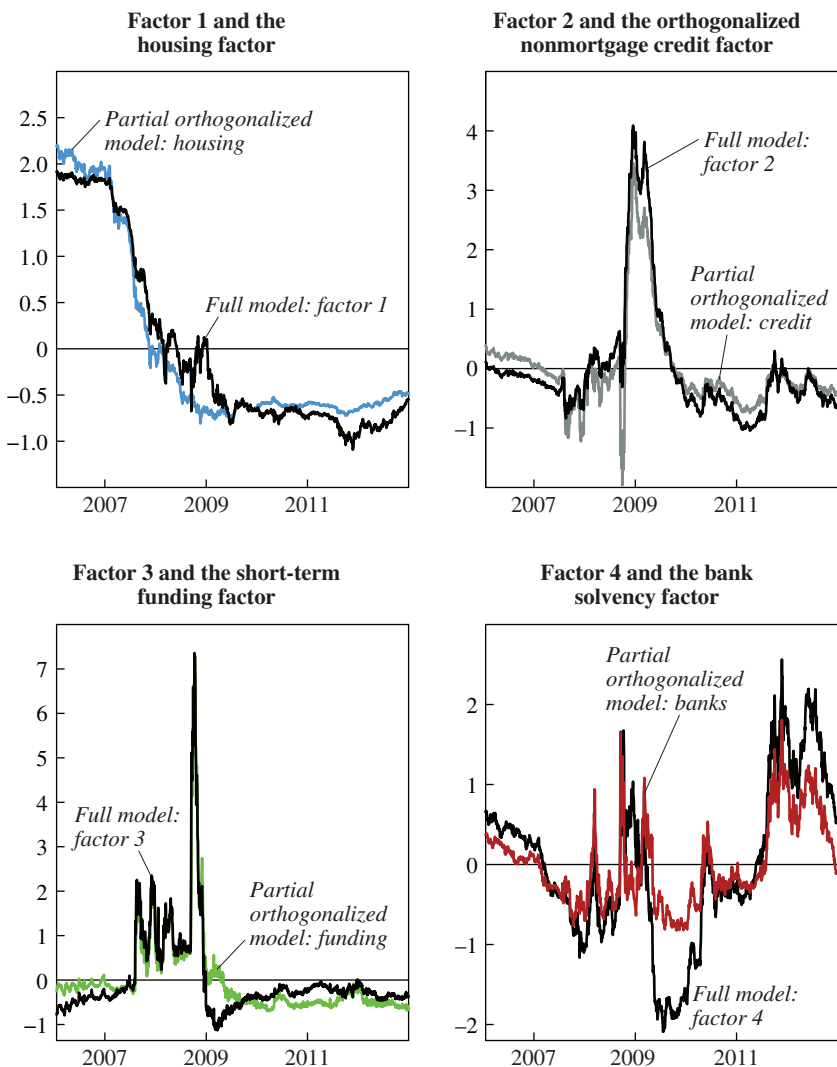
Table 5 shows the correlations of the full-sample factors with the orthogonalized subsample factors, and graphical comparisons of the full-sample factors with the orthogonalized subsample factors are shown in figure 12. The correlations of the first three full-sample factors with the housing, nonmortgage credit, and funding subsample factors respectively remain high, consistent with table 1. Interestingly, however, the fourth full-sample factor now lines up reasonably well with the factor estimated from the bank solvency subsample. (Recall that, in contrast, in table 1, factor 1 had the greatest correlation with the bank solvency factor.) Intuitively, the orthogonalization procedure appears to have isolated movements in bank balance sheets that are independent of housing and mortgage developments, and these movements in turn appear to constitute an independent (though relatively small) determinant of financial market outcomes during the crisis.

Table 6 reports the results of an exercise analogous to that shown in table 4, comparing the predictive power for monthly macroeconomic indicators of the two panic factors (funding and nonmortgage credit) and the two balance sheet factors (housing and bank solvency), except that here the orthogonalized subsample factors are used in place of the full-sample factors. Again, the predictive power of the panic factors is extremely strong, significant at the 1 percent level for all variables except housing starts and core inflation. The performance of the balance sheet factors is again much weaker.

For a second robustness check, I also considered proxies for the panic and balance sheet developments that make no use of factor analysis. Table 7 shows F statistics for prediction equations, constructed in analogy to tables 4 and 6, but using (in lieu of estimated factors) monthly values of the Federal Housing Finance Agency's housing price index, the three-month mortgage delinquency rate calculated by Fannie Mae (see figure 2), and the Gilchrist–Zakrajšek excess bond premium (see figure 1). The first

Figure 12. Estimated Factors: Full Sample versus Subsamples, 2006–12^a

Standard deviation from the mean



Source: Author's calculations.

a. This figure compares estimated full-sample factors with estimated factors from the subsamples, where the latter have been orthogonalized in this ordering: housing, funding, credit, and banks.

Table 6. *F* Statistics for Inclusion of Pairs of Factors in Prediction Equations^a

<i>Forecasted variable</i>	<i>Panic factors (orthogonalized)</i>	<i>Balance sheet factors (orthogonalized)</i>
GDP	3.25***	0.18
Industrial production	4.93***	0.85
Employment, excluding construction	4.60***	0.34
Unemployment	6.81***	2.82**
Real PCE	3.57***	0.84
Real PCE—durables	4.99***	0.23
Retail sales	8.45***	0.90
Housing starts	1.56	1.01
Capital goods orders	5.01***	0.91
ISM Manufacturing Index	15.67***	1.88*
Core PCE inflation	1.11	1.02
Degrees of freedom	(6;73)	(6;73)

Source: Author's calculations; see the appendix.

a. Panic and balance sheet factors are the orthogonalized partial factors. *F* statistics are for exclusion of pairs of factors, relative to an AR2 baseline. Statistical significance is shown for *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$. PCE = personal consumption expenditures.

Table 7. *F* Statistics for Exclusion of Alternative Crisis Measures in Prediction Equations^a

<i>Forecasted variable</i>	<i>House prices</i>	<i>Delinquencies</i>	<i>EBP</i> <i>(orthogonalized)</i>
GDP	2.62*	2.73**	7.85***
Industrial production	1.98	2.84**	11.12***
Employment, excluding construction	0.75	5.69***	8.44***
Unemployment	1.71	9.32***	15.24***
Real PCE	2.51*	2.95**	7.56***
Real PCE—durables	2.55*	2.05	6.1***
Retail sales	1.30	2.22*	8.93***
Housing starts	3.52**	3.14**	1.71
Capital goods orders	1.08	3.07**	7.91***
ISM Manufacturing Index	1.81	4.78***	15.47***
Core PCE inflation	1.01	1.81	1.86
Degrees of freedom	(3;76)	(3;76)	(3;76)

Source: Author's calculations, see the appendix.

a. *F* statistics are relative to an AR2 baseline. Statistical significance is shown for *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$. EBP = excess bond premium; PCE = personal consumption expenditures.

two variables capture developments in the housing and household balance sheets. Recall that the GZ excess bond premium is a measure of corporate bond interest rate spreads that controls for estimated default probabilities and thus reflects primarily investors' appetite for corporate credit, as reflected in the risk and liquidity premiums for this important category of private credit. We take this measure as a proxy for the panic; its sensitivity to the panic is evident in figure 1.

In table 7, the predictive power of (the log levels of) house prices and mortgage delinquencies are assessed (separately) in the first two columns, and that of the EBP in the third column. The fourth column shows the predictive power of the orthogonalized EBP—that is, the residual when the EBP is regressed against both house prices and delinquencies. This procedure has the effect of attributing any joint explanatory power of the EBP and the first two variables to the first two variables alone.

Table 7 shows that the EBP, even when orthogonalized, is a strong predictor of macro variables; its exclusion from the prediction equations is rejected at $p < .01$ for 9 of the 11 variables, with the exceptions (as in tables 4 and 6) being housing starts and core inflation. Interestingly, delinquencies also show some forecasting power in this exercise, although less than the EBP; I take this result as providing some support for the view that weak household balance sheets contributed to the broader economic decline. House prices are not very predictive, but, reasonably, both house prices and household delinquencies predict housing starts at the 5 percent level. Overall, the EBP's strong predictive power supports the conclusion obtained from the factor analysis: that the panic had powerful macroeconomic effects.

IV. Conclusions and Policy Implications

Ten years after the peak of the financial crisis, this paper has reviewed the role of credit factors in the crisis and in macroeconomics generally. A substantial body of evidence now suggests that such factors are important for the behavior of households, firms, and financial intermediaries. Macroeconomic modeling and analysis will need to consider such factors or risk substantial forecast misses, as were seen in 2008.

More specifically, the empirical portion of this paper has shown that the financial panic of 2007–09, including the runs on wholesale funding and the retreat from securitized credit, was highly disruptive to the real economy and was probably the main reason that the recession was so unusually deep. Presumably, the effects of the panic and the

associated disintermediation of credit were transmitted through a spike in the economy-wide EFP, together with sharp increases in risk aversion and liquidity preference. The results thus support the modeling done by Gertler and Kiyotaki (2015), among others. Again, the identification of the effects of the panic in this analysis is based on the evident discontinuities defining the key stages of the crisis. Although the panic was certainly not an exogenous event, its timing and magnitude were largely unpredictable, the result of diverse structural and psychological factors. Nor does it seem plausible that the panic happened because investors suddenly began to expect a severe deepening of the recession (that is, no reverse causality). Consequently, the fact that the panic preceded a broad-based downturn, and that the end of the panic preceded an improvement in macroeconomic conditions, is *prima facie* evidence that the panic had significant real effects.

Although variables related to housing and mortgages generally do not forecast well in my setup, it is worth reemphasizing that concluding these factors were unimportant is not justified, even putting aside their role as triggers for the panic. On balance, the cross-sectional evidence (and some more-limited time series evidence) supports the conclusion that the state of household balance sheets is an important determinant of spending decisions, both before and during the Great Recession. In particular, it seems plausible that the weakening of household balance sheets was a main reason for the slowing consumer spending in the period leading up to the crisis (Mian, Rao, and Sufi 2013), and that the need for household deleveraging and balance sheet repair was a significant headwind to recovery. Because balance sheet conditions usually evolve relatively slowly and continuously, however, identifying their macroeconomic effects by time-series methods (like mine) is difficult, particularly over a short period. Gertler and Gilchrist (2018), who combine time-series and cross-sectional data, find a larger role for household balance sheets in explaining the recession.

The findings concerning the importance of the panic have important policy implications, both retrospective and prospective. Retrospectively, policymakers (including the Federal Reserve and the Treasury) took aggressive and often highly unpopular measures to arrest the financial panic, including expanding lending well beyond the banking system and undertaking a series of interventions to recapitalize the banking system and to avoid the collapse of systemically important financial institutions. The stated rationale for these actions was policymakers' fears that, if not arrested, the panic would do severe and lasting damage to the economy,

perhaps resulting in a new Great Depression. The results of this paper provide some after-the-fact support for policymakers' claims.¹¹

Figure 13 provides a schematic of the panic and the policy response. The top two panels show the full-sample estimated factors corresponding to nonmortgage credit and to funding. These are the two panic factors, whose predictive power for the economy was shown above. Also shown in the top two panels are vertical lines indicating important policy initiatives undertaken by the Fed, the Treasury, and the Federal Deposit Insurance Corporation (FDIC). Box 1 briefly defines and describes these initiatives. As a metric of the policy response, the bottom panel of figure 13 shows the portion of the Federal Reserve's balance sheet associated with its various emergency lending programs (but excluding asset purchases associated with quantitative easing or the stabilization of Bear Stearns or AIG).


As can be seen in figure 13, in the first year or so of the crisis, from August 2007 to August 2008, policy mostly took the form of lender-of-last-resort activity, with the Federal Reserve extending its set of allowable counterparties beyond the banking system. Notably, the Fed provided liquidity to primary dealers—large broker-dealers that transact directly with the Fed—through its Term Securities Lending Facility and Primary Dealer Credit Facility programs. To overcome the stigma for banks of borrowing from the discount window, the Fed also started a program of auctioning term discount-window credit (Term Auction Facility). The Fed also reacted to global money market tensions by instituting currency swap programs with 14 foreign central banks, including 4 in emerging markets. These liquidity programs did not end the funding crisis but, as figure 13 suggests, stresses did not worsen significantly over the year.

However, funding problems intensified severely after the failure of Lehman Brothers and the rescue of AIG in September 2008.¹² After a money

11. Janice Eberly pointed out to me that these results also bear on the choice of policies to help distressed homeowners. Much debated at the time was whether it would be better to give lenders incentives to write down the principal of troubled mortgages or instead to focus on alleviating household liquidity constraints through government income support. If the problem was on the supply side of credit markets, as we have found here, then increasing net housing wealth or collateral through principal write-downs would not have led to higher spending, because households, not being able to borrow on almost any terms, could not liquefy their wealth. Instead, directly increasing current income would have been more effective at reducing financial distress, promoting spending, and enhancing welfare. See Eberly and Krishnamurthy (2014) and Ganong and Noel (2018). The latter find that, in fact, income supports during the Great Recession increased spending but principal reductions did not.

12. The government takeover of Fannie and Freddie in August 2008 is viewed by some as the seminal event of the crisis. Mishkin (2011) argues that the struggle to pass TARP in the weeks after the failure of Lehman Brothers also exacerbated market uncertainties.

Figure 13. Policy Interventions, 2007–9



Sources: Author's calculations; Federal Reserve Board.

a. TARP = Troubled Assets Relief Program; CPP = Capital Purchase Program; TALF = Term Asset-Backed Securities Loan Facility; PPIP = Public-Private Investment Program.

b. TAF = Term Auction Facility; TSLF = Term Securities Lending Facility; PDCF = Primary Dealer Credit Facility; AMLF, TGP = Asset-Backed Commercial Paper and Money Market Liquidity Facility, Temporary Guarantee Program; CPFF = Commercial Paper Funding Facility; TLGP = Temporary Liquidity Guarantee Program; MMIF = Money Market Investor Funding Facility; MBS = mortgage-backed securities.

Box 1. Policy Responses to the Panic

Federal Reserve and other U.S. government programs referred to in figure 13 include:

1. Discount window lending, including primary, secondary, and seasonal credit. Available to depository institutions only.
2. Term Auction Facility (TAF), under which the discount window credit was auctioned. See Armantier, Krieger, and McAndrews (2008) for details. McAndrews, Sarkar, and Wang (2017) find that TAF-related events were associated with downward moves in LIBOR.
3. Term Securities Lending Facility (TSLF). In this program, the Fed lent Treasury securities to primary dealers, taking mortgage-related securities as collateral. Fleming, Hrung, and Keane (2010) found that TSLF loans reduced repo spreads, but Wu (2008) reported that the TSLF and PDCF (see below) had negligible effects on interbank funding spreads when compared with the larger effects of the TAF.
4. Primary Dealer Credit Facility (PDCF), instituted after the near-failure of Bear Stearns, provided overnight credit to dealers. See Adrian and Schaumburg (2012) for a discussion.
5. Asset-Backed Commercial Paper and Money Market Liquidity Facility (AMLF) provided collateralized loans to depository institutions willing to purchase ABCP from money market funds. Duygan-Bump and others (2010) find that the program helped stabilize money market funds and improved liquidity in the ABCP market.
6. Swap lines by the Fed with foreign central banks. Goldberg, Kennedy, and Miu (2011) summarize the evidence on the effectiveness of the swap lines, finding that their establishment reduced funding pressures abroad and domestically.
7. Term Asset-Backed Securities Loan Facility (TALF). A joint Fed–Treasury operation, TALF involved Fed loans to holders of AAA-rated ABS. The Fed lent the market value of the ABS less a haircut, and received \$20 billion in credit protection through TARP from the Treasury. Covitz, Liang, and Suarez (2013) provide some evidence that TALF aided ABS market confidence.
8. Commercial Paper Funding Facility (CPFF). A vehicle through which the Fed purchased highly rated unsecured and asset-backed commercial paper, secured either via assets or issuer fees. Adrian, Kimbrough, and Marchioni (2011) describe the program and document associated declines in spreads for the classes of purchased paper.
9. Money Market Investor Funding Facility (MMIF). A complement to AMLF, the MMIF aimed to provide liquidity to the secondary money market. However, it was never drawn upon.
10. Temporary Guarantee Program for money market funds (TGP). To stop the run on MMFs, the Treasury Department guaranteed share prices of participating funds.
11. Temporary Liquidity Guarantee Program (TLGP). Under this program, the Federal Deposit Insurance Corporation insured new senior unsecured debt of depository institutions and their holding companies and guaranteed non-interest-bearing transactions accounts in full.
12. Troubled Asset Relief Program (TARP). Under TARP, Congress authorized up to \$700 billion to acquire troubled assets. Funds were used for capital injections in financial institutions, as well as for mortgage relief and to stabilize automobile companies.
13. Capital purchase program (CPP). Used TARP funds to put capital into both large and small banks.

Box 1. Policy Responses to the Panic (*Continued*)

14. MBS purchase program. A precursor to quantitative easing, under this program the Fed purchased mortgage-related securities issued or guaranteed by the government-sponsored enterprises. Hancock and Passmore (2010) found that the program lowered mortgage rates significantly in late 2008.

15. Stress tests (Supervisory Capital Assessment Program). A joint effort by the Fed, the Office of the Comptroller of the Currency, and the Federal Deposit Insurance Corporation, with the backing of the Treasury, this program evaluated the ability of large banks to withstand stress scenarios. Banks that failed the tests were required to raise private capital or accept capital from TARP. See Clark and Ryu (2015) for a description. Morgan, Peristiani, and Savino (2014) study the relationship between stress test announcements and bank stock returns.

16. Public–Private Investment Program (PPIP). In this program, the Treasury committed equity and debt financing to public–private funds that would acquire “legacy” residential and commercial MBS.

market fund holding Lehman Brothers commercial paper “broke the buck,” a broad-based run developed in the sector, to which the Treasury responded with a guarantee program and the Fed with new liquidity programs. Increasingly, however, funding concerns were morphing into solvency problems, with investors losing faith in a number of large institutions (Sarkar and Shrader 2010). The policy responses during this period evolved accordingly. Importantly, passage of the Troubled Assets Relief Program (TARP) legislation gave the Treasury the resources to put capital into the banking system, through its Capital Purchase Program; it would later use TARP funds also to support mortgage modifications and to prevent the failure of two large automobile companies. Two additional steps helped to stabilize the banking system: the guarantees of new senior bank debt by the FDIC, through its Temporary Liquidity Guarantee Program, and the stress tests of the banks conducted by the regulators (with the support of the Treasury) in the spring of 2009. The Fed and the Treasury also collaborated to support the ABS market through the Term Asset-Backed Securities Loan Facility (TALF) program.

A substantial body of literature has evaluated the various programs, in most cases finding that they worked as intended (see box 1 for selected references; also, for an overview, see Logan, Nelson, and Parkinson 2018). Many of these articles rely on event studies, however, which do not always give sharp results. In this vein, we matched up the dates of significant policy announcements or policy implementations with our estimated daily factors, looking for evidence that particular policies were linked to sharp

movements in one or more of the factors. We found some evidence of beneficial effects of some specific policies, including the Capital Purchase Program, the FDIC's loan guarantee program, the guarantee of money market funds, the announcement of stress test results, the Term Securities Lending Facility, and TALF. However, the results were not always robust, reflecting the usual difficulties in assessing the extent to which program announcements surprised markets, along with the fact that many programs were introduced at similar times and in the presence of confounding developments in financial markets.¹³ More research, preferably in the context of a consistent overarching framework, is needed to ascertain the relative importance and effectiveness of the various policies brought to bear during the crisis.

The gross fact, however, which is apparent in figure 13, is that the panic was brought under control relatively quickly. Funding conditions were substantially improved by the end of 2008, as is evident from the middle panel of figure 13. As the figure's top panel shows, stresses in nonmortgage credit markets continued into 2009, but following interventions—including the introduction of TALF and the successful stress testing of the banks—that aspect of the panic subsided as well. Given the results of this paper, which show the strong association of the panic factors and the economy, the suite of policies that controlled the panic likely prevented a much deeper recession than (the already very severe) downturn that we suffered.¹⁴

Looking forward, the findings of this paper argue for continued vigilance in ensuring financial stability. The costs of a financial crisis, particularly one that includes a sustained financial panic, are very high. Policymakers should err on the side of conservatism in ensuring that financial institutions are well capitalized, do not rely excessively on short-term funding, and have good systems for measuring and managing risk. Regulators should work to shine a light on the “dark corners” of the financial system and to take a systemic or macroprudential approach to thinking about risks. Although healthy debates continue—for example, on the appropriate level of bank capital—I think postcrisis reforms have significantly improved the resilience of the U.S. financial system to future shocks.

13. It is also sometimes difficult to identify when a program was “introduced”—for example, when it was announced, when it was implemented, or when its terms or size were changed.

14. Using a macroeconomic model with financial frictions, Del Negro and others (2017) conclude that the Fed's liquidity facilities in particular may have prevented a significantly worse economic collapse than occurred.

Even if financial crises are less likely than in the past, policymakers need to have appropriate tools to fight the next crisis, whenever it may occur. On this count, I am somewhat less sanguine. The orderly liquidation authority, created by the Dodd-Frank law, provides policymakers with important new authorities to help wind down a failing, systemic institution in an orderly way. These new liquidation authorities have not been tested, and some have doubts about their efficacy in the context of a systemic panic; but I think, nevertheless, that they are a significant improvement from the improvised authorities available during the last crisis. Other firefighting tools, however, have actually been cut back since the crisis. For example, the Treasury can no longer guarantee money market funds nor can the FDIC guarantee bank debt, as both did to very positive effect during the crisis. The Fed's emergency lending authorities have been limited to some degree; and, importantly, new disclosure requirements have probably stigmatized the discount window and other lending facilities to the point where they might prove useless in a crisis, as even troubled institutions would be reluctant to borrow.

The limitations on firefighting tools mostly reflect a fully understandable political reaction to some of the policy interventions made during the crisis. However, the evidence of this paper supports the view that these interventions were largely necessary to protect the broader economy. I hope that, as time passes, legislators will find it possible to conduct a balanced review and assessment of the tools available to fight the next crisis, to ensure that they will prove effective when needed.

ACKNOWLEDGMENTS I thank Olivier Blanchard, Janice Eberly, Mark Gertler, Gary Gorton, Raghuram Rajan, and James Stock for comments. Sage Belz and Michael Ng provided outstanding research assistance.

APPENDIX

Details on the Data

In the table that follows on the next page, data for the factor model are at a daily frequency, stripped of holidays (when present), and forward filled for missing values. Quarter-end dates for repo data are replaced with the preceding day's value to control for window dressing. All data used in the factor model are standardized over the period from 2006 to 2012 by the z score. The factor analysis estimates four factors using a varimax rotation, with a lower bound on uniqueness for optimization of 0.05. All citations of "Bloomberg" throughout the table, and also in the main text above, refer to "Bloomberg Finance LP." In the table, "Haver" refers to Haver Analytics.

<i>Series</i>	<i>Source</i>	<i>Raw source</i>	<i>Mnemonic form</i>	<i>Adjustment</i>	<i>Notes</i>
Standard & Poor's 500 Homebuilder and subprime lender stock prices	Haver Yahoo Finance; Bloomberg	Standard & Poor's None	SP500@DAILY DHL1, LEN, PHM, KBH, TOL, FNMA; FMCC, AHMI	Indexed to January 2007 Indexed to January 2007, taken as a percentage of S&P 500	In raw data, 1941-43 = 10; not used directly in factor model Using R function <code>quantmod::getSymbols</code> to pull directly from Yahoo Finance API. Subprime Lenders come from Bloomberg. Additional market-cap weighted index (as of 2006:Q4) constructed using CWD, AHMI, NEW
REIT prices	Yahoo Finance; Bloomberg	None	AVB, EQR, UDR	Indexed to January 2007, taken as a percentage of S&P 500	Using R function <code>quantmod::getSymbols</code> to pull directly from Yahoo Finance API. Additional Bloomberg Mortgage REIT index
ABX indexes	Bloomberg	IHS Markit ICE/BofA/ML	ABX [Rating] CDS1 S6-1 PRC Corp FMLSC@DAILY	None	ABX HE BBB 06-01 and ABX AAA HE 06-01 Effective yield
Home equity ABS yields	Haver	ICE	FLOD3@DAILY	For floating rate, spread over 10-year Treasury yield; LIBOR-OIS	
LIBOR	Haver	Federal Reserve	FLCPD@DAILY	Spread over 1-month OIS	
Commercial paper rates	Haver	None			Yields in percent per annum based on USD; include 3-month, 1-month, and 1-week Yields in percent per annum; include 1-day, 7-day, 15-day, 30-day for AA financial, AA asset-backed, and A2P2 nonfinancial
Treasury yields	Haver	Wall Street Journal	F30JON@DAILY	None	Yields in percent per annum; on-the-run Treasury; not used directly in factor model
OIS rate	Bloomberg	OTC composite (NY)	USSO[maturity]	None	USD overnight index fixed/float rate swap, where the index rate is the federal funds rate; Not used directly in factor model
GCF repo rates	DTCC	None	None	MBS and Agency rates taken as a spread over Treasury rate	MBS, Agency, and Treasury collateral; par weighted average based on overnight trades in dollars; http://www.dtc.com/charts/dtcc-gcf-repo-index
Bank CDS spreads	Bloomberg	IHS Markit	[Ticker] CDS USD SR 5Y D14 Corp	None	Senior 5-year D14 single name CDS for JPM, BofA, WFC, Citi, HBSC, GS, MS, DB, ML, UBS, CS
Bank Stock Prices	Bloomberg	None	JPM, BAC, C, HSB, WFC, GS, MS, UBS, DB, CS	Indexed to January 2007, taken as a percent of S&P 500	
Corporate bond yields	Haver	ICE/BofA/ML	FMLC3A@DAILY	Spread over 10-year Treasury yield	Effective Yield; using total corporate bonds by rating: AAA, AA, A, BBB, HY Master II, BB, B, CCC, BB/B
ABS yields	Haver	ICE/BofA/ML	FMLSC@DAILY	For fixed rate, spread over 10-year Treasury yield; for floating rate, spread over 3-month LIBOR	Effective yield; using fixed and floating rates; student loan (float only), credit card (fixed and floating), auto (fixed and floating)
ABS indexes	Bloomberg	Bloomberg/Barclays	LA [asset] TRU1 (raw index)	None	Bloomberg/Barclays Option-Adjusted Spread; using ABS auto and credit card total return indexes
CDX indexes	Bloomberg	IHS Markit	CDX [rating] CDS1 GEN 5Y PRC Corp	None	CDX IG in percent, CDX HY in index level, 5-year maturity
GDP	Haver	Macroeconomic Advisers	MGDP@USECON	YoY Percent Growth	
Industrial production	Haver	Federal Reserve	IP@USECON	YoY Percent Growth	Monthly, seasonally adjusted household data
Employment, excluding construction)	Haver	Bureau of Labor Statistics (BLS)	EA16@EMPL.; EAC16@EMPL	YoY Percent Growth	New residential construction, table 3; New, privately owned housing units started, seasonally adjusted at annual rates
Housing starts	Haver	Department of Commerce	HST@USECON	YoY Percent Growth	Household Survey, table A
Unemployment	Haver	BLS	LR@USECON	Percent (level)	Monthly retail sales and food services; adjusted for seasonal, holiday, and trading-day differences
Retail sales	Haver	Department of Commerce	NRST@USECON	YoY Percent Growth	Excluding food and energy
Core PCE inflation	Haver	Bureau of Economic Analysis (BEA)	JCXFEBM@USECON	YoY Inflation	Monthly, seasonally adjusted house price purchase-only index
Home prices	Haver	Federal Housing Finance Agency	USPHPM@USECON	YoY Percent Growth	
Real PCE	Haver	BEA	CBHM@USECON	YoY Percent Growth	Monthly, seasonally adjusted annual rates; personal income and outlays, tables 1 and 7
Real PCE—durables	Haver	BEA	CDBHM@USECON	YoY Percent Growth	Monthly, seasonally adjusted annual rates; personal income and outlays, tables 1 and 7
Capital goods orders (nondefense, excluding aircraft)	Haver	Census	NMOCNX@USECON	YoY Percent Growth	Monthly Advance Report on Durable Goods (NAICS), seasonally adjusted
ISM Composite Index	Haver	Institute for Supply Management	ISM@USECON	Percent (index level), seasonally adjusted	
Delinquencies	Haver	Fannie Mae	FNMSD@USECON	Percent (rate)	Values above 50 indicate growth Conventional Single-Family Serious Delinquency Rate (Fannie Mae Monthly Summary, table 8)

References

- Acharya, Viral V., and Nada Mora. 2015. "A Crisis of Banks as Liquidity Providers." *Journal of Finance* 70, no. 1: 1–43. <https://doi.org/10.1111/jofi.12182>.
- Adelino, Manuel, Antoinette Schoar, and Felipe Severino. 2015. "House Prices, Collateral, and Self-Employment." *Journal of Financial Economics* 117, no. 2: 288–306.
- Adrian, Tobias, Paolo Colla, and Hyun Song Shin. 2012. "Which Financial Frictions? Parsing the Evidence from the Financial Crisis of 2007–9." *National Bureau of Economic Research Macroeconomics Annual* 27 (August): 159–214.
- Adrian, Tobias, Karin Kimbrough, and Dina Marchioni. 2011. "The Federal Reserve's Commercial Paper Funding Facility." *FRBNY Economic Policy Review*. Federal Reserve Bank of New York. <https://www.newyorkfed.org/medialibrary/media/research/epr/11v17n1/1105adri.pdf>.
- Adrian, Tobias, and Ernst Schaumburg. 2012. "The Fed's Emergency Liquidity Facilities during the Financial Crisis: The CPFF." Federal Reserve Bank of New York. *Liberty Street Economics* (blog). August 20, 2012. <http://libertystreeteconomics.newyorkfed.org/2012/08/the-feds-emergency-liquidity-facilities-during-the-financial-crisis-the-cpff.html>.
- Afonso, Gara, Anna Kovner, and Antoinette Schoar. 2011. "Stressed, Not Frozen: The Federal Funds Market in the Financial Crisis." Federal Reserve Bank of New York Staff Report. https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr437.pdf.
- Aikman, David, Jonathan Bridges, Anil Kashyap, and Caspar Siegart. 2018. "Would Macroprudential Regulation Have Prevented the Last Crisis?" Bank of England Staff Working Paper 747 (August). <https://www.bankofengland.co.uk/working-paper/2018/would-macroprudential-regulation-have-prevented-the-last-crisis>.
- Aiyar, Shekhar. 2011. "How Did the Crisis in International Funding Markets Affect Bank Lending? Balance Sheet Evidence from the United Kingdom." Bank of England Staff Working Paper (April). <http://www.bankofengland.co.uk/research/Pages/workingpapers/2011/wp424.aspx>.
- . 2012. "From Financial Crisis to Great Recession: The Role of Globalized Banks." *American Economic Review* 102, no. 3: 225–30. <https://doi.org/10.1257/aer.102.3.225>.
- Aladangady, Aditya. 2014. "Homeowner Balance Sheets and Monetary Policy." Federal Reserve, Finance and Economics Discussion Series 9. <https://www.federalreserve.gov/econresdata/feds/2014/files/201498pap.pdf>.
- Albertazzi, Ugo, and Domenico J. Marchetti. 2010. "Credit Supply, Flight to Quality and Evergreening: An Analysis of Bank-Firm Relationships after Lehman." Economic Research and International Relations Area 756. Temi Di Discussione (Economic Working Papers), Bank of Italy. https://ideas.repec.org/p/bdi/wptemi/td_756_10.html.
- Alfaro, Laura, Manuel García-Santana, and Enrique Moral-Benito. 2018. "On the Direct and Indirect Real Effects of Credit Supply Shocks." SSRN Scholarly

- Paper ID 3132984. Rochester: Social Science Research Network. <https://papers.ssrn.com/abstract=3132984>.
- Almeida, Heitor, Murillo Campello, Bruno Laranjeira, and Scott Weisbenner. 2009. "Corporate Debt Maturity and the Real Effects of the 2007 Credit Crisis." Working Paper 14990. Cambridge, Mass.: National Bureau of Economic Research. <https://doi.org/10.3386/w14990>.
- Altavilla, Carlo, Matthieu Darracq Paries, and Giulio Nicoletti. 2015. "Loan Supply, Credit Markets and the Euro Area Financial Crisis." European Central Bank Working Paper 1861. European Central Bank. <https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp1861.en.pdf>.
- Arellano, Cristina, Yan Bai, and Patrick J. Kehoe. 2016. "Financial Frictions and Fluctuations in Volatility." Working Paper 22990. Cambridge, Mass.: National Bureau of Economic Research. <https://doi.org/10.3386/w22990>.
- Armantier, Olivier, Eric Ghysels, Asani Sarkar, and Jeffrey Shrader. 2015. "Discount Window Stigma during the 2007–2008 Financial Crisis." Federal Reserve Bank of New York Staff Report. https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr483.pdf.
- Armantier, Olivier, Sandy Krieger, and James McAndrews. 2008. "The Federal Reserve's Term Auction Facility." *Current Issues in Economics and Finance* 14, no. 5. https://www.newyorkfed.org/research/current_issues/ci14-5.html.
- Arteta, Carlos, Mark Carey, Ricardo Correa, and Jason Kotter. 2013. "Revenge of the Steamroller: ABCP as a Window on Risk Choices." *International Finance Discussion Papers* 1076 (April). <https://www.federalreserve.gov/pubs/ifdp/2013/1076/ifdp1076.pdf>.
- Bacchetta, Philippe, and Eric van Wincoop. 2016. "The Great Recession: A Self-Fulfilling Global Panic." *American Economic Journal: Macroeconomics* 8, no. 4: 177–98. <https://doi.org/10.1257/mac.20140092>.
- Baker, Scott R. 2018. "Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data." *Journal of Political Economy*, March, 1504–57. <https://doi.org/10.1086/698106>.
- Bao, Jack, Josh David, and Song Han. 2015. "The Runnables." *FEDS Notes* (blog). September 3, 2015. <https://www.federalreserve.gov/econresdata/notes/feds-notes/2015/the-runnables-20150903.html>.
- Bassett, William F., Mary Beth Chosak, John C. Driscoll, and Egon Zakrajšek. 2014. "Changes in Bank Lending Standards and the Macroeconomy." *Journal of Monetary Economics* 62 (March): 23–40.
- Benmelech, Efraim, Carola Frydman, and Dimitris Papanikolaou. 2017. "Financial Frictions and Employment during the Great Depression." Working Paper 23216. Cambridge, Mass.: National Bureau of Economic Research. <https://doi.org/10.3386/w23216>.
- Benmelech, Efraim, Ralf R. Meisenzahl, and Rodney Ramcharan. 2017. "The Real Effects of Liquidity During the Financial Crisis: Evidence from Automobiles." *Quarterly Journal of Economics* 132, no. 1: 317–65. <https://doi.org/10.3386/w22148>.

- Bernanke, Ben S. 1983. "Non-Monetary Effects of the Financial Crisis in the Propagation of the Great Depression." *American Economic Review* 73, no. 10: 257–76.
- . 1994. "The Macroeconomics of the Great Depression: A Comparative Approach." Working Paper 4814. Cambridge, Mass.: National Bureau of Economic Research. <https://doi.org/10.3386/w4814>.
- . 2010. "Implications of the Financial Crisis for Economics." Speech presented at the Conference Cosponsored by the Center for Economic Policy Studies and the Bendheim Center for Finance, Princeton, N.J., September 24. <https://www.federalreserve.gov/newsevents/speech/bernanke20100924a.htm>.
- . 2012. "Some Reflections on the Crisis and the Policy Response." Paper presented at Russell Sage Foundation and Century Foundation Conference on "Rethinking Finance," New York. <https://www.federalreserve.gov/newsevents/speech/bernanke20120413a.htm>.
- Bernanke, Ben S., and Mark Gertler. 1989. "Agency Costs, Net Worth, and Business Fluctuations." *American Economic Review* 79, no. 1: 14–31.
- . 1995. "Inside the Black Box: The Credit Channel of Monetary Policy Transmission." *Journal of Economic Perspectives* 9, no. 4: 27–48. <https://doi.org/10.1257/jep.9.4.27>.
- Bernanke, Ben S., Mark Gertler, and Simon Gilchrist. 1999. "The Financial Accelerator in a Quantitative Business Cycle Framework." *Handbook of Macroeconomics* 1, issue C: 1341–93.
- Bernanke, Ben S., and Harold James. 1991. "The Gold Standard, Deflation, and Financial Crisis in the Great Depression: An International Comparison." In *Financial Markets and Financial Crises*. University of Chicago Press for National Bureau of Economic Research. <http://www.nber.org/chapters/c11482.pdf>.
- Berrospide, Jose M., Lamont K. Black, and William R. Keaton. 2016. "The Cross-Market Spillover of Economic Shocks through Multi-Market Banks." *Journal of Money, Credit and Banking* 48, no. 5: 957–88.
- Boissay, Frédéric, Fabrice Collard, and Frank Smets. 2016. "Booms and Banking Crises." *Journal of Political Economy* 124, no. 2: 489–538. <https://doi.org/10.1086/685475>.
- Brunnermeier, Markus K. 2009. "Deciphering the Liquidity and Credit Crunch 2007–2008." *Journal of Economic Perspectives* 23, no. 1: 77–100.
- Brunnermeier, Markus K., and Yuliy Sannikov. 2014. "A Macroeconomic Model with a Financial Sector." *American Economic Review* 104, no. 2: 379–421. <https://doi.org/10.1257/aer.104.2.379>.
- Caballero, Ricardo J., Emmanuel Farhi, and Pierre-Olivier Gourinchas. 2017. "The Safe Assets Shortage Conundrum." *Journal of Economic Perspectives* 31, no. 3: 29–46.
- Calomiris, Charles W., and Charles M. Kahn. 1991. "The Role of Demandable Debt in Structuring Optimal Banking Arrangements." *American Economic Review* 81, no. 3: 497–513.
- Calomiris, Charles, and Joseph R. Mason. 2003. "Consequences of Bank Distress During the Great Depression." *American Economic Review* 93, no. 3: 937–47.

- Campello, Murillo, John R. Graham, and Campbell R. Harvey. 2010. "The Real Effects of Financial Constraints: Evidence from a Financial Crisis." *Journal of Financial Economics* 97, no. 3: 470–87. <https://doi.org/10.1016/j.jfineco.2010.02.009>.
- Carlson, Mark, and Jonathan D. Rose. 2015. "Credit Availability and the Collapse of the Banking Sector in the 1930s." *Journal of Money, Credit and Banking* 47, no. 7: 1239–71. <https://doi.org/10.1111/jmcb.12244>.
- Carlson, Mark A., Hui Shan, and Missaka Warusawitharana. 2013. "Capital Ratios and Bank Lending: A Matched Bank Approach." *Journal of Financial Intermediation* 22, no. 4: 663–87.
- Carlstrom, Charles T., and Timothy S. Fuerst. 1998. "Agency Costs, Net Worth, and Business Fluctuations: A Computable General Equilibrium Analysis." *American Economic Review* 87, no. 5: 893–910.
- Chaney, Thomas, David Sraer, and David Thesmar. 2012. "The Collateral Channel: How Real Estate Shocks Affect Corporate Investment." *American Economic Review* 102, no. 6: 2381–409. <https://doi.org/10.1257/aer.102.6.2381>.
- Chen, Brian S., Samuel G. Hanson, and Jeremy C. Stein. 2017. "The Decline of Big-Bank Lending to Small Business: Dynamic Impacts on Local Credit and Labor Markets." Working Paper 23843. Cambridge, Mass.: National Bureau of Economic Research. <https://doi.org/10.3386/w23843>.
- Chodorow-Reich, Gabriel. 2014. "Employment Effects of Credit Market Disruptions: Firm-Level Evidence from the 2008–9 Financial Crisis." *Quarterly Journal of Economics* 129, no. 1: 1–59.
- Christiano, Lawrence J., Martin S. Eichenbaum, and Mathias Trabandt. 2014. "Understanding the Great Recession." *American Economic Journal: Macroeconomics* 7, no. 1: 110–67.
- Clark, Tim P., and Lisa H. Ryu. 2015. "CCAR and Stress Testing as Complementary Supervisory Tools." Board of Governors of the Federal Reserve System, June 24. <https://www.federalreserve.gov/bankinforeg/ccar-and-stress-testing-as-complementary-supervisory-tools.htm>.
- Cohen, Jon, Kinda Cheryl Hachem, and Gary Richardson. 2017. "Relationship Lending and the Great Depression: New Measurement and Implications." Working Paper 22891. Cambridge, Mass.: National Bureau of Economic Research. <https://doi.org/10.3386/w22891>.
- Copeland, Adam, Antoine Martin, and Michael Walker. 2010. "The Tri-Party Repo Market Before the 2010 Reforms." Federal Reserve Bank of New York Staff Report. <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.363.5485&rep=rep1&type=pdf>.
- Cornett, Marcia Millon, Jamie John McNutt, Philip E. Strahan, and Hassan Tehrani. 2011. "Liquidity Risk Management and Credit Supply in the Financial Crisis." *Journal of Financial Economics* 101, no. 2: 297–312. <https://doi.org/10.1016/j.jfineco.2011.03.001>.
- Correa, Ricardo, Horacio Saprizza, and Andrei Zlate. 2013. "Liquidity Shocks, Dollar Funding Costs, and the Bank Lending Channel during the European

- Sovereign Crisis.” *International Finance Discussion Papers* 1059r (July). <https://www.federalreserve.gov/pubs/ifdp/2012/1059/ifdp1059r.pdf>.
- Covitz, Daniel M., Nellie Liang, and Gustavo A. Suarez. 2009. “The Evolution of a Financial Crisis: Panic in the Asset-Backed Commercial Paper Market.” Finance and Economic Discussion Series. Washington: Federal Reserve Board. <https://www.federalreserve.gov/pubs/feds/2009/200936/200936pap.pdf>.
- . 2013. “The Evolution of a Financial Crisis: Collapse of the Asset-Backed Commercial Paper Market.” *Journal of Finance* 68, no. 3: 815–48.
- Dagher, Jihad, and Kazim Kazimov. 2015. “Banks’ Liability Structure and Mortgage Lending During the Financial Crisis.” *Journal of Financial Economics* 116, no. 3: 565–82. <https://doi.org/10.1016/j.jfineco.2015.02.001>.
- Dang, Tri Vi, Gary Gorton, and Bengt Holmstrom. 2015. “The Information Sensitivity of a Security.” http://www.columbia.edu/~td2332/Paper_Sensitivity.pdf.
- . 2018. “Ignorance, Debt, and Financial Crises.” http://www.columbia.edu/~td2332/Paper_Ignorance.pdf.
- De Haas, Ralph, and Neeltje Van Horen. 2012. “International Shock Transmission after the Lehman Brothers Collapse: Evidence from Syndicated Lending.” *American Economic Review* 102, no. 3: 231–37. <https://doi.org/10.1257/aer.102.3.231>.
- Dell’Ariccia, Giovanni, Enrica Detragiache, and Raghuram Rajan. 2005. “The Real Effect of Banking Crises.” *Journal of Financial Intermediation* 17, no. 1: 89–112.
- Del Negro, Marco, Gauti Eggertsson, Andrea Ferrero, and Nobuhiro Kiyotaki. 2017. “The Great Escape? A Quantitative Evaluation of the Fed’s Liquidity Facilities.” *American Economic Review* 107, no. 3: 824–57.
- Del Negro, Marco, Raiden Hasegawa, and Frank Schorfheide. 2016. “Dynamic Prediction Pools: An Investigation of Financial Frictions and Forecasting Performance.” *Journal of Econometrics* 192: 391–405.
- Diamond, Douglas W., and Philip H. Dybvig. 1983. “Bank Runs, Deposit Insurance, and Liquidity.” *Journal of Political Economy* 91, no. 3: 401–19.
- Drechsler, Itamar, Alexi Savov, and Philipp Schnabl. 2018. “A Model of Monetary Policy and Risk Premia.” *Journal of Finance* 73, no. 1: 317–73. <https://doi.org/10.1111/jofi.12539>.
- Duchin, Ran, Oguzhan Ozbas, and Berk A. Sensoy. 2010. “Costly External Finance, Corporate Investment, and the Subprime Mortgage Credit Crisis.” *Journal of Financial Economics* 97, no. 3: 418–35. <https://doi.org/10.1016/j.jfineco.2009.12.008>.
- Duygan-Bump, Burcu, Alexey Lekov, and Judit Montoriol-Garriga. 2015. “Financing Constraints and Unemployment: Evidence from the Great Recession.” *Journal of Monetary Economics* 75 (October): 89–105.
- Duygan-Bump, Burcu, Patrick M. Parkinson, Eric S. Rosengren, Gustavo A. Suarez, and Paul S. Willen. 2010. “How Effective Were the Federal Reserve Emergency Liquidity Facilities? Evidence from the Asset-Backed Commercial Paper Money Market Mutual Fund Liquidity Facility.” QAU10-3. Federal Reserve

- Bank of Boston Working Paper. Federal Reserve Bank of Boston. <https://www.bostonfed.org/publications/risk-and-policy-analysis/2010/how-effective-were-the-federal-reserve-emergency-liquidity-facilities-evidence-from-the-asset-backed-commercial-paper-money-market-mutual-fund-liquidity-facility.aspx>.
- Eberly, Janice, and Arvind Krishnamurthy. 2014. "Efficient Credit Policies in a Housing Debt Crisis." *Brookings Papers on Economic Activity*, Fall: 73–118.
- Eggertsson, Gauti B., and Paul Krugman. 2012. "Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach." *Quarterly Journal of Economics* 127, no. 3: 1469–1513. <https://doi.org/10.1093/qje/qjs023>.
- Falato, Antonio, and Nellie Liang. 2016. "Do Creditor Rights Increase Employment Risk? Evidence from Loan Covenants." *Journal of Finance* 71, no. 6: 2545–90. <https://doi.org/10.1111/jofi.12435>.
- Favara, Giovanni, Simon Gilchrist, Kurt F. Lewis, and Egon Zakrajšek. 2016. "Updating the Recession Risk and the Excess Bond Premium," *FEDS Notes* (blog). October 6. <https://www.federalreserve.gov/econresdata/notes/feds-notes/2016/updates-the-recession-risk-and-the-excess-bond-premium-20161006.html>.
- Favilukis, Jack, Sydney C. Ludvigson, and Stijn Van Nieuwerburgh. 2010. "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium." Working Paper 15988. Cambridge, Mass.: National Bureau of Economic Research. <https://doi.org/10.3386/w15988>.
- Fazzari, Steven M., R. Glenn Hubbard, and Bruce C. Petersen. 1988. "Financing Constraints and Corporate Investment." *Brookings Papers on Economic Activity* 1988, no. 1: 141–206. <https://doi.org/10.2307/2534426>.
- Fleming, Michael J., Warren B. Hrung, and Frank M. Keane. 2010. "Repo Market Effects of the Term Securities Lending Facility." *American Economic Review* 100, no. 2: 591–96.
- Ganong, Peter, and Pascal Noel. 2018. "Liquidity vs. Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession." Working Paper 24964. Cambridge, Mass.: National Bureau of Economic Research.
- Geanakoplos, John. 2010. "The Leverage Cycle." *NBER Macroeconomics Annual* 2009 24 (April): 1–65.
- Gertler, Mark, and Simon Gilchrist. 2018. "What Happened: Financial Factors in the Great Recession." Working Paper 24746. Cambridge, Mass.: National Bureau of Economic Research. <https://doi.org/10.3386/w24746>.
- Gertler, Mark, and Peter Karadi. 2011. "A Model of Unconventional Monetary Policy." *Journal of Monetary Economics* 58, no. 1: 17–34.
- Gertler, Mark, and Nobuhiro Kiyotaki. 2015. "Banking, Liquidity, and Bank Runs in an Infinite Horizon Economy." *American Economic Review* 105, no. 7: 2011–43. <https://doi.org/10.1257/aer.20130665>.
- Gertler, Mark, Nobuhiro Kiyotaki, and Andrea Prestipino. 2017. "A Macroeconomic Model with Financial Panics." Working Paper 24126. Cambridge, Mass.: National Bureau of Economic Research. <https://doi.org/10.3386/w24126>.

- Gilchrist, Simon, Raphael Schoenle, Jae Sim, and Egon Zakrajšek. 2017. "Inflation Dynamics during the Financial Crisis." *American Economic Review* 107, no. 3: 785–823. <https://doi.org/10.1257/aer.20150248>.
- Gilchrist, Simon, and Egon Zakrajšek. 2012a. "Credit Spreads and Business Cycle Fluctuations." *American Economic Review* 102, no. 4: 1692–1720. <https://doi.org/10.1257/aer.102.4.1692>.
- . 2012b. "Credit Supply Shocks and Economic Activity in a Financial Accelerator Model." In *Rethinking the Financial Crisis*, edited by Alan S. Blinder, Andrew W. Loh, and Robert M. Solow. New York: Russell Sage Foundation.
- Giroud, Xavier, and Holger M. Mueller. 2017. "Firm Leverage, Consumer Demand, and Employment Losses During the Great Recession." *Quarterly Journal of Economics* 132, no. 1: 271–316. <https://doi.org/10.1093/qje/qjw035>.
- Gissler, Stefan, and Borghan Narajabad. 2017. "The Increased Role of the Federal Home Loan Bank System in Funding Markets, Part 1: Background." *FEDS Notes* (blog). October 18. <https://www.federalreserve.gov/econres/notes/feds-notes/the-increased-role-of-the-federal-home-loan-bank-system-in-funding-markets-part-1-background-20171018.htm>.
- Goetz, Martin, and Juan Gozzi. 2010. "Liquidity Shocks, Local Banks, and Economic Activity: Evidence from the 2007–2009 Crisis." SSRN Scholarly Paper ID 1709677. Rochester: Social Science Research Network. <https://papers.ssrn.com/abstract=1709677>.
- Goldberg, Linda S., Craig Kennedy, and Jason Miu. 2011. "Central Bank Dollar Swap Lines and Overseas Dollar Funding Costs." *FRBNY Economic Policy Review*. Federal Reserve Bank of New York. <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.422.11&rep=rep1&type=pdf>.
- Gorton, Gary B. 2008. "The Subprime Panic." Working Paper 14398. Cambridge, Mass.: National Bureau of Economic Research. <http://www.nber.org/papers/w14398.pdf>.
- Gorton, Gary B., and Andrew Metrick. 2012. "Securitized Banking and the Run on Repo." *Journal of Financial Economics* 104, no. 3: 425–51.
- Gorton, Gary, and George Pennacchi. 1990. "Financial Intermediaries and Liquidity Creation." *Journal of Finance* 45, no. 1: 49–71. <https://doi.org/10.2307/2328809>.
- Greenspan, Alan. 2005. "Reflections on Central Banking." Speech presented at Jackson Hole Economic Policy Symposium, Jackson Hole, Wyo., August 26. <https://www.kansascityfed.org/publicat/sympos/2005/pdf/Green-opening2005.pdf>.
- Guerrieri, Luca, and Matteo Iacoviello. 2017. "Collateral Constraints and Macroeconomic Asymmetries." *Journal of Monetary Economics* 90 (October): 28–49.
- Hall, Robert E. 2010. "Why Does the Economy Fall to Pieces after a Financial Crisis?" *Journal of Economic Perspectives* 24, no. 4: 3–20. <https://doi.org/10.1257/jep.24.4.3>.
- . 2011. "The High Sensitivity of Economic Activity to Financial Frictions." *Economic Journal* 121, no. 552: 351–78. <https://doi.org/10.1111/j.1468-0297.2011.02421.x>.

- Haltenhof, Samuel, Seung Jung Lee, and Viktors Stebunovs. 2014. "The Credit Crunch and Fall in Employment during the Great Recession." *Journal of Economic Dynamics and Control* 43 (June): 31–57.
- Hancock, Diana, and Wayne Passmore. 2010. "Did the Federal Reserve's MBS Purchase Program Lower Mortgage Rates?" *Finance and Economics Discussion Series* 2011, no. 1. <https://www.federalreserve.gov/pubs/feds/2011/201101/201101pap.pdf>.
- Hanson, Samuel, and Adi Sunderam. 2013. "Are There Too Many Safe Securities? Securitization and the Incentives for Information Production." *Journal of Financial Economics* 108, no. 3: 565–84.
- Hatzius, Jan. 2008. "Beyond Leveraged Losses: The Balance Sheet Effects of the Home Price Downturn." *Brookings Papers on Economic Activity* 2008, no. 2: 195–228.
- He, Zhiguo, In Gu Khang, and Arvind Krishnamurthy. 2010. "Balance Sheet Adjustments in the 2008 Crisis." Working Paper 15919. Cambridge, Mass.: National Bureau of Economic Research. <https://doi.org/10.3386/w15919>.
- He, Zhiguo, and Arvind Krishnamurthy. 2013. "Intermediary Asset Pricing." *American Economic Review* 103, no. 2: 732–70. <https://doi.org/10.1257/aer.103.2.732>.
- Huber, Kilian. 2018. "Disentangling the Effects of a Banking Crisis: Evidence from German Firms and Counties." *American Economic Review* 108, no. 3: 868–98. <https://doi.org/10.1257/aer.20161534>.
- Iacoviello, Matteo. 2014. "Financial Business Cycles." *Review of Economic Dynamics* 18, no. 1: 140–63.
- Irani, Rustom M., and Ralf R. Meisenzahl. 2014. "Loan Sales and Bank Liquidity Risk Management: Evidence from a U.S. Credit Register." *Finance and Economics Discussion Series* 2015, no. 001. <https://www.federalreserve.gov/econresdata/feds/2015/files/2015001pap.pdf>.
- Ivashina, Victoria, and David Scharfstein. 2009. "Bank Lending During the Financial Crisis of 2008." *Journal of Financial Economics* 97.
- Iyer, Rajkamal, Samuel Lopes, Jose-Luis Peydro, and Antoinette Schoar. 2014. "Interbank Liquidity Crunch and the Firm Credit Crunch: Evidence from the 2007–2009 Crisis." *Review of Financial Studies* 27, no. 1: 347–72. <https://academic.oup.com/rfs/article-abstract/27/1/347/1573768>.
- Jensen, Thais Lærholm, and Niels Johannesen. 2017. "The Consumption Effects of the 2007 & 2008 Financial Crisis: Evidence from Households in Denmark." *American Economic Review* 107, no. 11: 3386–414. <https://doi.org/10.1257/aer.20151497>.
- Jermann, Urban, and Vincenzo Quadrini. 2012. "Macroeconomic Effects of Financial Shocks." *American Economic Review* 102, no. 1: 238–71. <https://doi.org/10.1257/aer.102.1.238>.
- Jones, Callum, Virgiliu Midrigan, and Thomas Philippon. 2018. "Household Leverage and the Recession." Working Paper 16965. Cambridge, Mass.: National Bureau of Economic Research. <https://doi.org/10.3386/w16965>.

- Jordà, Òscar, Moritz Schularick, and Alan M. Taylor. 2016. "Macrofinancial History and the New Business Cycle Facts." Working Paper 22743. Cambridge, Mass.: National Bureau of Economic Research. <https://doi.org/10.3386/w22743>.
- Juselius, Mikael, and Mathias Drehmann. 2015. "Leverage Dynamics and the Real Burden of Debt." BIS Working Paper, Bank for International Settlements. <https://www.bis.org/publ/work501.htm>.
- Kacperczyk, Marcin, and Philipp Schnabl. 2010. "When Safe Proved Risky: Commercial Paper during the Financial Crisis of 2007–2009." *Journal of Economic Perspectives* 24, no. 1: 29–50.
- Kahle, Kathleen M., and René M. Stulz. 2013. "Access to Capital, Investment, and the Financial Crisis." *Journal of Financial Economics* 110, no. 2: 280–99. <https://doi.org/10.1016/j.jfineco.2013.02.014>.
- Kandrac, John. 2014. "Bank Failure, Relationship Lending, and Local Economic Performance." Finance and Economics Discussion Series 41. <https://www.federalreserve.gov/pubs/feds/2014/201441/201441pap.pdf>.
- Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante. 2016. "Non-Durable Consumption and Housing Net Worth in the Great Recession: Evidence from Easily Accessible Data." Working Paper 22232. Cambridge, Mass.: National Bureau of Economic Research. <https://doi.org/10.3386/w22232>.
- Kashyap, Anil, Jeremy C. Stein, and David W. Wilcox. 1993. "Monetary Policy and Credit Conditions: Evidence from the Composition of External Finance." *American Economic Review* 83, no. 1: 78–98.
- Keane, Frank M. 2013. "Securities Loans Collateralized by Cash: Reinvestment Risk, Run Risk, and Incentive Issues." *Current Issues in Economics and Finance* 19, no. 3. https://www.newyorkfed.org/medialibrary/media/research/current_issues/ci19-3.pdf.
- Kennickell, Arthur B., Myron L. Kwast, and Jonathan Pogach. 2015. "Small Businesses and Small Business Finance during the Financial Crisis and the Great Recession: New Evidence from the Survey of Consumer Finances." Finance and Economics Discussion Series 039. <https://www.federalreserve.gov/econresdata/feds/2015/files/2015039pap.pdf>.
- Kiyotaki, Nobuhiro, and John Moore. 1997. "Credit Cycles." *Journal of Political Economy* 105, no. 2: 211–48. <https://doi.org/10.1086/262072>.
- Kohn, Donald, and Brian Sack. 2018. "Monetary Policy During the Financial Crisis." Paper presented at conference sponsored by Hutchins Center on Fiscal and Monetary Policy at Brookings on "Responding to the Global Financial Crisis: What We Did and Why We Did It," Washington, September 11. <https://www.brookings.edu/wp-content/uploads/2018/08/13-Monetary-Policy-Prelim-Disc-Draft-2018.09.11.pdf>.
- Krishnamurthy, Arvind. 2010. "Amplification Mechanisms in Liquidity Crises." *American Economic Journal: Macroeconomics* 2: 1–30.
- Laeven, Luc, and Fabian Valencia. 2013. "The Real Effects of Financial Sector Interventions during Crises." *Journal of Money, Credit and Banking* 45, no. 1: 147–77.

- Levin, Andrew T., Fabio M. Natalucci, and Egon Zakrajšek. 2004. "The Magnitude and Cyclical Behavior of Financial Market Frictions." Finance and Economic Discussion Series 2004–70. Federal Reserve Board of Governors. <https://www.federalreserve.gov/pubs/feds/2004/200470/200470pap.pdf>.
- Logan, Lorie, William Nelson, and Patrick Parkinson. 2018. "Novel Lender-of-Last-Resort Programs." Paper presented at conference sponsored by Hutchins Center on Fiscal and Monetary Policy at Brookings on "Responding to the Global Financial Crisis: What We Did and Why We Did It," Washington, September 11. <https://www.brookings.edu/wp-content/uploads/2018/08/02-Novol-LOLR-Prelim-Disc-Draft-2018.09.11.pdf>.
- Longstaff, Francis A. 2010. "The Subprime Credit Crisis and Contagion in Financial Markets." *Journal of Financial Economics* 97, no. 3: 436–50. <https://doi.org/10.1016/j.jfineco.2010.01.002>.
- Mach, Traci L., and John D. Wolken. 2012. "Examining the Impact of Credit Access on Small Firm Survivability." Finance and Economics Discussion Series 10. <https://www.federalreserve.gov/pubs/feds/2012/201210/201210pap.pdf>.
- Manconi, Alberto, Massimo Massa, and Ayako Yasuda. 2012. "The Role of Institutional Investors in Propagating the Crisis of 2007–2008." *Journal of Financial Economics* 104, no. 3: 491–518. <https://doi.org/10.1016/j.jfineco.2011.05.011>.
- McAndrews, James, Asani Sarkar, and Zhenyu Wang. 2017. "The Effect of the Term Auction Facility on the London Interbank Offered Rate." *Journal of Banking & Finance* 83 (October): 135–52.
- McCabe, Patrick E. 2010. "The Cross Section of Money Market Fund Risks and Financial Crises." Finance and Economic Discussion Series. Federal Reserve Board of Governors. <https://www.federalreserve.gov/PUBS/feds/2010/201051/201051pap.pdf>.
- Merton, Robert C. 1974. "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates." *Journal of Finance* 29, no. 2: 449–70. <https://doi.org/10.2307/2978814>.
- Mian, Atif, Kamalesh Rao, and Amir Sufi. 2013. "Household Balance Sheets, Consumption, and the Economic Slump." *Quarterly Journal of Economics* 128, no. 4: 1687–1726. <https://doi.org/10.1093/qje/qjt020>.
- Mian, Atif, and Amir Sufi. 2010. "Household Leverage and the Recession of 2007 to 2009." *IMF Economic Review* 58, no. 1: 74–117. <https://doi.org/10.3386/w15896>.
- . 2014a. *House of Debt*. Chicago: University of Chicago Press. <https://www.press.uchicago.edu/ucp/books/book/chicago/H/bo20832545.html>.
- . 2014b. "What Explains the 2007–2009 Drop in Employment?" *Econometrica* 82, no. 6: 2197–223.
- . 2018a. "Credit Supply and Housing Speculation." Working Paper 24823. Cambridge, Mass.: National Bureau of Economic Research. <https://doi.org/10.3386/w24823>.
- . 2018b. "Finance and Business Cycles: The Credit-Driven Household Demand Channel." Working Paper 24322. Cambridge, Mass.: National Bureau of Economic Research. <https://doi.org/10.3386/w24322>.

- Mishkin, Frederic S. 2011. "Over the Cliff: From the Subprime to the Global Financial Crisis." *Journal of Economic Perspectives* 25, no. 1: 49–70. <https://doi.org/10.1257/jep.25.1.49>.
- Mitchener, Kris James, and Gary Richardson. 2016. "Network Contagion and Interbank Amplification during the Great Depression." Working Paper 22074. Cambridge, Mass.: National Bureau of Economic Research. <https://doi.org/10.3386/w22074>.
- Morgan, Donald P., Stavros Peristiani, and Vanessa Savino. 2014. "The Information Value of the Stress Test." *Journal of Money, Credit and Banking* 46, no. 7: 1479–1500.
- Nolan, Charles, and Christoph Thoenissen. 2009. "Financial Shocks and the U.S. Business Cycle." *Journal of Monetary Economics* 56, no. 4: 596–604.
- Peek, Joe, and Eric Rosengren. 2000. "Collateral Damage: Effects of the Japanese Bank Crisis on Real Activity in the United States." *American Economic Review* 90, no. 1: 30–45. <https://doi.org/10.1257/aer.90.1.30>.
- . 2016. "Credit Supply Disruptions: From Credit Crunches to Financial Crisis." SSRN Scholarly Paper ID 2870864. Rochester: Social Science Research Network. <https://papers.ssrn.com/abstract=2870864>.
- Puri, Manju, Jörg Rocholl, and Sascha Steffen. 2011. "Global Retail Lending in the Aftermath of the U.S. Financial Crisis: Distinguishing Between Supply and Demand Effects." *Journal of Financial Economics* 100, no. 3: 556–78. <https://doi.org/10.1016/j.jfineco.2010.12.001>.
- Rajan, Raghuram G. 2005. "Has Financial Development Made the World Riskier?" Working Paper 11728. Cambridge, Mass.: National Bureau of Economic Research. <https://doi.org/10.3386/w11728>.
- Ramcharan, Rodney, and Raghuram Rajan. 2014. "Financial Fire Sales: Evidence from Bank Failures." Finance and Economics Discussion Series 67. <https://www.federalreserve.gov/pubs/feds/2014/201467/201467pap.pdf>.
- Ramcharan, Rodney, Skander Van den Heuvel, and Stephane H. Verani. 2016. "From Wall Street to Main Street: The Impact of the Financial Crisis on Consumer Credit Supply." *Journal of Finance* 71, no. 3: 1323–56.
- Reinhart, Carmen M., and Kenneth S. Rogoff. 2009. *This Time Is Different: Eight Centuries of Financial Folly*. Princeton University Press.
- Rose, Jonathan. 2015. "Old-Fashioned Deposit Runs." SSRN Scholarly Paper ID 2705730. Rochester: Social Science Research Network. <https://papers.ssrn.com/abstract=2705730>.
- Sahm, Claudia R., Matthew D. Shapiro, and Joel Slemrod. 2015. "Balance-Sheet Households and Fiscal Stimulus: Lessons from the Payroll Tax Cut and Its Expiration." Federal Reserve, Finance and Economics Discussion Series 037. <https://www.federalreserve.gov/econresdata/feds/2015/files/2015037pap.pdf>.
- Santos, João A. C. 2011. "Bank Corporate Loan Pricing Following the Subprime Crisis." *Review of Financial Studies* 24, no. 6: 1916–43. <https://doi.org/10.1093/rfs/hhq115>.

- Sarkar, Asani, and Jeffrey Shrader. 2010. "Financial Amplification Mechanisms and the Federal Reserve's Supply of Liquidity during the Crisis." Federal Reserve Bank of New York Staff Report 431 (March). https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr431.pdf.
- Schroth, Enrique, Gustavo A. Suarez, and Lucian A. Taylor. 2014. "Dynamic Debt Runs and Financial Fragility: Evidence from the 2007 ABCP Crisis." *Journal of Financial Economics* 112, no. 2: 164–89. <https://doi.org/10.1016/j.jfineco.2014.01.002>.
- Scott, Hal S. 2016. *Connectedness and Contagion: Protecting the Financial System from Panics*. Cambridge, Mass.: MIT Press.
- Shiller, Robert J. 2007. "Understanding Recent Trends in House Prices and Homeownership." In *Housing, Housing Finance, and Monetary Policy*. Federal Reserve Bank of Kansas City. https://www.kansascityfed.org/publicat/sympos/2007/PDF/Shiller_0415.pdf.
- Shin, Hyun Song. 2011. "Global Banking Glut and Loan Risk Premium." Mundell-Fleming Lecture. Washington: International Monetary Fund. <https://www.imf.org/external/np/res/seminars/2011/arc/pdf/hss.pdf>.
- Shleifer, Andrei, and Robert Vishny. 2010. "Asset Fire Sales and Credit Easing." *American Economic Review Papers and Proceedings* 100, no. 2: 46–50.
- Siemer, Michael. 2014. "Firm Entry and Employment Dynamics in the Great Recession." Finance and Economics Discussion Series 56. <https://www.federalreserve.gov/pubs/feds/2014/201456/201456pap.pdf>.
- Stock, James H., and Mark W. Watson. 2012. "Disentangling the Channels of the 2007–09 Recession." *Brookings Papers on Economic Activity* 2012, no. 1: 81–156.
- Sunderam, Adi. 2015. "Money Creation and the Shadow Banking System." *Review of Financial Studies* 28, no. 4: 939–77. <https://doi.org/10.1093/rfs/hhu083>.
- Van den Heuvel, Skander. 2002. "Does Bank Capital Matter for Monetary Transmission?" *Economic Policy Review*, May: 259–65.
- Woodford, Michael. 2010. "Financial Intermediation and Macroeconomic Analysis." *Journal of Economic Perspectives* 24, no. 4: 21–44. <https://doi.org/10.1257/jep.24.4.21>.
- Wu, Tao. 2008. "On the Effectiveness of the Federal Reserve's New Liquidity Facilities." Federal Reserve Bank of Dallas Working Paper 8. <https://www.dallasfed.org/~media/documents/research/papers/2008/wp0808.pdf>.
- Zhang, Lu, Arzu Uluc, and Dirk Bezemer. 2017. "Did Pre-Crisis Mortgage Lending Limit Post-Crisis Corporate Lending? Evidence from U.K. Bank Balance Sheets." Bank of England Staff Working Paper 651 (March). <http://www.bankofengland.co.uk/research/Pages/workingpapers/2017/swp651.aspx>.

Comments and Discussion

COMMENT BY


OLIVIER BLANCHARD Ben Bernanke has written an important and information-rich paper, with three separate but related sections. Each section is likely to become required reading for anyone who wants to understand financial crises in general, and the 2008 global financial crisis in particular.

The first section of Bernanke's paper focuses on the interactions between the financial system and the real economy, and the various mechanisms behind a financial crisis. The second section reviews the explosion of micro-economic and macroeconomic research on these different mechanisms, research that was largely triggered by the global crisis and that has built on the numerous quasi-natural experiments it has generated. And the third section looks at the data through the lens of this conceptual framework, and reaches a strong conclusion: It is the panic aspects of the crisis that explain its very large macroeconomic effects.

In this comment, building on the first section of Bernanke's paper, I offer a five-level typology of financial crises, extending his analysis to take into account what happened in Europe. But first, I try to narrow down what he refers to as "panics" in terms of multipliers versus multiple equilibria. I conclude with policy implications.

Should we think of the large effects of financial shocks on the real economy in terms of large multipliers or multiple equilibria? Does this distinction make sense, at least in theory? And does it have important policy implications? My answer to both questions is that it does, even if distinguishing empirically between the two is not straightforward. It is useful to sketch an example.

Consider the two-way interaction between solvency and activity. Think of solvency as standing for variables such as the capital ratio of financial

Figure 1. Activity and Solvency

Source: Author's calculations.

institutions, or the distance to bankruptcy of firms; for my purposes, I do not need to be more specific.


A decrease in activity will decrease solvency. The effect is likely to be nonlinear, becoming stronger the larger the decrease in activity: Most financial institutions and firms are likely to have a sufficient cushion to avoid insolvency for small decreases in activity. Larger decreases are likely, however, to eliminate this cushion, leading to proportionately larger decreases in solvency. This relation is represented by the concave schedule $S(Y)$ in my figure 1.

Similarly, a decrease in solvency will decrease activity. Again, the effect is likely to be nonlinear, becoming stronger the larger the decrease in solvency. Limited bankruptcies may have little effect on activity; widespread bankruptcies are likely to create proportionately larger disruptions and a larger decline in activity. This relation is represented by the convex schedule $Y(S)$ in my figure 1.

The initial equilibrium is at point A in my figure 1. Suppose that, for any reason, solvency decreases at a given level of activity, so $S(Y)$ shifts down to $S'(Y)$, and so the equilibrium moves from A to A' . Given the slopes of $Y(S)$ and $S(Y)$ at this point, the decrease in solvency has a small effect on

Figure 2. Multipliers versus Multiple Equilibria

Solvency, S



Source: Author's calculations.

activity, and the small decline in activity in turn has a small effect back on solvency. Put another way, the multiplier associated with a decrease in solvency is small. Given the concavity and the convexity of the two loci, the larger the shift, however, the stronger the feedback effects and the larger the multiplier.

As shown in my figure 2, for a large enough adverse shift, the two loci now cross twice. There is a “good” (or perhaps, more accurately, a not-so-good) equilibrium at point B , where activity and solvency are low, and a “bad” one at point C , where activity and solvency are much lower: Very low activity leads to very low solvency, which in turn leads to very low activity. Multipliers are large, and explain a decrease to point B . But if the economy moves to the bad equilibrium, the effects are much larger, as the economy moves to point C .

Are “panics” the result of large multipliers or evidence of multiple equilibria? I believe the sharp movements in rates documented by Bernanke strongly suggest multiple equilibria. The answer has obvious policy implications. If the outcome is the result of large multipliers, and it is a case of the good equilibrium becoming less good, then policy measures must deal with fundamentals in order to shift one or both loci back, and improve the

good equilibrium. If it is, instead, a move to the bad equilibrium, policy measures must return the economy to the good equilibrium. One way is to improve fundamentals and shift the locus back so as to eliminate the bad equilibrium. But this may be difficult. On paper, an easier way is to eliminate the bad equilibrium without necessarily improving fundamentals. In principle, this can be done by, for example, providing liquidity to the financial institutions or the firms that are in trouble.

Both types of measures are likely to be used; working on the fundamentals helps, but—and I read this as one of the messages from Bernanke’s paper—it is essential to focus on eliminating the panic—that is, the emergence of the bad equilibrium—by solving the coordination problem in some way. I would put it even more strongly: Given the asymmetric outcomes, it is better to provide liquidity at the risk of finding out that it was largely a solvency problem than it is to limit liquidity provision and allow for a panic, that is, a move to the bad equilibrium.

With this distinction between multipliers and multiple equilibria in mind, let me turn to interactions between the financial sector and the real economy, and offer a tentative five-level typology of financial crises. The first three levels parallel those in the first section of Bernanke’s paper, so I cover them only briefly. Given his focus on the U.S., he did not mention the other two levels; but they were essential in shaping the European version of the financial crisis.

Let us start from an adverse shock to some asset prices—for example, a drop in housing prices—as was the case in the United States, Spain, or Ireland. The first level focuses on the effect on borrowers’ balance sheets. As the drop in asset prices deteriorates their balance sheet, the value of their collateral decreases, forcing them to decrease borrowing. Lower borrowing leads to lower spending and lower output. Lower output in turn leads to lower asset prices, lower collateral, and lower borrowing.

This mechanism was well understood before the 2008 global financial crisis; and indeed, it was already integrated in some macroeconomic models (for example, that of Bernanke, Gertler, and Gilchrist 1999). It was not seen, however, as a likely trigger for a major financial crisis, but rather as implying both a stronger effect of other shocks on activity and, on its own, as a minor source of output fluctuations. The multiplier associated with the borrowers’ balance sheet effect was thought not to be very large. Is this right? The size of the multiplier clearly depends on the initial degree of leverage, and this explains, for example, the difference between the effects of the bursting of the high-tech bubble in the 2000s and the decline in housing prices during the 2008 crisis. How much this mechanism

exacerbated the crisis is the subject of some disagreement, but I agree with Bernanke that it is clearly not enough to explain the depth of the effect on output that we observed.

The second level of my typology focuses on the effect of lenders' balance sheets, in particular the balance sheets of financial intermediaries. Either directly because they hold the lower-priced assets, or indirectly because lower activity deteriorates the borrowers' balance sheets and by implication their own, intermediaries decrease lending, leading to lower spending and lower output. Lower output leads in turn to a deterioration of their own balance sheet.

This effect was also well understood before the crisis, and, as Bernanke shows in his review of the research on the microeconomic evidence, has been thoroughly documented for the crisis. The multiplier is larger than in the first level because the degree of leverage of financial institutions is much higher than that of the typical borrower. It does not take too much of an adverse shock to have a substantial effect on the capital ratio of banks. Still, I agree with Bernanke that more was at work. This is where the third level comes in.


The third level focuses on liquidity runs. These come under different names, from sudden stops to rollover crises to market freezes or, as Bernanke calls them, panics. At some point, some investors, worried about solvency, decide to cut their losses, stop lending to some institutions, or exit some asset markets altogether. Uncertainty about the state of the financial institutions or the underlying value of particular assets plays a central role here: The more uncertain investors become, the more likely they will be to sell or cut their lending. As they do, fire sale prices will reinforce concerns about solvency, further retrenchment, and further effects on asset prices, on lending, and, in turn, on output. The so-called run on repo, which was documented by Gary Gorton and Andrew Metrick (2012), clearly played a central role in the crisis. What we learned from the crisis is that this was not just the province of traditional bank runs but could also happen in the shadow banking system and in asset markets. But just as in the textbook bank run case, this is clearly fertile ground for multiple equilibria and very large effects on output.

Two other interactions, which did not play a role in the United States and so are not mentioned by Bernanke, are similarly fertile ground and played a central role in Europe. These constitute the fourth and fifth levels of my typology.

The fourth level involves the state, and focuses on the interactions between the balance sheet of the state and those of financial institutions.

Figure 3. CDS Rates on Irish Banks and Government Bonds, 2006–18

Basis points



Source: Bloomberg.

These interactions again have been given many names, from doom loops to diabolical loops to deadly embrace. The weak balance sheets of banks lead to the expectation that the state will bail them out. This leads investors to question whether the state's balance sheet is strong enough to absorb these bailouts. And this in turn leads to decreases in the price of government bonds on the banks' balance sheets, and thus to weaker balance sheets (see, for example, Farhi and Tirole 2018).

A particularly salient example is that of Ireland, shown in my figure 3, which plots the evolution of the credit default swap (CDS) rates for banks and for government bonds from October 2006 to April 2008. Starting in the fall of 2007, investors started to worry about Irish banks' balance sheets, leading to increases in the CDS rates of different banks (the variable in the figure is an average, with the shortcoming that, as some banks are closed, the average can jump down a lot). In September 2008, under pressure, the Irish government extended a general guarantee to banks' creditors. The immediate effect was a decrease in CDS rates on banks and an increase in CDS rates on government bonds. But thereafter, the two rates moved very much together in a "deadly embrace," both reaching very high levels, until they stabilized at a low level in 2014.

How should we think of these episodes? There is little question that they involve high multipliers, and that even the good equilibrium implies some solvency risk. But the same mechanisms as those described above can lead to multiple equilibria. Fears, initially justified or not, about the solvency of banks or the state can again become self-fulfilling, leading to a bad equilibrium.


The fifth level involves foreign investors. If a country operates under flexible exchange rates, worries about solvency—whether of the state, of firms, or of financial institutions—are likely to lead to capital outflows and an exchange rate depreciation. To the extent that a large proportion of bonds is denominated in foreign currency, the depreciation further deteriorates balance sheets, whether those of firms and banks, and those of the state. These in turn lead to lower spending, lower solvency, and further worries. This has been a standard scenario in many emerging markets.

If countries instead operate under a peg or within a common currency zone, then the effect comes from worries that the country will give up the peg—or, in the case of the common currency zone, that the country will exit the zone. The effect is an even larger spread on bonds, reflecting not only the risk of default but also the risk of depreciation following exit. This is shown in my figures 4 and 5, which portray the evolution of the CDS on Irish and Portuguese government bonds, together with the market-based euro-exit probabilities, from 2010 to 2018.¹ The similarity of movements is striking between CDS prices, which reflect default risk, and the computed euro-exit probability. Again, these interactions imply potentially large multipliers. But they also can lead to multiple equilibria. Fears of default, justified or not, can lead to expectations of exit, high spreads, and default becoming self-fulfilling.

Finally, I return to panics, whether they reflect high multipliers or multiple equilibria, and ask when they are more likely to arise. There are, I believe, two views. I have called the first one “dark corners” (Blanchard 2015). Although panics will always be hard to predict, the argument is that they are much more likely to arise under some configurations of the financial system, say, when leverage or uncertainty is high. A particularly strong form of this view is the notion of “Minsky cycles,” where steady increases in risk lead, nearly mechanically, to a crash and a financial crisis. At the other extreme is the “hidden mines” view. To a first approximation,


1. These exit probabilities were kindly provided to me by Fathom Consulting. These are computed by using government bond rates and CDS prices (on the assumption that CDS prices, for CDS issued under the pre-2014 rules of the International Swaps and Derivatives Association, do not reflect denomination risk, only credit risk).

Figure 4. CDS Rates on Government Bonds, and Euro Exit Probabilities, Ireland, 2010–18



Sources: Bloomberg; Fathom Consulting.

Figure 5. CDS Rates on Government Bonds, and Euro Exit Probabilities, Portugal, 2010–18



Sources: Bloomberg; Fathom Consulting.

under this view, panics are nearly unpredictable, and thus are about as likely (or as unlikely) to happen under nearly any configuration. These two views lead to different policy conclusions. Under the first, policy should be used aggressively to reduce the risk. Under the second, not much can be done, beyond trying to maintain good fundamentals in general, to specifically prevent the panic; if and when it happens, measures must be taken to limit or eliminate it. This is not an abstract debate: There are those who argue that the Federal Reserve should have seen the risks building up in 2007 and 2008, and thus taken much more aggressive measures to avoid dark corners. There are those who argue, instead, that it was next to impossible to predict the panics that started in late 2008, and that the best that could be done was to deal with them when they came, which the Fed indeed did. More research on this strikes me as a high priority. To be more specific, think of the following relation:

$$P(\text{panics}) = f(\text{financial state}) + \varepsilon.$$

We need to know much more about $f'(\cdot)$, the relevant variables going into “financial state,” and the R^2 equivalent of the regression.

REFERENCES FOR THE BLANCHARD COMMENT

- Bernanke, Ben, Mark Gertler, and Simon Gilchrist. 1999. “The Financial Accelerator in a Quantitative Business Cycle Framework.” In *Handbook of Macroeconomics*, volume 1, part C, edited by J. Taylor and M. Woodford. Amsterdam: Elsevier.
- Blanchard, Olivier. 2015. “Dark Corners.” Lecture given at American Economic Association–American Finance Association joint luncheon, at Annual Meeting of the American Economic Association, Boston, January 3–5. <https://www.aeaweb.org/webcasts/2015/luncheon>.
- Farhi, Emmanuel, and Jean Tirole. 2018. “Deadly Embrace: Sovereign and Financial Balance Sheets Doom Loops.” *Review of Economic Studies* 85, no. 3: 1781–1823.
- Gorton, Gary, and Andrew Metrick. 2012. “Securitized Banking and the Run on Repo.” *Journal of Financial Economics* 104, no. 3: 425–51.

COMMENT BY

RAGHURAM RAJAN It is very difficult to discuss a paper like this one by Ben Bernanke. Usually, great scholars who write on a subject can be told how their academic speculations could benefit from more awareness

of the way the world actually functions. Our colleagues who have their feet more squarely planted in the real world can benefit from counseling on their academic methods. But with Bernanke, we have a first-rate scholar who understands the financial sector and its workings better than most practitioners. He makes the life of a commenter hard by eliminating the scope for easy observations. It is also hard to fault the actions of the Federal Reserve in response to the 2008 global financial crisis, even with the benefit of hindsight. Collectively, the actions paid off and saved the world from a second Great Depression. Bernanke and his fellow central bankers at that time are heroes, something that those who lived through those panic-filled days in 2008 and 2009, when the system was on the brink of failure, fully appreciate. So instead of offering criticism of Bernanke's paper or assessing the Fed's crisis response, I do three things in this comment. First, I describe what the paper attempts to do. Second, I speculate on the paper's political economy. And third, I focus on a question the paper does not ask, but I wish it had.

The paper starts by noting that as the crisis hit, most forecasters, including the Federal Reserve, constantly underestimated its depth. The paper suggests that the triggers for the crisis may well have been the weakness in housing and in household balance sheets, and the way they infected financial institutions' balance sheets. The real blow, however, was the financial panic, resulting from a loss of investor confidence in financial intermediaries, which choked off the supply of credit. Through factor analysis, the paper suggests that the crisis can be attributed to an initial weakening of households' and mortgage lenders' balance sheets, and then to the increasing reluctance of wholesale funders to continue rolling over their debt (epitomized by a blowout of the spread between the London Interbank Offered Rate and the Overnight Indexed Swap Rate), followed by a full-scale panic as yields on even nonmortgage securitizations jumped, and finally to a weakening of the commercial banking system.

However, when it comes to explaining the depths of the recession and the strength of the subsequent recovery, funding and credit factors have the greatest explanatory power. The paper concludes that though balance sheet factors, including the debacle in the housing market, may have been important as triggers, the unanticipated panic transformed it into the 2008 global financial crisis. The paper does offer the caveat that perhaps "both household and bank balance sheets evolve too slowly and (comparatively) smoothly for their effects to be picked up in the type of analysis presented in this paper." Indeed, David Aikman and others (2018) suggest that "the 11.5 percentage point increase in U.S. household debt to GDP that occurred

between 2004 and 2007 can explain between 3.5 and 5.8 percentage points of the decline in GDP. This is around half the GDP shortfall relative to trend.” In other words, tests over longer horizons can indeed produce different interpretations.

The point of this exercise, then, seems not so much to play down balance sheet explanations of the prolonged recession—they were important but perhaps not useful in explaining some of the short-term macroeconomic fluctuations—but to argue that the U.S. authorities were right in focusing on fixing the illiquidity and mistrust in the financial system, because these prevented a much more prolonged and devastating downturn. I agree. To further our learning, it would be nice to try and disentangle the effects of the various interventions. For instance, I believe the bank stress tests in the spring of 2009 (accompanied by government capital infusion for banks that fell short and could not raise any from the markets, ensuring that they did not shrink their balance sheets further to meet capital requirements) did a lot to build confidence in the system. However, the success of such interventions suggests that it is hard to separate panic factors from balance sheet factors.

Why, then, undertake this exercise? Clearly, the perception that the authorities were much more interested in bailing out the banks than in restructuring household debt—they were focused on Wall Street rather than Main Street, as the parlance goes—was widespread. It is not clear that there was an easy way to write down the debt of households that were in over their heads. Nevertheless, no one really explained why enormous amounts of taxpayer funds were put at risk in saving banks. The public perception was that few bankers paid the price for taking the system to the brink. The widespread conclusion was that this was the elite looking after their own, the rest be damned. This paper should then be seen as explaining why the interventions in the financial sector were needed, though it will not assuage those who believe that the authorities should have done more for households.

Let us turn now to the question I wish Bernanke’s paper had asked: What caused the crisis? If indeed it was not just housing but also the entire financial sector that was fragile, then what responsibilities do the Fed and other regulators have? The paper does not tackle this question directly, but there are hints strewn around, perhaps unintentionally. The paper starts by pointing out how no one really saw the entire contours of the crisis, and how forecasters, including in the Fed, constantly underestimated the size of the downturn. Then, in emphasizing the word “panic” and the classic paper by Douglas Diamond and Philip Dybvig (1983), it might appear that


the paper attributes the panic to an unforecastable “sunspot.” Although banks and bank-like institutions are indeed vulnerable to sunspot-based runs, more updated versions of the research done by Diamond and Dybvig relying on global games emphasize that bank runs are set off by underlying bad fundamentals (see, for example, Morris and Shin 2000). Indeed, any postmortem (including the current paper) would also point to the growing leverage across the financial sector, and the clear evidence of exploding risk-taking before the crisis as contributors. Moreover, the then-often-articulated claim by the Fed that it could not deflate a bubble or stop financial excess, but could pick up the pieces when the bubble collapsed, is rarely heard now. How much did that claim lead to complacency, both in the private sector and among authorities? The spark that set off the conflagration may have been hard to forecast, but the dry timber was well and truly piled up and ready to burn.

Put differently, the paper’s emphasis on the supply of credit during the crisis as an important constraint is no doubt right; but why was it so fragile? A number of papers suggest that anticipated easy financing conditions cause an increase in asset prices, in leverage, and in financial fragility. For instance, several researchers—Emmanuel Farhi and Jean Tirole (2012); Diamond and Rajan (2012); and Itamar Drechsler, Alexi Savov, and Philipp Schnabl (2017)—suggest that anticipation of easy liquidity leads to more short-term borrowing by financial institutions and more investment in illiquid assets. Diamond, Yunzhi Hu, and Rajan (2018) argue that this can also make borrowers rely more on continuing liquidity to support their borrowing capacity, and to neglect other sources of debt capacity such as better governance. What, then, causes expectations of easy liquidity? Sustained accommodative monetary policy may be a contributor. A large number of papers now document the link between easy monetary policy and risk-taking by banks (for example, Ioannidou, Ongenga, and Peydró 2015; Jiménez and others 2014).

There is an ongoing debate about whether the Fed was behind the curve—for example, whether Taylor rule residuals were negative well before and after the Fed started raising rates in 2004, with John Taylor suggesting that they were, and Bernanke arguing the opposite.¹ Be that as it may, regardless of whether the policy rate matched economic conditions, it may have been too accommodative from the perspective of financial fragility. The Taylor rule may not be the best measure of the appropriateness

1. See John Taylor’s views at Taylor (2007) and Ben Bernanke’s response at Bernanke (2010).

Figure 1. The Average Lending Distance between a Corporate Borrower and the Nearest Bank Branch from Which It Borrowed, 1996–2016




Source: Taken from Granja, Leuz, and Rajan (2018).

of policy for this purpose.² For instance, A. Maddaloni and J.-L. Peydró (2011) note that the net percentage of loan officers reporting a tightening of credit standards continued to be negative until early 2007.

That there was generalized financial excess can be seen in a number of areas. Consider, for example, small business lending. My figure 1 is illustrative. It shows the average distance between a corporate borrower and the nearest bank branch from which it borrowed. Although this distance has been increasing because of technological change, which enables business to be done at a distance (Petersen and Rajan 2002), it departed from this trend in about 2003, only to come crashing back down during the global crisis. Distant loans turned out to be much more prone to defaulting, and the higher risk involved was not compensated with higher interest rates. Banks therefore seemed to be stretching to lend, and taking on additional risk, before the crisis.

The jury is still out on how much easy monetary conditions exacerbated financial excess before the crisis. What is not in dispute is that regulators

2. However, for an attempt at correlating Taylor rule deviations with financial excess, see Kahn (2010).

Figure 2. Covenant-Lite Leveraged Loans in the United States, 2004–17

Source: Taken from Diamond, Hu, and Rajan (2018).

and supervisors could have done more, if nothing else to stand in the way of the leveraging of the financial sector. This, then, leads to the most important question for our current times: Have we absorbed the lessons of the past, and are we doing enough to prevent financial excess? For instance, starting in 2013, the Federal Reserve offered guidance to banks, counseling them against making highly leveraged loans. The political establishment has since pushed back against this guidance, with the Government Accountability Office suggesting that the guidance is subject to congressional review, and the comptroller of the currency allowing well-capitalized banks to transgress limits (Berlin 2018). Banks have started ignoring the guideline at a time when the market is already frothy. The expansion in so-called covenant-lite loans suggested by my figure 2 gives some cause for concern.

The broader point is that even from a political economy perspective, it is not enough to explain why the authorities acted the way they did when the global crisis hit. We also need a much more detailed study of what went wrong in the period leading up to the crisis. Otherwise, there is a great danger that we will repeat the mistakes of the past.

In this light, it is worrisome that the United States still does not have a regulatory structure that is able to assess the need for macroprudential intervention as well as a set of relevant macroprudential tools (Aikman and others 2018). Such tools could possibly bridge the policy gap between a monetary policy that is set for the needs of the real economy but turns out to be too easy for the financial sector (again, whether such gaps exist is another area that needs more research). For now, the only tool the United States has is the countercyclical capital buffer, which is a blunt tool. Some regulators have suggested that its use be explored. Clearly, there is much to be learned if we are to avoid the terrible experiences that Ben Bernanke documents so well in this paper.

REFERENCES FOR THE RAJAN COMMENT

- Aikman, David, Jonathan Bridges, Anil Kashyap, and Caspar Siebert. 2018. "Would Macroprudential Regulation Have Prevented the Last Crisis?" Staff Working Paper 747, Bank of England.
- Berlin, Andrew. 2018. "Regulated Banks Soften Stance on Leveraged Lending Guidance." Reuters, April 19. <https://www.reuters.com/article/us-lev-regulation/regulated-banks-soften-stance-on-leveraged-lending-guidance-idUSKBN1H Q2XV>.
- Bernanke, Ben. 2010. "Monetary Policy and the Housing Bubble." Paper given at Annual Meeting of American Economic Association, Atlanta, January 3. <https://www.federalreserve.gov/newsevents/speech/files/bernanke20100103a.pdf>.
- Diamond, Douglas W., and Philip H. Dybvig. 1983. "Bank Runs, Deposit Insurance, and Liquidity." *Journal of Political Economy* 91, no. 3: 401–19.
- Diamond, Douglas, Yunzhi Hu, and Raghuram Rajan. 2018. "Pledgeability, Industry Liquidity, and Financing Cycles." NBER Working Paper 23055. Cambridge, Mass.: National Bureau of Economic Research. <http://www.nber.org/papers/w23055>.
- Diamond, Douglas, and Raghuram Rajan. 2012. "Illiquid Banks, Financial Stability, and Interest Rate Policy." *Journal of Political Economy* 120, no. 3: 552–91. DOI:10.1086/666669.
- Drechsler, Itamar, Alexi Savov, and Philipp Schnabl. 2018. "A Model of Monetary Policy and Risk Premia." *Journal of Finance* 73, no. 1: 317–73.
- Farhi, Emmanuel, and Jean Tirole. 2012. "Collective Morality Hazard, Maturity Mismatch, and Systemic Bailouts." *American Economic Review* 102, no. 1: 60–93.
- Granja, João, Christian Leuz, and Raghuram Rajan. 2018. "Going the Extra Mile: Distant Lending and Credit Cycles." NBER Working Paper 25196. Cambridge, Mass.: National Bureau of Economic Research. <http://www.nber.org/papers/w25196>.

- Ioannidou, Vasso, Steven Ongena, and José-Luis Peydró. 2015. "Monetary Policy, Risk-Taking, and Pricing: Evidence from a Quasi-Natural Experiment." *Review of Finance* 19, no. 1: 95–144.
- Jiménez, Gabriel, Steven Ongena, José-Luis Peydró and Jesús Saurina. 2014. "Hazardous Times for Monetary Policy: What Do Twenty-Three Million Bank Loans Say About the Effects of Monetary Policy on Credit Risk-Taking?" *Econometrica* 82, no. 2: 463–505.
- Kahn, George A. 2010. "Taylor Rule Deviations and Financial Imbalances." *Economic Review, Federal Reserve Bank of Kansas City*, no. Q II, 63–99. <https://www.kansascityfed.org/Publicat/EconRev/PDF/10q2Kahn.pdf>.
- Maddaloni A., and J.-L. Peydró. 2011. "Bank Risk-Taking, Securitization, Supervision, and Low Interest Rates: Evidence from Euro-Area and U.S. Lending Standards." *Review of Financial Studies* 24: 2121–65.
- Morris, Stephen, and Hyun Song Shin. 2000. "Rethinking Multiple Equilibria in Macroeconomic Modeling." In *NBER Macroeconomics Annual 2000*, edited by Ben S. Bernanke and Kenneth Rogoff. MIT Press.
- Petersen, Mitchell A., and Raghuram G. Rajan. 2002. "Does Distance Still Matter? The Information Revolution in Small Business Lending." *Journal of Finance* 57, no. 6: 2533–70.
- Taylor, John. 2007. "Housing and Monetary Policy." NBER Working Paper 13682. Cambridge, Mass.: National Bureau of Economic Research. https://web.stanford.edu/~johnhtayl/Onlinepaperscombinedbyyear/2007/Housing_and_Monetary_Policy.pdf.

GENERAL DISCUSSION Robert Hall began by remarking that although financial panics originate in credit markets, it is interesting to consider how much of the decline in the stock market during the 2008 global financial crisis represented a panic—as people perhaps applied higher personal discount rates in the stock market—or how much it represented the fact that the stock market reflected changes in fundamentals in the economy and credit markets. He proposed that an extension of the author's analysis that includes evolutions in the stock market might be of value.

Robert Gordon proposed constructing an analysis of the 1929 panic in a fashion analogous to the author's. Fundamental factors in the run-up to the 1929 crash included the fact that real housing had peaked in 1927 and had begun to rapidly decline, in part because of the Federal Reserve's attempts to dampen stock market speculation with higher interest rates. The panic in the stock market was fueled in part by very-low-margin requirements on stocks at the time, which resulted in self-perpetuating downward movements in stock prices as investors lost confidence. The emergence

of the panic in 1929 very much mirrored that in 2008. But after the panic, the policy responses—the shrinkage of the money supply; and a fiscal contraction that was particularly strong at the state and local levels, the tax increase of 1932, and Smoot-Hawley tariffs—were remarkably different. Gordon argued that the divergence of the two policy responses should give pause in using the Great Depression as a counterfactual to the Great Recession.

Stanley Fischer commended the paper. He remarked that the political system in the United States has largely decided that the government will not play an active role in future financial crises, and that lender-of-last-resort authorities should be limited because they encourage moral hazard. In contrast, Britain's system has made efforts to extend its capacity to deal with credit disruptions, across the economy. He remarked on current political sentiment that favors deregulation for banks, because the Dodd-Frank Act had differential effects on both small and large banks, and even these effects will result in lower supervision on average. Fischer expressed concern that the conditions for a crisis are developing faster than they did after the Great Depression.

Ben Friedman commented that in a financial panic, the central bank's distinction between liquidity and solvency problems in determining whether to lend to institutions is not operative, because the value of institutions' assets is endogenous to whether the central bank chooses to intervene and in what securities it does so. Similarly, on the liability side of institutions' balance sheets, whether depositors return to the institution in the future depends on current intervention or the lack of it. Friedman proposed that recognizing the extent to which solvency and liquidity are endogenous with respect to the central banks' actions would help reduce opposition to the banks' interventions in a crisis.

Eric Rosengren commended the paper, and noted that he agrees with the author that the academic literature has underemphasized the role of credit supply in influencing real economic activity. He suggested that the author incorporate volumes of transactions in credit markets, along with the spreads on which the author focused. Doing so would shed light on the persistence of the effects from credit market disruptions, given that many key credit markets had almost no activity for some time after the crisis. In addition, he remarked that some credit markets revived in direct response to Fed programs, particularly after the results of the Fed's stress tests were published. Another example is the Fed's Asset-Backed Commercial Paper Money Market Mutual Fund Liquidity Facility, which largely revived asset-backed commercial paper markets.

Jay Shambaugh commented that he found the paper to be very interesting, and that he appreciated the author's choice to reassert that the analysis is not an attempt to entirely dismiss the role of household and bank balance sheets in the recession. He remarked, however, that the design of the analysis is unlikely to produce a strong role for household balance sheets, because they do not have the time series variation that would allow for strong identification. In parallel, the author's credit market variables would be insignificant in a cross-sectional study like that done by Atif Mian and Amir Sufi, who find a significant role for household balance sheets.¹ In short, the two channels—the “balance sheet” and “panic” channels—will look relatively more important depending on how one constructs the analysis.

Nellie Liang remarked that in recent research, she has found that the United States' recovery since the financial crisis was in fact quicker than the recoveries of most other countries. She noted that this difference could help to further identify the panic's effects, given that most other countries also experienced financial panics but many did not also experience the deterioration in balance sheets that was central in the United States.

Frederic Mishkin said that the paper has significant implications for policy design during a panic—and in particular, for the design of policies like the Troubled Asset Relief Program (TARP). At its beginning, TARP intended to offload bad assets from banks, and was designed so that healthy and distressed institutions would participate to avoid the stigma attached to participating in a relief program. As a result, Mishkin noted, a large portion of TARP funds went to banks not facing restrictions on how they would use the funds. That TARP had a provision allowing the Treasury to recapitalize banks was a fortunate turn of events, but Mishkin argued that its poor initial design had political ramifications that could dampen the Fed's ability to act in future crises.

Athanasios Orphanides affirmed Eric Rosengren's assessment of the role of monetary policy in setting the stage for the financial crisis, observing that the Fed's policy rate in the years leading up to 2007 matched what would have been prescribed by a Taylor-type rule. In retrospect, there is no evidence that the Fed's rate deviated from a systematic approach to policy-making in a way that would have fueled the crisis.

Orphanides reaffirmed Ben Friedman's point that the distinction between solvency and liquidity is endogenous to the central bank's interventions,

1. Atif Mian and Amir Sufi, “What Explains the 2007–2009 Drop in Employment?” *Econometrica* 82, no. 6 (2014): 2197–2223.

and observed that it relates to Olivier Blanchard's framework of multiple equilibria in the financial sector. The central bank may decide which equilibria are met. But the fact that central banks have no systematic rules that prescribe actions during a crisis prevents this intervention from being entirely effective. Orphanides remarked that policymakers ought to consider ahead of time where losses will be absorbed, and how to avoid moral hazard while providing a credible backstop to lending.

Mark Gertler responded to Jay Shambaugh's remarks, noting that in recent research done with Simon Gilchrist, he has considered both the state-level, cross-sectional variation and the time series aggregate movements of household balance sheets and banking distress.² They have found, like Bernanke, that the financial variable is most important for explaining real economic developments. Gertler noted that, as Shambaugh's intuition would suggest, they find a more important role for balance sheets in the cross-sectional analysis than Bernanke does in the time series, however.

Ben Bernanke began by thanking the participants for great comments. He agreed with Olivier Blanchard that the international ramifications of the crisis would be an interesting direction for future research.

Regarding models of the financial panic, Bernanke proposed the model put forth by Mark Gertler and Nobuhiro Kiyotaki, in which some fundamental variable—for example, bank capital—varies over time and if it falls below a certain threshold, a bank run becomes a possible equilibrium.³ Gertler and Kiyotaki assume that if a run equilibrium exists, a run in fact occurs. Bernanke also acknowledged, as many commenters also noted, that policy interventions may determine whether the run equilibrium arises. However, he concluded that his interest in the paper is on the effects of the panic, and does not require that he take a position on whether the panic was a sunspot or a fundamental crisis.

Responding to Rajan's comment paper, Bernanke agreed that basic arbitrage breaks down when liquidity is very tight. He noted that bank credit default swaps and other bond spreads that would get at this do correlate closely with his panic-related variables. He observed that market volumes would also be a useful indicator of arbitrage failure, as Eric Rosengren had suggested.

2. Mark Gertler and Simon Gilchrist, "What Happened: Financial Factors in the Great Recession," NBER Working Paper 24746 (Cambridge, Mass.: National Bureau of Economic Research, 2018).

3. Mark Gertler and Nobuhiro Kiyotaki, "Banking, Liquidity, and Bank Runs in an Infinite Horizon Economy," *American Economic Review* 105, no. 7 (2015): 2011–43.

Bernanke commented that Rajan's discussion of monetary policy is not directly relevant to the paper. But, in any case, the Federal Open Market Committee's rate decisions in 2004 are not strong evidence that the Fed deviated from normal policymaking practice, because the recovery from the 2001 recession had been fairly weak until that point. In 2003, job creation was sluggish, the Federal Open Market Committee members were worried about the possibility of deflation, and the only sector performing well was housing. By many measures, the Fed's interest rate decisions did not deviate from Taylor rule predictions. In retrospect, the challenge of the crisis was in fact a large and run-vulnerable sector that was neither appropriately overseen nor capitalized. Thus, the proper policy response in 2004 was more likely to have been regulatory rather than a monetary tightening in the face of a weak economy.

Bernanke agreed with Robert Hall's remarks about examining the stock market. He referred to work by John Campbell, Stefano Giglio, and Christopher Polk finding that the decline in the stock market during the crisis was mostly due to lower expectations about future profits, not to increased discount rates.⁴ This suggests that the stock market was internalizing the effects of the crisis on the economy, rather than functioning as an additional factor that drove the real downturn.

Bernanke responded to comments that the design of the analysis may not have captured variation in household balance sheets by reasserting the last analysis conducted in the paper, which includes house prices, delinquencies, and additional panic variables in prediction equations for real economic variables. The analysis has the same results as the main findings of the paper—that the panic variables explain much more of the variation in economic activity than do the housing variables. Moreover, house prices and delinquencies have quite a bit of time-series variation in the aggregate. Bernanke remarked that Gertler and Gilchrist's approach to studying the time series and the cross section is commendable and that more work should be done in this spirit.⁵

Bernanke concluded by stating that he simply hopes to reject the hypothesis that the panic was irrelevant to the recession. He acknowledged that the stress on housing was obviously enormous, and that policymakers likely wish they had done more to relieve that distress, but to entirely ignore the panic would have been fatal.

4. John Campbell, Stefano Giglio, and Christopher Polk, "Hard Times," *Review of Asset Pricing Studies* 3, no. 1 (2013): 95–132.

5. Gertler and Gilchrist, "What Happened."

OLIVIER COIBION

University of Texas at Austin

YURIY GORODNICHENKO

University of California, Berkeley

MAURICIO ULATE

University of California, Berkeley

The Cyclical Sensitivity in Estimates of Potential Output

ABSTRACT The fact that declines in output since the Great Recession have been parlayed into equivalent declines in measures of potential output is commonly interpreted as implying that output will not return to previous trends. We show that real-time estimates of potential output for the United States and other countries respond gradually and similarly to both transitory and permanent shocks to output. Observing revisions in measures of potential output therefore tells us little about whether changes in actual output will be permanent. Some alternative methodologies to estimate potential output can avoid these shortcomings. These approaches suggest a much more limited decline in potential output since the Great Recession.

The Great Recession was characterized not only by large declines in economic activity in most advanced economies but also by ones that have persisted for a decade, with no sign of these affected economies catching up to previously expected trend levels. If anything, trends are now being revised down in light of these economies' continuing inability to close the output gaps first generated in 2008. As illustrated in figure 1 for the United States (see below, in section I), estimates of potential output

Conflict of Interest Disclosure: The authors did not receive financial support from any firm or person for this paper or from any firm or person with a financial or political interest in this paper. They are currently not officers, directors, or board members of any organization with an interest in this paper. No outside party had the right to review this paper before circulation.

have been systematically revised downward since the Great Recession, such that all the current deviations of output from past estimates of potential are now being reinterpreted as permanent declines in the economy's productive capacity. These large downward revisions imply that the output gap appears closed, and this absence of any remaining slackness in the economy is a primary motivation for the Federal Reserve's progressive tightening of monetary policy.

However, before we take these dynamics in the estimates of potential output at face value, we should understand their properties and what determines revisions in these estimates. In this paper, we focus on how real-time estimates of potential output respond to different economic shocks in the United States, and also across a wide range of countries. Using a variety of institutional sources for estimates of potential GDP, we find that real-time estimates of this variable respond to cyclical shocks that have no long-run effects on the economy and underrespond to shocks that do. In all cases, adjustments in real-time estimates of potential GDP are extremely gradual, much like a moving average of past output changes. In fact, given their gradual pace of adjustment to shocks and the fact that these real-time estimates fail to differentiate between shocks that do and do not affect the productive capacity of the economy, there seems to be little value added in estimates of potential GDP relative to simple measures of statistical trends. At a minimum, the fact that estimates of potential GDP are revised, either upward or downward, should *not* be taken as a sign that future changes in GDP will in fact be more or less persistent than usual but rather indicates little more than that the prior changes in GDP have been persistent.

Because estimates of potential GDP are not necessarily created in the same fashion across institutions, we consider estimates from the Federal Reserve Board and from the Congressional Budget Office (CBO) for the United States as well as estimates from the International Monetary Fund and the Organization for Economic Cooperation and Development (OECD) for a broader cross section of countries. We complement this with long-term forecasts of output growth from the professional forecasting firm Consensus Economics. Most public or international organizations follow production function approaches, in which estimates of the potential productive capacity of an economy reflect estimates of the capital stock, potential labor force sizes combined with estimates of human capital, and measures of total factor productivity (TFP). Hence, estimates of potential output should change when the technological capacity of the economy improves but not in response to purely cyclical variations in employment, such as those arising from monetary policies.

To test these propositions, we bring to bear not just a wide range of estimates of potential output but also a range of shock measures. Somewhat surprisingly, given the short samples, we find several clear patterns in the data that should give one pause before interpreting changes in estimates of potential output as indicators of permanent changes in output. First, and perhaps most strikingly, though we reproduce the common and well-documented finding that monetary shocks have only transitory effects on GDP, we then document the startling feature that these shocks are followed by a gradual change in estimates of potential GDP. This finding occurs not just in the United States but across other countries as well and is true for a range of sources of estimates of potential GDP.

We find a similar set of results when we focus on government spending shocks. Regardless of the identification strategy, increases in government spending have transitory effects on GDP, but estimates of potential GDP again display a delayed response to these shocks, ultimately responding to the shock in the same direction as the short-run response of GDP. As with the effects of monetary shocks, the fact that estimates of potential GDP respond so unambiguously to these shocks strongly suggests that real-time estimates of potential GDP are failing to adequately distinguish between permanent and transitory shocks. In this respect, estimates of potential GDP are *sensitive to cyclical fluctuations in GDP* originating from demand shocks.

Turning to supply shocks that should affect potential GDP, the results are more mixed. With productivity shocks, which have immediate and persistent effects on GDP, we find that estimates of potential GDP again respond only very gradually but, after several years, fully incorporate the effects of new productivity levels. With tax shocks, we similarly observe that, after a long delay, estimates of potential GDP eventually catch up to actual changes in GDP. Hence, these two supply shocks provide evidence that real-time estimates of potential output ultimately embody some changes in potential GDP. However, the very slow rate at which information about these shocks is incorporated into estimates of potential GDP points to *an insufficient sensitivity of these estimates in response to supply shocks*. With oil price shocks, however, an even more severe problem arises. We observe persistent declines in GDP after these shocks, but estimates of potential GDP actually go in the opposite direction. As with demand shocks, this specific type of supply shock therefore also presents a challenge to the view that estimates of potential GDP are actually capturing what they are meant to.

Furthermore, we can consistently reproduce the way in which estimates of potential GDP respond to shocks by applying a one-sided Hodrick-Prescott (HP) filter to real-time GDP data. In the U.S. as well as in the cross-country data, this approach generates impulse responses to shocks that are nearly indistinguishable from those found using the actual estimates of potential GDP from all organizations, including the countercyclical behavior of measured potential GDP after oil supply shocks. The HP filter is effectively just a weighted moving average of recent GDP changes, and by construction it does not differentiate between the underlying sources of changes in GDP, be they monetary, technological, or others. Thus, a reliance on simple statistical filters like HP by official agencies could readily rationalize why one might observe a gradual response by real-time measures of potential output to any economic shock, even those that have only transitory effects on GDP and that should presumably be stripped out of estimates of potential GDP.

Fortunately, other approaches to identifying potential output can do better. For example, the approach taken by Olivier Blanchard and Danny Quah (1989) to identify supply and demand shocks can successfully generate real-time estimates of potential output that are consistent with theoretical predictions. Indeed, when Blanchard and Quah's approach is applied to real-time data to recover potential output measured as the historical contribution of shocks with permanent effects on output, the resulting real-time estimate of potential output reacts strongly to identified supply shocks (TFP, taxes, and oil price shocks), and it does not respond significantly to identified demand shocks (monetary policy and government spending shocks). Hence, it does not suffer from the problems associated with most other measures of potential output. Furthermore, this approach yields a starkly different interpretation for changes in U.S. potential output since the Great Recession. Our estimates imply that the gap between potential and actual output in the U.S. increased by about 5 log percentage points between 2007:Q1 (when the gap was likely close to zero) and 2017:Q1, leaving ample room for policymakers to close this gap through demand-side policies, if they chose to do so.

We find similar evidence of a large output gap using other methods to calculate measures of potential output, such as the ones proposed by Jordi Galí (1999), which uses information from labor productivity and hours, or by John Cochrane (1994), which brings in additional information from consumption. Using information from inflation to make inferences about potential output through an estimated Phillips curve also points toward significant slackness. All these methodologies give similar results, pointing to

an increase in the gap of 5 to 10 percentage points between 2007:Q1 and 2017:Q1. This assures us that this result is not an artifact of the Blanchard-Quah approach and instead is a feature that is robust to different identification schemes. The idea that significant slackness remained in the U.S. economy through 2017 is also consistent with the low levels of capacity utilization, contained wage growth, and the evolution of labor force participation since the Great Recession.

This paper touches on several bodies of literature. It is most directly tied to recent work since the Great Recession focusing on the possibility of hysteresis—that is, cases where demand shocks lead to permanent effects on the level of economic activity. Though many mechanisms can generate such effects—for example, less research and development during periods of low investment, as shown by Diego Anzoategui and others (2016), Gianluca Benigno and Luca Fornaro (2018), and Patrick Moran and Albert Queralto (2018)—empirical evidence on hysteresis remains scant, as emphasized by Blanchard (2017), with most estimates of monetary and government spending shocks being consistent with the null hypothesis that these shocks have no permanent effects on GDP (for reviews of the literature on monetary and government spending shocks, see Nakamura and Steinsson 2017; Ramey 2016). Recent research has focused on the degree to which the sustained declines in output since the Great Recession have ultimately been interpreted as reflecting declines in potential GDP and therefore can be expected to be long-lasting. Laurence Ball (2014) documents that for most advanced economies, much of the declines in output since the Great Recession have been matched with declines in estimates of potential output. Antonio Fatas and Lawrence Summers (2018) focus on the degree to which fiscal consolidations map first into output changes and then into changes in estimates of potential GDP, with the latter being an indicator that GDP changes will be permanent. *Our results suggest that one should draw little inference from the evolution of estimates of potential GDP about the persistence of GDP changes*; these estimates fail to exclusively identify supply shocks that should drive potential GDP and instead also respond to transitory demand shocks. The fact that most of the output declines observed since the Great Recession are now attributed to declines in potential GDP implies little, other than that these declines have been persistent because estimates of potential GDP fail to adequately distinguish between the underlying sources of changes in GDP.

Our paper also relates to research on news shocks and beliefs about long-run productivity. A strand of the literature studies how news about future productivity can have contemporaneous effects on economic activity

long before the productivity changes actually occur (for example, Beaudry and Portier 2006; Barsky and Sims 2011, 2012). In this spirit, Blanchard, Guido Lorenzoni, and Jean-Paul L’Huillier (2017) show that revisions in estimates of future potential output are correlated with contemporaneous changes in consumption and investment. If estimates of future potential output were invariant to transitory shocks, then one could entertain a causal interpretation of these correlations as reflecting the effect of news about the future on current economic decisions. But our results call for caution with this type of interpretation; estimates of potential GDP display sensitivity to demand shocks, and this sensitivity calls into question the basis for causal inference of the type made by Blanchard, Lorenzoni, and L’Huillier (2017).

A third strand of the literature on which we build focuses on the implications of real-time measurement of the output gap for monetary policy. Athanasios Orphanides and Simon van Norden (2002), for example, illustrate how real-time estimates of potential GDP can, in short samples, be sensitive to the method used to measure either the trend or deviations from it. Orphanides (2001, 2003, 2004) argues that the Federal Reserve’s mismeasurement of the output gap in the 1970s was one of the primary reasons why inflation was allowed to rise so sharply in the 1970s. We are similarly interested in the difficulties with measuring potential output and the output gap; but rather than studying how sensitive estimates of potential output can be to the different statistical techniques used to identify it, we instead characterize whether the historical estimates of potential output from public and international organizations respond to the “correct” shocks. Our estimates imply that just as the Federal Reserve likely overstimulated the economy in the 1970s because of mismeasurement of potential output, it is now at risk of understimulating the economy by underestimating its productive capacity.

Finally, by comparing actual responses of output after economic shocks to the predictions of agents about these variables, our paper is closely related to recent work studying the expectations formation process of economic agents. Coibion and Gorodnichenko (2012) study the forecast errors of agents to economic shocks and find that these errors are persistent after shocks, consistent with models where agents are not fully informed about the state. By comparing the long-run response of GDP with estimates of potential GDP, this paper similarly provides some insight about how these potential GDP estimates are formed.

The paper is organized as follows. Section I presents information about the estimates of potential output used in the paper. Section II presents our baseline estimates, using U.S. data, of how measures of potential GDP

respond to economic shocks. Section III extends these results to a broader range of countries. Section IV presents examples of how estimates of potential output can be improved. And section V concludes.

I. How Estimates of Potential Output Are Created and Used

A seminal description of potential output is in Arthur Okun's (1962) presidential address to the American Statistical Association. Although the notion of potential or natural levels of output had been discussed as far back as research done by Knut Wicksell (1898) and John Maynard Keynes (1936), Okun (1962) provided a sharper definition than had been previously utilized as well as guidance about how to estimate potential output (Hauptmeier and others 2009). Okun emphasized that potential output is a "supply concept, a measure of productive capacity." But it is not designed to represent the maximum amount that an economy could produce. Instead, Okun defines it as the amount that could be produced without generating inflationary pressure. Hence, though potential GDP is related to the nonaccelerating inflation rate of unemployment (NAIRU), potential output provides a more comprehensive assessment of how much an economy can produce without triggering above-normal inflation. This interpretation of potential output advocated by Okun serves as the foundation for most approaches to estimating potential output.

Although Okun proposed to estimate potential output through a combination of knowing the NAIRU and applying what subsequently became known as Okun's law, few organizations follow the specific approach suggested by Okun. As classified by Frederic Mishkin (2007), there are three broad classes of methods to construct a measure of potential output: statistical, production function, and structural (based on dynamic stochastic general equilibrium, DSGE). We first review these methods and then discuss how various agencies measure potential output.

Statistical methods typically impose little theoretical structure on the properties of potential output and interpret low-frequency variation in output series as potential output. One example of this approach is to use univariate time series methods, such as autoregression (AR) models or different types of filters, on actual output to extract a statistical trend component, which is then identified with potential output. Another example is given by methods using several variables—such as output, unemployment, and inflation—to obtain potential output via an unobserved components model and a Phillips curve (Kuttner 1994; Staiger, Stock, and Watson 1997).

In the production function approach, independent estimates of the different inputs that go into the aggregate production function (for example, labor, capital, and multifactor productivity) are plugged into the production function to obtain potential output. Because the objective is to obtain potential output and not actual output, the estimates of the different inputs must correspond to the concept of the maximum (or “normal”) amount of each variable that could be used for production without leading to an acceleration of inflation (for example, the labor force participation rate and a level of natural unemployment should be used instead of the cyclical level of employment). In the latter sense, this approach to estimating potential output remains in the spirit suggested by Okun. This approach is also related to growth accounting, because after log-differentiation of a Cobb–Douglas production function, the growth of potential output can be expressed as the weighted average of the growth rates of the different inputs (for an application of this approach to the dynamics of output in the post–Great Recession period, see Fernald and others 2017).

Finally, structural approaches use DSGE models, typically with a New Keynesian structure, to back out potential output. This requires calibrating or estimating the parameters of the model to the relevant economy so that the different shocks hitting the economy can be identified. Once this stage is completed, potential output can be obtained from the solution of the model when certain shocks and frictions are turned off (for example, Andres, Lopez-Salido, and Nelson 2005). This methodology is particularly dependent on models and relies heavily on the estimation of a sophisticated model, which, given limited variation in macroeconomic data, may be a challenge for identification of structural parameters and shocks. Furthermore, because estimated DSGE models have only been used in recent years, no historical, real-time data are available to assess their properties.

The implicit assumptions about the nature of potential output are not identical across methods. The production function approach, for example, explicitly tries to strip out cyclical factors from estimates of potential output. Statistical filters similarly try to separate cyclical fluctuations in output from changes in the trend, with the latter being equivalent to potential. In contrast, with a New Keynesian DSGE model, where the potential level of output reflects counterfactual outcomes under flexible prices, transitory “demand” shocks like temporary changes in government spending can affect the level of potential output for some time, whereas they would be excluded from estimates of potential under the other two approaches (see Blanchard 2017). Because our empirical strategy involves studying the

response of real-time estimates of potential output to supply (long-lived) versus demand (transitory) shocks, we are adopting an interpretation of potential output that hews most closely to the production function and statistical filtering approaches, in part because this is precisely the conceptual framework that is most often used by statistical and other agencies when they construct estimates of potential.

1.A. The Congressional Budget Office

The CBO uses the production function approach for estimating potential output. As described by the CBO (2001, 2014), this institution estimates potential output with different methods for five sectors in the economy. The main one is the nonfarm business (NFB) sector, which represents about 75 percent of the U.S. economy. The remaining four smaller sectors are agriculture and forestry, households, nonprofit organizations serving households, and government.

In each of these sectors, the CBO projects the growth of each input by estimating a trend growth rate for it during the previous and current business cycles (as dated by the National Bureau of Economic Research) and by extending that trend into the future. This implies that the trend growth for inputs depends on recent history and on business cycle dating, with possibly large changes in trends when a new business cycle begins. The CBO tries to remove the cyclical component of the growth rate of different variables by estimating the relationships between those variables and a measure of the unemployment rate gap, the difference between the actual unemployment rate and the natural rate of unemployment.

For the NFB, the CBO uses a production function with three inputs: potential labor, services from the stock of capital, and the sector's potential TFP. For the agriculture and forestry sectors, and for nonprofits serving households, potential output is estimated using trends in labor productivity for those sectors. For the household sector, potential output is obtained as a flow of services from the owner-occupied housing stock. Finally, for the government sector, potential output is estimated using trends in labor productivity and depreciation of government capital. The CBO's real-time estimates of potential output have been available since 1991 at an annual frequency and since 1999 at a semiannual frequency.

Estimates of potential output by the CBO play an important role in fiscal policy discussions in the United States. When new tax or spending policies are under review by the U.S. Congress, their implications for future tax revenues, government expenditures, and deficits are assessed under assumptions about the long-run future path of the economy, as

captured by estimates of potential GDP (although some policies require the CBO to make inferences about how these policies themselves may change potential output over time, for example, via “dynamic scoring”). How these estimates are formed and how well they separate cyclical from permanent shocks therefore matters for how well these policy measures are scored.

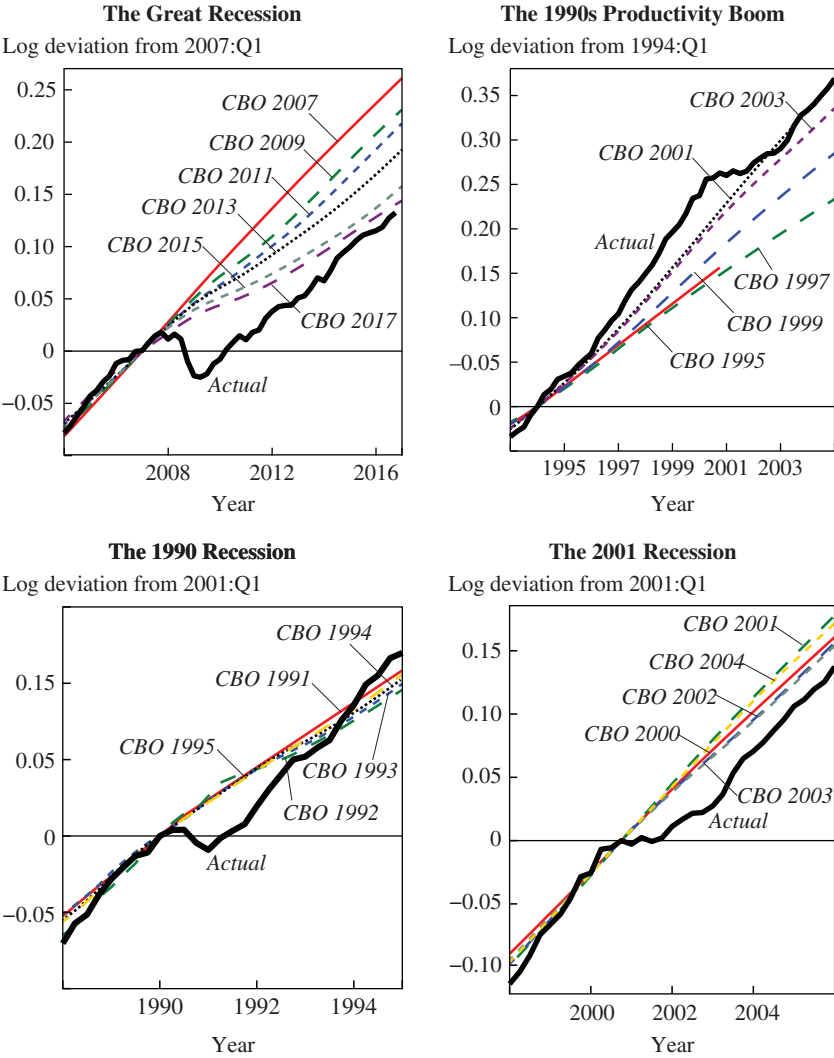
These estimates of potential output are sometimes subject to very large revisions. Preceding the revisions over the course of the Great Recession, for example, the CBO had similarly made a sequence of large *upward* revisions to the projected path of potential output over the course of the 1990s, as illustrated in panel B of figure 1. These upward revisions were tied to the higher-than-expected productivity growth in the U.S. over this period.¹ Other episodes reveal less dramatic sequences of revisions. For example, panels C and D of figure 1 illustrate the CBO’s revisions during the two previous U.S. recessions. In both cases, the CBO first started reducing its predicted path of potential output during the recession but then ultimately raised them back up again. In the case of the 1990 recession, GDP ultimately overtook estimates of potential output, whereas over the same time horizon of three years after the start of the recession, the CBO continued to estimate a large output gap after the 2001 recession. But in neither case do we observe a systematic pattern of downward revisions toward the path of actual GDP such as that which was observed after the Great Recession.

1.B. The Federal Reserve

While preparing macroeconomic projections (historically known as Greenbook forecasts) for meetings of the Federal Open Market Committee (FOMC), the staff members of the Federal Reserve Board construct

1. Although it is true that some of these revisions were not related to productivity changes—such as the ones coming from the shift to chained GDP, the addition of software, or revisions to the National Income and Product Accounts—CBO (2001, 2) summarized one of the larger revisions as follows, “CBO also altered its method to address changing economic circumstances. In particular, labor productivity has been growing much faster since 1995 than its post-1973 trend. Because that acceleration has coincided with explosive growth in many areas of information technology (IT), . . . many observers have speculated that the U.S. economy has entered a new era, characterized by more rapid productivity growth. . . . After analyzing the data and the relevant empirical literature, CBO has concluded that elements of the so-called IT revolution . . . explain much of the acceleration in the growth of labor productivity during the late 1990s. CBO has incorporated many of those elements into its economic projections.”

Figure 1. Historical Revisions in the CBO's Estimates of U.S. Potential Output^a



Source: Congressional Budget Office.

a. This figure plots estimates of U.S. potential output from the Congressional Budget Office made at different time periods (that is, at the beginning of the corresponding year). The heavy solid line represents real GDP in the U.S. In each panel, each series is normalized to zero—for 2007, in panel A; for 1994, in panel B; for 1990, in panel C; and for 2000, in panel D.

a measure of the output gap (that is, the difference between actual and potential output) to assist the FOMC's members in their decisionmaking. As pointed out by Rochelle Edge and Jeremy Rudd (2016, 785), from the Board of Governors of the Federal Reserve System, the estimate of the output gap from the Greenbook "is judgmental in the sense that it is not explicitly derived from a single model of the economy. In particular, the staff's estimates of potential GDP pool and judgmentally weight the results from a number of estimation techniques, including statistical filters and more structural model-based procedures."

While describing the evolution of measuring potential output by the Fed, Orphanides (2004, 157) mentions that in the Greenbook estimates, "the underlying model for potential output was a segmented/time-varying trend. The specific construction methods and assumptions varied over time. During the 1960s and until 1976, the starting point was Okun's (1962) analysis. From 1977 onward, the starting point was Clark's (1979) analysis and, later, the related methods explained in Clark (1982) and Braun (1990). Throughout, these estimates of potential output were meant to correspond to a concept of noninflationary 'full employment.' However, judgmental considerations played an important role in defining and updating of potential output estimates throughout this period, so the evolution of these estimates cannot be easily compared to that of estimates based on a fixed statistical methodology."

More recently, Charles Fleischman and John Roberts (2011) describe a methodology to compute potential output using a multivariate unobserved components model that is taken into account by the Federal Reserve Board when producing its judgmental estimates of potential output. Its procedure embeds some parts of many of the methodologies described above; it uses multivariate statistical methods, trend estimation, growth accounting (as in the production function approach), and the relationship between cyclical fluctuations in output and unemployment (as in Okun's law). The authors use data on nine macroeconomic series: real GDP; real gross domestic income; the unemployment rate; the labor force participation rate; aggregate hours for the NFB; a measure of NFB sector employment; two measures of NFB sector output (measured on the product side and on the income side); and inflation as measured by the Consumer Price Index, excluding food and energy. The common cyclical component of the economy is constrained to follow an AR(2) process, and trends in the series are related to each other via structural equations (for example, Okun's law, production function) to obtain a final measure of the trend of output, which is associated with potential output.

Real-time estimates of potential output can be computed from the estimates of actual output and the output gap reported in Greenbooks since 1987.² Real-time estimates for the same variables in the 1969–87 period are provided by Orphanides (2004). For this earlier period, the quality of the estimates is likely to be worse because the estimates sometimes had to be obtained from a variety of sources (for example, the Council of Economic Advisers) other than the Federal Reserve. As a result, we take the 1987–2011 series as the benchmark and explore the longer time series in robustness checks. Because the Greenbooks only forecast potential output growth for up to a few years, we cannot reproduce figure 1 (the evolution of real-time forecasts of potential GDP during the Great Recession) for Greenbook forecasts.

Estimates of potential output play an immediate role in decisionmaking by the Federal Reserve. One of the objectives of the FOMC is to stabilize output around potential, and whether output is below or above potential is commonly interpreted as having implications for inflation, the other objective targeted by the Federal Reserve. Potential mismeasurement of the output gap (the difference between actual output and potential) is mentioned (for example, Orphanides 2001) as a reason why the Federal Reserve allowed inflation to rise during the 1970s, and Fed chairman Alan Greenspan's perception that potential output was growing unusually rapidly in the 1990s explains why, over this period, monetary policymakers were less concerned about inflation than they normally would have been, given the low unemployment rates (Gorodnichenko and Shapiro 2007).

1.C. The International Monetary Fund

The IMF provides estimates of potential output for a wide range of countries. There is considerable methodological variation across countries in how the IMF generates estimates of potential output. As summarized by Carlos de Resende (2014, 24), in a study conducted by the IMF's Independent Evaluation Office, "Interviews with staff showed that the use of the macro framework is country-specific and varies greatly in detail and sophistication, ranging from the use of 'satellite' models to simply entering numbers based on judgment." In this respect, the IMF's approach to measuring

2. This series is available from the Real-Time Data Research Center at the Federal Reserve Bank of Philadelphia. There is a five-year delay period for the release of Greenbook projections.

potential output is methodologically similar to measures reported in the Greenbooks, in the sense that they use a combination of different methods to compute potential output and then aggregate them using a great deal of judgment. At the same time, the IMF staff often uses the Hodrick–Prescott filter and/or multivariate methods such as the ones described in Patrick Blagrove and others (2015) to construct measures of potential output. The IMF provides potential output estimates for 27 countries.³ Nowcasts and one-year-ahead forecasts are available for the period 2003–16. Since 2009, the IMF has also provided up to five-year-ahead forecasts for potential output.

Estimates of potential output can play an important role in the IMF’s policy decisions. To assess the sustainability of countries’ fiscal policies, tax and spending levels are commonly evaluated at the level of potential GDP to control for the cyclical changes in revenues and expenditures that are expected to be transitory, thereby helping to gauge any “structural” fiscal imbalances. These imbalances are then the primary focus of policy reforms undertaken by those countries receiving funds from the IMF during times of crisis.

I.D. The Organization for Economic Cooperation and Development

The OECD’s estimates of potential output are based on a production function approach. In particular, the OECD uses a Cobb–Douglas production function with constant returns to scale that combines physical capital, human capital, labor, and labor-augmenting technological progress. Each of these inputs is projected using a trend, and TFP is assumed to converge to a certain degree among different countries in the medium run. As pointed out by the OECD (2012, 195): “The degree of convergence in total factor productivity depends on the starting point, with countries farther away from the technology frontier converging faster, but it also depends on the country’s own structural conditions and policies.” Note that when forecasting potential output in the medium term,

3. These countries are Australia, Austria, Belgium, Canada, Switzerland, Cyprus, the Czech Republic, Germany, Denmark, Spain, Estonia, Finland, France, the United Kingdom, Greece, Hungary, Ireland, Iceland, Italy, Japan, South Korea, Luxembourg, Malta, the Netherlands, Norway, New Zealand, Poland, Portugal, Slovak Republic, Slovenia, Sweden, Turkey, and the United States. For more information on the time periods for the data on these countries, see online appendix table 1. The online appendixes for this and all other papers in this volume may be found at the *Brookings Papers* web page, www.brookings.edu/bpea, under “Past BPEA Editions.”

the OECD assumes that output gaps close over a period of 4 to 5 years, depending on their initial size. Therefore, one should expect to see above average future growth for countries with large output gaps. Relative to the IMF, the OECD covers more countries and has longer time series (see the online appendix). For many countries, nowcasts and one-year-ahead forecasts have been available since 1989. Since 2005, the OECD has also reported five-year-ahead forecasts for potential output. As with the IMF, estimates of potential output in the OECD are commonly used to assess cyclically adjusted fiscal balances and to characterize the need for structural reforms.

1.E. Consensus Economics

Consensus Economics, a global survey firm of professional forecasters, does not provide estimates of potential output, but it does report forecasts for the growth rate of actual output from 1 to 10 years into the future. Because estimates made for several years into the future (for example, years 6 through 10) are likely to be independent of business cycle conditions, we use these long-run estimates as an approximation of the growth rate of potential output at the same horizon. These data are available for 12 countries, and the starting date varies across countries from 1989 to 1998 (see online appendix table 1). Given the wide range of forecasters included in the Consensus Economics forecasts, one cannot readily summarize how these forecasts are made. Private forecasts, however, are widely used in both public and international organizations for comparison purposes with in-house forecasts.

1.F. Comparison of Potential Output Measures

Table 1 documents some basic moments for estimates of the potential output growth rate (nowcasts) produced by the IMF and OECD, as well as the forecasted long-term output growth rate from Consensus Economics. We work with growth rates of potential output rather than levels because the definition of output varies across time (base year) and agencies. The growth rate series are highly correlated and generally have similar moments across sources. This is especially true for the IMF and OECD forecasts, which conceptually are measuring the same objects (nowcasts of potential GDP). The Consensus Economics forecasts, in contrast, are at a different horizon and are for actual rather than potential GDP. These strong correlations are not driven by outliers. Indeed, there are few large differences across sources, and these tend to be concentrated in a handful of countries and periods (see online appendix figure 1).

Table 1. Comparison Output Measures from the IMF, OECD, and Consensus Economics^a

<i>Basis for comparison and correlation</i>	<i>Institution and output measure</i>		
	<i>IMF: Potential output growth rate (nowcast)</i>	<i>OECD: Potential output growth rate (nowcast)</i>	<i>Consensus Economics: 6- to 10-year-ahead forecast for actual output growth rates</i>
Observations	607	1,358	581
Mean	1.64	2.30	2.22
Standard deviation	1.10	1.25	0.54
Correlation			
IMF	1.00		
OECD	0.87	1.00	
Consensus Economics	0.72	0.78	1.00


a. This table reports moments of measures of potential output from the IMF and OECD across the countries described in online appendix table 1 and listed in footnote 3 in the text, as well as moments of forecasted growth rates of GDP 6 to 10 years ahead from Consensus Economics. See subsection I.F for details.

Figure 2 illustrates that this strong correlation across series is not restricted to differences in growth rates across countries. Time series for the growth rate of U.S. potential output across the different institutions that produce estimates (Fed Greenbooks, CBO, IMF, OECD, and Consensus Economics long-term forecasts of actual output) track each other closely as well. There are nonetheless occasional differences across estimates. After the 1990–91 recession, for example, the CBO reduced its estimate of potential GDP growth significantly more than the staff of the Federal Reserve Board, whereas private forecasters hardly changed their long-term forecasts of growth at all. After the Great Recession, the IMF and OECD both lowered their estimates of potential GDP growth far more than the Fed Greenbooks or the CBO, but then revised them back up while the CBO continued to progressively revise its estimates of potential GDP growth down.

Figure 3 plots a longer time series of estimates of potential GDP available from the Fed Greenbooks, as extended backward by Orphanides (2004). In addition, we plot several statistical approaches to estimating potential GDP, including a one-sided, five-year moving average of real-time GDP and a one-sided HP-filter ($\lambda = 500,000$) of real-time GDP. The HP filter tracks the Greenbooks' estimates of potential output quite closely, especially since the mid-1980s, while the moving-average approach tends to display larger fluctuations. All series co-move relatively closely with a moving average of capacity-adjusted TFP changes as measured by John Fernald (2012).

Figure 2. Estimates of Potential U.S. Output Growth Rate and Forecasted Long-Term Growth for Actual Output^a

Growth rate of potential output (annualized percentage)



Sources: International Monetary Fund; Congressional Budget Office; Organization for Economic Cooperation and Development; Federal Reserve Greenbooks; Consensus Economics.

a. All series in the figure are real-time data at the semiannual frequency. The potential outputs for the IMF, OECD, and CBO are reported for the current calendar year. Potential output for Greenbooks is the semiannual average of quarterly growth rates of potential output for the quarters in a given semester. Series for Consensus Economics show the 6- to 10-year-ahead forecast for the actual output growth rate (per year).


The persistence in revisions of potential GDP shown in figures 2 and 3 suggests that some of these revisions might be predictable from recent changes. We evaluate this formally by regressing revisions of potential GDP on lags of itself:

$$(1) \quad (\Delta \log Y_{it}^* - \Delta \log Y_{it-1}^*) = \alpha + \beta (\Delta \log Y_{it-1}^* - \Delta \log Y_{it-2}^*) + error_t$$

where $\Delta \log Y_{it}^*$ is the growth rate of potential output in time t according to a projection made at time s . We find (table 2) a mild amount of predictability in the Greenbooks' revisions of potential GDP. With the CBO, the coefficient on lagged revisions is similar but not significantly

Figure 3. Real-Time Estimates of U.S. Potential Output Growth Rate and Trends in Actual Output Growth Rate^a

Percent per year



Sources: Orphanides (2004); Federal Reserve Greenbooks; Federal Reserve Bank of Philadelphia.

a. All series are real time at the quarterly frequency. Potential output for the pre-1987 period is from Orphanides (2004). Potential output for 1987–2011 is from the Federal Reserve Bank of Philadelphia. Potential output is measured as the growth rate of potential output between a given quarter and the next three quarters. *HP-filtered actual output* (HP = Hodrick–Prescott) is calculated as the value of the one-sided HP filter trend for the quarter given the first vintage of GDP data that covers the given quarter, with HP filter smoothing parameter of 500,000. *MA (20) actual output* is calculated as the 20-quarter moving average over the current and preceding 19 quarters reported in the first vintage of GDP data that covers the given quarter. *MA (20) total factor productivity* for a given quarter is calculated as the 20-quarter moving average running on the current quarter and the preceding 19 quarters.

different from zero. The results are different for international data, with coefficients on past OECD revisions being not different from zero and with those on past IMF and Consensus Economics revisions exhibiting negative predictability.

II. How Estimates of U.S. Potential Output Are Adjusted after Economic Shocks

Although a limited unconditional predictability is a desirable attribute of estimates of potential GDP, it does not imply that there is no predictability in estimates of potential output *conditional* on different economic shocks.

Table 2. Predictability of Revisions in Estimates of Potential GDP^a

<i>Dependent variable:</i> ($\log Y_{it}^* - Y_{it-1}^*$)	<i>Source</i>				
	<i>CBO</i> (1)	<i>Fed Greenbooks</i> (2)	<i>OECD</i> (3)	<i>IMF</i> (4)	<i>Consensus Economics</i> (5)
($\log Y_{it-1}^* - \log Y_{it-2}^*$)	0.204 (0.132)	0.294*** (0.086)	-0.066 (0.040)	-0.154*** (0.044)	-0.355*** (0.045)
Observations	42	96	1,282	548	566
R^2	0.065	0.085	0.163	0.351	0.288
Number of countries			31	27	12

a. This table presents regressions of the revision in estimates of potential GDP on the previous revisions in estimate of potential GDP (equation 1). Newey–West standard errors are in parentheses. “Source” indicates where estimates of potential output come from the Congressional Budget Office (CBO), Greenbooks of the Federal Reserve Board, the Organization for Economic Cooperation and Development (OECD), the International Monetary Fund (IMF), and Consensus Economics. For Consensus Economics, revisions are for the growth rate of GDP at horizons of 6 to 10 years. Columns 3–5 are across countries and include time and country fixed effects. Within R^2 is reported for columns 3–5.

To assess how estimates of potential output respond to economic shocks, we combine the estimates described in the previous section with identified measures of economic or policy shocks.

II.A. Measures of Economic Shocks

There is an extensive literature on identifying shocks that potentially drive business cycle and longer-term fluctuations, particularly for the United States (for a survey, see Ramey 2016). Following this literature, we employ several measures of both “demand” and “supply” shocks for the U.S. Our use of the terms “supply” and “demand” reflects a certain abuse of terminology. All the shocks we consider have both supply and demand effects in modern business cycle models. Our classification instead primarily relies on whether these shocks appear to have permanent or transitory effects on GDP. We define demand shocks as those whose real effects appear to be transitory and therefore should not affect estimates of potential output.⁴

For supply shocks, we consider changes in TFP, oil price shocks, and tax shocks. TFP changes are measured as by Fernald (2012), who adjusts Solow residuals for time-varying utilization of inputs. Although these data are somewhat sensitive to vintage (see Kurmann and Sims 2017), we rely on the final vintage of the data because the data by vintage are available for relatively recent times. For oil price shocks, we use oil supply shocks as

4. Because the units of these shocks vary, we normalize all shocks to be mean zero and have unit variance.

identified by Lutz Kilian (2009).⁵ For tax shocks, we use Christina Romer and David Romer's (2010) narrative measure of exogenous tax changes. To be clear, tax shocks have both demand and supply effects. We denote them here as "supply" shocks because Romer and Romer (2010) document that they have permanent effects on output, and therefore should be captured by estimates of potential GDP.

We consider three identified demand shocks, all related to policy. The first are monetary policy shocks. For the United States, our baseline measure of these shocks follows the quasi-narrative approach of Romer and Romer (2004). They use the narrative record to construct a consistent measure of policy changes at FOMC meetings since 1969, then orthogonalize these policy decisions to the information available to policymakers at each FOMC meeting, as captured by the Greenbook forecasts prepared by the staff of the Federal Reserve Board before each FOMC meeting. The unexplained policy changes are then defined as the monetary shocks. We use the updated version of these shocks from Coibion and others (2017) and set values after the onset of the zero lower bound equal to zero.⁶

The second type of demand shock we consider are the military spending news shocks given by Valerie Ramey (2016). Using real-time measures of the expected future path of defense spending in the United States, Ramey constructs a measure of the present discounted value of future defense expenditures for each quarter. Changes in these measures from one quarter to the next thus reflect changes in either current or future defense spending.

Finally, we consider a broader measure of government spending shocks, namely, differences between ex-post government spending and ex-ante forecasts of this spending following Alan Auerbach and Gorodnichenko (2012b). Unlike the Ramey news measure, this measure captures unanticipated short-run changes in government spending but is broader in that it includes more than just military spending.

All three types of demand shocks have repeatedly been found to have only transitory effects on GDP (see Nakamura and Steinsson 2017; Ramey 2016), so there is little evidence supporting the hysteresis hypothesis that transitory shocks have long-lived effects on output (and therefore potential) through endogenous productivity or tax responses. As emphasized by

5 We also tried using the oil shocks identified by Baumeister and Hamilton (2015) in place of the ones identified by Kilian (2009). The results were very similar and are available from the authors upon request.

6. We also experimented with monetary policy shocks identified via recursive ordering of vector autoregression residuals, as done by Bernanke and Blinder (1992), and we found similar results, as documented in online appendix figure 5.

Blanchard (2017), these transitory shocks could still affect potential GDP in a transitory fashion in the presence of physical or human capital. As a result, we study not just the response of nowcasts of potential GDP to these shocks but also of long-run forecasts of potential from the CBO as well as long-run forecasts of GDP growth from private forecasters. The latter two should unambiguously not respond to these transitory shocks. Finally, even if the real world were characterized by hysteresis, monetary policymakers explicitly rule out this channel and emphasize that, in their view, monetary policy has only transitory effects on GDP.⁷ Their estimates of potential GDP should therefore be invariant to monetary shocks.

II.B. Effects of Shocks on Actual Output and Estimates of Potential Output in the United States

To provide a benchmark for how we might expect estimates of potential output to respond to economic shocks, we first characterize the response of actual output to these shocks. Specifically, we regress ex-post changes in output on current and past values of a shock, as follows:

$$(2) \quad \Delta \log Y_t = \alpha + \sum_{k=0}^K \phi_k \epsilon_{t-k} + \text{error}_t,$$

where t indexes time (quarters), $\Delta \log Y_t$ is the growth rate of real GDP, ϵ is an identified shock, and error is the residual. A key advantage of this moving-average specification is that it allows us to handle data with mixed frequencies and gaps in the time series as well as correlations of the error term. For consistency, we run these regressions at the same time frequency as what is available for estimates of potential output, namely, quarterly when comparing with Greenbook forecasts, and semiannually otherwise. Because Greenbook forecasts of potential output begin in 1987, we run the regression for output over the same time sample. Given the limited number of observations available, we include only one shock at a time (the shocks are roughly uncorrelated). Because the error term is not necessarily white noise, we use Newey–West standard errors everywhere.⁸ Impulse responses

7. For example, in a speech on March 3, 2017, Janet Yellen stated that “monetary policy cannot, for instance, generate technological breakthroughs or affect demographic factors that would boost real GDP growth over the longer run or address the root causes of income inequality. And monetary policy cannot improve the productivity of American workers. Fiscal and regulatory policies—which are of course the responsibility of the Administration and the Congress—are best suited to address such adverse structural trends” (Yellen 2017).

8. Because the null hypothesis we are testing is that of zero response of output and potential output, the fact that shocks are estimated does not constitute an issue for standard errors and tests of the null hypothesis, as shown by Pagan (1984).

Figure 4. Responses of U.S. Output and Greenbook Estimates of Potential U.S. Output to Shocks^a

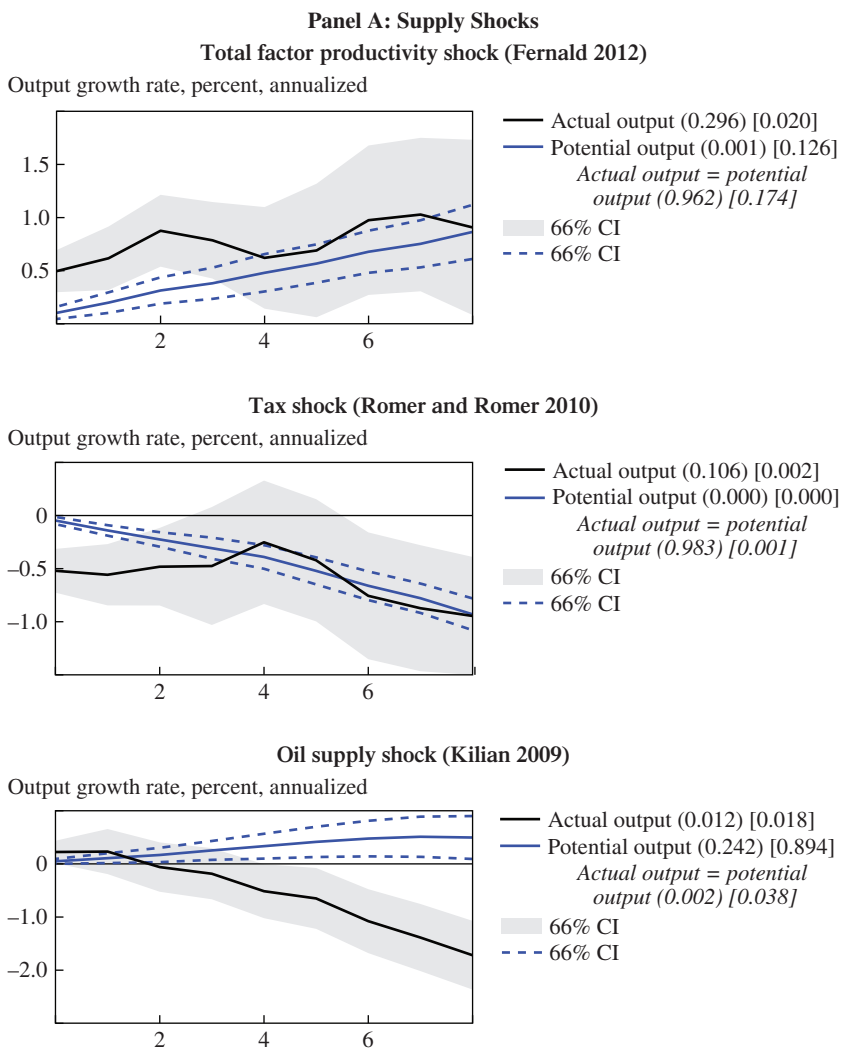
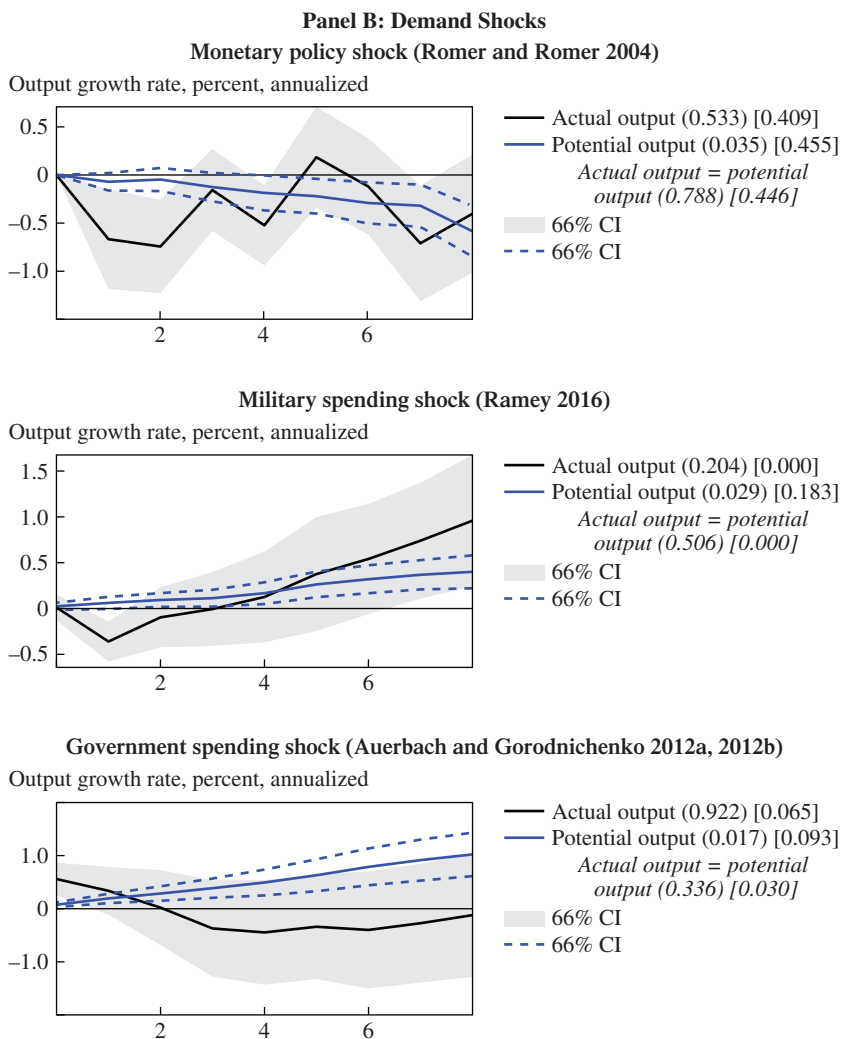


Figure 4. Responses of U.S. Output and Greenbook Estimates of Potential U.S. Output to Shocks^a (*Continued*)

Sources: Authors' calculations, with potential output from Federal Reserve Greenbooks and identified shocks from Fernald (2012); Romer and Romer (2004, 2010); Kilian (2009); Ramey (2016); Auerbach and Gorodnichenko (2012a, 2012b).

a. This figure reports impulse response functions (IRFs) estimated using equations 2 and 3. The estimation sample covers the longest possible period with nonmissing observations for shocks and potential output (output gap) available at the Federal Reserve Bank of Philadelphia. In parentheses, we report the p values for a test of whether the response of actual (potential) output is different from zero at the maximum horizon (eight quarters). In square brackets, we show the p values for a test of whether the path of the response of actual (potential) output is different from zero over the entire duration of the IRF. The last row of the legend—for which there is no line in the graphs—reports p values for a test of equality of responses of actual and potential output at the maximum horizon (parentheses) and a test of equality of the paths of the responses for actual and potential output are equal across horizons. CI = confidence interval.

come directly from the estimates of ϕ . To recover responses of the *level* of output, we cumulate ϕ_k up to a given horizon. For example, the level responses are ϕ_0 for $h = 0$, $\phi_0 + \phi_1$ for $h = 1$, $\phi_0 + \phi_1 + \phi_2$ for $h = 2$, and so on.⁹

For each impulse response, we include 66 percent confidence intervals and the legend of each associated graph reports the p values for two types of tests. In parentheses we report the p value for a test of whether the response of actual output is different from zero at the maximum horizon (eight quarters), while in square brackets we show the p value for a test of whether the path of the response of actual output is different from zero over the entire horizon of the impulse response. These p values are also included in panel A of online appendix table 2, together with more information that we describe later in the paper.

We plot the responses of actual output to each type of shock in figure 4, which appears on the previous two pages. Panel A of the figure focuses on the three supply shocks. In response to a TFP shock, output immediately rises about 0.5 percentage point and remains persistently higher by about this magnitude. Hence, these TFP shocks appear to have permanent effects on output. Tax increases have a (negative) contemporaneous effect on output that is similarly sustained over the entire impulse response horizon. In contrast, negative oil supply shocks have a more delayed effect on output, but are associated with a long-lived decline in GDP. In short, all three supply shocks have the expected long-lived effects on GDP. As a result, we would expect them to be captured by high-quality measures of potential GDP.

Turning to demand-side shocks (panel B of figure 4), we again find the expected responses of output. Contractionary monetary policy shocks push output down. The point estimates are much less precise than those of Romer and Romer (2004), reflecting the shorter time sample, the fact that monetary shocks are smaller over this limited sample, and the different approach to estimating impulse responses. Increases in expected military expenditures have a delayed positive effect on GDP (which reflects the fact that the expenditures themselves are also generally delayed).¹⁰ Immediate spending shocks, as given by Auerbach and Gorodnichenko (2012b), have transitory, short-run effects on GDP and no long-run effects. Demand-side shocks therefore generally deliver cyclical variation in output but no long-run effects on GDP. As a result, we would expect high-quality measures of potential GDP to be insensitive to these shocks.

9. For monetary policy shocks, we constrain $\phi_0 = 0$ to capture the minimum delay restriction.

10. Although our horizon of impulse responses is too short to illustrate this, Ramey (2016) shows that news about future military spending has only transitory effects on GDP.

To characterize the effects of these economic shocks on estimates of potential output, we run equivalent specifications:

$$(3) \quad \Delta \log Y_{it}^* = \alpha + \sum_{k=0}^K \phi_k \epsilon_{t-k} + error_t,$$

where $\Delta \log Y_{it}^*$ is the (nowcast) estimated growth in potential in quarter t given information in quarter t at an annualized rate. We first consider Greenbook estimates of potential output and extend our results to alternative estimates of potential in subsequent sections. Responses of the implied *level* of potential output are constructed in the same way as before. For comparison, we plot the responses of potential output in the same graphs as the responses of actual output, and we also include 66 percent confidence intervals and the p values for the same tests mentioned above (now for the responses of potential output instead of actual output). Finally, we also include the p values for a test of whether paths of the responses for actual and potential output are equal over the entire duration of the impulse response (in square brackets) and the p values of a test of whether the responses are equal at the maximum horizon (in parentheses). The p values are also included in panel A of online appendix table 2.

Looking first at TFP shocks, we find that estimates of potential GDP respond very gradually but in the same direction as actual GDP. The shock has little immediate impact on estimates of potential; but after two years, the responses are overlapping and estimates of potential GDP have caught up to actual GDP. Very similar results are obtained with tax shocks: Estimates of potential GDP are unchanged immediately after the shock, but gradually converge to the path of actual GDP. Hence, with both TFP and tax shocks, one would ultimately attribute the decline in output to a decline in potential output, but only with some delay. One possible reason for delayed responses of forecasts is information rigidity, as suggested by Coibion and Gorodnichenko (2012, 2015a). However, the fact that estimates of potential GDP evolve very gradually after tax shocks (which occur only for large legislative tax changes of which staff members at the Fed would be well aware) suggests that other mechanisms must be at play to explain the inertia in real-time estimates of potential output.

Turning to the response to oil price shocks, we find a starkly different response: Estimates of potential GDP *increase* over time while actual GDP *falls*. In contrast to TFP and tax shocks, in which the long-run response of output is ultimately matched by the response of potential, contractionary oil price shocks are associated with sharply falling measured output gaps (Y_t/Y_t^*) in the long run, as estimates of potential are progressively increased

while output itself is falling. Policymakers facing a trade-off between stabilizing inflation (which rises after a negative oil supply shock, thereby calling for higher interest rates) and closing the output gap (which is falling, calling for lower interest rates) are therefore perceiving an even starker trade-off because the rise in the estimate of potential output makes the output gap seem even more negative.¹¹ This result is not driven by the specific measure of oil supply shocks (we find a similar result with Kilian's 2008 measure of OPEC supply shocks) or by the sample period (we find similar results for alternative periods).

There are several potential explanations for this finding. One is that policymakers are confounding oil supply and demand shocks: If they observe a supply-driven increase in oil prices that they incorrectly attribute to stronger global demand for oil from, for example, improved technology, then this might lead them to revise their estimates of potential GDP upward, even as actual GDP is falling. An alternative explanation is that higher oil prices might be perceived as inducing greater investment in new energy sources and alternative energy technologies, which could then raise potential GDP in the long run, even as short-run GDP falls, though there is little evidence that GDP ultimately responds in a positive manner. The available data unfortunately do not enable us to identify the underlying explanation. If nothing else, this result provides a surprising example of how estimates of potential GDP can move in the direction opposite to that of actual GDP.

Turning to demand shocks, we again observe important deviations from what one would expect of estimates of potential GDP. With monetary and both types of fiscal shocks, estimates of potential respond little on impact to these shocks but progressively respond in the same manner as the short-run response of GDP. The transitory decline in GDP after a contractionary monetary shock is followed by a persistent decline in the real-time estimates of potential GDP, while the transitory increase in output after an increase in government spending is followed by a persistent rise in estimates of potential GDP. Hence, these *cyclical* fluctuations in output lead to the perception among forecasters that they are *permanently* affecting output, as if they were TFP or tax shocks, despite the fact that their effects on income are actually short-lived.

11. The pronounced decline in the *perceived* output gap after oil supply shocks is consistent with the view that monetary policymakers were too willing to accommodate these shocks with lower interest rates, and that this accommodation may have exacerbated the Great Inflation of the 1970s.

Our results are not limited to these specific examples of identified shocks. For example, we can identify supply and demand shocks jointly, as was done by Blanchard and Quah (1989), by running a vector autoregression (VAR) with output growth and unemployment and restricting demand shocks to have no long-run effects on output. When we use these supply and demand shocks to characterize the response of real-time estimates of potential output over the same period, we again find that real-time estimates respond very gradually to both shocks, moving in the direction of the change in output (online appendix figure 2). Importantly, because this identification explicitly imposes the fact that only supply shocks have permanent effects on GDP, it addresses the possibility that some demand shocks might have hysteresis effects and therefore should be incorporated into estimates of potential GDP. In short, across identification schemes, we find an overresponse of real-time estimates of potential GDP to demand shocks and an underresponse to supply shocks.

II.C. The Robustness of Baseline Results for the United States

Because of the relatively short samples involved, we want to verify that our results are robust to a range of reasonable variations. Our first check is on the empirical method used to estimate impulse responses. As an alternative to equations 2 and 3, we reproduce impulse responses of actual output and nowcasts of potential GDP to each of the shocks using autoregressive distributed lag specifications to estimate impulse response functions (IRFs), as done by Romer and Romer (2004), namely:

$$(4) \quad \Delta \log Y_t = \alpha + \sum_{j=1}^J \delta_j \Delta \log Y_{t-j} + \sum_{k=0}^K \phi_k \epsilon_{t-k} + error_t,$$

using $J = 4$ and $K = 8$. The results are presented in online appendix figure 3. By and large, the results are very similar. With productivity and tax shocks, we continue to find persistent but delayed effects on estimates of potential GDP that are ultimately converging to the responses of actual GDP. Similarly, with all three demand shocks, we find the same qualitative patterns as with the previous empirical specification. The only difference lies in the response to oil supply shocks, where we no longer observe a pronounced rise in estimates of potential GDP. Instead, our estimates indicate no response of the nowcasts of potential, suggesting some sensitivity in this result.

One potential source for this empirical sensitivity is the limited time sample. As a result, we replicate our baseline results over an extended time period, where for each shock we now use the maximum time sample

Figure 5. Responses of U.S. Output and Greenbook Estimates of Potential U.S. Output to Shocks: Extended Sample^a

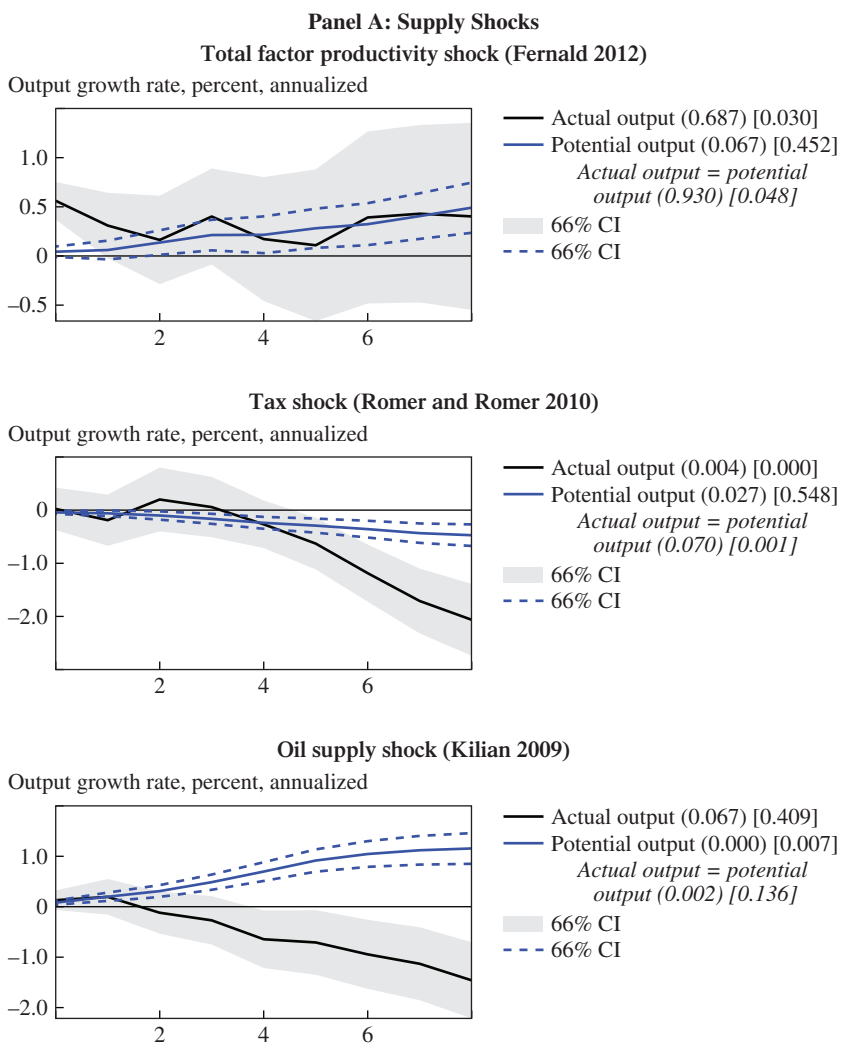
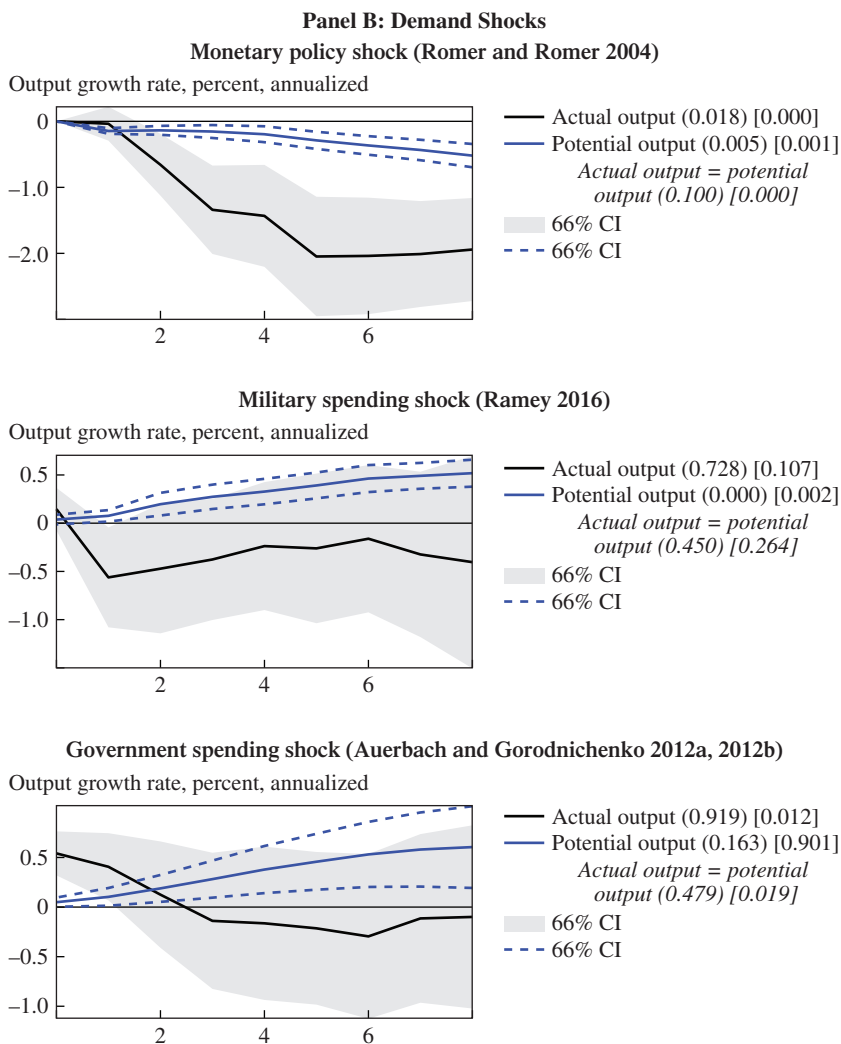


Figure 5. Responses of U.S. Output and Greenbook Estimates of Potential U.S. Output to Shocks: Extended Sample^a (*Continued*)

Sources: Authors' calculations, with potential output from Federal Reserve Greenbooks and identified shocks from Fernald (2012); Romer and Romer (2004, 2010); Kilian (2009); Ramey (2016); Auerbach and Gorodnichenko (2012a, 2012b).

a. This figure reports impulse response functions (IRFs) estimated using equations 2 and 3. The estimation sample covers the longest possible period with nonmissing observations for shocks and potential output (output gap) using output gap data starting in 1970. In parentheses, we report the p values for a test of whether the response of actual (potential) output is different from zero at the maximum horizon (eight quarters). In square brackets, we show the p values for a test of whether the path of the response of actual (potential) output is different from zero over the entire duration of the IRF. The last row of the legend—for which there is no line in the graphs—reports p values for a test of equality of responses of actual and potential output at the maximum horizon (parentheses) and a test of equality of the paths of the responses for actual and potential output are equal across horizons. CI = confidence interval.

available across both the shocks and the Greenbook estimates of potential GDP (1969–2011). The results, presented in figure 5 (which appears on the previous two pages), confirm our baseline findings: There is a delayed but persistent response of the estimates of potential GDP to all shocks. In every case but oil supply shocks, the nowcasts evolve in the direction of the short-run changes in GDP. With oil supply shocks, the estimates of potential GDP rise in an even more pronounced fashion, while actual output falls.¹² Hence, the baseline results are not specific to the period since 1987.

We also consider whether our results are sensitive to relying on nowcasts of potential GDP growth. Because the Fed Greenbooks also include forecasts and backcasts of potential GDP growth (two years in each direction), we can characterize how the perceived *path* of potential GDP evolves after each shock. We find very little difference relative to nowcasts, implying that Federal Reserve staff members raise or lower the entire path of projected and past potential GDP growth in response to shocks (online appendix figure 4).

Another potential issue with these results is our reliance on estimates of potential GDP from a single source: the staff of the Federal Reserve Board. In figure 6 (which appears two pages down from here), we reproduce our results using estimates of potential GDP from the Congressional Budget Office. One advantage of CBO estimates is they are available at longer horizons. As a result, we consider both “nowcasts” of potential GDP (equivalent to Greenbook estimates) as well as five-year-ahead forecasts (that is, the growth rate of potential output in five years from the date when a forecast is made). A disadvantage of CBO estimates, as discussed in subsection I.A, is that the sample for these is more limited and the time frequency at which forecasts are available is reduced. Not surprisingly, the effects of each shock on GDP are therefore considerably less precisely estimated. However, the responses of the estimates of potential GDP are still quite precise. Qualitatively, we find that the CBO’s estimates of current potential GDP respond much like those from the Greenbooks: gradually but persistently to all shocks. Long-run forecasts of potential GDP generally respond by less than those of current potential GDP. However, they still ultimately respond to demand shocks, implying that the CBO implicitly interprets cyclical shocks as having permanent effects on GDP.

12. When we apply the autoregressive distributed lag specification to oil supply shocks over the whole sample, we find the same result.

The fact that CBO forecasts of *long-run* potential respond similarly to nowcasts of potential GDP addresses one possible issue raised by Blanchard (2017), namely, that demand shocks might have transitory effects on potential output. This can occur even in standard models through a number of channels, such as lower levels of physical capital following periods of disinvestment or lower levels of human capital after extended unemployment stretches. But in these models, demand shocks would still have only transitory effects on potential output, so forecasts of *long-run* potential output should remain unchanged after demand shocks, even if contemporaneous levels of potential were responding to these shocks. The fact that both nowcasts and long-run forecasts of potential respond to demand shocks suggests that the mechanism emphasized by Blanchard (2017) is not driving these results.

In short, we document a systematic response of estimates of potential GDP to shocks that have only cyclical effects on GDP. Furthermore, even some supply shocks have contradictory effects on estimates of potential GDP, in the sense that changes in the latter after oil supply shocks speak little to actual long-run changes in output. Thus, seeing ex-post that declines in GDP seem to be accounted for by changes in potential GDP, as has been the case in the U.S. since the Great Recession, says little about whether the decline in output is likely to persist or can be reversed by standard countercyclical policies.

II.D. Explaining Patterns in Impulse Responses

Why are estimates of potential GDP responding to shocks that only have cyclical effects, such as monetary policy and government spending shocks? One possibility is that policy institutions and statistical agencies perceive these shocks as affecting current levels of potential output (for example, if they affect current capital stocks) but not long-run levels of potential output (as would be implied by, for example, monetary neutrality). This is unlikely to be the case, however, because the long-horizon CBO forecasts of potential GDP respond about as much as their nowcasts of potential GDP.

An alternative possibility is that these estimates are relying to a large extent on simple statistical methods to measure trend (potential) levels from actual GDP. As illustrated in figure 3, one can come close to replicating the real-time Greenbook estimates of potential GDP growth by using a one-sided HP-filter on real-time GDP data available each quarter or by taking a simple, one-sided moving average of recent GDP outcomes. Because these types of methods fail to identify the different potential sources of

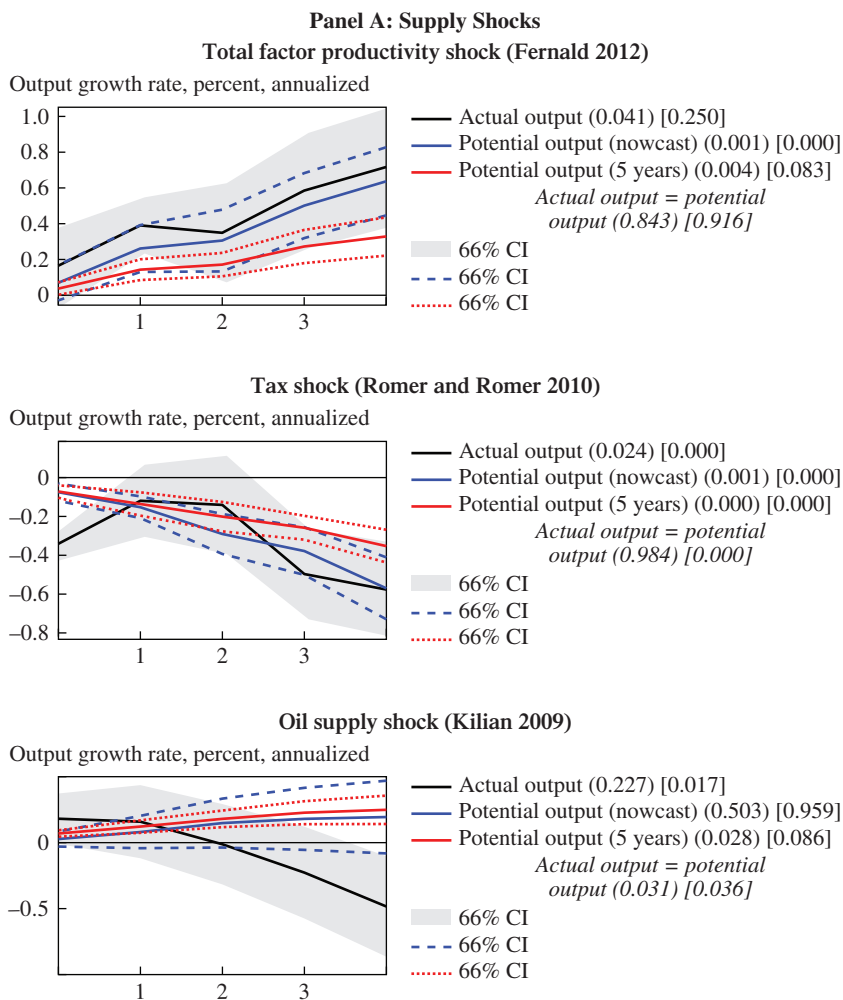
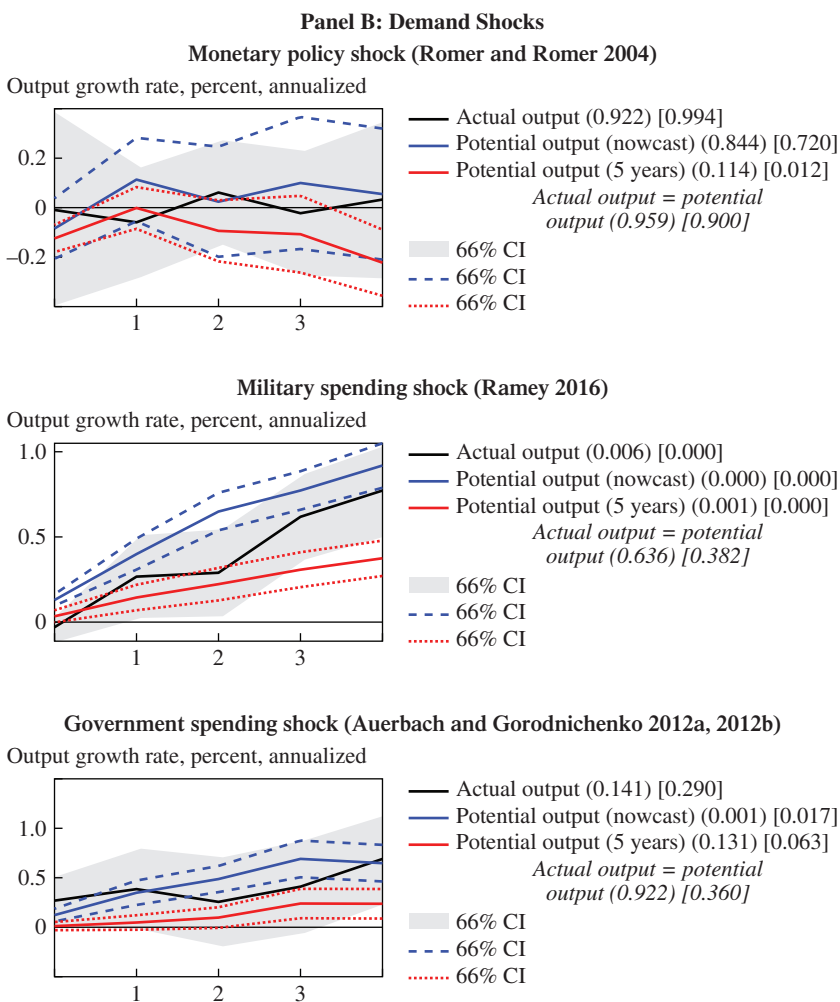
Figure 6. Responses of U.S. Output and the CBO's Estimates of Potential U.S. Output to Shocks^a

Figure 6. Responses of U.S. Output and the CBO's Estimates of Potential U.S. Output to Shocks^a (*Continued*)

Sources: Authors' calculations, with potential output from Federal Reserve Greenbooks and identified shocks from Fernald (2012); Romer and Romer (2004, 2010); Kilian (2009); Ramey (2016); Auerbach and Gorodnichenko (2012a, 2012b).

a. This figure reports impulse response functions (IRFs) estimated using equations 2 and 3. The estimation sample covers the longest possible period with nonmissing observations for shocks and potential output (output gap) available from the Congressional Budget Office. In parentheses, we report the p values for a test of whether the response of actual (potential) output is different from zero at the maximum horizon (8 quarters). In square brackets, we show the p values for a test of whether the path of the response of actual (potential) output is different from zero over the entire duration of the IRF. The last row of the legend—for which there is no line in the graphs—reports p values for a test of equality of responses of actual and potential output (nowcast) at the maximum horizon (parentheses) and for a test of equality of the paths of the responses for actual and potential (nowcast) output are equal across horizons. CI = confidence interval.

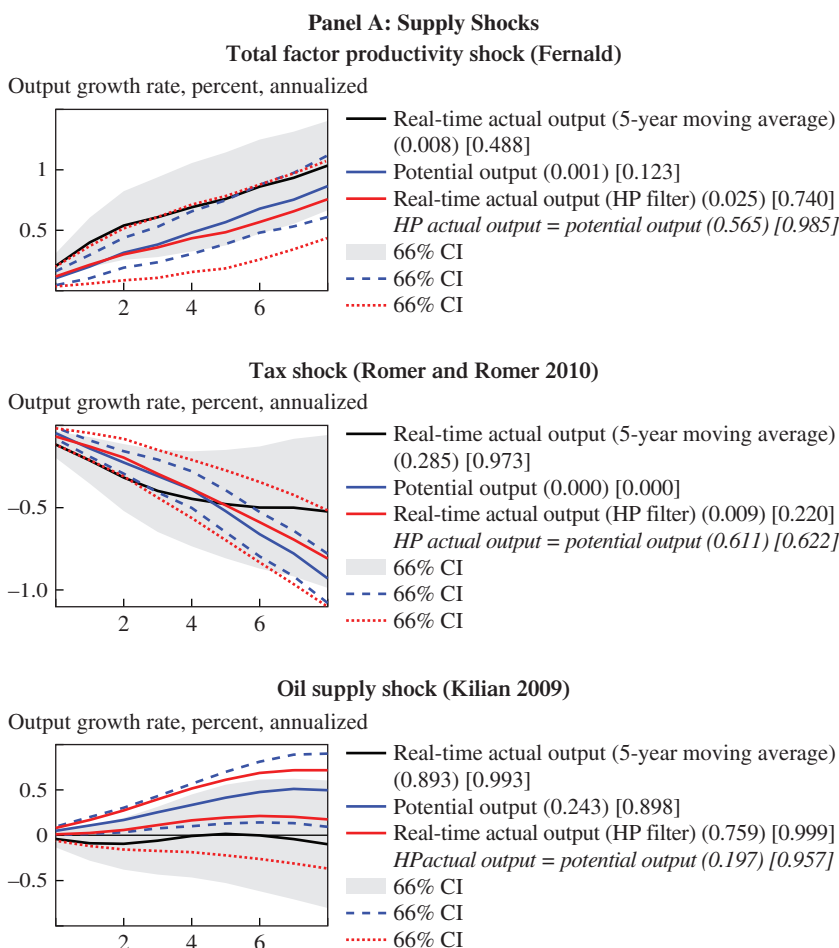
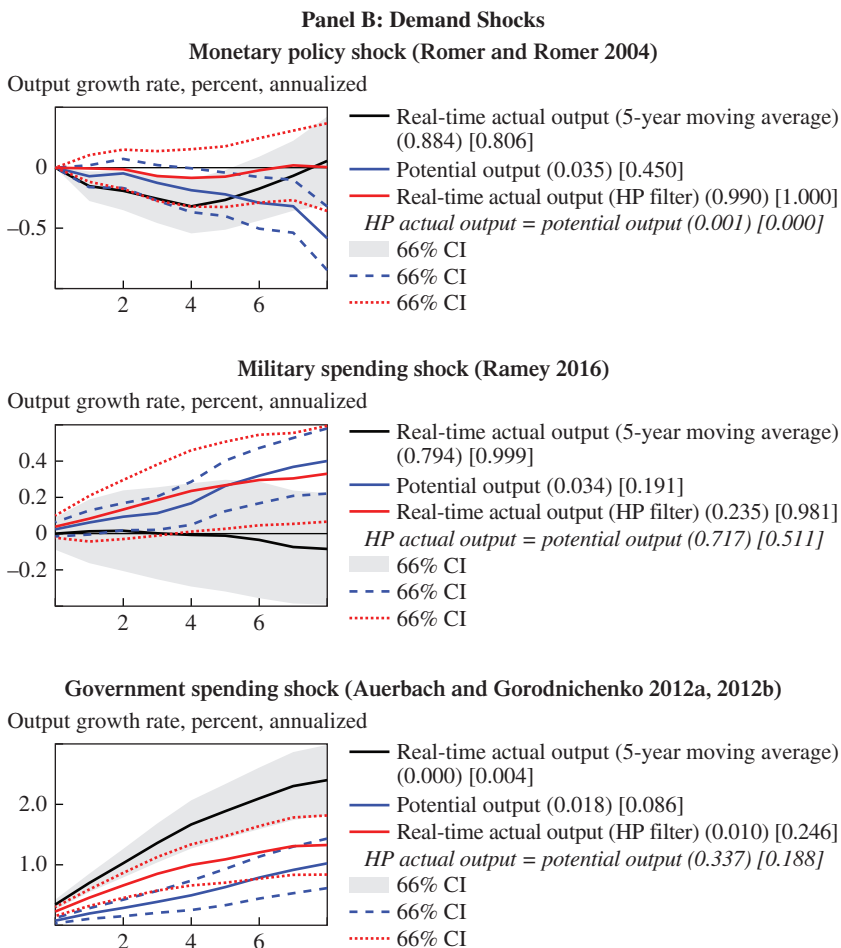
Figure 7. Responses of Greenbook Estimates of Potential U.S. Output and HP-Filtered Output to Shocks^a

Figure 7. Responses of Greenbook Estimates of Potential U.S. Output and HP-Filtered Output to Shocks^a (*Continued*)

Sources: Authors' calculations, with potential output from Federal Reserve Greenbooks and identified shocks from Fernald (2012); Romer and Romer (2004, 2010); Kilian (2009); Ramey (2016); Auerbach and Gorodnichenko (2012a, 2012b).

a. This figure reports impulse response functions (IRFs) estimated using equations 2 and 3. The estimation sample covers the longest possible period with nonmissing observations for shocks and potential output (output gap) available at the Federal Reserve Bank of Philadelphia. HP-filtered (HP = Hodrick-Prescott) actual output for a given quarter is calculated as the value of the HP-filter trend for the quarter given the first vintage of GDP data that covers the given quarter. The smoothing parameter for the HP filter is set at 500,000. The five-year moving average (MA) actual output for a given quarter is calculated as the 20-quarter MA running on the current quarter and the preceding 19 quarters reported in the first vintage of GDP data that covers the given quarter. In parentheses, we report the p values for a test of whether the response of actual (potential) output is different from zero at the maximum horizon (eight quarters). In square brackets, we show the p values for a test of whether the path of the response of actual (potential) output is different from zero for all horizons of the IRF. The last row of the legend—for which there is no line in the graphs—reports p values for a test of equality of responses of actual and potential output at the maximum horizon (parentheses) and for a test of equality of the paths of the responses for actual and potential output are equal across horizons. CI = confidence interval.

changes in economic activity, they would naturally lead to slow-moving dynamic responses to *all* economic shocks that move actual output.

To assess this possibility, we replicate our baseline impulse responses using the same two statistical approaches to estimating potential GDP as in figure 3. In the first case, we apply a one-sided HP filter with smoothing parameter $\lambda = 500,000$ to real-time data on GDP. In the second, we take a five-year moving average of real GDP using real-time data. We present the results, along with the responses of potential GDP as measured by the Greenbooks in figure 7 (which is on the two preceding pages; also see the p values included in panel C of online appendix table 2). When using the HP-filtered series, we can very closely replicate the response of estimated potential GDP after every shock.¹³ With the moving average, the fit is not as strong. The very close fit of the impulse responses using the HP filter, as well as how closely one can reproduce the unconditional time series of historical estimates of potential GDP in figure 3 with an HP-filtered series, suggests that Greenbook estimates of potential GDP incorporate little additional information relative to this purely statistical approach to estimating potential GDP.¹⁴ It is then quite natural for these series to respond to all shocks that affect GDP, even if these movements are transitory. But this endogenous response to cyclical shocks should not be interpreted as reflecting permanent effects of these shocks on output but rather as a mechanical reaction based on how estimates of potential GDP are constructed. Equivalently, observing a downward revision in real-time estimates of potential GDP is *not* informative about whether the associated declines in actual GDP are likely to be sustained.

Another way to see how closely the HP filter can mimic real-time estimates of potential GDP, along with the potential dangers of doing so, is

13. The fact that we can match the increase in estimated potential output after an oil supply shock with the HP filter points toward a possible identification issue with these shocks. They are identified from a three-variable VAR of oil production, global economic activity (measured using an index of shipping prices), and oil prices. If oil prices are disproportionately sensitive to U.S. output (rather than global output) or shipping prices are an otherwise imperfect measure of global activity, then one might observe identified oil supply shocks disproportionately happening after sustained U.S. economic expansions (because oil prices and production are endogenous). This could lead an HP filter of real GDP to rise after an oil supply shock.

14. The best match of HP-filtered series comes with high values of λ (we use $\lambda = 500,000$). This high value is consistent with a low pass filter that allows only low frequencies with periods of about 10 years and higher. Lower values do not replicate Greenbook measures of potential GDP as closely, as can be seen in online appendix figure 6. Similarly, with moving-average measures, we can better replicate the dynamic response of Greenbook estimates of potential when averaging over long periods (10–20 years) than over shorter horizons (3–5 years), as illustrated in online appendix figure 5.


illustrated in figure 8. In the top panel, we plot the time path of potential GDP that would have been estimated in real time using the HP filter during the Great Recession. Specifically, for each quarter, we apply an HP filter to the available data and extract the trend level for that period. We then plot the sequence of these estimates over time, thereby showing the evolution of the implied real-time trend level of GDP during this historical episode for different values of the smoothing parameter. Regardless of the smoothing parameter, estimates of real-time trend output from an HP filter exhibit a significant downward revision (the magnitude of the revision declines in λ), much like the real-time estimates of official organizations in the United States, providing another illustration of how closely one can reproduce historical real-time estimates of potential output using a simple statistical filter. The danger of doing so is illustrated in the bottom panel of figure 8, which replicates this exercise for the Great Depression using data from Ramey and Sarah Zubairy (2018). The use of an HP filter to estimate potential GDP in real time over the course of the Great Depression would have implied that the output gap closed sometime between 1934 and 1936, depending on the smoothing parameter. But as illustrated in figure 8, GDP surged thereafter and real-time estimates of potential GDP began to climb back up. Unless one is prepared to entertain the idea that the Great Depression reflected negative supply shocks that were offset by positive supply shocks in the middle to late 1930s, we interpret this experience as illustrating the potential pitfalls of relying on simple statistical filters to make inferences about potential output during long-lived downturns.¹⁵

III. Cross-Country Evidence on the Incorporation of Shocks into Estimates of Potential

The Great Recession was of course not limited to the United States, and the persistence of output declines in most major advanced economies has also been associated with declines in their potential output, as documented by Ball (2014). Indeed, despite widespread lackluster growth by historical

15. Papell and Prodan (2012) analyze large recessions in the United States and other countries using long samples. Consistent with our analysis of the Great Depression, they find that actual output eventually catches up with prerecession projections of potential output. Gordon and Krenn (2010) document that using a bandpass filter to estimate potential GDP during the Great Depression would similarly imply implausible declines in potential between 1929 and the mid-1930s.

Figure 8. Real-Time Estimates of U.S. Potential GDP from an HP Filter during the Great Recession and Great Depression^a



Sources: The data in the top panel are authors' calculations using underlying GDP data from FRED, the database of the Federal Reserve Bank of Saint Louis. The data in the bottom panel are from Ramey and Zubairy (2018).

a. This figure reports estimates of trend (potential output) generated by the one-sided Hodrick-Prescott (HP) filter for various values of the smoothing parameter λ . The filter is recursively applied to the final vintage of the data. For example, an estimate for 2008:Q1 uses data only up to 2008:Q1, an estimate for 2008:Q2 uses data only up to 2008:Q2, and so on.

standards since the Great Recession, the World Bank recently estimated that advanced economies have on average an output gap of zero, indicating that the large downward revisions to potential output estimated by the CBO for the U.S. since 2007 also extend to other advanced economies (World Bank 2018). To what extent can the cyclical patterns documented above in estimates of potential GDP be generalized to other countries? In this section, we turn to cross-country estimates of potential GDP, from both international organizations and professional forecasters. Using international data gives us many more observations and thus more statistical precision and power.

III.A. IMF and OECD Estimates of Potential GDP

We consider first estimates of potential GDP from two international organizations, the IMF and the OECD. Both provide estimates of the level of potential GDP for a wide range of countries.¹⁶

We follow the same strategy as with the U.S. and compare impulse responses of actual GDP and estimates of potential GDP from each of these two organizations to different economic shocks. However, because time samples are much shorter for most countries, we pool data across all countries in our sample. In short, for each identified shock ϵ , we estimate the following specifications:

$$(5) \quad \Delta \log Y_{j,t} = \alpha_j + \gamma_t + \sum_{k=0}^K \phi_k \epsilon_{j,t-k} + error_{t,j}, \text{ and}$$

$$(6) \quad \Delta \log Y_{j,t}^* = \delta_j + \kappa_t + \sum_{k=0}^K \psi_k \epsilon_{j,t-k} + error_{t,j},$$

where j indicates the country and α_j , δ_j and γ_t , κ_t respectively denote country and time fixed effects. The time frequency is semiannual, as determined by the frequency of real-time estimates of potential GDP by both the IMF and OECD.

Because of more limited data availability across countries, we cannot identify as many shocks and in the same way as was done for the United States. For productivity, we use innovations in labor productivity, after conditioning on past changes in labor productivity as well as country and

16. We exclude Norway from our analysis because this country relies heavily on energy exports.

time fixed effects.¹⁷ For oil shocks, we continue to use Kilian's measure of oil supply shocks but interact it with a country-specific measure of oil sufficiency, from the International Energy Agency's World Energy Statistics and Balances (IEA 2017) to distinguish it from the time fixed effects.¹⁸ For monetary policy shocks, we run a VAR for each country on GDP growth, unemployment, inflation, and the interest rate and apply a Choleski decomposition on this ordering to recover country-specific interest rate shocks. The VAR has four lags using quarterly data from 1980:Q1 until 2016:Q4 or as available.¹⁹ Finally, fiscal shocks are differences between ex-post government spending and ex-ante forecasts of government spending from the OECD, following Auerbach and Gorodnichenko (2012a).

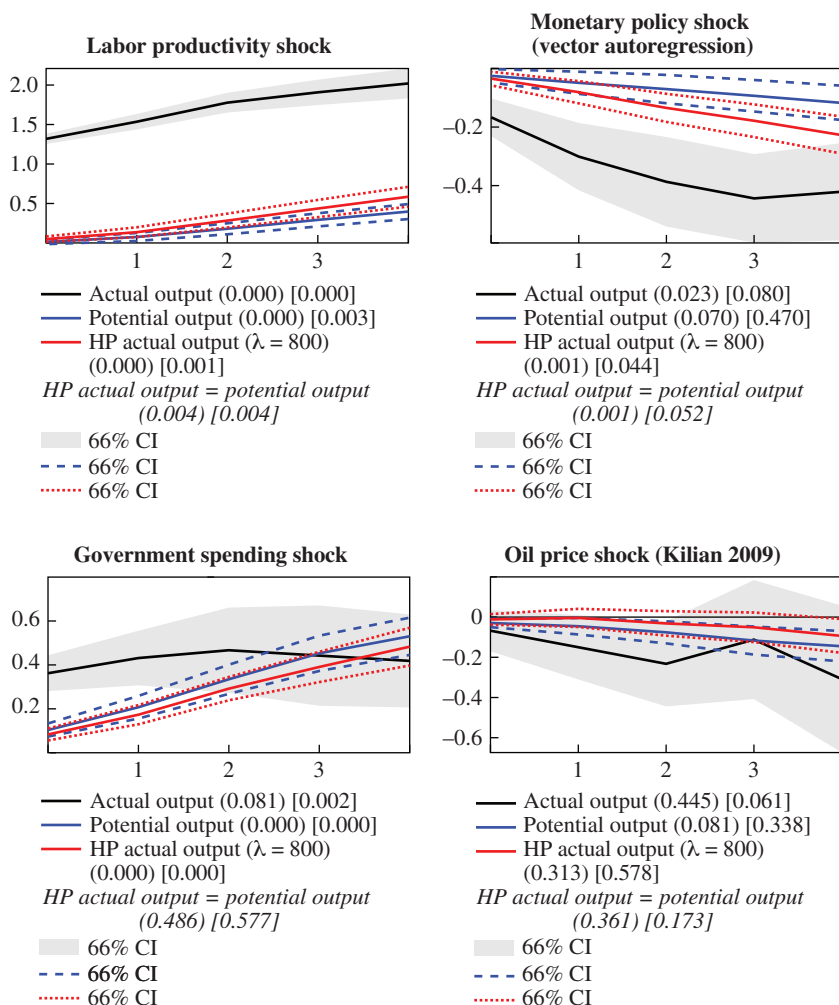
Turning first to the OECD sample of countries and estimates of potential GDP, figure 9 presents responses of both GDP and potential to each of the four shocks (the p values for the same tests discussed in section II are included in the figure and summarized in online appendix table 3). All four shocks yield the expected changes in GDP. Productivity shocks have an immediate and permanent effect on output, while oil supply shocks have a negative albeit delayed persistent effect on output. Both demand shocks have transitory effects on GDP which start dissipating in about one or one and a half years and are mostly gone after three years (we only show IRFs up to four semesters in the figure).

The effects of these shocks on potential GDP are consistent with those obtained for the United States. In response to productivity shocks, estimates of potential GDP evolve gradually in the direction of actual changes in output. After oil supply shocks, estimates of potential GDP decrease

17. Specifically, we use a measure of labor productivity at the semiannual frequency taken from the OECD and then regress it on lags of itself in a panel regression with country and time fixed effects, allowing coefficients on the lags of labor productivity to vary over countries, as well as a dummy for Ireland in 2015 due to its very big outliers in productivity changes. It is important to notice that this OECD measure of labor productivity is highly correlated with other measures of productivity, such as multifactor productivity from the OECD or productivity from EU-KLEMS data.

18. Oil sufficiency measures what percentage of total oil usage can be satisfied from each country's supply. Hence it ranges from 0 (if the country has no oil supply at all—for example, Belgium), passing through 1 (if the country can exactly satisfy its oil demand—for example, Australia), up to high numbers like 20 (if the country is a net exporter of oil).

19. A group of countries is in the euro zone after 1999. For these countries, we construct monetary policy shocks as follows. For the pre-euro period, we run a country-specific VAR and obtain monetary policy as described in the text. For the euro period, we run a VAR with variables measured at the level of the euro zone. From this VAR, we obtain monetary policy shocks, which we append to the shocks identified in the pre-euro period.

Figure 9. Response of the U.S. Growth Rate for Actual Output and the OECD's Measure of Potential Output (Nowcast)^a

Sources: Authors' calculations, with potential output from the OECD and identified shocks described in the text.

a. The figure shows impulse response functions (IRFs) for growth rates of actual and potential output (nowcast). IRFs are estimated using equations 5 and 6. The horizontal axis measures time in semesters (six months). The vertical axis measures growth rate of output per year. In parentheses, we report the p values for a test of whether the response of actual (potential) output is different from zero at the maximum horizon (eight quarters). In square brackets, we show the p values for a test of whether the path of the response of actual (potential) output is different from zero across all horizons of the IRF. The last row of the legend—for which there is no line in the graphs—reports p values for a test of equality of responses of actual and potential output at the maximum horizon (parentheses) and for a test of equality of the paths of the responses for actual and potential output are equal across horizons. HP = Hodrick-Prescott filter. CI = confidence interval.

slightly, but this response is very weak. After both demand shocks, estimates of potential GDP gradually and persistently evolve in the same direction as the short-run changes in GDP even though these changes in GDP are transitory. Thus, we observe both the undercyclicality after productivity shocks and the overcyclicality after demand shocks documented in the United States.

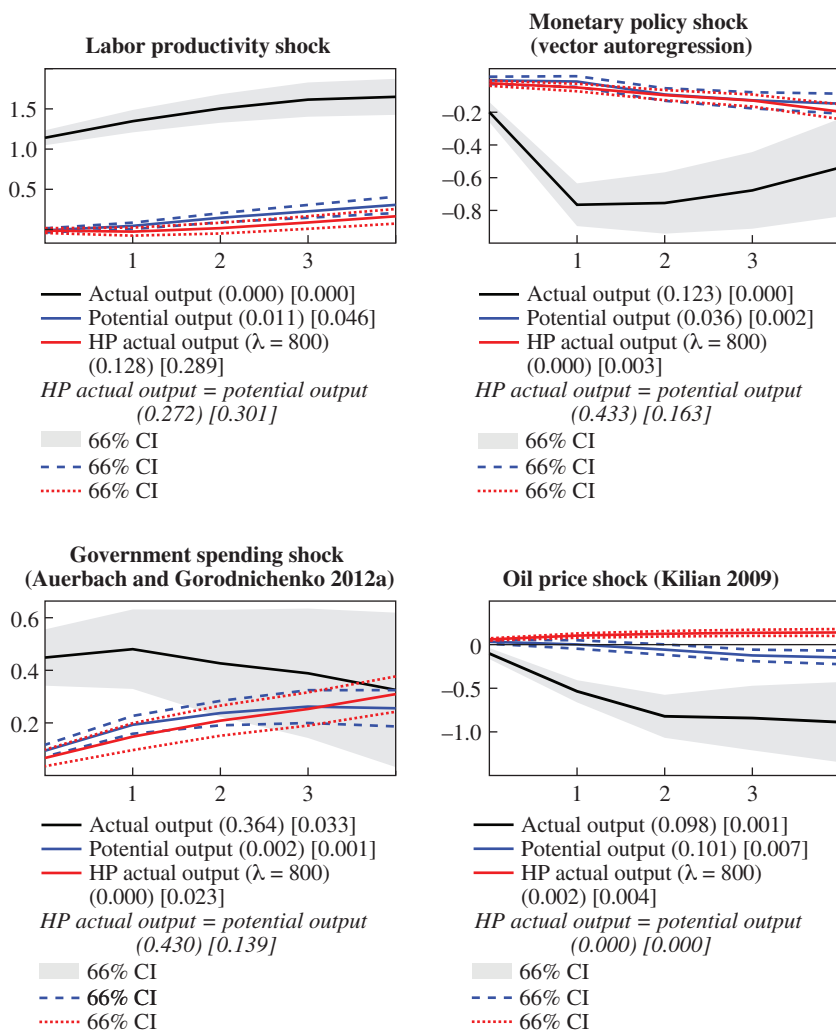
Furthermore, we include in the figure the impulse response of HP-filtered real GDP (constructed for each country using real-time data and a one-sided filter) to each shock. As was the case with the United States, we find that HP-filtered GDP responds almost identically to each shock as the OECD's estimates of potential GDP. As was the case with the Greenbook estimates of potential GDP, OECD estimates do not appear to capture much more information than what is embodied in a simple univariate filter of real-time actual GDP growth rates, which can account for why their estimates of potential GDP growth rates therefore respond to shocks that have only cyclical effects on GDP.

In figure 10, we produce equivalent results for the IMF sample of countries and IMF estimates of potential GDP. Despite the different countries in the sample, the estimated effects of the shocks on actual GDP are very similar to those found in the OECD sample. The responses of the IMF's estimated levels of potential GDP respond similarly to those from the OECD: They rise inertially after productivity shocks, and also respond inertially after monetary and fiscal shocks, in the same direction as the short-run response of GDP. Their response after oil supply shocks is equally weak. For comparison, we also again include responses of real-time, HP-filtered output and find, as with the OECD, that these very closely track the IMF estimates of potential output after shocks, with the only exception again being oil supply shocks.

Overall, the evidence from these two international organizations closely aligns with previous evidence from the United States: Their estimates of potential GDP are well approximated by an HP filter applied to real-time data and therefore seem to respond mechanically to short-run changes in GDP, regardless of the underlying source of economic variation. This suggests that observing revisions in one of these organizations' estimates of potential GDP in a country tells us little about how persistent the concurrent changes in GDP are likely to be.

III.B. Private Long-Horizon Forecasts of the GDP Growth Rate

In addition to forecasts from international policy organizations, we consider how private forecasters adjust their beliefs about the long-run GDP

Figure 10. Response of the U.S. Growth Rate for Actual Output and the IMF's Measure of Potential Output (Nowcast)^a

Sources: Authors' calculations, with potential output from the IMF and identified shocks described in the text.

a. The figure shows impulse response functions (IRFs) for growth rates of actual and potential output (nowcast). IRFs are estimated using equations 5 and 6. The horizontal axis measures time in semesters (six months). The vertical axis measures growth rate of output per year. In parentheses, we report the p values for a test of whether the response of actual (potential) output is different from zero at the maximum horizon (eight quarters). In square brackets, we show the p values for a test of whether the path of the response of actual (potential) output is different from zero across all horizons of the IRF. The last row of the legend—for which there is no line in the graphs—reports p values for a test of equality of responses of actual and potential output at the maximum horizon (parentheses) and for a test of equality of the paths of the responses for actual and potential output are equal across horizons. CI = confidence interval.

growth rate in response to shocks. Although forecasts of potential GDP are not readily available, Consensus Economics provides forecasts of GDP at long horizons on a semiannual basis. To the extent that cyclical fluctuations in GDP should be complete within five years or so, these long-horizon forecasts should be equivalent to forecasts of potential GDP growth at the same horizon.

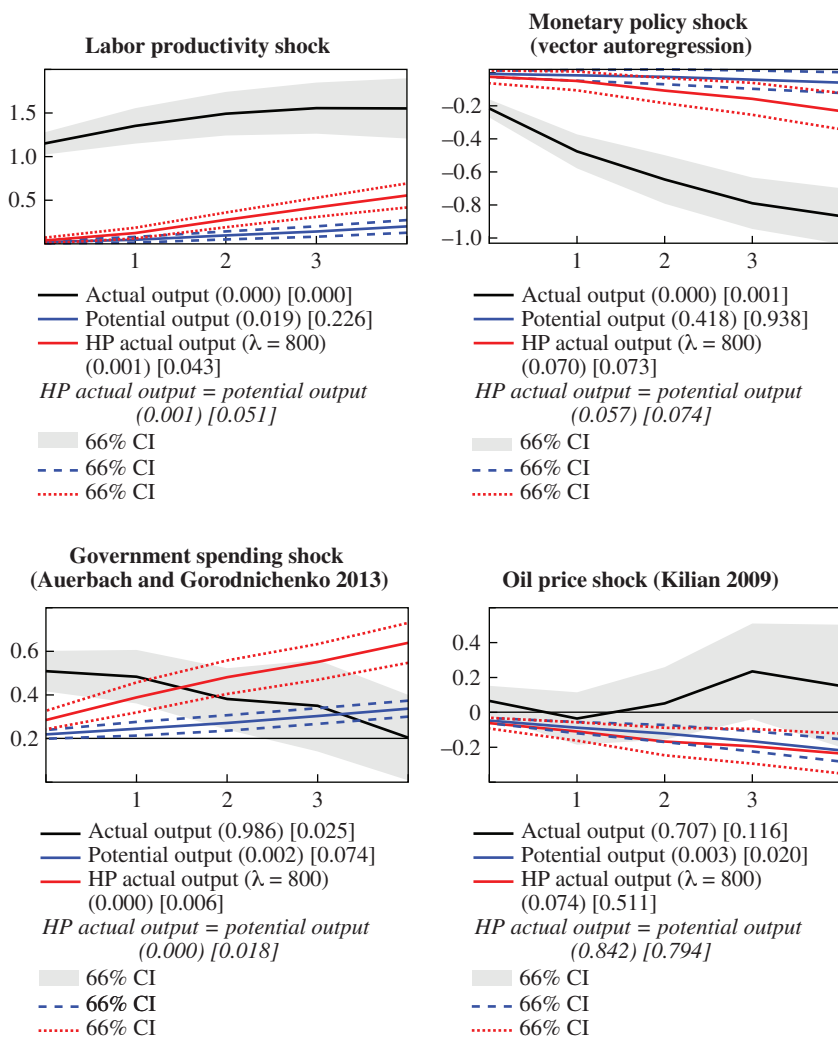
Using the same shocks as those used with the OECD and IMF samples, we replicate our previous results using private forecasts of long-run GDP for the 12 countries for which we have these forecasts (see online appendix table 1 for the countries and periods included in this sample). With the different sample of countries and time periods, the impulse responses of actual GDP are broadly similar (figure 11), although the output responses to monetary shocks are more persistent while the response to oil supply shocks is much less precise.

After productivity shocks, private forecasts gradually evolve in the same direction as actual output, therefore replicating the pattern observed with forecasts from public and international organizations. After the two demand shocks, the private sector forecasts also gradually evolve in the direction of the short-run movements in GDP, although the response after monetary shocks is not significant at standard levels. With respect to oil supply shocks, private forecasts of long-run GDP decline gradually.

For comparison, we also plot the implied response of HP-filtered levels of output to the same shocks and countries. For all shocks, HP-filtered forecasts evolve in the same direction as private forecasts but more rapidly. This is in contrast to what was found with estimates of potential from public and international organizations, when the estimates of potential GDP were almost identical in the impulse responses to those of an HP-filtered level of output. The more inertial response of private forecasters could reflect less rapid information updating or a difference in forecasting horizon (private forecasts are for long-run levels of GDP rather than current estimates of potential GDP).

IV. Alternative Approaches to Estimating Potential Output

The apparent inability of available estimates of potential output to differentiate between shocks that have permanent effects and those with only transitory effects raises the question of whether alternative approaches might be better. Obviously, this is a challenging task, and developing a single satisfactory method is beyond the scope of the paper. However, we can utilize available tools to get a glimpse of what may constitute

Figure 11. Response of the U.S. Growth Rate for Actual Output and Consensus Economics' 6- to 10-Year-Ahead Forecast for Actual Output^a

Sources: Authors' calculations, with potential output from the IMF and identified shocks described in the text.

a. The figure shows impulse response functions (IRFs) for growth rates of actual and potential output (nowcast). IRFs are estimated using equations 5 and 6. The horizontal axis measures time in semesters (six months). The vertical axis measures growth rate of output per year. In parentheses, we report the p values for a test of whether the response of actual (potential) output is different from zero at the maximum horizon (eight quarters). In square brackets, we show the p values for a test of whether the path of the response of actual (potential) output is different from zero across all horizons of the IRF. The last row of the legend—for which there is no line in the graphs—reports p values for a test of equality of responses of actual and potential output at the maximum horizon (parentheses) and for a test of equality of the paths of the responses for actual and potential output are equal across horizons. CI = confidence interval.

a basis for a satisfactory method to estimate potential output. Specifically, we first use Blanchard and Quah's (1989) approach, designed specifically to separately identify supply and demand shocks, to show that long-run restrictions may provide a practical solution to some of the issues we have identified above. We show that this approach implies significantly different estimates of potential output during the Great Recession, and that alternative approaches yield similar conclusions.

IV.A. Blanchard and Quah Approach to Estimating Potential Output

In this simple, proof-of-concept exercise, we follow Blanchard and Quah (1989; henceforth, BQ) and estimate a bivariate VAR(8), where the variables are output growth and the unemployment rate. The identifying restriction of this model is as follows: Supply-side shocks are the structural shocks that have permanent effects on the level of output, and demand-side shocks are restricted to have zero effect on the level of output in the long run. We then interpret predicted movements in output driven by supply-side shocks as capturing potential output. The restriction that only supply-side shocks have permanent effects on output is broadly consistent with the responses of output observed in figure 4 and other results in the literature, namely, that monetary and government spending shocks do not seem to have permanent effects on output (for example, Romer and Romer 2004; Ramey 2016).

Because BQ and others emphasize the importance of structural breaks, we use a rolling window of 120 quarters.²⁰ When applying the BQ approach, we use real-time data to ensure that our results are not driven by information that is not available to the econometrician. In a particular quarter (say 1995:Q1), we use the vintages of real output growth and unemployment rate that were available at that point in time (obtained from the Federal Reserve Bank of Philadelphia's real-time database for macroeconomists), estimate the structural vector autoregression with long-run restriction using these series, and then perform the historical decomposition on these data to recover the component of the growth rate of actual output due to supply-side shocks for the given quarter. That is, we keep only the data point that corresponds to the last quarter in a rolling-window sample. The next quarter's (1995:Q2) historical decomposition data point is going to

20. We would like the rolling window to be big for the long-run identifying restriction to work well, but at the same time we would like it to be small to minimize exposure to structural breaks. We compromise by using a rolling window of 120 quarters, but results are similar when we use alternative rolling windows, such as 80, 100, 140, or 160 quarters.

use vintages that were not available yet in 1995:Q1, and the previous quarter's (1994:Q4) historical decomposition data point used vintages that contained less information and stopped in 1994:Q4. This approach therefore uses no more information than what was available to agents in real time, making our estimates comparable to real-time estimates of U.S. potential GDP growth.

After we recover the time series of the growth rate of output due to supply shocks (that is, our estimate of potential output), we estimate regression equations 2 and 3 on actual output and our estimate of potential output. Figure 12 shows the resulting impulse responses. We find that, in contrast to the conventional estimates of potential output, our estimate strongly reacts to supply shocks and exhibits no significant sensitivity to demand shocks. Interestingly, the reaction of our estimate for potential output to a TFP shock is stronger at short horizons than the reaction of actual output. This pattern is consistent with theoretical responses in New Keynesian models, where frictions prevent actual output from an immediate adjustment to a productivity shock so that a productivity shock creates a negative output gap in the short run. Despite its simplicity, the BQ approach can therefore make progress toward resolving puzzles in the reaction of conventional estimates of potential output to identified shocks.

It is notable that real-time estimates of potential output coming from BQ do not suffer from the same issues as those found from official estimates of potential output. One interpretation of how the latter respond to shocks is that they represent the optimal outcome in the presence of noisy information; if agents cannot differentiate between supply and demand shocks in real time, then their estimates of potential should slowly respond to each kind of shock. But the fact that the BQ methodology can, *in real time*, successfully distinguish between the two kinds of shocks suggests that this is not a binding constraint on real-time analysis but rather reflects the specific methodologies used by each organization to create measures of potential output.²¹

We can also use the BQ decomposition to revisit how potential output may have changed over the course of the Great Recession. In generating

21. Another piece of evidence consistent with this interpretation is that even final (2017) estimates of potential output respond to historical supply and demand shocks in the same qualitative manner as in figure 6 (see online appendix figure 8). Despite a long delay, revised estimates of potential GDP from official agencies do not successfully distinguish between transitory and permanent shocks, suggesting that this reflects a feature of how these estimates are constructed, not an inability to distinguish between these shocks in real time.

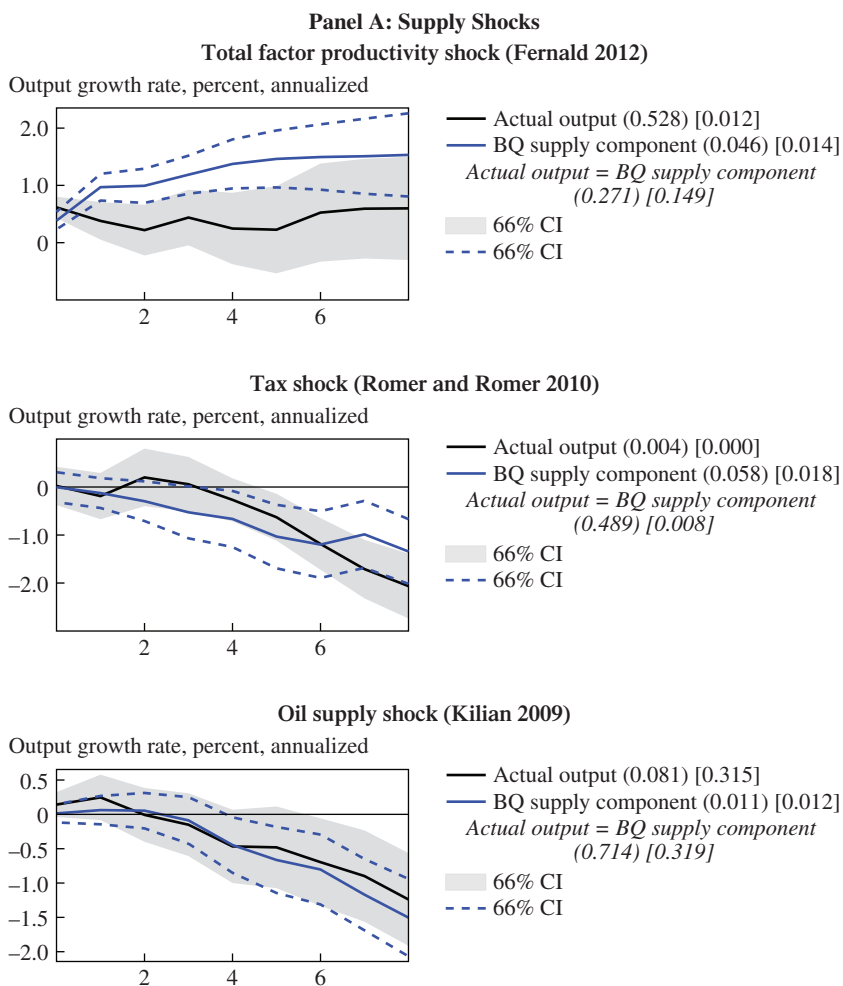
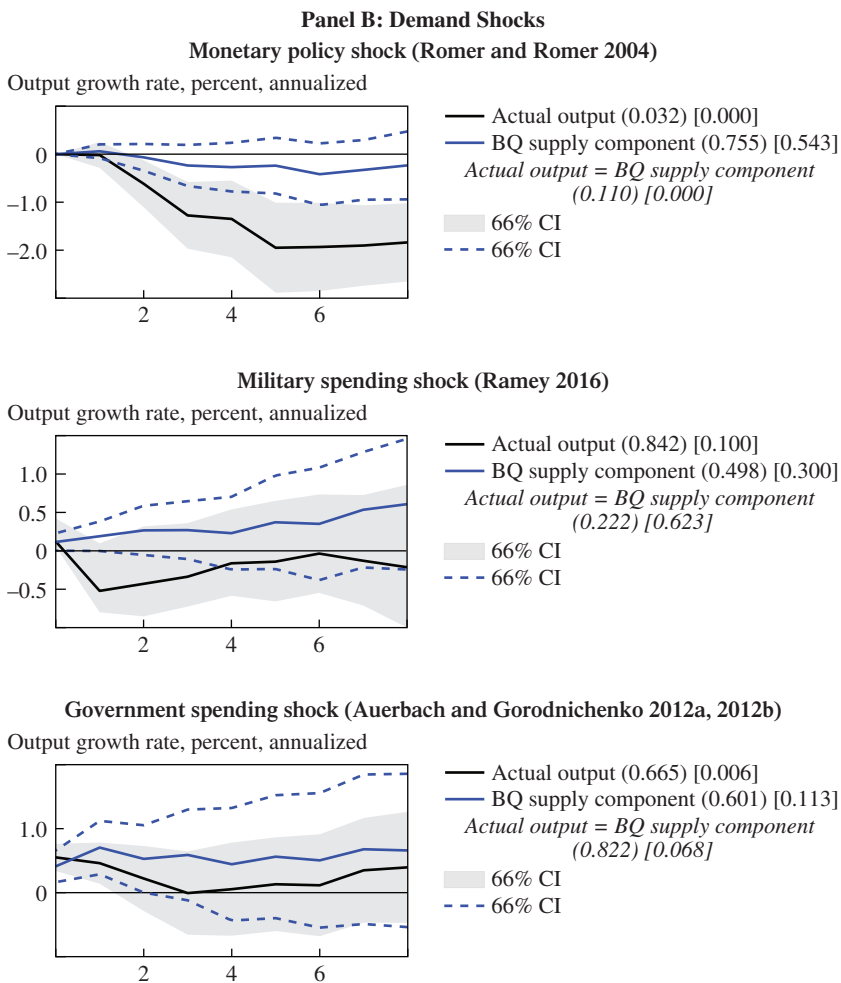
Figure 12. Response of the U.S. Growth Rate for Actual Output and the SVAR Identified Historical Supply Component of Actual Output^a

Figure 12. Response of the U.S. Growth Rate for Actual Output and the SVAR Identified Historical Supply Component of Actual Output^a (*Continued*)

Sources: Authors' calculations, with potential output from Federal Reserve Greenbooks and identified shocks from Fernald (2012); Romer and Romer (2004, 2010); Kilian (2009); Ramey (2016); Auerbach and Gorodnichenko (2012a, 2012b).

a. SVAR = structural vector autoregression. This figure reports impulse response functions (IRFs) estimated using equations 2 and 3. The "BQ supply component" is the historical contribution of supply-side shocks—as identified by Blanchard and Quah (1989)—to the output growth rate. The estimation sample covers the longest possible period with nonmissing observations for shocks and potential output (output gap) using output gap data starting in 1970. In parentheses, we report the p values for a test of whether the response of actual (potential) output is different from zero at the maximum horizon (eight quarters). In square brackets, we show the p values for a test of whether the path of the response of actual (potential) output is different from zero across all horizons of the IRF. The last row of the legend—for which there is no line in the graphs—reports p values for a test of equality of responses of actual and potential output at the maximum horizon (parentheses) and for a test of equality of the paths of the responses for actual and potential output are equal across horizons. CI = confidence interval.


real-time estimates and forecasts of potential output using the BQ methodology, it is important to note that one must take a stand on the long-run growth rate of the economy. Heuristically, we can decompose the growth rate of output as $\Delta \log Y_t = g + \Delta \log Y_t^p + \Delta \log Y_t^c$, where g is the long-run rate of output, $\Delta \log Y_t^p$ is the growth rate of output due to “supply” shocks with permanent effects on the level of output, and $\Delta \log Y_t^c$ is the growth rate of output due to transitory “demand” shocks. We define the growth rate of potential output as $\Delta \log Y_t^* \equiv g + \Delta \log Y_t^p$. By iterating VAR coefficients from BQ forward, we construct forecasts for $\Delta \log Y_{t+h|t}^* = g + \Delta \log Y_{t+h|t}^p$, given the history of supply shocks up to period t . Then we cumulate $\Delta \log Y_{t+h|t}^*$ over $0, \dots, H$ to compute the response of the level of potential output to a shock. Note that in this calculation, we follow BQ and assume that shocks do not influence g , the growth rate of output in the long run. Although this assumption is consistent with the fact that the growth rate of output per capita in the United States has been remarkably stable, at 2 percent a year over the last 150 years (Jones 2016), it is nonetheless an important assumption. In the context of using BQ for the Great Recession, we apply the long-run growth rate of GDP from the 1977–2007 period (3.1 percent) and assume that it remains invariant to the Great Recession.

The resulting real-time revisions in potential output using the BQ methodology during the Great Recession are plotted in the top panel of figure 13. Like official estimates, we find that there are declines in potential output during the Great Recession that take some time to uncover; the first significant downward revisions for 2009 potential output occur using the 2013 estimates. But there is little predictability in subsequent revisions; they all closely track the 2013 estimates of the path of output. And unlike the official estimates, the BQ approach points to a large and continuing gap between actual output and potential. By 2016, we estimate U.S. potential output to have grown by about 5 log percentage points more than actual output since 2007, a difference that could potentially be closed through the use of demand-side policies.

Furthermore, it is likely that BQ estimates represent an overestimate of the decline in potential output. This is because, since the onset of the zero bound on interest rates, even transitory demand shocks should be expected to have more persistent effects than they normally would, given the absence of offsetting monetary policy actions. Because the BQ approach is estimated over a long period, more persistent demand shocks during the zero lower bound are likely to be in part attributed to “supply shocks” in the BQ decomposition. Some of the estimated decline in potential output since the


Figure 13. Revisions in Potential GDP during the Great Recession from the Blanchard-Quah Methodology^a

Assuming Long-Run Growth Rate Equal to Average 1977–2007 Value (3.1 Percent)
Log deviation from 2007:Q1



Using Alternative Long-Run Growth Rates, BQ 2017 Vintage

Log deviation from 2007:Q1



Sources: Authors' calculations, following the structural vector autoregression methodology of Blanchard and Quah (1989) and various measures of long-run growth between 2007 and 2017 from Consensus Economics, Macroeconomic Advisers, Survey of Professional Forecasters, the Congressional Budget Office, and Blue Chip Economic Forecasts.

a. The top panel plots the real-time estimates and forecasts of potential GDP, following Blanchard and Quah (1989), for different rolling windows. YYYY in "BQ YYYY" shows the last year of the rolling window. See section IV.A for details. The bottom panel plots BQ 2017 for different values of g , which are taken from the sources indicated in the legend: Macro Adv = Macroeconomic Advisers; Blue Chip = Blue Chip Economic Forecasts; SPF = Survey of Professional Forecasters; CBO = Congressional Budget Office. See section IV.B for details.

Great Recession attributed to supply-side factors is therefore likely to be transitory, making the output gap even larger than our estimates suggest.

Because of the possible sensitivity of BQ estimates of potential GDP to assumptions about the long-run growth rate, we consider a number of other values for the long-run growth rate of output that were suggested before the Great Recession. We view it as important to restrict our attention to pre–Great Recession estimates because these already include predictable deterministic changes in growth after 2007 (such as from the retirement of the Baby Boomers) but are not contaminated by the persistent changes in output since the Great Recession. Indeed, as we documented using long-run projections of professional/official forecasters in subsection III.B, real-time estimates of long-run growth respond to shocks that have only transitory effects, so we should expect these estimates to have been significantly reduced since the Great Recession (as most in fact have been), but this is not informative about whether these changes should be expected to persist.²²

Given the difficulty inherent in making forecasts about future productivity growth, the main driver of long-run GDP growth, there was significant uncertainty about the long-run future growth rates of the United States' GDP before the Great Recession. For example, Macroeconomic Advisers, a prominent economic forecasting firm, was predicting a relatively high long-run growth rate of 3.3 percent. Many other professional forecasters were similarly optimistic, with forecasters in both the Blue Chip Economic Forecasts and the Survey of Professional Forecasters predicting long-run growth rates of 3.0 percent, just under the postwar average of 3.1 percent. Other forecasters were somewhat more pessimistic. For example, forecasters at Consensus Economics were predicting an average long-run growth rate of 2.8 percent (there was much disagreement across forecasters; the standard deviation is 0.6 percent). The CBO was even more pessimistic, predicting an average growth rate of just 2.6 percent in the long run. We show the implications of each of these assumptions for BQ decompositions since the Great Recession in the bottom panel of figure 13. Depending on the source of long-term projections, the output gap has fallen anywhere between 15 percent (Macroeconomic Advisers) to 2 percent (CBO) since the Great Recession.

22. We find similar results when we adjust output by the size of the civilian population (online appendix figure 9).

IV.B. Alternative Estimates of Potential Output since the Great Recession

Although these different estimates from the BQ methodology all imply significant remaining slackness, they also point to the difficulty of pinning down the output gap using a single procedure. In this subsection, we consider several alternative theory-based approaches to investigate the robustness of this finding.

One approach closely related to BQ is from Galí (1999). He proposes identifying technology shocks in a VAR through long-run restrictions by assuming that these shocks change labor productivity in the long run while other shocks do not. We apply the same two-variable VAR as used by Galí (1999) to real-time data and define the real-time level of potential output as that coming only from the identified technology shocks. As illustrated in figure 14, this approach points to even smaller changes in potential output over the course of the Great Recession, perhaps due to the narrower interpretation of the types of shocks that affect potential output than in BQ. The 2017 level of potential output is only 5 log percentage points lower when estimated using 2017 data than forecasted from 2006 data, yielding a growth in the output gap by 2017 of well over 10 log percentage points relative to 2007.²³


Cochrane (1994) proposes an alternative approach to identifying permanent changes in GDP by exploiting the consumption/output ratio. Under Milton Friedman's (1957) Permanent Income Hypothesis, consumption changes reflect permanent changes in income, so adding information about consumption can help decompose transitory from permanent changes in income. Applying his methodology to real-time data on consumption and GDP, and identifying potential GDP as those changes associated with changes in consumption, yields a surprisingly similar path of revisions in potential output over the Great Recession as the BQ approach, as illustrated in figure 14. As with Galí's (1999) approach, the implied output gap in 2017 is therefore more than 10 log percentage points bigger than in 2007 when applying the same long-run growth rate as in BQ estimates (3.1 percent).²⁴

23. One could also follow King and others (1991), Gonzalo and Ng (2001), and others to consider VARs that include more than two variables or use other permanent/transitory decompositions.

24. We report results for different vintages of the Galí and Cochrane approaches in online appendix figure 10.

Figure 14. Alternative Approaches to Estimating Potential GDP in Real Time during the Great Recession^a

Log deviation from 2007:Q1



Sources: Authors' calculations of potential output, following various methodologies (Blanchard and Quah 1989; Galí 1999; Cochrane 1994), or using information in the Phillips curve or from the Congressional Budget Office.

a. The figure plots the 2017 estimates of the path of potential GDP from these approaches, as well as the Blanchard and Quah (1989, "Blanchard") approach, the Phillips curve, the CBO's estimates for 2017, and the 2007 precrisis estimate. "Actual" denotes the path of real GDP. See section IV.B for the details.

Importantly, the Cochrane approach is immune to concerns about hysteresis, because it does not try to distinguish between supply and demand shocks based on their long-run effects. If hysteresis is present, then even transitory shocks should have effects on consumption due to their long-lived effects on income. As a result, they would be incorporated into the resulting estimates of potential output. Furthermore, this approach is also likely to overstate the decline in potential output over this time period. If some households are credit-constrained ("hand-to-mouth") and adjust their consumption to transitory income changes, then we will measure declines in potential GDP even from some transitory shocks, thereby overstating the

change in potential GDP since the Great Recession and understating the current amount of economic slackness.

Closer in spirit to Okun's (1962) approach is to infer information about potential output from the inflation rate. In New Keynesian models, nominal rigidities generate an expectations-augmented Phillips curve that relates inflation to expected inflation and the output gap (or the deviation of unemployment from the natural rate of unemployment). Conditional on observing inflation, expected inflation, and real GDP, one can then use the Phillips curve to infer the potential level of GDP (under the assumption of no markup shocks). Following Coibion and Gorodnichenko (2015b), we estimate an expectations-augmented Phillips curve during the pre-Great Recession period using inflation expectations from the Michigan Survey of Consumers. As shown by Coibion and Gorodnichenko (2015b), conditioning on household forecasts of inflation yields a stable Phillips curve since the 1960s and eliminates the puzzle of the "missing disinflation" during the early years of the Great Recession. We then apply this Phillips curve to the period since the Great Recession to infer what path of potential output is required to account for inflation dynamics during this period.

A key advantage of this approach is that it does not rely on long-run restrictions, which may be sensitive to structural breaks (Fernald 2007). We plot a smoothed version of 2017 estimates of potential GDP over the period of the Great Recession in figure 14, along with the 2017 estimates from other approaches for comparison.²⁵ The implied potential GDP from the Phillips curve does not decline much until 2011, significantly later than other approaches. However, by 2017, the resulting estimate of potential GDP is close to that of the BQ approach, pointing to an output gap of about 5 log percentage points. In a related paper (Coibion, Gorodnichenko, and Ulate 2019), we do more extensive work using the expectations-augmented Phillips curve to back out potential output. We show that this approach works systematically across countries and that the measures of potential output that it delivers paint a similar picture to the ones obtained in this paper using the BQ approach.

In short, bringing additional information to bear on the identification of potential output—whether from labor productivity, consumption, or inflation—combined with theoretical predictions regarding how these variables relate to potential GDP, largely confirms the findings of the BQ

25. We plot a smoothed version because sampling uncertainty in inflation expectations measured by the Michigan Survey of Consumers (500 households participate in the survey in a typical month) generates high-frequency noise in estimates of potential GDP.

approach. Each approach points to nontrivial revisions in potential output since the Great Recession, but not nearly as large as those coming from the official organizations. This implies that current U.S. output likely remains significantly below potential output, and therefore that further stabilization policies could be warranted.

IV.C. Can the Output Gap Be Large When Unemployment Is Low?

Our view that a significant output gap likely remains in the United States a decade since the start of the Great Recession may seem at odds with the conclusion one might reach from looking at recent U.S. unemployment rates. For example, an output gap of 5 percent would, using Okun's law, require a negative unemployment rate gap of about 1.5 percent.²⁶ With the U.S. unemployment rate having fallen below 4 percent in April 2018, this would imply a natural rate of unemployment of about 2.5 percent. In contrast, typical estimates of the NAIRU point toward much higher values (the 2018 CBO estimate is 4.6 percent). Is it possible to reconcile recent labor market dynamics with our estimates of potential output? In this subsection, we argue that the answer is unambiguously yes, and that it is the alternative view—namely, that labor markets are currently very tight—that seems at odds with other economic dynamics.

First, the evidence from a number of other macroeconomic variables is consistent with the view that much economic slackness remains. Consumption dynamics, for example, suggest that permanent declines in income have been quite limited since the recession, as shown in subsection IV.B, which also documents that the behavior of inflation relative to inflation expectations is consistent with significant economic slackness remaining. Other variables point toward a very similar conclusion. For example, capacity utilization is a commonly used measure of the state of the business cycle. By the end of 2017, utilization was at 77 percent, well below its average value of 81 percent over the 1977–2007 period, with only 14 percent of quarters during that period having utilization rates of less than 77 percent. By historical standards, such low utilization rates are hard to reconcile with output being at or above its normal productive capacity. Wages also paint a picture of a labor market that remains slack: Annual nominal and real wage growth in the last quarter of 2017 were at the 21st and 6th percentiles, respectively, of the distribution of their historical

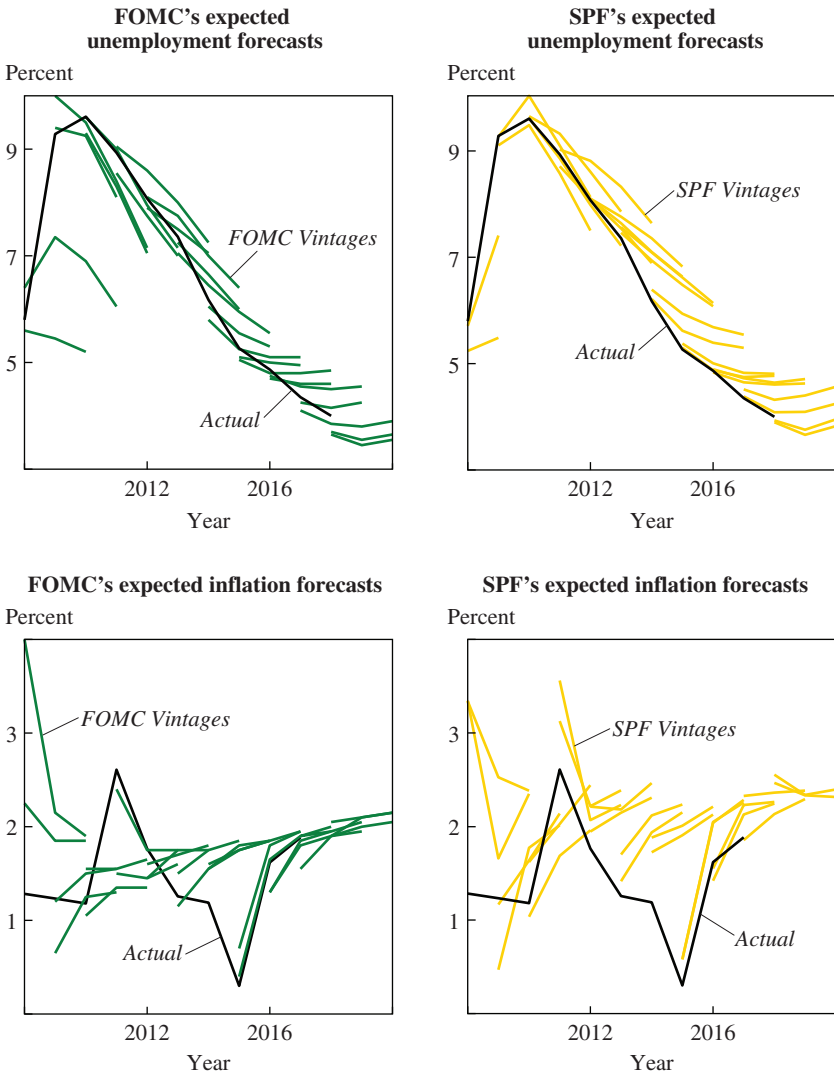
26. For all Okun's law calculations, we use a coefficient of 3, such that each change in the unemployment gap of 1 percentage point is associated with a change in the output gap of 3 percentage points (for a range of estimates of Okun's law, see Knotek 2007).

values from 1977 to 2007. Also by historical standards, it is difficult to reconcile tight labor markets with such low growth rates in wages.

Second, any statement about the natural rate of unemployment must be tentative at best, given the conceptual and measurement issues involved. Indeed, many of the same challenges as those associated with estimating the potential level of GDP are also present in estimating the natural rate of unemployment, so there is little reason to expect one to be more accurately measured than the other. Consistent with this, we observe similar patterns of systematic revisions in estimates of the natural rate of unemployment as we do in estimates of potential GDP. For example, these revisions tend to be in the direction of actual changes in unemployment, much as we observed with potential GDP. The top right panel of figure 15 plots projected unemployment rates of professional forecasters at different moments during the recovery, and their estimates of the natural rate of unemployment over time are given in figure 16. When unemployment first began to decline after its peak during the Great Recession, professional forecasters expected a gradual decline in unemployment toward a natural rate that was estimated to be nearly 6 percent. But as unemployment rates fell over time, professionals also continuously revised their estimates of the natural rate downward, with their current estimates being just above 4 percent. Importantly, professional forecasters have been consistently too pessimistic in their unemployment projections since 2011. The CBO's estimates of the natural rate of unemployment have followed an identical pattern, albeit with smaller changes (figure 16). The top left panel of figure 15 shows that FOMC members have similarly adjusted downward the levels toward which they project unemployment rates will converge, though they do not publicly provide explicit forecasts of the natural rate of unemployment.


Third, predictions about nominal variables based on perceptions of a tightening labor market have been significantly off target in recent years. As described in subsection IV.B, an expectations-augmented Phillips curve requires a significant output gap to account for inflation dynamics since the Great Recession. But even without imposing an expectations-augmented Phillips curve, forecasts based on tight labor markets have failed to adequately predict inflation. For example, the bottom right panel of figure 15 plots inflation forecasts from the Survey of Professional Forecasters over the course of the Great Recession; these have repeatedly overpredicted inflation since 2013, consistent with professionals overestimating the tightness in labor markets. A similar pattern is visible using inflation forecasts from the FOMC members over the same period (the bottom left panel of figure 15). The degree of overestimation of inflation is more limited in

Figure 15. Unemployment and Inflation Forecasts since the Great Recession^a



Sources: FRED, the database of the Federal Reserve Bank of Saint Louis; Federal Open Market Committee; Survey of Professional Forecasters.

a. This figure plots actual unemployment rates (top panels) and inflation rates (bottom panels), as well as projected rates reported by the Survey of Professional Forecasters (SPF) and from surveys of members of the Federal Open Market Committee (FOMC).


Figure 16. Estimates of the Natural Rate of Unemployment^a

Sources: Congressional Budget Office; Survey of Professional Forecasters.

a. This figure shows time series of the nonaccelerating inflation rate of unemployment (NAIRU), estimated by the Congressional Budget Office (CBO), and the consensus real-time estimate of the equilibrium rate of unemployment from the Survey of Professional Forecasters (SPF).

FOMC forecasts, but this likely reflects the institutional nature of these forecasts: Policymakers need to present forecasts of inflation that converge to the 2 percent target or risk casting doubt on their credibility (Tarullo 2017).

The issues with measuring tightness in labor markets extend beyond the difficulties associated with estimating the natural rate of unemployment and extend to the challenge of using the unemployment rate as a measure of slackness. In an environment where labor force participation exhibits clear business cycle variation, the unemployment rate may not be a sufficient metric of business cycle conditions. And this issue is not new; over the course of the late 1990s, for example, Federal Reserve chairman Alan Greenspan allowed unemployment to fall significantly below the then-estimated natural levels of unemployment (the “Greenspan gamble”). Instead of generating a rise in inflation, the result was an increase in labor force participation (from 66.5 percent in January 1996 to 67.3 percent

Figure 17. Labor Force Participation in Canada and the United States^a

Sources: Statistics Canada; Congressional Budget Office; U.S. Bureau of Labor Statistics.

a. The figure plots time series of actual and projected labor force participation rates. The Canadian series is from Statistics Canada. The U.S. actual series is from the Bureau of Labor Statistics. The 10-year-ahead projection—as of 2007—for the U.S. is from the Congressional Budget Office.

in April 2000), which led the CBO to later revise downward its estimate of the 1999 natural rate of unemployment from 5.6 to 4.8 percent. This endogeneity of the labor force participation rate appears to have become increasingly pronounced since the Great Recession. It is well known that labor force participation in the U.S. has declined significantly since the start of the Great Recession relative to 2007 projections (figure 17). How much of this decline is likely to reflect an endogenous decision by some to abandon the labor force because of limited job prospects? One way to gauge this is to compare labor force participation in the United States with that in Canada, which has a similar demographic structure and trends and thus is a frequent benchmark for comparison (see, for example, Card and Freeman 1993). But Canada is also a country that did not experience a serious financial crisis or a recession anywhere near the size of what was experienced in the U.S. As illustrated in figure 17, labor force

participation in Canada also declined since 2007, but by far less than in the U.S.—1.7 versus 3.2 percent. In fact, the decline in labor force participation in Canada since 2007—2.0 percent—corresponds almost exactly to the decline in participation that the CBO predicted would happen in the U.S. in 2007, before the start of the Great Recession. If we measured the 2017 U.S. unemployment rate relative to a labor force size consistent with a declining participation rate of 2.0 percent instead of 3.2 percent, we would have an estimated unemployment rate in 2017 of 5.3 percent (instead of 4.4 percent) and an output gap of 5 percent, which would imply, via Okun's law, a natural rate of unemployment of 3.7 percent.

Christopher Erceg and Andrew Levin (2014) provide another way to gauge the cyclical sensitivity of labor force participation during the Great Recession by exploiting the cross-state variation in employment outcomes. They find that states experiencing larger increases in unemployment during the Great Recession also experienced larger declines in participation in subsequent years, a feature we verify over a longer time span in online appendix figure 11. They find that each 1 percentage point of higher unemployment is associated with a 0.3 percent decline in the participation rate. Extrapolating this to the aggregate economy, the increase in the national unemployment rate by 5 percentage points between 2007 and 2009 should therefore be expected to generate a decline in participation of about 1.5 percentage points. Hence, endogenous participation can account for all the unexpected decline in the participation rate observed since the Great Recession.²⁷ Accounting for this change in the participation of the unemployed yields an adjusted unemployment rate of 5.8 percent for 2017 and, via Okun's law and an estimated output gap of 5 percent, a natural rate of unemployment of 4.1 percent.

This sensitivity of both the measured unemployment rate and the estimated natural rate of unemployment should give one pause when thinking about the cyclical state of the economy based on the labor market. The endogeneity of labor force participation puts typical values of both in question. Because estimates of potential output are not being normalized by

27. Erceg and Levin (2014) focus on the labor force participation rate for prime age adults. In online appendix figure 11, we present equivalent results using changes in total labor force participation from 2007 to 2017 across states. We find that an increase of 1 percentage point in the unemployment rate between 2007 and 2009 is associated with a decline of 0.15 percentage point in the labor force participation rate through 2017, or half the sensitivity found by Erceg and Levin (2014). Hence, our estimates imply that the aggregate rise in unemployment from 2007 to 2009 can account for three-fourths of the unpredictable component of the decline in participation.

an endogenous variable the way unemployment rates are, this provides another reason to focus on measuring output gaps rather than unemployment gaps. However, estimating potential output is no panacea for the measurement problems associated with labor market variables. As Okun (1962, 1) observed, “The quantification of potential output is at best an uncertain estimate and not a firm, precise measure.” Indeed, estimating potential output is hard because statistical issues are magnified by sensitivity to economic assumptions. For instance, forecasts of actual output are routinely associated with wide confidence bands (for example, standard errors for the Fed and private one-year-ahead forecasts are often greater than 1 percentage point). Because potential output is aimed to project long-run dynamics, sampling uncertainty is amplified in these projections. This uncertainty is further exacerbated by using long-run restrictions, as in the BQ approach and similar methods, in relatively short samples. Structural breaks and low-frequency variation in the data add another layer of complexity.

The sensitivity of potential output estimates to variation in economic assumptions is equally humbling. For example, BQ and similar approaches assume that g , the long-run growth rate of potential output, does not respond to economic shocks; but conceivably, g may persistently react to these shocks. Because even small differences in growth rates are compounded into large magnitudes over time, a weak sensitivity of g to shocks can translate into significant variation in potential output estimates. Concretely, if we overstate g by 0.1 percent a year, over 10 years we can overstate the output gap by 1 percentage point.²⁸ As a result, because estimating potential output is inherently so challenging, one should interpret our estimates in this section, and indeed all estimates of the potential level of output, as tentative. This uncertainty surrounding estimates of potential output and the natural rate of unemployment implies that risk management should be a primary consideration in policymakers’ decisionmaking processes.

V. Conclusion

Our results speak to two distinct but related questions. The first is how real-time estimates of potential output respond to transitory versus permanent economic shocks and therefore how we should interpret revisions in

28. The degree of uncertainty about what value to use for g is large. Gordon (2014), for example, argues that g is likely to be only 1.6 percent a year between 2014 and 2020, well under the CBO’s forecast of 2.2 percent a year, and far below the historical average of 3.1 percent (1947–2017 sample).

estimates of potential output observed in the data. The second is how high-quality, real-time estimates of potential *should* react to economic shocks.

With respect to the first question, we provide robust evidence that real-time estimates of potential output respond to *all* identified economic shocks, whether transitory or permanent. Observing a sequence of revisions in estimates of potential output, like those since the start of the Great Recession, therefore tells us little about whether declines in GDP are likely to be permanent or transitory. Instead, approaches like those of Blanchard and Quah (1989), who explicitly distinguish between temporary and long-lived shocks, are much more successful in this respect. Importantly, they suggest that current U.S. GDP is significantly below its longer-run potential and therefore that the U.S. economy remains in need of ample stimulus from monetary and fiscal authorities.

In terms of how high-quality estimates of potential *should* respond to shocks, the answer is sensitive to the concept of potential output one has in mind and the purpose that it is supposed to serve. For an agency like the International Monetary Fund that is concerned with constructing cyclically adjusted balances and long-run fiscal trends, the relevant measure of potential output is precisely one that strips out cyclical variation in GDP and identifies long-run changes. Our results suggest that the current methods used by this and similar agencies are largely unsuccessful in this respect; their revisions are contaminated by transitory shocks and respond too slowly to long-lived shocks. For example, tax cuts that have immediate and permanent effects on output are not fully reflected in official estimates of potential output for several years, suggesting that the effects of tax changes on projected revenues are likely overstated. In this sense, our results are related to the research of Blanchard and Daniel Leigh (2012), who argue that the IMF underestimates the fiscal multipliers of austerity measures.

At the same time, it is important to bear in mind the severe constraints that hamper the ability of both public and private organizations to estimate potential GDP in real time. Not only are there profound statistical and economic challenges involved, as described in subsection IV.C, but tight budgetary restrictions also make the systematic creation and updating of these estimates in real time a significant challenge for public institutions. The political implications of the estimates of potential GDP created by these agencies also present additional constraints on officials' ability to experiment with alternative procedures. The objective of our paper should therefore not be interpreted as criticizing these particular organizations but rather as highlighting the limitations of the methods that are currently

being relied upon for both fiscal and monetary policymaking, as well as proposing potential alternatives.

The approaches that we consider here, either because they explicitly distinguish between transitory and permanent shocks like Blanchard and Quah (1989) or incorporate additional information like consumption or inflation, can help address some of the limitations of currently used methods and lead to improved estimates of cyclically adjusted levels of GDP. It is likely that there remains much room for further improvement in the real-time measurement of potential output. One strategy would be to combine some of the different approaches used in this paper (as well as others), in the hope that combining different sources of information could augment the precision of the resulting estimates. A complementary approach might be to consider the dynamics of potential GDP jointly with the natural rate of unemployment and the natural rate of interest, concepts that are closely related but typically are estimated separately. Because theory implies a tight link between these different measures, considering their joint determination might also lead to more precise estimates. But until new research provides more refined and reliable estimates of potential GDP, we should likely heed Okun's (1962) warning that "meanwhile, the measure of potential must be used with care."

ACKNOWLEDGMENTS An earlier version of this paper was prepared for the second annual conference of the National Bank of Ukraine. Yuriy Gorodnichenko thanks the National Science Foundation for financial support. We are grateful to seminar participants at the Brookings Institution, the College of William & Mary, the University of Texas at Austin, the Central Bank of Chile, the Federal Reserve Bank of San Francisco, and the National Bureau of Economic Research's Summer Institute. We thank Laurence Ball, Olivier Blanchard, Aaron Flaaen, Serena Ng, Valerie Ramey, and James Stock for helpful comments, and Alexander Kopoin for sharing the OECD's data on potential output.

References

- Andres, Javier, David Lopez-Salido, and Ed Nelson. 2005. "Sticky-Price Models and the Natural Rate Hypothesis." *Journal of Monetary Economics* 52, no. 5: 1025–53.
- Anzoategui, Diego, Diego Comin, Mark Gertler, and Joseba Martinez. 2016. "Endogenous Technology Adoption and R&D as Sources of Business Cycle Persistence." NBER Working Paper 22005. Cambridge, Mass.: National Bureau of Economic Research.
- Auerbach, Alan, and Yuriy Gorodnichenko. 2012a. "Fiscal Multipliers in Recession and Expansion." In *Fiscal Policy after the Financial Crisis*, edited by Alberto Alesina and Francesco Giavazzi. University of Chicago Press.
- . 2012b. "Measuring the Output Responses to Fiscal Policy." *American Economic Journal—Economic Policy* 4, no. 2: 1–27.
- Ball, Laurence. 2014. "Long-Term Damage from the Great Recession in OECD Countries." NBER Working Paper 20185. Cambridge, Mass.: National Bureau of Economic Research.
- Barsky, Robert, and Eric Sims, 2011. "News Shocks and Business Cycles." *Journal of Monetary Economics* 58, no. 3: 273–89.
- Baumeister, Christiane, and J. D. Hamilton. 2015. "Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks." Working paper. <https://econweb.ucsd.edu/~jhamilto/bh2.pdf>.
- Beaudry, Paul, and Franck Portier. 2006. "Stock Prices, News and Economic Fluctuations." *American Economic Review* 96, no. 4: 1293–1307.
- Benigno, Gianluca, and Luca Fornaro. 2018. "Stagnation Traps." *Review of Economic Studies* 85, no. 3: 1425–70.
- Bernanke, Ben S., and Alan S. Blinder. 1992. "The Federal Funds Rate and the Channels of Monetary Transmission." *American Economic Review* 82, no. 4: 901–21.
- Blagrove, Patrick, Roberto Garcia-Saltos, Douglas Laxton, and Fan Zhang. 2015. "A Simple Multivariate Filter for Estimating Potential Output." IMF Working Paper 2015–79. <https://www.imf.org/external/pubs/ft/wp/2015/wp1579.pdf>.
- Blanchard, Olivier. 2017. "Should One Reject the Natural Rate Hypothesis?" Working paper, Peterson Institute of International Economics. <https://piie.com/publications/working-papers/should-we-reject-natural-rate-hypothesis>.
- Blanchard, Olivier, and Daniel Leigh. 2013. "Growth Forecast Errors and Fiscal Multipliers." *American Economic Review* 103, no. 3: 117–20.
- Blanchard, Olivier, and Danny Quah. 1989. "The Dynamic Effects of Demand and Supply Disturbances." *American Economic Review* 79, no. 4: 655–73.
- Blanchard, Olivier, Guido Lorenzoni, and Jean-Paul L'Huillier. 2017. "Short-Run Effects of Lower Productivity Growth: A Twist on the Secular Stagnation Hypothesis." NBER Working Paper 23160. Cambridge, Mass.: National Bureau of Economic Research.

- Braun, Steven. 1990. "Estimation of Current-Quarter Gross National Product by Pooling Preliminary Labor-Market Data." *Journal of Business and Economic Statistics* 8: 293–304.
- Card, David, and Richard B. Freeman. 1993. "Introduction." In *Small Differences That Matter: Labor Markets and Income Maintenance in Canada and the United States*, edited by David Card and Richard B. Freeman. University of Chicago Press.
- CBO (Congressional Budget Office). 2001. *CBO's Method for Estimating Potential Output: An Update*. August.
- . 2014. *Revisions to CBO's Projection of Potential Output Since 2007*. February.
- Clark, Peter. 1979. "Potential GNP in the United States 1948–80." *Review of Income and Wealth* 25: 141–65.
- . 1982. "Okun's Law and Potential GNP." Board of Governors of the Federal Reserve System, October.
- Cochrane, John. 1994. "Permanent and Transitory Components of GNP and Stock Prices." *Quarterly Journal of Economics* 109, no. 1: 241–65.
- Coibion, Olivier, and Yuriy Gorodnichenko. 2012. "What Can Survey Forecasts Tell Us About Information Rigidities?" *Journal of Political Economy* 120, no. 1: 116–59.
- . 2015a. "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." *American Economic Review* 105, no. 8: 2644–78.
- . 2015b. "Is the Phillips Curve Alive and Well After All? Inflation Expectations and the Missing Disinflation." *American Economic Journal—Macroeconomics* 7, no. 1: 197–232.
- Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia. 2017. "Innocent Bystanders? Monetary Policy and Inequality in the U.S." *Journal of Monetary Economics* 88: 70–89.
- Coibion, Olivier, Yuriy Gorodnichenko, and Mauricio Ulate. 2019. "Is Inflation Just Around the Corner? The Phillips Curve and Global Inflationary Pressures." NBER Working Paper 25511. Cambridge, Mass.: National Bureau of Economic Research.
- de Resende, Carlos. 2014. "An Assessment of IMF Medium-Term Forecasts of GDP Growth." Background Paper, Independent Evaluation Office of the International Monetary Fund.
- Edge, Rochelle, and Jeremy Rudd. 2016. "Real-Time Properties of the Federal Reserve's Output Gap." *Review of Economics and Statistics* 98, no. 4: 785–91.
- Erceg, Christopher, and Andrew Levin. 2014. "Labor Force Participation and Monetary Policy in the Wake of the Great Recession." *Journal of Money, Credit and Banking* 46, no. S2: 3–49.
- Fatas, Antonio, and Lawrence Summers. 2018. "The Permanent Effects of Fiscal Consolidations." *Journal of International Economics* 112: 238–50.
- Fernald, John G. 2007. "Trend Breaks, Long-Run Restrictions, and Contractionary Technology Improvements." *Journal of Monetary Economics* 54, no. 8: 2467–85.

- . 2012. "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity." Federal Reserve Bank of San Francisco Working Paper 2012–19.
- Fernald, John G., James H. Stock, Robert E. Hall, and Mark W. Watson. 2017. "The Disappointing Recovery of Output after 2009." *Brookings Papers on Economic Activity*, Spring: 1–58.
- Fleischman, Charles, and John Roberts. 2011. "From Many Series, One Cycle: Improved Estimates of the Business Cycle from a Multivariate Unobserved Components Model." Finance and Economics Discussion Series 2011–46. Washington: Divisions of Research and Statistics and Monetary Affairs of the Federal Reserve Board.
- Friedman, Milton. 1957. *A Theory of the Consumption Function*. Princeton University Press.
- Gali, Jordi. 1999. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review* 89, no. 1: 249–71.
- Gonzalo, Jesus, and Serena Ng. 2001. "A Systematic Framework for Analyzing the Dynamic Effects of Permanent and Transitory Shocks." *Journal of Economic Dynamics and Control* 25, no. 10: 1527–46.
- Gordon, Robert. 2014. "A New Method of Estimating Potential Real GDP Growth: Implications for the Labor Market and the Debt/GDP Ratio." NBER Working Paper 20423. Cambridge, Mass.: National Bureau of Economic Research.
- Gordon, Robert, and Robert Krenn. 2010. "The End of the Great Depression 1939–41: Policy Contributions and Fiscal Multipliers." NBER Working Paper 16380. Cambridge, Mass.: National Bureau of Economic Research.
- Gorodnichenko, Yuriy, and Matthew D. Shapiro. 2007. "Monetary Policy When Potential Output Is Uncertain: Understanding the Growth Gamble of the 1990s." *Journal of Monetary Economics* 54, no. 4: 1132–62.
- Hauptmeier, Sebastian, Friedrich Heinemann, Marcus Kappler, Margit Kraus, Andreas Schrimpf, Hans-Michael Trautwein, and Qingwei Wang. 2009. *Projecting Potential Output: Methods and Problems*. ZEW Economic Study. Heidelberg: Physica-Verlag.
- IEA (International Energy Agency). 2017. "World Indicators." IEA World Energy Statistics and Balances (database). <http://dx.doi.org/10.1787/data-00514-en>.
- Jones, Charles I. 2016. "The Facts of Economic Growth." In *Handbook of Macroeconomics, Volume 2*, edited by John B. Taylor and Harald Uhlig. Amsterdam: Elsevier.
- Keynes, John M. 1936. *A General Theory of Employment, Interest and Money*. New York: Palgrave Macmillan.
- Kilian, Lutz. 2008. "Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?" *Review of Economics and Statistics* 90, no. 2: 216–40.
- . 2009. "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market." *American Economic Review* 99, no. 3: 1053–69.

- King, Robert G., Charles I. Plosser, James H. Stock, and Mark W. Watson. 1991. "Stochastic Trends and Economic Fluctuations." *American Economic Review* 81, no. 4: 819–40.
- Knotek, Edward S., II. 2007. "How Useful Is Okun's Law?" *Federal Reserve Bank of Kansas City Economic Review* 92, no. 4: 73–103.
- Kurmann, Andre, and Eric Sims. 2017. "Revisions in Utilization-Adjusted TFP and Robust Identification of News Shocks." NBER Working Paper No. 23142. Cambridge, Mass.: National Bureau of Economic Research.
- Kuttner, Kenneth. 1994. "Estimating Potential Output as a Latent Variable." *Journal of Business and Economic Statistics* 12, no. 3: 361–68.
- Mishkin, Frederic. 2007. "Estimating Potential Output." Speech given at Conference on Price Measurement for Monetary Policy, Federal Reserve Bank of Dallas, Dallas, May 24.
- Moran, Patrick, and Albert Queralto. 2018. "Innovation, Productivity and Monetary Policy." *Journal of Monetary Economics* 93: 1–20.
- OECD (Organization for Economic Cooperation and Development). 2012. "Medium and Long-Term Scenarios for Global Growth and Imbalances." In *OECD Economic Outlook, Volume 2012/1*. Paris: OECD Publishing.
- Okun, Arthur. 1962. "Potential GNP: Its Measurement and Significance." *American Statistical Association—Proceedings of the Business and Economics Statistics Section*, 98–104.
- Orphanides, Athanasios. 2001. "Monetary Policy Rules Based on Real-Time Data." *American Economic Review* 91, no. 4: 964–85.
- . 2003. "The Quest for Prosperity Without Inflation." *Journal of Monetary Economics* 50: 633–63.
- . 2004. "Monetary Policy Rules, Macroeconomic Stability, and Inflation: A View from the Trenches." *Journal of Money, Credit and Banking* 36, no. 2: 151–75.
- Orphanides, Athanasios, and Simon van Norden. 2002. "The Unreliability of Output-Gap Estimates in Real Time." *Review of Economics and Statistics* 84, no. 4: 569–83.
- Pagan, Adrian. 1984. "Econometric Issues in the Analysis of Regressions with Generated Regressors." *International Economic Review* 25, no. 1: 221–47.
- Papell, David, and Ruxandra Prodan. 2012. "The Statistical Behavior of GDP after Financial Crises and Severe Recessions." *B.E. Journal of Macroeconomics* 12, no. 3: 1–31.
- Ramey, Valerie. 2016. "Macroeconomic Shocks and Their Propagation." In *Handbook of Macroeconomics, Volume 2*, edited by John B. Taylor and Harald Uhlig. Amsterdam: Elsevier.
- Ramey, Valerie A., and Sarah Zubairy. 2018. "Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data." *Journal of Political Economy* 126, no. 2: 850–901.
- Romer, Christina, and David Romer. 2004. "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review* 94, no. 4: 1055–84.

- . 2010. “The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks.” *American Economic Review* 100, no. 3: 763–801.
- Staiger, Douglas O., James H. Stock, and Mark W. Watson. 1997. “How Precise Are Estimates of the Natural Rate of Unemployment?” In *Reducing Inflation: Motivation and Strategy*, edited by Christina Romer and David Romer. University of Chicago Press.
- Tarullo, Daniel K. 2017. “Monetary Policy Without a Working Theory of Inflation.” Brookings Hutchins Center Working Paper 33.
- Wicksell, Knut. 1898. *Interest and Prices*. English translation, New York: Sentry Press, 1936.
- World Bank. 2018. *Global Economic Prospects: Broad Based Upturn, but for How Long?* Washington: World Bank. <https://openknowledge.worldbank.org/handle/10986/28932>.
- Yellen, Janet L. 2017. “From Adding Accommodation to Scaling It Back.” Speech to the Board of Governors of the Federal Reserve System, Washington, March 3. <https://www.federalreserve.gov/newsevents/speech/yellen20170303a.htm>.

Comments and Discussion

COMMENT BY

SERENA NG This paper by Olivier Coibion, Yuriy Gorodnichenko, and Mauricio Ulate is motivated by uncertainty over the state of the economy due to diverse estimates of potential output. As documented by George Perry (1977), there were similar concerns after the 1974–75 recession. Interest in how to measure potential output is just as strong now as it was then.

The premise of this paper is that a reasonable estimate of potential output should have certain key properties, among which are dynamic responses to shocks. Using estimates of potential output and shocks collected from a variety of sources, Coibion and colleagues find that potential output tends to respond too slowly to permanent shocks and too much to transitory shocks. This finding is robust across measures of potential output for the United States and other countries. The authors attribute the problem to a strongly persistent component that gets embedded into the potential output estimates. Because of this, the authors warn against making inference about GDP from revisions in potential output. The paper then proposes a new measure of potential output that depends only on permanent shocks. The methodology yields an output gap ranging from 5 to 10 percent in 2017. Explanations are then given for why the output gap might differ from the unemployment rate.

I enjoyed reading this creative paper, and appreciate the immense amount of work put into synthesizing the data from so many sources. I particularly like the idea of using auxiliary information (here, shocks constructed elsewhere) to validate estimates of latent variables. In what follows, I first provide a framework to help understand why measures of potential output might respond to shocks in the manner documented in the paper. I then turn to some issues with the proposed methodology. I suggest

that the maintained assumption of a linear trend is largely responsible for the large output gap of 5 to 10 percent, which is at odds with an unemployment rate that is near a historical low.

In what follows, I let Y_t be log GDP, Y_t^* be the level of potential output, and $\hat{Y}_{(t|t)}^*$ be a nowcast estimate of Y_t^* that can be obtained from the Congressional Budget Office, the Greenbook of the Federal Reserve Bank of Philadelphia, or other sources. The proposed nowcast estimate is $\tilde{Y}_{(t|t)}^*$. I let $\hat{Y}_t^* = \hat{Y}_{(t|T)}^*$ be the full sample (smoothed) estimate. In this notation, the paper considers the regression

$$(1) \quad \Delta Y_{it}^* = a + \sum_{k=0}^K \phi_k \hat{\epsilon}_{t-k} + error$$

where $\hat{\epsilon}_t$ is one of the shock series collected from a variety of sources and $K = 8$. The base case uses the Greenbook's $\Delta \hat{Y}_{(t|t)}^*$ over the sample 1987:Q1–2016:Q4. With the responses of ΔY_t as a benchmark, the paper by Coibion and colleagues reports (1) insufficient sensitivity of $\Delta \hat{Y}_{(t|t)}^*$ to some permanent shocks and incorrect response to other permanent shocks; (2) excess sensitivity of $\Delta \hat{Y}_{(t|t)}^*$ to monetary policy and government spending shocks that should only have transitory effects; and (3) a pattern of dynamic responses that can be replicated using a one-sided Hodrick–Prescott (HP) filter with $\lambda = 500K$ as Y_t^* , leading to the suggestion that “there seems to be little value added in estimates of potential GDP relative to simple measures of statistical trends.”

A quick remark on the third result. Even though the dynamic responses of $\hat{Y}_{(t|t)}^*$ are similar to those of an HP trend with $\lambda = 500K$, there is little in Coibion and colleagues' analysis to suggest that the level of such an HP trend is similar to the level of any of the $\hat{Y}_{(t|t)}^*$ that were carefully constructed. Furthermore, an HP trend is mean-squared optimal only if the cycle is white noise. As James Hamilton (2017) points out, $\lambda = 1600$ is already larger than the data-determined value. Thus, an HP trend with $\lambda = 500K$ cannot be seen as optimal in any meaningful sense.

UNDERSTANDING EXCESS AND INSUFFICIENT SENSITIVITY Results 1 and 2 hold for all five U.S. estimates of Y_t^* as well as for international data. It is thus useful to dig deeper into this. The shocks used in the analysis are themselves estimated, and their properties can in principle be questioned, but I abstract from this possibility. I suggest below that the first result is generic of filtering integration of order one [I(1)] processes, while the second result is symptomatic of a cyclical component that is strongly persistent.

Potential output plays a prominent role in policy work, but the variable is latent. Different estimates are obtained under different assumptions that are often not made explicit. To make sense of Coibion and colleagues' analysis, potential output must have a unit root (stochastic trend) component. Thus, let Y_t be the sum of a trend component τ_t and a stationary cyclical component c_t . The trend is itself the sum of two components: a deterministic trend d_t , a stochastic trend s_t :

$$(2) \quad Y_t = \tau_t + c_t$$

$$(3) \quad d_t + s_t + c_t.$$

I broadly define potential output as

$$Y_t^* = \mathbb{E}[\tau_t | \mathcal{R}] = \mathbb{E}[d_t + s_t | \mathcal{R}].$$

The challenge in identifying Y_t^* is that d_t , s_t , and c_t are latent, so we impose statistical and/or economic restrictions represented by \mathcal{R} . Different values of \mathcal{R} lead to different estimates. As shown by Arthur Okun (1983), Y_t^* cannot generate inflationary pressure. The HP filter constrains ΔY_t^* to change slowly, while a production function approach requires Y_t^* to be consistent with full employment. As is seen below, Coibion and colleagues restrict the transitory shocks to have no long-run effect on Y_t^* and produce an estimate of s_t under the maintained assumption that d_t is a linear trend.

To understand results 1 and 2, suppose that

$$(4) \quad s_t = s_{t-1} + e_t^s$$

$$(5) \quad a(L)c_t = e_t^c.$$

I assume, for simplicity, that (e_t^s, e_t^c) are serially and mutually uncorrelated innovations to the trend and the cycle, respectively. For understanding results 1 and 2, there is also no loss in assuming that $d_t = a + gt$ is a linear trend, where the growth rate g is known. The assumption that s_t is a random walk is without loss of generality. However, it is important that c_t is stationary ergodic: $\alpha(z) = 1 - \phi_1 z - \phi_2 z^2 - \phi_p z^p \neq 0$ for $|z| \leq 1$,

In the simplest case, when $c_t = \alpha c_{t-1} + e_t^c$, ΔY_t can be represented by

$$(6) \quad \begin{aligned} \Delta Y_t &= \Delta d_t + \Delta s_t + \Delta c_t \\ &= g + \Psi_y^s(L)e_t^s + \Psi_y^c(L)e_t^c. \end{aligned}$$

Under my assumptions, the true long-run effect of a unit permanent shock e_t^s on ΔY_t is $\psi_y^s(1) = \psi_s^s(1) = 1$, while the long-run effect of a unit transitory shock on ΔY_t is $\psi_y^c(1) = 0$, because $\psi_y^c(L) = \psi_c^c(1) = \frac{1-L}{1-\alpha L}$ is zero, evaluated at $L = 1$.

Suppose in addition to Y_t , we observe e_t^s , e_t^c , d_t , and s_t . Note, first, that Δs_t is white noise and c_t is serially correlated with a first-order autocorrelation coefficient of α . My table 1 reports results for four values of α : 0, 0.25, 0.8, and 0.95—with $T = 200$. The sample autocorrelation coefficients when s_t and c_t are observed correctly reflect the dynamics of the two stochastic processes, as seen from the row denoted (ρ_s, ρ_c) in the table's fourth and fifth columns. Least squares regressions should recover the dynamic effects of e_{t-j}^s and e_{t-j}^c on Δs_t . Let $(\psi_{sj}^s, \psi_{sj}^c)$ denote these regression coefficients for $j = 1, \dots, J$. The impact effect is given by $j = 0$ and the cumulative effect is captured by summing the coefficients from $j = 0$ to J . Because s_t is a random walk, one would expect $\psi_{s,0}^s = 1$ and $\psi_{sj}^s = 0$ for $j > 1$. Furthermore, ψ_{sj}^c should be zero for all j because e_{t-j}^c should have no effect on s_t at any lag.¹ As shown in the fourth and fifth columns of table 1, the estimates of ψ_{sj}^s and ψ_{sj}^c have the values that we expect.

Now suppose we regress ΔY_t instead of Δs_t on the shocks and denote the coefficients by $(\psi_{yj}^s, \psi_{yj}^c)$. The second and third columns of my table 1 indicate that the cumulative effects are close to the true values of $(1, 0)$ only when α is small. At $\alpha = 0.95$, the estimates of $(0.909, 0.229)$ are biased. Though the downward bias of $\hat{\psi}_y^s$ is largely gone when $T = 2,000$, $\hat{\psi}_y^s$ remains biased at 0.146. This suggests, on one hand, that the problem is not just a finite sample issue, but also that $\hat{\psi}_y^s$ and $\hat{\psi}_y^c$ may not be reliable benchmarks because they are biased at precisely the parameter region of interest.

We do not observe Y_t^* or its components d_t and s_t . Applying a one-sided filter $H(L)$ to Y_t gives the $\hat{Y}_t^* = H(L)Y_t$

$$\begin{aligned}\Delta \hat{Y}_t^* &= \Delta H(L)Y_t = H(1)g + H(L)e_t^s + \frac{H(L)(1-L)}{1-\alpha L}e_t^c \\ &= g^* + \psi_*^s(L)e_t^s + \psi_*^c(L)e_t^c.\end{aligned}$$

1. Coefficients with t statistics with less than 1 in absolute value are set to zero to get a more precise estimate of the long-horizon effect.

Table 1. Simulations to Illustrate Excess Sensitivity and Excess Smoothness^a

DGP: $Y_t = d_t + s_t + c_t = \tau_t + c_t$								
$\Delta s_t = e_t^s, \quad e_t^s \sim N(0, .03)$								
$(1 - \alpha L)c_t = e_t^c, \quad e_t^c \sim N(0, .3), \quad \text{corr}(e_t^s, e_t^c) = 0.$								
Regression: $Z_t = a + \psi_{sj}^s e_{t-j}^s + \psi_{sj}^c e_{t-j}^c + \text{error}.$								
j	$\hat{\Psi}_{sj}^s$	$\hat{\Psi}_{sj}^c$	$\hat{\Psi}_{sj}^s$	$\hat{\Psi}_{sj}^c$	$\hat{\Psi}_{MMj}^s$	$\hat{\Psi}_{MMj}^c$	$\hat{\Psi}_{HPj}^s$	$\hat{\Psi}_{HPj}^c$
$T = 200, \alpha = 0.0$								
0	1.003	0.999	1.000	-0.000	0.057	-0.001	0.103	0.101
1	-0.003	-1.000	0.000	0.001	0.042	-0.000	0.070	-0.022
2	-0.007	-0.000	-0.000	-0.000	0.044	-0.000	0.060	-0.010
0-12	0.009	0.001	0.000	0.001	0.043	-0.000	0.045	-0.005
0-40	1.013	-0.004	1.000	-0.003	0.381	-0.064	0.984	0.021
(ρ_s, ρ_c)			-0.005	-0.003	0.096	0.613	-0.015	0.422
$T = 200, \alpha = 0.25$								
0	0.989	0.999	1.000	-0.001	0.059	-0.001	0.100	0.099
1	0.017	-0.750	-0.000	-0.000	0.043	-0.000	0.071	0.003
2	-0.025	-0.187	0.000	0.001	0.043	-0.001	0.058	-0.008
0-12	0.007	-0.011	0.000	0.000	0.043	-0.001	0.047	-0.007
0-40	0.933	-0.002	1.000	-0.003	0.398	-0.091	0.984	0.022
(ρ_s, ρ_c)			-0.002	0.240	0.123	0.691	0.058	0.548
$T = 200, \alpha = 0.8$								
0	0.993	0.999	1.000	-0.001	0.057	0.008	0.102	0.103
1	-0.006	-0.201	0.000	-0.001	0.042	0.002	0.071	0.052
2	0.005	-0.161	0.000	-0.001	0.042	0.001	0.058	0.032
0-12	0.010	-0.103	-0.000	-0.000	0.040	0.000	0.047	0.010
0-40	1.057	0.014	1.000	-0.007	0.390	-0.289	0.985	0.086
(ρ_s, ρ_c)			-0.002	0.780	0.268	0.880	0.236	0.811
$T = 200, \alpha = 0.95$								
0	0.992	1.000	1.000	-0.001	0.055	0.033	0.104	0.103
1	0.002	-0.049	-0.000	-0.001	0.041	0.025	0.066	0.068
2	-0.002	-0.044	0.000	0.001	0.041	0.024	0.056	0.055
0-12	0.006	-0.040	0.000	0.000	0.041	0.022	0.047	0.038
0-40	0.909	0.229	1.000	-0.004	0.313	-0.314	0.952	0.422
(ρ_s, ρ_c)			-0.006	0.929	0.391	0.938	0.308	0.912
$T = 200, \alpha = 0.95$								
0	1.000	1.000	1.000	-0.000	0.049	0.032	0.042	0.043
1	0.001	-0.050	0.000	0.000	0.049	0.029	0.039	0.038
2	0.002	-0.048	0.000	-0.000	0.049	0.028	0.038	0.034
0-12	0.002	-0.043	0.000	0.000	0.049	0.025	0.035	0.029
0-40	1.011	0.146	1.000	-0.002	0.541	-0.240	0.973	0.359
(ρ_s, ρ_c)			0.001	0.948	0.754	0.949	0.611	0.931

Source: Author's calculations.

a. Results are based on the mean over 1,000 replications. The parameter ψ_{sj}^s is the effect of e_{t-j}^s on Z_t . Results are reported for $j = 0, 1$, and 2 , as well as the sum of the coefficients up to lag 12 and 40. The (ρ_s, ρ_c) reports the first-order autocorrelation of variable Δs_t and c_t when they are observed, or the estimated trend and cycle when they are not.

I consider two choices of $H(L)$. The first is a 20-period, moving-mean filter (denoted *MM* in the sixth and seventh columns of my table 1). The second is the one-sided HP filter with $\lambda = 500K$ (denoted *HP* in the table's eighth and ninth columns). As seen from the first-order autocorrelation coefficient of $\Delta\hat{Y}_t^*$, the series exhibits correlation, even though Δs_t is white noise by design, and the gap estimate is more persistent than the true cycle. As seen from (ρ_s, ρ_c) for the last two columns of my table 1, $\Delta\hat{Y}_t^*$ is less persistent than Δs_t , while \hat{c}_t is much more persistent than c_t . This is merely echoing the findings of Timothy Cogley and James Nason (1995), and many others, about the consequences of filtering.

What do we get when we regress $\Delta\hat{Y}_t^*$ instead of Δs_t on the shocks? Evidently, $\hat{\psi}_*^s(L)$ (with $*$ = *MM* or *HP*) is severely biased for the true value of 1, while the estimated effects of e_{t-j}^s on $\Delta\hat{Y}_t^*$ all differ from the value of zero. This arises because $H(L)$ spreads out the effect of the permanent shock over time. Result 1 documented in the paper by Coibion and colleagues—that $\Delta\hat{Y}_{t/t}^*$ reacts insufficiently to the permanent shock—is consistent with the simulation results. This bias is largely invariant to the dynamics of c_t .

In contrast, result 2 depends on the dynamics of c_t . Though the effect of the transitory shock e_t^c on the Δs_t is zero, the effect on ΔY_t^* is not. As seen from my table 1, the *MM* estimate of ΔY_t^* yields estimates of ψ_*^c that are similar in magnitude each period, a reflection of the constant weights in the *MM* filter. Though the bias at each lag is small, the bias in the cumulative effect is not. The bias in the *HP* estimates reflect the declining pattern of the *HP* weights. For both filters, the bias in the cumulative effect grows as α increase. When $\alpha = 0.95$, the cumulative effect of 40 lags is -0.314 for *MM* and 0.422 for *HP*. The bias is opposite in sign; hence, the choice of filter matters. This finding of excess sensitivity to transitory shocks is robust to changing alternative trend specifications as long as c_t has an autoregressive root local to unity.

The excess sensitivity result can be traced to the discontinuity of $\psi_*^c(L) = \frac{H(L)(1-L)}{1-\alpha L}$ at $\alpha = 1$. When α is far from 1, the long-run effect of e_t^c as measured by $\psi_*^c(1)$ is zero, because $1-L=0$ evaluated at $L=1$. However, the term is of order 1 when $\alpha=1$. When α is close to 1, there is a near-cancellation of $(1-L)$ in the numerator, with $(1-\alpha L)$ in the denominator, and e_t^c will appear as if it has permanent effects in finite samples. In other words, when c_t is highly persistent, it can be mistaken for s_t . The simulations in my table 1 bear this out; when c_t is highly persistent, transitory shocks have effects on $\Delta\hat{Y}_t^*$ that persist even after 40 periods.

One may ask if a cyclical component that is highly persistent is realistic. My own estimation of an unobserved components model (not reported) finds that the largest autoregressive root in c_t is bigger than 0.95. James Morley, Charles Nelson, and Eric Zivot (2003), and, more recently, Angelia Grant and Joshua Chan (2017) reported similar estimates for α . The autoregressive root in the unemployment series (more on this below) is also suggestive of a strongly persistent c_t . Now the half-life of a shock when $\alpha = 0.95$ is $\log(0.5) \div \log(0.95) = 14$ periods and increases to $\log(0.5) \div \log(0.98) = 34$ periods when $\alpha = 0.98$. Thus, though the observation by Coibion and colleagues that $\hat{\Delta Y}_{(t|t)}^*$ tends to respond to transitory shocks cannot be disputed, the finding that the responses are nonzero after eight quarters is not informative as to whether the long-run responses will be zero.

Coibion and colleagues have identified an interesting feature of many nowcast estimates of potential output that surprisingly has gone unnoticed. I conjecture that the finding will also hold for the smoothed estimates of potential output because the root problem is a cyclical component in GDP that is highly persistent, not data revisions.

THE PROPOSED ESTIMATE OF Y_t^* Coibion and colleagues' premise is that the $\hat{Y}_{(t|t)}^*$ measures are contaminated by persistent cyclical variations. But how to remove the "nearly permanent" cyclical component from these estimates? Bias-adjusting $\hat{\Delta Y}_{(t|t)}^*$ is difficult because we know how to estimate the purely permanent purely transitory quantities, but we are not very good at dealing with nearly permanent ones. Coibion and colleagues do not bias-adjust existing estimates but try something different. For $\Delta Y_t = g + \theta_i^s(L)e_t^s + \theta_i^c(L)e_t^c$, their idea is to define potential output growth as

$$(7) \quad \begin{aligned} \Delta Y_{(t|t)}^* &= \mathbb{E}_t [\Delta Y_t | e_t^s, \dots, e_0^s, \dots] \\ &= g + \theta_i^s(L)e_t^s \end{aligned}$$

where e_t^s, \dots, e_0^s are permanent shocks up to period t . Given real-time estimates of these shocks from a rolling-window application of the Blanchard-Quah (BQ) method, Coibion and colleagues then construct

$$(8) \quad \Delta \tilde{Y}_{(t|t)}^* = g + \hat{\theta}_i^s(L)\hat{e}_{(t|t)}^s.$$

The proposed potential output $\tilde{Y}_{(t|t)}^*$ is the cumulative sum of $\Delta \tilde{Y}_{(t|t)}^*$, with $g = 0.031$, which is the average growth rate of Y over the sample period 1977–2007. This notion of trend output is similar in spirit to Beveridge

and Nelson's (1981) trend. Both methods aim to produce an estimate of s_t , assuming d_t is a linear trend. But Coibion and colleagues use the nowcast (instead of forecast) of output growth, their analysis is bivariate instead of univariate, and they take the extra step to make a permanent/transitory decomposition of the shocks.²


Estimation of the permanent shocks using the BQ method comes with some caveats. First, it depends on one variable (output) being $I(1)$ and one variable (the unemployment rate) being $I(0)$. But the unemployment rate has an $AR(1)$ coefficient of 0.97. This is computed over the sample 1961:Q1–2016:Q4. Furthermore, the methodology depends on the choice of variables. Using capacity utilization in lieu of the unemployment rate, for example, will give a different estimate of potential output. Instead of two variables, Coibion and colleagues could have used a bigger vector autoregression, as is done by Anders Warne (1991) and by Jesús Gonzalo and me (2001), or they could have identified the stochastic trend directly—as is done by James Stock and Mark Watson (1988), Robert King and others (1991), and Gonzalo and Clive Granger (1995)—without going through the step of identifying the underlying shocks. A larger cause of concern is that all permanent/transitory decompositions implicitly or explicitly rely on estimates of the spectral density at frequency zero from the data.³ If all carefully constructed $\hat{Y}_{(t|t)}^*$ series considered in Coibion and colleagues' paper have failed to isolate the pure trend component, one cannot be overly optimistic that the BQ methodology can succeed in doing so with 20 years of data. The standard errors around $\tilde{Y}_{(t|t)}^*$ must be unacceptably large.

THE IMPORTANCE OF D_t Although having a measure of Y_t^* that is not affected in the long run by transitory shocks is desirable, the level of $Y_t^* = \mathbb{E}[d_t + s_t | \mathcal{R}]$ is of interest, not the counterfactual response of s_t to shocks. For this, the assumption on d_t becomes important. I will suggest that their implied output gap of 5 to 10 percent is due to the questionable assumption of the linear deterministic trend.

Even though Coibion and colleagues performed a rolling window estimation, $\tilde{Y}_{(t|t)}^*$ is still based on the assumption that g is a constant 0.031

2. Assuming that $d_t = \alpha + gt$ and s_t is a random walk, Beveridge and Nelson (1981) define $s_t^{BN} = \lim_{k \rightarrow \infty} \mathbb{E}_t(Y_{t+j} - kg) = Y_t + \lim_{k \rightarrow \infty} \sum_{k=1}^{\infty} \mathbb{E}_t(\Delta Y_{t+k} - g)$. Beveridge and Nelson's (stochastic) trend adjusts the current level Y_t by deviations of future growth from mean g .

3. In particular, the BQ and Galí (1999) methods essentially impose zero restrictions on the spectral density at frequency zero, while the Cochrane (1994) and Gonzalo and Ng's (2001) methods rely on cointegration arguments and still need some restriction on the spectral density at the zero frequency.

Figure 1. Smoothed GDP Growth, 1960–2020^a

Source: Author's calculations.

a. The mw line is the low-frequency component of annualized GDP growth, as constructed by Mueller and Watson (2008). The med5 line is the five-year moving median, and the mean5 line is the five-year moving mean of the series.

throughout. Over the sample 1948:Q1–2018:Q2 the residuals from linear detrending—that is, from regressing Y_t on a constant and a trend—have been -12 percent on average since 2012 and have been becoming more negative. The magnitude is in line with the -5 to -10 percent gap implied by $\tilde{Y}_{(t|t)}^*$. But a linear trend is monotone and cannot adapt to changes in demographics, technology, or any other structural aspects that have evolved over time. A quadratic trend that bends, for example, yields a gap of about -2 percent, and the second-order term is strongly statistically significant. The peak-to-peak method considered by Bradford De Long and Lawrence Summers (1988) gives a gap of about -1 percent. This is not to say that these methods are optimal, but rather that the linear trend is too rigid and leaves too much predictable variation unexplained to be desirable. My figure 1 shows the five-year moving mean and moving median of annualized GDP growth, along with the low-frequency component in the series


estimated by the procedure of Ulrich Mueller and Watson (2008). It is then clear that the issue is not whether g should be 3.1 percent or 2.6 percent, but rather that it is assumed to be the same constant throughout. Because growth is well below the overall mean of 3.1 percent in the last 10 years, the constant growth assumption will overestimate Y_t^* during this period, giving a large output gap. Pierre Perron (1989) finds that misspecifying d_t can lead one to conclude that a unit root is present when the data are actually trend stationary. Misspecification of d_t will also affect the estimation of Y_t^* .

It is not a trivial task to disentangle s_t from c_t when c_t is strongly persistent, even if d_t is known. When d_t is itself of unknown form, as seen from my figure 1, the exercise becomes a formidable task. A different way to see the problem is that any s_t that is a unit root process is $Op(T)$. Any polynomial time trend is at least $O(T^{\frac{3}{2}})$, so d_t dominates s_t when both are present. We cannot model s_t without first removing d_t . Given this difficulty, some may prefer to use the unemployment rate as a guide to the state of the economy. Coibion and colleagues provide compelling economic arguments for why the output gap is still a variable of interest. But from an econometric point of view, extracting a c_t from the unemployment rate (UR) is a more manageable exercise because it does not show a trend over time. As such, the d_t component of UR is just a constant and the only possible source of nonstationarity in UR is s_t . Identification of s_t is then much simpler, at least within the framework of unobserved components.

This, then, raises the question that perhaps the unobserved components model is asking too much of the data, and we should be content with being able to separate variations above and below certain frequencies. Mueller and Watson (2008) suggests a procedure (hereafter MW) that consists of projecting the series of interest on $K = 12$ cosine functions and taking the residuals as the cycle. After the low-frequency component is removed, the MW/UR gap still has an autoregressive root of about 0.92, similar to the one in the MW/output gap of about 0.89. My figure 2 plots the MW/output gap along with the MW/UR gap, but renormalized and centered to have the same mean and variance as the MW/output gap. We see that the two series match up remarkably closely over the last six decades. In 2016, which is the end of Coibion and colleagues' sample, there indeed appears to be more slackness in output than in the labor market, but much smaller than the 5 to 10 percent suggested by Coibion and colleagues. Both gaps suggest that the economy is near capacity in 2017.

In summary, Coibion and colleagues' results are consistent with a cyclical component in GDP that is strongly persistent. When the GDP data alone are uninformative about the trend component, using auxiliary information

Figure 2. The Unemployment Rate versus GDP: The Mueller–Watson Method, 1945–2020^a



Source: Author's calculations.

a. The unemployment rate line is the cyclical component constructed by removing the low-frequency component of the unemployment rate. The GDP line is the cyclical component of log GDP. Both low-frequency and cyclical components are computed using the method of Mueller and Watson (2008).

to help with identification is potentially useful. These variables Z can be thought of “external” instruments. The question is how to use this information. Coibion and colleagues use shocks as Z and require that the sum of coefficients on the temporary shocks in equation 1 sum to zero. But their restriction only gives us a better estimate of s_t for a given d_t , while the level of Y_t^* is largely determined by d_t . The exercise is incomplete without a careful modeling of d_t .

REFERENCES FOR THE NG COMMENT

Beveridge, Stephen, and Charles R. Nelson. 1981. “A New Approach to Decomposition of Economic Time Series into Permanent and Transitory Components,

- with Particular Attention to Measurement of the Business Cycle.” *Journal of Monetary Economics* 7, no. 2: 151–74.
- Cochrane, John H. 1994. “Permanent and Transitory Components of GNP and Stock Prices.” *Quarterly Journal of Economics* 109, no. 1: 241–65.
- Cogley, Timothy, and James M. Nason. 1995. “Effects of the Hodrick-Prescott Filter on Trend and Difference Stationary Time Series: Implications for Business Cycle Research.” *Journal of Economic Dynamics and Control* 19, nos. 1–2: 253–78.
- De Long, J. Bradford, and Lawrence H. Summers. 1988. “How Does Macroeconomic Policy Affect Output?” *Brookings Papers on Economic Activity*, no. 2: 433–80.
- Galí, Jordi. 1999. “Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?” *American Economic Review* 89, no. 1: 249–71.
- Gonzalo, Jesús, and Clive Granger. 1995. “Estimation of Common Long-Memory Components in Cointegrated Systems.” *Journal of Business and Economic Statistics* 13, no. 1: 27–35.
- Gonzalo, Jesús, and Serena Ng. 2001. “A Systematic Framework for Analyzing the Dynamic Effects of Permanent and Transitory Shocks.” *Journal of Economic Dynamics and Control* 25, no. 10: 1527–46.
- Grant, Angelia L., and Joshua C. C. Chan. 2017. “Reconciling Output Gaps: Unobserved Components Model and Hodrick-Prescott Filter.” *Journal of Economic Dynamics & Control* 75: 114–21.
- Hamilton, James D. 2017. “Why You Should Never Use the Hodrick-Prescott Filter.” *Review of Economics and Statistics* 100, no. 5: 831–43.
- King, Robert G., Charles I. Plosser, James H. Stock, and Mark W. Watson. 1991. “Stochastic Trends and Economic Fluctuations.” *American Economic Review* 81, no. 4: 819–40.
- Morley, James C., Charles R. Nelson, and Eric Zivot. 2003. “Why Are the Beveridge-Nelson and Unobserved-Components Decompositions of GDP So Different.” *Review of Economics and Statistics* 85, no. 2: 235–43.
- Mueller, Ulrich K., and Mark W. Watson. 2008. “Testing Models of Low-Frequency Variability.” *Econometrica* 76, no. 5: 979–1016.
- Okun, Arthur M. 1983. “Potential GNP: Its Measurement and Significance.” In *Economics for Policymaking: Selected Essays of Arthur M. Okun*, edited by Joseph A. Pechman. MIT Press.
- Perron, Pierre. 1989. “The Great Crash, the Oil Price Shock, and the Unit Root Hypothesis.” *Econometrica* 57, no. 6: 1361–1401.
- Perry, George L. 1977. “Potential Output and Productivity.” *Brookings Papers on Economic Activity*, no. 1: 11–47.
- Stock, James H., and Mark W. Watson. 1988. “Testing for Common Trends.” *Journal of the American Statistical Association* 83, no. 404: 1097–1107.
- Warne, Anders. 1993. “A Common Trends Model: Identification, Estimation, and Asymptotics.” Working paper. Stockholm: Institute for Economic Studies.

COMMENT BY

VALERIE A. RAMEY This paper by Olivier Coibion, Yuriy Gorodnichenko, and Mauricio Ulate presents surprising new results showing that the leading real-time estimates of potential GDP for the United States and other industrialized countries react to temporary demand shocks. Potential GDP is intended to be an estimate of the maximum sustainable level of output that does not generate inflationary pressure. Because it is a supply-side concept, potential output should not react to demand shocks with temporary effects but should react fully to supply shocks with permanent effects. Coibion and colleagues present convincing evidence that none of the leading estimates of potential GDP satisfies this dichotomy.

Coibion and colleagues have three goals for their paper. Their first goal is to demonstrate that estimates of potential GDP by the various governmental and nongovernmental institutions in the U.S. and other industrialized countries overreact to shocks that have temporary effects on actual GDP and underreact to shocks that have permanent effects on actual GDP. The authors carefully construct real-time databases and use a variety of methods for estimating shocks to show convincingly that leading institutions, such as the Congressional Budget Office (CBO), revise their estimates of potential GDP in response to shocks that are easily identified, even in real time, as temporary. The authors estimate a variety of standard demand shocks, such as monetary and fiscal shocks, first showing that the impulse responses of actual GDP imply temporary effects and then showing that the estimates of potential GDP are revised in response to those shocks. They then estimate supply shocks, showing first that they have permanent effects on actual GDP and then that estimates of potential GDP are not revised sufficiently in response. Achieving this first goal constitutes two-thirds of the paper, and is its heart. These sections of the paper make a substantial contribution: the demonstration is very convincing, and the results are important because estimates of potential GDP are central to numerous quantitative models and are also important guides for policymakers. Perhaps one of the most surprising details in their findings is that estimates of potential GDP by the Federal Reserve's army of Ph.D. economists are virtually indistinguishable from a simple Hodrick–Prescott filter trend and that the Federal Reserve's own estimates of potential GDP are revised based on estimated monetary policy shocks. That is, the Federal Reserve's estimates of potential GDP behave as if monetary policy shocks have permanent supply-side effects, even though the impulse responses of actual GDP show no permanent effects of monetary policy shocks.

The paper's second goal is to explore alternative methods for estimating potential GDP that overcome the problems highlighted in the authors' demonstration. Their main suggested alternative is Olivier Blanchard and Danny Quah's (1989) decomposition of GDP shocks into demand and supply shocks using long-run restrictions, known as the BQ method. Coibion and colleagues show that this measure of potential GDP does not suffer from the same weaknesses as standard measures documented in the earlier sections of the paper. In addition, they explore a variety of other methods based on economic-theory with either alternative long-run restrictions based on theory or Phillips curves.

Finally, Coibion and colleagues' third goal is the production of an alternative measure of the current output gap. Using their implementation of the BQ method, they offer an alternative estimate of current potential GDP and conclude that actual GDP was still more than 5 percent below potential GDP in 2017.

I believe that Coibion and colleagues are very successful in achieving their first goal. Their careful demonstration of the weaknesses of current methods makes it clear that estimates of potential GDP can be improved. Regarding their second goal, their explorations of alternative methods are very promising. I believe that their choice of alternatives is very good. However, as I make clear below, there are remaining challenges with the implementation of their preferred alternative, so more work needs to be done. I demonstrate that key assumptions in their implementation lead to their implausible conclusion that current GDP is significantly below potential GDP. As a result, I do not think their estimates are ready for use by policymakers.

THE PAPER'S ALTERNATIVE METHODS FOR ESTIMATING POTENTIAL GDP To address the weaknesses of the standard estimates of potential GDP, Coibion and colleagues explore alternative methods for estimating potential GDP that can distinguish between shocks that have temporary versus permanent effects on actual output. The main alternative method they explore is the BQ decomposition method. This method uses a bivariate time series model with real GDP and the unemployment rate, and it identifies supply shocks as those shocks that have long-run effects on GDP and demand shocks as all other shocks that have temporary effects. Even if one does not agree with BQ's supply shock-versus-demand shock dichotomy, their method is still useful for separating out temporary from permanent shocks to GDP, which is the key to improving estimates of potential GDP.

Coibion and colleagues also explore other alternatives. For example, they use Jordi Galí's (1999) long-run restriction to identify permanent

shocks to technology; John Cochrane's (1994) permanent income hypothesis-motivated method for using the behavior of consumption to identify permanent shocks to GDP; and a Phillips curve model to infer potential GDP from inflation dynamics. The authors' implementation of all these methods implies much larger current output gaps—that is, actual GDP is farther below potential GDP than those implied by the CBO's estimates and others.

I focus on Coibion and colleagues' implementation of the BQ method because that is their favored method, and that method actually gives a more conservative estimate of the gap relative to their other alternatives. Nevertheless, the authors' particular implementation of the BQ method implies a large gap. Their estimate of potential GDP leads them to conclude that “the gap between potential and actual output in the U.S. increased by about 5 log percentage points between 2007:Q1 (when the gap was likely close to zero) and 2017:Q1, leaving ample room for policymakers to close this gap through demand-side policies if they chose to do so.” Thus, their estimates can be seen as an encouragement for policymakers to undertake more demand-side stimuli, even when the unemployment rate is below 4 percent.

IMPLICATIONS OF THE AUTHORS' POTENTIAL GDP ESTIMATES I now demonstrate that Coibion and colleagues' alternative estimates of potential, while avoiding the weaknesses they highlighted for the standard estimates, have a number of implications ranging from questionable to implausible. I argue, however, that the problem is that their estimates are based on questionable auxiliary identifying assumptions that are relatively easy to fix.

Implication 1: Coibion and colleagues' estimates of potential GDP decline as much as the CBO's estimates after the Great Recession. One of Coibion and colleagues' main critiques of the CBO revisions of potential GDP is that they lowered them too much from 2007 to 2017, in response to cyclical fluctuations. Figure 1 of their paper shows how the CBO's estimates of potential GDP changed from 2007 to 2017, and figure 13 shows how their own BQ estimates changed in real time from 2007 and 2017. Consider the revision for the estimate of potential output at the end of their sample, 2016:Q4. Using their data and programs, I calculated that the CBO revised down its estimate of potential GDP in 2016:Q4 by about 0.12 log points, whereas Coibion and colleagues' BQ estimate was revised down by about 0.11 log points over the same period. Thus, both methods lead to the same downward revision in potential GDP. If we believe that Coibion and colleagues' method is accurately capturing only permanent shocks, then their method *validates* the CBO revisions.


Implication 2: The implied natural rate of unemployment is implausibly low. We can combine Coibion and colleagues' estimate of the output gap with Okun's law to calculate the implied natural rate of unemployment. In their paper, Coibion and colleagues conduct this exercise in subsection IV.C. However, they use the older historical estimates of -3 for the parameter on the unemployment gap term rather than the more up-to-date estimates of -2 (Ball, Leigh, and Loungani 2017). Furthermore, they use their estimate of the output gap in 2016:Q4. Because the unemployment rate has fallen so much since then, adding more recent data is instructive.

Thus, I update Coibion and colleagues' BQ estimates through 2018:Q2, using their same programs and the same rolling window over the previous 30 years. I find that actual output is about 6.6 percent below their estimate of potential GDP in 2018:Q2. Thus, using their method, I find that actual GDP is farther below potential GDP in 2018:Q2 than it was in 2016:Q4.

The unemployment rate in 2018:Q2 was 3.9 percent. Using Okun's law with a modern unemployment gap coefficient of -2 implies that the natural rate of unemployment in 2018:Q2 was about 0.6 percent. This unemployment rate is below any level ever achieved in the United States, including World War II, and is completely implausible.

Coibion and colleagues argue, however, that the usual Okun's law relationship no longer applies because the employment-to-population ratio in the U.S. fell so much during the Great Recession. Though this is an intriguing possibility, I show below that there is a much simpler explanation for why they estimate such a large output gap and implied low natural rate of unemployment: one of their auxiliary identifying assumptions leads potential GDP to have a significantly higher growth rate than actual GDP in the long run.

Implication 3: Coibion and colleagues' implied output gap has a strong upward trend. As mentioned briefly in discussing the last point, Coibion and colleagues' method for estimating potential GDP implies a bigger output gap in 2018 than at the end of 2016, which seems odd given the fast pace of growth of the U.S. economy and the significant decline in the unemployment rate. This feature led me to inspect Coibion and colleagues' implied output gap for the last 30 years more closely, because they use 30-year rolling regressions to counter possible breaks in trends. In my figure 1, I show the output gap estimated by the CBO and by Coibion and colleagues, where the gap is defined as log actual output minus log potential output so that the gap should be negative at the end of a recession. The CBO's gap behaves as expected, varying cyclically but with no trend. In contrast, the dominant feature of the Coibion and colleagues' implied gap

Figure 1. Estimated Output Gaps, $Y - Y^*$, 1988–2018^a

Sources: Author's estimates, using programs from Coibion and colleagues' paper and updated data from the FRED database of the Federal Reserve Bank of Saint Louis.


a. Y = log actual GDP; Y^* = log potential GDP.

is a strong downward trend—the estimated gap declines at a rate of about 0.6 percent per year. According to their estimates, the output gap was very positive in 1988, implying that *actual* output was almost 12 percent above *potential*. However, over time, this gap has narrowed and has become negative. According to the authors' estimates, the output gap is wider now, at about –6.6 percent, than it was at the end of the Great Recession, when it was about –2.2 percent.

This result is a direct consequence of Coibion and colleagues' estimated potential GDP having a much higher trend than actual GDP over the last 30 years. My figure 2 shows the path of both series. Even in the second half of the 1990s, when the growth of total factor productivity surged because of the information technology revolution, they estimate that actual GDP was significantly above potential GDP. The two series cross in 2007, and then the gap becomes negative and widens over time because their estimate of potential GDP grows more quickly than actual GDP. The next section explains which of the assumptions made by Coibion and colleagues lead to this implausible behavior.

Figure 2. Actual GDP versus Potential GDP, as Estimated by Coibion and Colleagues, 1988–2018

Logarithms




Sources: Author's estimates, using programs from Coibion and colleagues' paper and updated data from the FRED database of the Federal Reserve Bank of Saint Louis.

THE BQ METHOD IS NOT ENOUGH TO IDENTIFY POTENTIAL GDP Coibion and colleagues recognize that in order to implement the BQ method to derive a path of potential GDP, they must take a stand on the long-run growth rate of the economy. To see the identification problem, consider the intuitive equation they offer in subsection IV.A of their paper:

$$\Delta \log Y_t = g + \Delta \log Y_t^p + \Delta \log Y_t^c,$$

where $\Delta \log Y_t$ is the growth rate of actual GDP, g is the long-run growth rate of GDP, $\Delta \log Y_t^p$ is the growth rate of output due to permanent shocks, and $\Delta \log Y_t^c$ is the growth rate of output due to temporary shocks. The BQ method assumes that permanent shocks can permanently affect the *level* of GDP, but not the *growth rate* of GDP. Therefore, the BQ method identifies only *deviations* from a long-run path; hence, neither the slope (g) of this path nor the intercept is identified.

Thus, Coibion and colleagues are forced to make two additional assumptions to identify the path. To identify the *slope* of the path, they assume a value of g of 3.1 percent, which equals both the average growth rate of

Figure 3. Estimated Output Gaps, $Y - Y^*$, 1950–2020^a

Sources: Author's estimates, using programs from Coibion and colleagues' paper and updated data from the FRED database of the Federal Reserve Bank of Saint Louis.

a. The CBO's estimate versus the Blanchard–Quah method's estimate on the full sample.

real GDP from 1977 to 2007 and for the entire post–World War II period. To identify the *intercept* of the path, they assume that potential GDP was equal to actual GDP in 2007:Q1. Also, the CBO's estimated gap is then only about -0.3 percent, so this assumption is close to the CBO's estimates. However, as my figure 2 shows, the slope estimate for g leads the authors' estimate of potential GDP to grow much faster than actual GDP from 1988 to 2018. It is this divergence in growth rates that leads directly to their estimate that output is currently 6.6 percent below potential GDP.

The problem of different growth rates for actual and potential GDP would not occur if g were set equal to the actual growth rate of GDP over the sample used in the estimation. To demonstrate this, I updated the authors' data and reestimated their BQ model back to 1948 and created output gap estimates. These are shown in my figure 3, along with the CBO's estimates. As the figure shows, there is no longer a trend in the gap estimate. However, the two estimates do not move in lockstep. The correlation between the CBO's gap estimate and the BQ gap estimate is about 0.5, suggesting that much could be learned from the differences in the implied gaps.

CONCLUSION Overall, this is an important paper that effectively demonstrates that standard measures of potential GDP overreact to temporary shocks and underreact to permanent shocks. It makes a convincing argument that we can do better, even in real time. The alternative methods explored are promising, but the methods still need work, so any implied gap estimates are “not yet ready for prime time.” For now, I think I will stick with the CBO’s estimate of the gap, which indicates no slackness in the U.S. economy.

REFERENCES FOR THE RAMEY COMMENT

- Ball, Laurence, Daniel Leigh, and Prakash Loungani. 2017. “Okun’s Law: Fit at 50?” *Journal of Money, Credit and Banking* 49, no. 7: 1413–41.
- Blanchard, Olivier J., and Danny Quah. 1989. “The Dynamic Effects of Aggregate Demand and Supply Disturbances.” *American Economic Review* 79, no. 4: 655–73.
- Cochrane, John H. 1994. “Permanent and Transitory Components of GNP and Stock Prices.” *Quarterly Journal of Economics* 109, no. 1: 241–65.
- Gali, Jordi. 1999. “Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?” *American Economic Review* 89, no. 1: 249–71.

GENERAL DISCUSSION James Stock began by drawing a firmament analogy, wondering if among the stars in the firmament, potential output— Y^* —had any contributions beyond the natural rate of unemployment— U^* . He postulated that in principle, the answer could be yes, because potential output can incorporate capital accumulation, total factor productivity growth, changes in underlying population growth, and changes in the labor force participation rate. This can provide additional information and help explain measures of slackness in the economy and, therefore, thinking about monetary and fiscal policy. However, each of these additional factors has many problems in practice. He acknowledged that perhaps it is plausible to forecast population growth or put aside immigration issues, but there are still ongoing challenges in understanding the labor force participation rate and total factor productivity growth. Although, in principle, it might be possible to get these things right—such as determining the underlying growth rate, and thus making measures of potential output more informative than the natural rate of unemployment—whether this can be pulled off in practice remains doubtful.

Stock stated that based on evidence presented by the authors and other evidence that he had seen, it was not clear to him if the additional challenges of moving to potential output from the natural rate of unemployment are worth it. It was also not clear to him if much more can be done than just estimating an output (Y) gap as an unemployment (U) gap times a rolling Okun's law coefficient (or something along those lines), and if there is any value in that. A plot of the output gap estimated by the Congressional Budget Office (CBO) against its estimation of the unemployment gap reveals that the two gaps are almost same, with a slight time variation in the Okun's law coefficient between the two plots. Therefore, in practice, the CBO's methodology states that it is tough to learn anything about an output gap that is not already observed in an unemployment gap, although this fact remains buried in the methodology. Stock thought that while the unemployment gap is challenging to measure, the scope for challenge is a little bit less than that for measuring the output gap. Going back to his firmament analogy, he concluded that despite the potential for learning more by looking at an output gap and potential output, it made sense for him to pull that star, Y^* , out of the firmament and to continue focusing on the natural rate of unemployment, U^* .

Jonathan Pingle asked the authors and commenter Valerie Ramey if they thought that a greater consideration of the 2005–7 period might be useful. Considering the paper's analysis of growth rate shocks, Pingle noted that the stepping off point was that the output gap in the last expansion barely seemed to close, although that is a little obscured by indexing in many of the estimates. However, the subsequent level of the gap matters a lot for policy. Pingle observed that in the 2005–7 period, inflation was running above the central bank target, despite globalization pressures, and there were imbalances in the economy, including overactivity in housing, a sector that is sensitive to the interest rate. This would imply an output gap that was more than just closed, or is inconsistent with an appropriate equilibrium target for monetary policy. He asked whether the level and stepping off point deserve more consideration, or whether a level of potential output that is too high is simply being carried forward.

Olivier Blanchard had two comments. The first was about the Blanchard–Quah approach, which is mentioned in the paper and the discussions. There are two conceptually separate steps to this approach. The first is a statistical decomposition of output between the part due to shocks that have a permanent effect on output, and shocks that only have a temporary effect. The second is the reference to the shocks with permanent effects

as “supply shocks,” and to the shocks with temporary effects as “demand shocks.” The first step is simply data description and should be uncontroversial. The second step is controversial. One can buy the first step and construct a series for output due to shocks with permanent effects, and call it potential output, without accepting the second step.

Blanchard, drawing on his experience having seen the construction of potential output in various institutions, suggested that it often suffered from two problems. The first is intellectual laziness and the ease of using a simple statistical method, such as the Hodrick–Prescott (HP) filter, without thinking hard about the implicit assumptions behind it. The second includes political factors. Looking at retrospective revisions of the output gap in Greece is revealing: In real time, the assessment was that there was not much of an output gap. However, the output gap at the start of the crisis is now viewed as having been a large positive number, which makes the fall look less bad than it would otherwise be viewed as being. Blanchard suspected that though this is much less relevant for the CBO, it is nonetheless worth considering.

Turning to the question of what to do about measuring potential output and the output gap, Blanchard thought that economists need to look for signals of whether the economy is overheating or underheating. Inflation is far from a perfect signal, but it is the most natural one and the first one that should be considered. Blanchard did not think there was enough consideration of inflation by the authors, other than mentions of the Phillips curve. If economists really think that the inflation signal is becoming worse, which many do, then it is important to look at many dimensions of the labor market—such as the degree of labor force participation relative to a reasonable trend, the ratio of vacancies to unemployment, and the degree of involuntary part-time work—and then to use all these variables, together with inflation, to get to the natural rate of unemployment.

John Haltiwanger was struck by how few data go into producing potential output and real-time output, and he thought they are somewhat related. Starting with real output, he noted that statistical agencies continue to use very crude methods for their benchmark revisions, both methodologically and in timing. For example, the “birth-death model” used for the benchmark revisions from the payroll survey is really outdated in many ways and has always been crude. The Bureau of Labor Statistics is sitting on top of administrative data and could do a much better job. Haltiwanger acknowledged that the bureau is strapped for resources,

so he did not want to pick on it, but that the general question is whether much more progress could be made in solving some of these problems with real-time data.

Moving on to potential output, Haltiwanger observed that there has been a lot of thinking about changes in potential output, building up from microeconomic evidence. A recent Jackson Hole conference was about these topics, such as changing market structure, concentration, and the changing role of start-ups.¹ Many classes of models suggest that the way to understand this is via heterogeneous firm and worker models and associated data. Economists are getting closer to real-time data on these factors. For example, the census has published a real-time business formation index that could be used in such contexts. Haltiwanger concluded that it might be possible to make progress on some problems by paying attention to the push toward heterogeneous agent models and the data that go along with them.

Eswar Prasad observed that one of the crucial issues in the literature was figuring out the right benchmark for evaluating different measures of output gaps. Much of the paper focuses on defining statistical benchmarks, which are very important. However, building on earlier discussions, Prasad argued that economic benchmarks are potentially far more important because of their implications for inflation or variables that the output gap may eventually affect. For instance, taking the Blanchard–Quah approach, which Prasad described as formidable, one could use the same model with inflation instead of the unemployment rate, because the identification restrictions would work very similarly if one made assumptions about how supply and demand shocks affect inflation. The right approach would be to slim down the number of economically meaningful variables and add more variables. Trying to infer what is happening with the output gap by looking at as many indicators as economic models might suggest would have some relationship with the output gap.

Prasad also stated that the univariate filter seemed to work very well. He recalled that back in 1991, when he was at the International Monetary Fund and was responsible for calculating measures of potential output, the HP filter was new and fresh, and so it seemed appropriate. However, even then there were concerns about whether the filter was too sensitive to observations

1. “Changing Market Structures and Implications for Monetary Policy,” Conference, Jackson Hole, Wyo., August 23–25, 2018, <https://www.kansascityfed.org/publications/research/escp/symposiums/escp-2018>.

toward the end of the sample. The problem would be exacerbated by increasing the smoothing parameter, which is very sensitive to what happens at the end of the sample. Prasad recommended considering alternatives, like the King–Baxter band-pass filter, that have slightly better end-of-sample properties, and measure univariate filters against those slightly more robust alternatives, which are not as sensitive to end-of-sample problems.

Steven Braun stated that the one thing that he would have changed about the paper would be not using GDP, but instead using gross domestic output (GDO), which is the average of GDP and gross domestic income (GDI), because GDO has a higher correlation with the unemployment rate than either GDP or GDI does individually. Braun recalled having observed this in a past Brookings paper by Jeremy Nalewaik, and thought that that advice had been neglected.² Including Okun's law in the list of ways to calculate potential GDP given by the authors is the second thing that Braun would have changed. He noted that Okun's law, which is much simpler than the production function, is the first item that he would have listed. Yet even while using the production function, Okun's law comes in through the back door because of the adjustment from the actual labor force to the potential labor force. In addition, the Phillips curve must also be used, because a natural rate of unemployment is required to use Okun's law. Braun observed that the basic problem is that the Phillips curve has stopped working. An estimation of the Phillips curve that is restricted to the past 25 years shows that a zone of two-sigma uncertainty now includes plus or minus infinity. He concluded that it is difficult or impossible to estimate potential output without a natural rate of unemployment.

Robert Hall agreed with James Stock and reiterated that labor is the most important input to the economy. A measure like the output gap, which describes slackness or the lack of slackness, should be mapped into the labor market. Put differently, the economy is at potential when there is full employment. Today, every measure of the labor market screams tight, with no exception. Labor force participation has been considered an exception by some; however, those who have looked carefully at participation, including Hall himself, have concluded that there was a steep decline in participation that has not been erased by the restoration of full

2. Jeremy J. Nalewaik, "The Income- and Expenditure-Side Estimates of U.S. Output Growth," *Brookings Papers on Economic Activity*, Spring 2010, <https://www.brookings.edu/bpea-articles/the-income-and-expenditure-side-estimates-of-u-s-output-growth/>.

employment.³ It would therefore be a mistake to incorporate changes in participation into a measure of labor market tightness. Hall noted that the labor market is as tight as it has ever been since the Current Population Survey was created under its present name in 1947–48, and all the reasonable measures, including those from the employer side, from the Job Openings and Labor Turnover Survey (the average duration of vacancies), are at an all-time high. All suggestions that unemployment is a bad measure of tightness have faded from influence, as has the notable persistence of long-term unemployment. Since 1948, there has been no trend in unemployment; it is a remarkably stable indicator, and this is reflected in the stability of Okun’s law. Hall concluded that the measures of the output gap that track unemployment very closely are right, and any paper that says otherwise should be questioned.

Athanasios Orphanides applauded the paper, as he recalled earlier panels with Arthur Okun, Bob Hall, and George Perry presenting work on exactly the question that the paper tried to answer.⁴ He noted the difficulty in identifying temporary effects as distinct from permanent effects, and he wondered how this translated into estimates of potential output and the corresponding implications for policy. The difficulty of separating temporary and permanent shocks is evidence for the need to identify robust ways of formulating countercyclical policy, monetary policy, and fiscal policy. He thought that the CBO is doing a very good job of this, considering the difficulties. In the case of monetary policy, economists have been making progress in recognizing that output gaps are mismeasured by downplaying the role of output gap measurement and taking more signals from inflation and inflation expectations. Employment gaps are useful; however, their measurement also has the issue of trying to evaluate the natural rate of unemployment. Although monetary policy has drawn these policy conclusions, the next item on the research agenda is finding out what advice can be drawn for fiscal policy. Orphanides wondered how uncertainty about long-term estimates of potential output can be incorporated into fiscal projections, taking into account the sensitivity from one-sided political pressures. Everybody is happy to raise estimates of potential, and using these

3. John G. Fernald, James H. Stock, Robert E. Hall, and Mark W. Watson, “The Disappointing Recovery of Output after 2009,” *Brookings Papers on Economic Activity*, Spring 2017, 1–58, <https://www.brookings.edu/wp-content/uploads/2017/08/fernaldtextsp17bpea.pdf>.

4. George L. Perry, “Labor Force Structure, Potential Output, and Productivity,” *Brookings Papers on Economic Activity*, no. 3, 1971, <https://www.brookings.edu/bpea-articles/labor-force-structure-potential-output-and-productivity/>.

estimates for policy. Conversely, however, everybody is unhappy when the estimates of potential output growth are reduced.

Robert Gordon agreed with James Stock and commented on how there had been two papers over two days—one on monetary policy,⁵ and one on potential output,⁶ both of which did not mention the unemployment rate. He defined potential output as a situation in which inflation is neither accelerating nor decelerating, which is exactly the same as the definition of the natural rate of unemployment. Therefore, by definition, the output gap would be zero when the unemployment gap is zero. Gordon reiterated Valerie Ramey's point, made while giving her comment, that the paper's conclusion that the output gap is currently 6 to 10 percent is implausible.

Regarding the questions of estimating potential output and the CBO's method of doing it implicitly, Gordon thought that James Stock had come close to the answer through his suggestion of aligning the output gap with the unemployment gap. A Kalman filter can be applied that extracts anything that is correlated with the unemployment gap from the cycles in output, using the unemployment gap as information. The result is a series of potential output data that is much more stable in comparison with that generated using the HP filter. This series does not respond to the decline in actual output during the 1981–82 recession and behaves similarly in the years 2007–9. It slows down radically after 2009, not in response to the demand decline but because of the underlying decline in the growth rate of productivity and the decline in labor force participation. Therefore, potential output backed out from the unemployment gap is radically slow growing and suggests a zero output gap in the current economy. Gordon also responded to Steve Braun's comments. He agreed with Braun that using the GDO in studying the response of productivity and output per hour produces very stable results. Regarding Braun's comments on the range of uncertainty of infinity and the disappearance of the Phillips curve, Gordon recommended waiting. He noted that the core personal consumption expenditure inflation had risen from about 1.4 to 2.0 percent in the previous year and that the Federal Reserve had forecasted continued 2.1 percent inflation over the next two years without any upward movement in inflation. One

5. Emmanuel Farhi and François Gourio, "Accounting for Macroeconomic and Finance Trends: Market Power, Intangibles, and Risk Premia," in this issue of *Brookings Papers on Economic Activity*.

6. Olivier Coibion, Yuriy Gorodnichenko, and Mauricio Ulate, "The Cyclical Sensitivity in Estimates of Potential Output," in this issue of *Brookings Papers on Economic Activity*.

would need to wait for two years to know whether the Phillips curve has truly disappeared.

Steven Davis agreed with Robert Hall's comments about the natural rate of unemployment but saw the current unemployment picture as murky. In particular, the frictional rate of unemployment is lower today than it was 10, 20, or 30 years ago—for two reasons. First, the labor force has aged considerably since the 1980s. Older workers have fewer short-term unemployment spells, which leads to a lower frictional rate of unemployment. Second, the frictional rate of unemployment has also fallen because of a trend decline in business volatility and job reallocation rates since the 1980s. Davis then referred to Valerie Ramey's estimate of a 3.5 percent natural rate of unemployment using the Blanchard–Quah methodology. Although Davis does not see 3.5 percent as his point estimate for the natural rate of unemployment, he does not find it outside the realm of plausibility, given the forces driving the decline in frictional unemployment. He concluded by noting that there is a fair degree of uncertainty about the current natural rate of unemployment and, hence, about the implied output gap.

Mark Gertler addressed two issues that he thought are being conflated. The first concerns what the output gap is, and the second is whether the current potential output is in part a response to demand contraction during the Great Recession. Regarding the first issue, Gertler agreed with Robert Hall and Valerie Ramey that the current output gap is low. Regarding the second issue, Gertler said that he was sympathetic to the view of the paper's authors, and he pointed out that the Great Recession looked similar to a financial crisis in an emerging market, with permanent deviations from output trends and permanent declines in productivity growth.

Glenn Rudebusch disagreed with those who considered the “laziness” of government economists to be an important factor responsible for the excess sensitivity of real-time estimates and potential output. He agreed with James Stock that measures of the unemployment gap are invariably quite persistent and smooth. If the output gap were to be set equal to the unemployment gap, then it would also be fairly smooth. However, there is notable noise or transitory variation in quarterly aggregate output. Therefore, transitory variation in potential output is a convenient mechanical offset that results in a smooth output gap. A more transparent accounting of measurement error and noise in measured output would help resolve this problem.

Wendy Edelberg, who was working at the CBO at the time of her remarks, discussed the CBO's experience of projecting potential output. She noted that the CBO has endeavored to not be overly influenced by

recent movements in the weakness of total factor productivity when projecting potential output growth over the forthcoming decade. For its 10-year projections, the CBO has put more weight on the growth of total factor productivity before the last few years than its normal procedure would suggest. Edelberg discussed figure 1 of the paper, which shows the CBO's downward revisions to potential output since 2007. The line extended for potential output in the 2007 projection in an ocular regression looks as if it is a continuing trend from the data before 2007. But this is an illusion. The CBO's projection for 2007 predicted steeper potential output growth than since 2004. In 2007, the CBO was projecting a sizable pickup in hours growth that in retrospect seemed implausible and inconsistent with the demographic data. Therefore, for the first few years of downward revisions, the CBO had been incorporating the fact that its projections of hours growth were too strong, and they had little to do with the weaknesses in output growth that were being witnessed.

Edelberg also noted that the CBO projects potential output by building it up from data on total factor productivity, the labor market, and capital. One of the reasons it was being marked down was because the growth of capital services was weak. However, in relatively recent years, the persistent weakness of total factor productivity growth had been the major reason for downward revisions to potential output, with which the CBO had been grappling. The CBO projects that current potential total factor productivity growth is weak, consistent with recent incidences, but that potential total factor productivity will revert upward in the future, in line with long-term trends. Therefore, the CBO is projecting an improvement in potential output growth, which, Edelberg acknowledged, was based more on long-term trends than on developments in recent data. And the CBO was projecting potential output growth to improve from about 1.7 percent at the end of 2017 to almost 2 percent over the year. Discussing the paper, she noted that it would be hard to reconcile that the output gap is big and negative given all the other indicators in the economy. However, perhaps the real question is whether current estimates of the output gap are a good indicator of the behavior of potential output over the forthcoming 5 or 10 years. Edelberg referred to Glenn Rudebusch's hurricane analogy, wondering how much weight should be applied to temporary factors that hold down potential output for a short period, when considering output growth for longer periods.⁷

7. Glenn D. Rudebusch, Daniel Wilson, and Tim Mahedy, "The Puzzle of Weak First-Quarter GDP Growth," *Economic Letter 2015–16* (Federal Reserve Bank of San Francisco), 2015.

Alan Blinder agreed with James Stock and picked up on Wendy Edelberg's comments, noting that though there are uncertainties in everything, those in capital and labor are relatively easy to handle. Total factor productivity is the real challenge. There are three things to consider about total factor productivity growth. The first is that it is not constant and changes over time. The second is that these changes in total factor productivity growth are completely unpredictable. And the third is that it is very hard to recognize the changes when they happen. It took a long time to catch on to the productivity deceleration in the 1970s and to the productivity acceleration in the 1990s. It is therefore pretty much impossible to forecast productivity. Blinder observed that, for monetary policy purposes, potential GDP growth is forecasted for the next three years. He suggested that the authors consider whether anything beats the forecast that says that total factor productivity growth in the next three to five years will be similar to what it was in the last three to five years.

Kristin Forbes asked the authors if they had looked at past estimates of different agencies in real time to track the most accurate ones in hindsight, assuming that the authors' estimate of potential output is the best one. She wondered which estimates should be used to make a set of potential output estimates, if there is no time to replicate the authors' technology.

Olivier Coibion thanked the organizers; the commenters, Valerie Ramey, and Serena Ng; and the participants for their insightful comments. He stated that the paper does two things. First, it evaluates how existing real-time estimates of potential respond to shocks; and second, it asks if measures of potential can be created that do better along this metric. Coibion observed that almost all the comments focused on the second aspect, so he would do the same. Responding to a common comment about using information from inflation, he stated that though the authors did not cover it in their presentation, they discuss this extensively in the paper. One view of recent inflation dynamics, as suggested by Steve Braun, is that it reflects a broken or very flat Phillips curve, in which case inflation is uninformative about the output gap. In the paper, Coibion and his colleagues consider a second view, which is an expectation-augmented Phillips curve using household inflation expectations. As they have shown in previous work, this provides a stable Phillips curve with no missing disinflation. That Phillips curve can therefore successfully be used to infer an output gap. Because inflation remains well below inflation expectations, this Phillips curve implies an output gap in the same range as the other measures imply. Coibion stated that their results therefore were also consistent with inflation dynamics.

Coibion also agreed with the broader point about the usefulness of combining information from other sources. For example, consumption information can be used, and long-run restrictions can be combined with information about inflation to get more precise estimates of the output gap. Coibion noted that as Serena Ng emphasized, the lack of precision in the estimates is a major concern and, therefore, combining additional information would be useful. He stated that he and the other authors had attempted to understand the implications of off-the-shelf methods relative to the CBO's estimates. He concluded that they were surprised to find that, by and large, all the methods gave a similar answer about the evolution of the output gap relative to the start of the Great Recession.

JEFF FUHRER

Federal Reserve Bank of Boston

ERIC S. ROSENGREN

Federal Reserve Bank of Boston

GIOVANNI P. OLIVEI

Federal Reserve Bank of Boston

GEOFFREY M.B. TOOTELL

Federal Reserve Bank of Boston

Should the Federal Reserve Regularly Evaluate Its Monetary Policy Framework?

ABSTRACT Would a more open and regular evaluation of the monetary policy framework improve policy in the United States? Even when considering a relatively short time frame spanning the 1960s to the present, it is possible to point to many significant changes to the framework. Some of the changes were precipitated by acute economic conditions; others were considered and implemented only gradually, as a response to long-standing problems with the framework. But to date, the process for evaluating and changing frameworks has not always been transparent, and changes have not always been timely. Could a more formal, and open, review process improve how well we adhere to our current framework? Could transitions to a new framework be made more effectively? We conclude that such a review might indeed be beneficial, and outline one possible review process.

From the inception of central banking, policymakers have adjusted their monetary policy frameworks in light of the economics profession's evolving understanding of monetary economics, changes in the structure of the economy, and the obvious failures of previously used regimes. The lineaments of the current framework for the Federal Reserve are outlined in the most recent January "Statement on Longer-Run Goals and Monetary

Conflict of Interest Disclosure: The authors are officers and employees of the Federal Reserve Bank of Boston. Beyond this shared affiliation, the authors did not receive financial support from any firm or person for this paper or from any firm or person with a financial or political interest in this paper, and they are currently not officers, directors, or board members of any organization with an interest in this paper. No outside party had the right to review this paper before circulation. The views expressed in this paper are those of the authors, and do not necessarily reflect those of the Board of Governors of the Federal Reserve System or the Federal Open Market Committee.

Policy Strategy.” In addition to specifying the 2 percent numerical inflation objective and the specific price index that the Federal Open Market Committee (FOMC) will target, the document emphasizes the symmetry of the inflation goal and the role that communication plays in anchoring longer-term inflation expectations. The document also articulates the symmetry of the FOMC’s loss function with respect to deviations of inflation from target and employment from the FOMC’s assessment of its long-run level, noting also that in circumstances when these dual objectives are in conflict, the FOMC “follows a balanced approach in promoting them.” Finally, and particularly important for this paper, the document notes the FOMC’s intent “to reaffirm these principles and to make adjustments as appropriate in its annual organizational meeting each January.” Exploring whether a formal process might help reduce any obstacles to making these adjustments more effective is the aim of this paper.

How often have such adjustments been required? As this paper illustrates, changes have occurred quite frequently. Almost none of the elements in the current framework existed when the Federal Reserve System was founded in 1913, and most of them have been codified only very recently. These changes, though sometimes significant, did not require enabling legislation, but simply the FOMC’s agreement. In short, the history of the United States’ monetary policy framework is one of nearly continuous changes, both minute and momentous. Broadly, over the past 100 years, the monetary framework has progressed from the Gold Standard, to the Bretton Woods monetary system, to the Treasury Accord, to goal-and-instrument independence, to just instrument independence, to the formal adoption of an explicit numerical objective for price stability, to the use of balance sheet policy to augment conventional policy during the Great Recession, to the 2012 adoption of an explicit framework document that evolved to outline a symmetric and equally weighted emphasis on both aspects of the Fed’s congressionally given dual mandate.

Thus, in reality, the question is not whether the framework can or should change, but what are the appropriate triggers for such changes and what process might best aid the central bank in considering how to change it. As a point of comparison, on February 26, 1991, the Bank of Canada announced an inflation-control target framework, which was extended in 1993, along with a pledge to review it again in 1998. In 2001, the Bank of Canada established a five-year cycle to ensure that its monetary policy framework would remain effective as the economy, and the central bank’s understanding of it, evolve. Every five years, the Bank of Canada conducts a formal review of the goals of monetary policy as well as alternative

approaches to attaining these goals, considerations such as lowering or raising the 2 percent inflation target, whether to target the price level, and the role of financial stability. The process of reevaluation includes staff research on key topics in the years leading up to the decision date, and invites feedback from the public, the government, and academics (see Murray 2018).

Would a more comprehensive and regular evaluation of the framework improve monetary policy in the United States? To answer this question, one needs to understand why framework changes have occurred; how and how quickly the shortcomings of earlier frameworks were recognized, and thus whether one can reasonably expect to improve the way in which framework changes occur; and whether a regular review process could be part of this improvement. It is important to recognize that framework reassessment is not entirely episodic and event-driven. Staff and principals within the Federal Reserve System are involved in a continuous reassessment of the framework. And there is also considerable interaction between the Fed, academics, other central banks, and other policy institutions, and this interaction provides some opportunity to consider emerging ideas about how to improve the conduct of policy.

However, U.S. monetary history—certainly including the Great Depression, and possibly the Great Inflation of the 1970s, the Great Recession, and the 2008–9 financial crisis of very recent history—might suggest that the existing combination of internal processes and external interactions does not always produce the optimal framework. Would more focused internal analysis at regular intervals be helpful? Would a more formal incorporation of external analysis from academics and others improve the Fed’s performance at key junctures? This paper tentatively concludes that such a process may help the Fed more effectively make needed framework changes. Whether the source of any problems is Fed errors or the profession’s understanding, a regular reevaluation process, both external and internal, may help to more efficiently change the framework when needed.

We wish to emphasize that the framework changes we have in mind are not ones that would require amendments to the existing legislation, which would obviously fall under the purview of Congress. Rather, they are largely technical changes meant to improve the conduct of monetary policy to better achieve the congressionally mandated goals.

The paper begins by defining, in section I, what is meant by a monetary policy framework. To anticipate, the definition necessarily entails “gray areas.” In section II, we consider the history of changes in the monetary policy framework for the U.S. central bank in the modern era.

From this review, we hope to get a sense of the frequency of changes, the motivation for changes, and a sense of which measures we might use to gauge the success or failure of the historical frameworks. Section III discusses a host of practical questions about the process for reevaluating the Fed's monetary policy framework. Section IV discusses whether the Fed might consider a change in the framework given the current circumstances, and section V concludes.

I. How Do We Define a Monetary Policy Framework?

Broadly, a monetary policy framework may be defined as the set of tools and processes by which the central bank attempts to define and attain its high-level economic goals. The central bank might be allowed to choose some components of this process, such as the precise inflation target and the transparency of the policy. But some elements of the framework are strongly influenced by other factors outside the central bank's control, such as the structure of the economy and the desires of the public. Given this definition of the framework, it follows that the changes in the framework in which we are most interested are those that significantly affect the central bank's ability to achieve its high-level goals on behalf of the public. More specifically, a monetary policy framework will include eight main elements.

The first element is the governance structure of the central bank. This paper largely abstracts from how the central bank fits into the country's governmental structure—for example, whether it is, statutorily, an instrument independent of the executive branch and the Treasury. Such considerations have been shown to importantly affect the efficacy of central bank actions, but we assume in this paper that the Fed, both legislatively and practically, has a high degree of independence.¹ A related high-level concept is that of accountability: the responsibility delegated to the central bank by Congress to deliver acceptable economic outcomes to the country's citizens. Many efforts to improve transparency have been rooted in a desire to provide the public with explanations for why the Fed does what it does, an essential component of accountability.

The second element is a set of ultimate goals for the central bank. Today, we have a congressionally mandated set of goals—the so-called dual mandate, which comprises “stable prices” and “maximum employment”—phrases that have been modified in common usage to “price stability” or

1. For issues pertaining to the Federal Reserve's governance, see, among others, Binder and Spindel (2016) and Conti-Brown (2016).

“low and stable inflation” and “maximum *sustainable* employment.”² The goals have changed through time. When the Fed was founded in 1913, its goals focused primarily on stability of the banking system. The gold standard demanded fixing the dollar price of gold. Obviously, both sets of goals differed dramatically from the Fed’s responsibilities in the current framework. In the long run, even the dual mandate might be altered. For example, recurring bouts of financial instability might prompt the Fed to be more explicit about the role of monetary policy in preventing and offsetting such disruptions.³

The third element is a loss function. An articulation of goals is not sufficient. Unless the framework entails a single, rigidly defined goal, it needs to include a loss function (or the equivalent) that describes how the central bank weights its (sometimes competing) goals.⁴ For example, the Fed needs to say whether it considers losses on either side of the target inflation or employment goals symmetrically; how it weights deviations from its inflation goal versus deviations from its employment goal; whether it chooses a point target for inflation or a band; whether it allows the operational inflation goal to move somewhat over time; over what horizon it intends to bring inflation back to its goal; and so on. As an institution accountable to the public, the Fed might also explain where the targets come from, and what determines their values. Such an explication would also help illuminate why the framework might change through time.

The fourth element is instruments. The instruments that the central bank has directly controlled in attempting to achieve its key goals have, over time, included the dollar price of gold, the volume of various monetary and reserve aggregates, the level of short-term interest rates, and the size and composition of its balance sheet. Multiple instruments can and have been part of the same framework; for example, during the financial crisis, the Fed used the federal funds rate until it fell to its effective lower bound, at which point it pursued balance sheet policies in an attempt to better achieve its mandated goals.

The fifth element is the central bank’s operational targets. The operational “target” (or targets) that the central bank sets so that it can achieve its

2. Federal Reserve Act, as amended, Section 2A–12 USC 225a, as added by act of November 16, 1977 (91 Stat. 1387), and as amended by acts of October 27, 1978 (92 Stat. 1897), August 23, 1988 (102 Stat. 1375), and December 27, 2000 (114 Stat. 3028).

3. See Peek, Rosengren, and Tootell (2015).

4. In some cases, a model’s explicit microeconomic foundations allow one to derive a model-consistent loss function (see Rotemberg and Woodford 1997; Woodford 2002), and could inform the central bank’s choice of a specific loss function.

primary goals sometimes completely overlap its ultimate goals. A central bank that is a pure inflation targeter can use inflation as both its ultimate and operational targets. However, a central bank with both inflation and output as its ultimate goals could choose to use, for example, nominal GDP as the operational target to achieve its goals. Targeting nominal GDP imposes specific weights on deviations of prices from the desired price-level path and deviations of real output from potential, the ultimate goals of monetary policy.⁵

The sixth element is transparency. Transparency is often an important part of the framework, especially when it is enhanced to improve the efficacy of policy actions. It is also an important element of the framework because it improves a central bank's accountability to the public. One goal of transparency is to make monetary policy more predictable. For example, the release of FOMC or staff forecasts, which signal future policy, and publishing alternative scenarios could (at least in theory) help the public to understand the current and expected setting of policy, which might lead to a more predictable and efficient transmission of policy actions into other asset prices. In fact, one motivation for revisiting the monetary policy regime every few years is to ensure that the regime is clear to the public. The costs of opacity can be high. For example, the profession has struggled with understanding the poor performance of the economy in the 1970s and early 1980s (and this paper is no exception), in part because of this lack of clarity about the framework. In the extreme, the Fed may want to set expectations with clear forward guidance. The ability of the central bank to affect expectations is a topic of active discussion, and much has been written about the wisdom of attempting it, as well as the efficacy of historical attempts (see, for example, King, Lu, and Pastén 2008).

The seventh element includes rules and discretion—the systematic component of monetary policy. Given a set of goals, an articulation of the loss function, a set of instruments, and perhaps an intermediate target, a central bank should generally aim to conduct monetary policy in a systematic, and thus predictable, fashion. As a consequence, even if transparency is minimal, one may be able to discern with some accuracy the policy rule implications of a framework. That rule will not capture all features of the framework—in particular, a simple rule would fail to capture asymmetries in the uncertainty about the outlook—but it can reflect, in a compact way,

5. The desire to use nominal GDP as an operational target might arise from equal weights on price and output deviations in the loss function, or from other practical considerations that suggest it would deliver desirable outcomes relative to other operational targets.

many aspects of the framework. To the extent that a central bank's behavior can be well described by a policy rule, whether this rule is articulated by the bank or can be accurately inferred by the public from the bank's actions, policy predictability will be enhanced, and the transmission channel will be more effective. At the same time, under certain conditions, truly optimal policy may deviate noticeably from simple rules, and thus discretion may be an important component of the policy framework.

One important element of discretion is risk management. Most discussion to this point has abstracted from how the evaluation of and response to risk might fit into the monetary policy framework. This is not a trivial omission; indeed, then-Fed chairman Alan Greenspan (2004, 2005) often described the business of monetary policy as in large part an exercise in risk management. The evaluation of risk—or, more specifically, the consideration of asymmetry in the distribution of policy-relevant outcomes, along with the possibility of abnormally large tail risks—has clearly played a role in FOMC deliberations over the years. Most notably, financial stability risks have risen in prominence in the FOMC's discussions. Providing a precise analytical framework for the Fed's, or any other central bank's, systematic response to such risks is beyond the scope of this paper. But in attempting a definition of the monetary policy framework, the response to and management of risk is a nontrivial element.

Finally, the eighth element is the central bank's depiction of the economy—"the model." Broadly speaking, the model that the central bank uses to describe the economy's evolution and the interactions between policy and the real and financial economies can both constrain and influence the regime chosen by the central bank.⁶ In a committee such as the FOMC, different members can base their policy recommendations on different models while still sharing the same elements of the framework we have already outlined. Nevertheless, common features across models are crucial inputs to the policy process—the equilibrium real rate of interest, the natural rate of unemployment, and the slope of the Phillips curve. Post-war U.S. history appears to have experienced quite persistent and significant fluctuations in most if not all of these key parameters, as illustrated in section II below. Such changes in economic structure can also spur modifications to the monetary policy framework, although not all changes will require a shift in the framework. For example, when changes in economic

6. For a given economic structure or model, one can entertain any number of monetary policy frameworks that might work within it. In this sense, the model is not part of the framework, although it can clearly influence the choice and efficacy of frameworks.

structure *constrain* the framework—such as when a drop in the equilibrium rate makes it more likely that the effective lower bound will bind—then a framework change may be needed.

In addition, the current instantiation of a central bank’s economic model reflects current economic wisdom as accepted (and perhaps modified) by the central bank. One can take for granted, in present circumstances, the importance of explicit expectations; of macroeconomic behavior that is grounded to some extent in microeconomic behaviors; of the importance of accounting identities, budget constraints, and adding up constraints; and of the absence of a long-run trade-off between unemployment and inflation, given that most modern models reflect such concepts to varying degrees. But these have not always been features of the models used by central banks in the conduct of monetary policy, and several of them have changed the way banks think about conducting monetary policy, and thus about what are viewed as better and worse frameworks.

There is no widely agreed-upon definition of what constitutes a monetary policy framework, but the eight key elements described here should be useful as guideposts as we consider both the history of the U.S. monetary policy framework and its possible evolution. Again, it is important to recognize that though these elements appear as distinct components in the description above, in practice there will be both gray areas in the definitions and overlap among the components as they are used in any specific framework.

II. A Review of Monetary Policy Frameworks since the 1960s

We now provide our assessment of framework changes over time, using the previous section’s taxonomy. Many of the changes we identify are discussed in Allan Meltzer’s history of the Federal Reserve (2002, 2010, 2014) and by Christina Romer and David Romer (2002, 2013). These works are based on a thorough reading and interpretation of the minutes of the FOMC’s meetings. Here, we complement some empirical evidence with a word count of specific phrases used at FOMC meetings that may indicate a change in focus on key elements of the policy regimes at the time. When a given framework is operative, one would expect certain words related to this framework to arise more frequently. Our analysis is also organized around specific elements of the framework. As such, it is not necessarily exhaustive, but it is meant to highlight the fundamental issues and provide explanations for the reasons and processes that led to or hindered changes in the monetary policy framework.

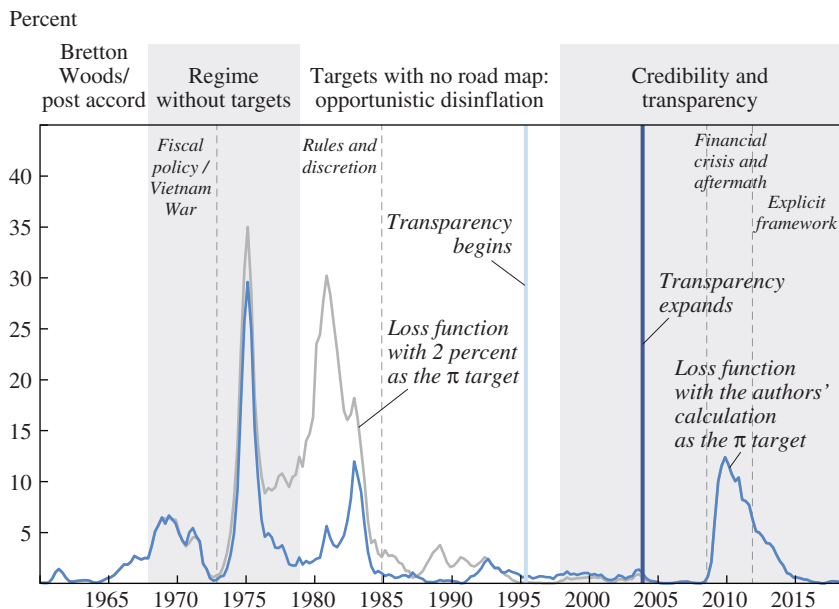
As Fed insiders, we also wish to emphasize that the framework is to some extent always under discussion and debate. The staffs of the Federal Reserve Board and the regional Fed banks are constantly working on memos and papers that examine possible changes to the framework. Another element of the work conducted at the Board and the regional banks concerns potential technical changes in key parameters of the economy. Such work is ongoing in the system and might need less coverage at a regularly scheduled public meeting, such as the one under consideration here. The deeper issues that correspond to the framework debates may be more appropriate for such gatherings.

It is useful to frame the discussion first in terms of realized outcomes and the policy frameworks in which they occurred. Figure 1 presents our version of the frameworks that have existed since the 1960s. The regimes are drawn in figure 1 with very broad brushstrokes. Still, we attempt to show some of the finer strands of the tapestry of monetary policy that run throughout the past 60 years, such as independence and transparency. It is important to note that, for the most part, lessons were not forgotten over this time period, so in many ways the regime changes are really an accumulation of knowledge. The regime names attempt to emphasize the added pieces in the puzzle acquired over a given period.

The realized outcomes are also presented in figure 1 by means of a quadratic loss function that weights inflation and unemployment equally. These losses could capture the costs of using the wrong framework, along with adverse shocks not related to monetary policy. Inflation and unemployment are taken as deviations from an estimate of the inflation target (when the target was not explicit), and the estimate by the Congressional Budget Office (CBO) of the natural rate of unemployment, respectively.⁷ The largest losses appear in the second half of the 1960s, the 1970s and early 1980s, and with the recent Great Recession.⁸ The “Volcker disinflation” occurred after about 15 years of large welfare losses, and required a very costly recession to alter the course of inflation and inflation expectations. Although not all the large economic losses represented in the figure were the direct consequence of FOMC policies, it is relevant to ask whether a

7. We measure inflation with the latest vintage of the Q4/Q4 change in the core personal consumption expenditures deflator. Details about the estimation of the time-varying inflation target are provided in section II of the paper.

8. A time-varying target for inflation reduces the loss during the 1970s and early 1980s, but the qualitative results in the figure continue to hold even with an inflation target fixed at 2 percent.

Figure 1. Loss Functions and Regimes, 1960:Q1–2018:Q1

Sources: Authors' calculations; Bureau of Economic Analysis; Bureau of Labor Statistics; Congressional Budget Office; Federal Reserve Board.

more systematic evaluation of the framework might reduce such losses, whether they resulted from delaying actions, adherence to a broken framework, misperception of key aspects of economic structure, or discretionary deviations from an otherwise well-functioning framework.

II.A. Regimes without and with Explicit Targets

With the demise of the Bretton Woods monetary system, and the demands of financing the Vietnam War, the Fed's mandate became less clear. To examine this issue, this subsection explores the Fed's inflation model. In so doing, we also comment on recent developments that have a bearing on the policy framework. The Fed grapples constantly with its model of inflation. Here, we infer the evolution of the FOMC's views about inflation from the inflation predictions made by the staff of the Federal Reserve Board and published in the Greenbook/Tealbook (GB/TB). This analysis is related to and extends the work of Romer and Romer (2002). The GB/TB inflation

forecast for a particular quarter is modeled as a function of lagged inflation and the unemployment rate:


$$(1) \quad E_t \pi_{t+i} = \beta_{0,t} + \beta_{1,t} E_t \pi_{t+i-1}^4 + \beta_{2,t} E_t u_{t+i-1} + v_{i,t}, \quad i = 1, 2, 3$$

where π_{t+i} denotes the annualized rate of inflation in quarter $t + i$, π_{t+i-1}^4 is the average rate of inflation prevailing over the four quarters from $t + i - 1$ to $t + i - 4$, and u_{t+i-1} is the level of the unemployment rate at $t + i - 1$. The operator E_t denotes a forecast made in quarter t . We consider forecasts of inflation one, two, and three quarters out, as indexed by i . At each of the three forecasts of quarterly inflation horizons, the relationship is augmented by an error term, $v_{i,t}$, which captures other factors that influence the inflation forecast besides past inflation and the unemployment rate. An important feature of equation 1 is time variation in the β coefficients, which is assumed to occur as a random walk.


Details about the data and estimation are provided in the appendix to this paper. Figure 2 reports the unsmoothed, time-varying estimates of the coefficients over the period 1966:Q4 to 2017:Q4.⁹ It is apparent that the weight given to lagged inflation, as measured by $\beta_{1,t}$, was low in the late 1960s and started to rise noticeably in the early 1970s. The first few estimates in the sample need to be interpreted with caution, because the available forecasts in the GB/TB often did not extend out four quarters. Initial conditions also matter, but it can be shown that the qualitative result of an increase in the importance of lagged inflation in the 1970s relative to the late 1960s is robust. This strand of the framework has recently regained importance, as there has been a noticeable decline in the weight placed on lagged inflation. As concerns the assessment of the short-run trade-off between inflation and unemployment, $\beta_{2,t}$ in equation 1, the estimates are again noisy at the beginning of the sample, but views about the trade-off appear to have changed in the 1970s and 1980s. More recently, there has been a gradual but steady decline in the estimated impact of economic slack on inflation. The intercept term, $\beta_{0,t}$, also exhibits noticeable variation, and we comment on these fluctuations below.

In all, though admittedly simple, this exercise points to changes in the inflation model. Some of these changes have had a significant impact on

9. Given that the staff's forecasts are made public with a five-year lag, for the period from 2013 to the present, our analysis uses the FOMC's economic projections. The appendix provides details on how the forecasts from the "Summary of Economic Projections" are used in the analysis.

Figure 2. Unsmoothed Estimated Coefficients for Equation 1, 1966:Q4–2017:Q4

Sources: Authors' calculations; Federal Reserve Board.

Figure 3. Mentions of Inflation as a Problem, 1966–79^a

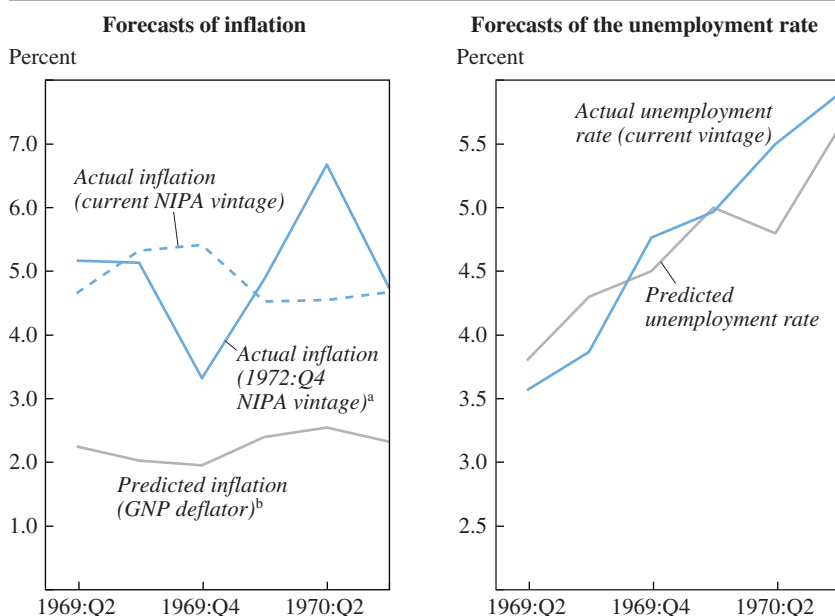
Source: Federal Reserve Board.

a. Four-meeting moving-average term counts as a percentage of total words in FOMC transcripts, memorandums of discussions, and historical minutes.

the policy framework. By the end of the 1960s, the need to design monetary policy to account for the requirements of fiscal policy and the executive branch—particularly given the increases in spending on the Vietnam War—had produced disappointing inflation outcomes.¹⁰ The FOMC's nervousness about the inflation situation at that time can be inferred from an increase in the mentions of inflation, as shown in figure 3.¹¹ Policy tightening in 1969 was seen as an opportunity to reduce inflation. However, the realized decline was noticeably less than expected. Figure 4 shows that the persistent miss in the inflation forecast at the time cannot

10. At the January 1969 FOMC meeting, Chairman William McChesney Martin mentioned that “to Mr. Nixon, he had expressed his view that inflation was the primary economic problem now facing the nation, and that the new Administration would have to deal with it effectively from the beginning if inflation were not to get out of control. He had done his best to emphasize the seriousness of the problem” (Minutes, January 14, 1969).

11. It is important to note that the terms used to discuss similar topics have changed throughout U.S. monetary history. Thus, the use of “inflationary psychology” was fairly common in the 1960s but is less common today. More commonly used terms included “inflation expectations” and the “anchoring” of expectations. For this reason, we must take care in interpreting the frequency with which specific phrases are used.

Figure 4. Federal Reserve Greenbook Forecasts, 1969:Q2–1970:Q3

Sources: Bureau of Economic Analysis; Bureau of Labor Statistics; Federal Reserve Board; Haver Analytics.

a. NIPA = National Income and Product Accounts.

b. GNP = gross national product.

be attributed to a persistent downward bias in the unemployment rate forecast.¹²

The inflation underprediction appears to have led to a reconsideration of the Fed's inflation model. The increase in the weight given to past inflation in equation 1 in the early 1970s signals a move toward an accelerationist view of inflation.¹³ The estimates of $\beta_{1,t}$ do not reach unity in equation 1 because, with time-varying coefficients, some of the persistence in the inflation process is shifted from lagged inflation to the time-varying intercept. But

12. The horizontal axis in the graphs denotes the quarter in which the Greenbook forecast was made. The forecast is given by the value of inflation or the unemployment rate expected to prevail on average three and four quarters into the forecast. We use only the third quarter of the forecast whenever the fourth quarter is not available. The exercise stopped with the 1970:Q3 forecast because the Nixon wage and price controls were enacted in 1971:Q3.

13. Sargent (2001) attributes the run-up in inflation in the late 1960s and 1970s to slow learning about the true process of inflation.


a fixed-coefficients estimation of equation 1 over a period spanning from the 1970s to the early 2000s would yield a coefficient on lagged inflation very close to unity, which is consistent with the Friedman-Phelps natural rate framework. This change, coupled with a decline in the short-run trade-off between inflation and unemployment, entailed a significant increase in the perceived sacrifice ratio during the mid-1970s. Note that an increase in $\beta_{1,t}$ from 0.2 to 0.8 and a decline in the absolute value of $\beta_{2,t}$ from 0.4 to 0.25—which is roughly the magnitude of the movements that occurred from the late 1960s to the mid-1970s—imply that bringing inflation down from 5 to 4 percent over the course of eight quarters would require, other things being equal, an unemployment rate gap of about 1 percent on average, up from 0.2 percent.

This pessimism about the cost/benefit trade-off of using monetary policy to lower inflation has been documented before (for example, by Romer and Romer 2013). As a consequence, beginning in the early 1970s, price and wage controls were advocated as an alternative means for controlling inflation, and figure 5 highlights how FOMC members were discussing such fiscal solutions to the inflation problem. The reluctance to engineer large employment losses as a way of reducing inflation had notable implications for another aspect of the framework, the inflation target. In the context of equation 1, it is possible to infer the FOMC operational inflation target from the time-varying intercept, which can be written as

$$(2) \quad \beta_{0,t} = (1 - \beta_{1,t})\pi_t^* - \beta_{2,t}\bar{u}_t,$$


where π_t^* and \bar{u}_t are time-varying measures of longer-run inflation and the natural rate of unemployment, respectively. Together, equations 1 and 2 provide a representation of the Phillips curve, which is now often used to parsimoniously describe inflation. Although this is a “modern” view of the inflation process, a looser interpretation in terms of a reduced form where inflation has a tendency to revert over the forecast horizon to the π_t^* objective—after controlling for an activity gap and supply shocks—is still valid and likely to have informed the Federal Reserve’s inflation forecast consistently over time.

Figure 6 depicts a derivation of π_t^* according to equation 2, given our estimated time-varying β s under the assumption that the natural rate of unemployment, \bar{u}_t , evolves as in the most recent vintage of the CBO’s estimate, over the period 1969–2007. The current vintage of the CBO’s natural rate of unemployment differs from real-time estimates, and such a

Figure 5. Mentions of Fiscal Solutions to the Inflation Problem, 1960–2012^a

Source: Federal Reserve Board.

a. Four-meeting moving-average term counts as a percentage of total words in FOMC transcripts, memorandums of discussions, and historical minutes.

Figure 6. Estimated Longer-Run Inflation, 1968:Q4–2007:Q4

Sources: Authors' calculations; Congressional Budget Office; Federal Reserve Board.


difference will introduce biases in the estimate of π_t^* , a point to which we return below. Because the estimated time-varying parameters are noisy, figure 6 depicts the minimum and median values of π_t^* over a centered moving window of nine quarters. We report the minimum value to provide a conservative assessment of the time variation in π_t^* .

The main takeaway from this exercise is that the attainable rate of inflation in the medium term was subject to profound reevaluations in the late 1960s and the 1970s. Furthermore, the tolerable level of inflation was also subject to reevaluations that continued until the late 1990s, when π_t^* finally settled at about 2 percent. The figure depicts a steady increase in the implicit inflation goal over the course of the 1970s. Because the CBO's current view of the natural rate of unemployment in the 1970s is likely higher than most real-time assessments, our estimate of the rise in the inflation goal over this period is conservative.¹⁴ It is possible that policymakers' long-run aspirations were always for low inflation, but in practice their perception of the attainable rate of inflation in the medium run was subject to frequent reevaluations in the face of adverse supply shocks.¹⁵ Without a clear mandate, the costs of returning to a lower target were considered too high to be paid directly, a topic that is further addressed when we discuss "opportunistic disinflation." Needless to say, the lack of explicit targets was a significantly important missing piece to the monetary policy framework in the 1970s. But it is also important to note that the target was subject to, admittedly milder, revisions in the 1980s and most of the 1990s. By then, the Federal Reserve had regained credibility in its stance toward inflation; but as we discuss below, it was not yet transparent about its inflation goal.

From an inflation model perspective, the most recent period also stands out. The role for past inflation and economic slack in determining inflation has diminished, and more emphasis is placed on long-run inflation expectations. Figure 7 shows that discussions about "well-anchored expectations" increasingly appear in the transcripts starting in 2004. With a stable inflation goal at 2 percent, the focus was to maintain inflation near

14. In deriving our estimate of the inflation goal from equation 2, we have purposely chosen the CBO's estimate because it averages about 6.0 percent in the 1970s, with relatively little variation over the decade. This estimate is likely on the high side of the range of real-time estimates of the natural rate of unemployment, and therefore makes the reported π_t^* in the 1970s a conservative estimate.

15. Ireland (2007) reaches similar conclusions about time variation in the inflation objective using a different approach based on estimating a small-scale, dynamic stochastic general equilibrium model on actual data.

Figure 7. Mentions of Well-Anchored Inflation Expectations, 1960–2012^a

Source: Federal Reserve Board.

a. Four-meeting moving-average term counts as a percentage of total words in FOMC transcripts, memorandums of discussions, and historical minutes.

this target rather than to achieve lower inflation. The notion here is that insofar as *long-run* inflation expectations are “well-anchored,” inflation will deviate only modestly from the inflation goal in proportion to the deviation of the unemployment rate from its equilibrium level. In this context, the role of the central bank is to ensure that long-run expectations are centered on the inflation goal, and to stabilize the economy at full employment, at which point inflation will equal its target.

As is discussed later in this paper, the anchoring power of long-run expectations and the small effect of the unemployment rate gap on realized inflation have important repercussions for the monetary policy framework. The inflation costs of deviating from full employment are small in this setup. As a result, the cost/benefit analysis of probing for better labor market outcomes (in the form of a lower equilibrium unemployment rate) may be more favorable now. Issues surrounding the shape of the loss function are also coming into better focus. With small inflation costs, what are the welfare costs of overshooting full employment? Are the losses symmetric to undershooting full employment, as the current statement on monetary policy strategy implies?

II.B. Changes in Policy Rules and Opportunistic Disinflation

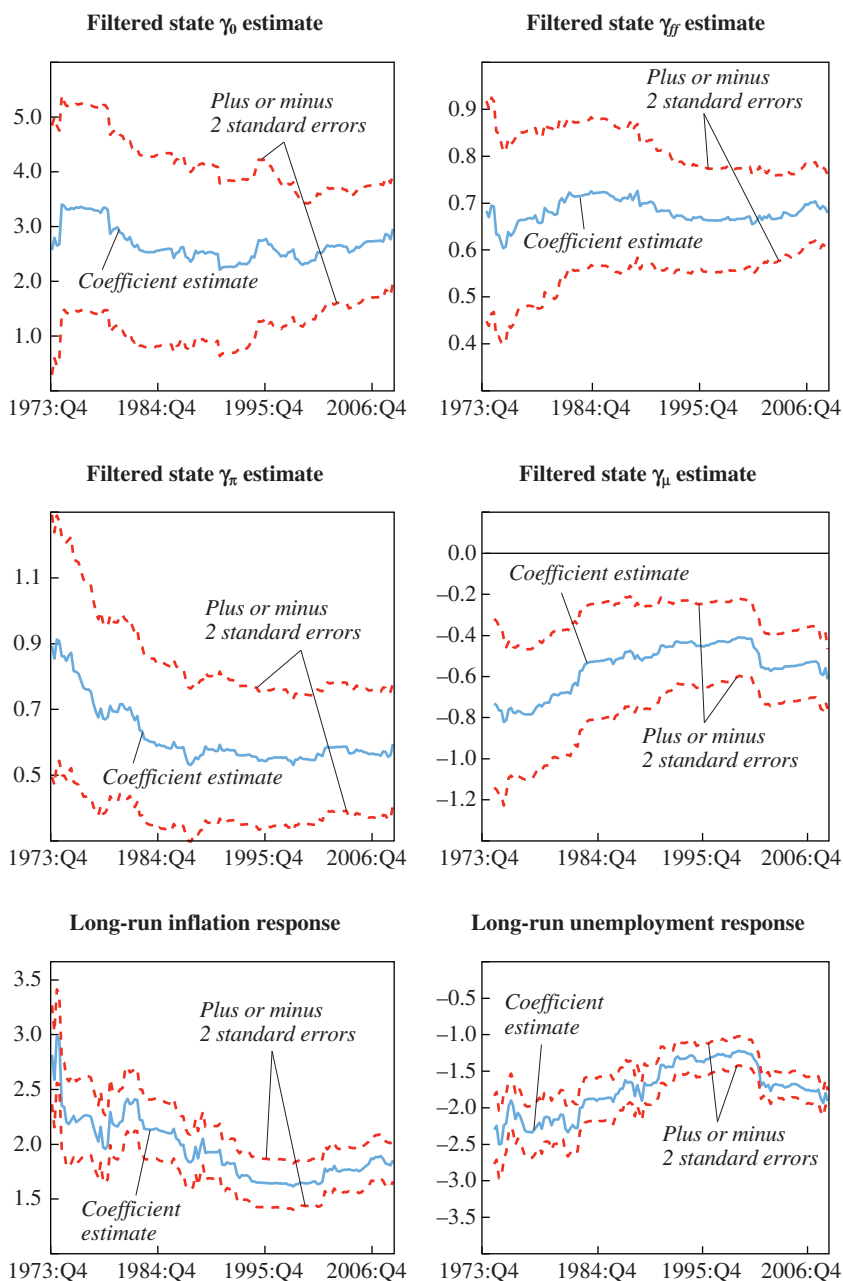
As the target-less regime began to crumble by the end of the 1970s, the passage of the Humphrey-Hawkins Act of 1978 provided the Federal Reserve with a mandate to pursue targets—but not a road map for how to get there. Along with calling for semiannual reports to Congress, the Humphrey-Hawkins Act amended the monetary policy objectives contained in the Federal Reserve Act, and thus directed the FOMC to “maintain long run growth of the monetary and credit aggregates commensurate with the economy’s long run potential to increase production so as to promote effectively the goals of maximum employment, stable prices, and moderate long-term interest rates.” The 1979 monetary aggregates experiment could be looked at in the context of just such a road map or rule. More generally, the topic of how a policy rule for the FOMC has evolved over time has been widely debated in the literature. To capture time variation in the policy rule, we consider this reaction function:

$$(3) \quad \hat{ff}_t = \gamma_{0,t} + \gamma_{\hat{ff},t} \hat{ff}_{t-1} + \gamma_{\pi,t} E_t \pi_{t+3} + \gamma_{u,t} E_t u_{t+3} + v_t,$$

where \hat{ff} is the federal funds rate and the other variables are defined as above, with v denoting the error term.¹⁶ The rule is forecast-based, with the forecasts being given again by the Fed staff’s projections as published in the GB/TB. In addition to the forecasts of inflation and unemployment, the rule allows for smoothing interest rates. As before, the coefficients in the rule are time-varying, with their evolution assumed to follow a random walk.

Equation 3 is estimated over the sample from 1969:Q1 to 2008:Q4; more details about the data and estimation are provided in the appendix. Figure 8 reports the unsmoothed filtered estimates starting in 1973:Q4. We omit the earlier period because the estimates might be affected by the choice of initial conditions, for which we do not hold strong priors. The figure also reports the long-run responses to inflation and unemployment, computed by dividing the contemporaneous responses by $1 - \gamma_{\hat{ff},t}$. In all,

16. The specification is similar to that given by Boivin (2006). Our exercise, however, is conducted at a quarterly rather than at a Greenbook frequency. Another important difference is that we let the time-varying intercept capture not just potential changes in the equilibrium federal funds rate but also changes in policymakers’ assessment of the natural rate of unemployment.

Figure 8. Unsmoothed Estimated Coefficients for Equation 3, 1973:Q4–2008:Q4

Sources: Authors' calculations; Federal Reserve Board.

there has been some variation over time in the degree of interest rate smoothing and in how the FOMC has reacted to inflation and unemployment rate forecasts.

It is interesting to note that the weight given to unemployment relative to inflation (in absolute terms) was, overall, at its largest in the 1970s, and then declined in the 1980s and the 1990s. The relative emphasis on unemployment deviations in the 1970s is consistent with the previous discussion about the FOMC being unwilling to generate large employment losses in order to reduce inflation. Such a focus on employment stabilization also raises the much-discussed issue of the Fed's independence. Still, once taking into account changes in the operational inflation target (which, in the context of the policy rule given in equation 3, are subsumed in $\gamma_{0,t}$), the tenet that in the 1970s the FOMC was violating the "Taylor principle," whereby policy rates move more than one-for-one with inflation, is far from settled.¹⁷ In our exercise, the long-run response of the federal funds rate to inflation projections is always estimated to be above unity.¹⁸

The estimated parameters in equation 3 signal a greater emphasis placed over the course of the 1980s and 1990s on deviations of inflation from target relative to deviations from full employment. After the sharp decline in inflation achieved by 1984, it is notable how the strategy over most of the years 1984–86 was one with a strong resemblance to a gradualist approach to driving inflation lower. During those years, the unemployment rate was stable, but at levels near 7 percent, above the natural rate. A variant of this strategy was later undertaken under Alan Greenspan in the late 1980s and early 1990s, with the FOMC pursuing "opportunistic disinflation" to reduce inflation below its average of 3.5 to 4 percent. The strategy accomplished a reduction in inflation by allowing some slack to remain in the economy following the 1990–91 economic downturn, avoiding the arguably larger costs of initiating another recession. It is possible to motivate such a strategy by assuming an unconventional loss function in employment and inflation (see Orphanides and Wilcox 2002), a notion that hints at the flexible

17. For contrasting views about the FOMC rule's consistency with the Taylor principle in the 1970s, see, for example, Clarida, Galí, and Gertler (2000); and Orphanides (2003). Boivin (2006) reaches different conclusions from ours, likely as a result of the differences in the specifications that we have already discussed.

18. In this regard, it is relevant to note that the effective federal funds rate increased substantially in 1973 before the oil price shock. And monetary policy had tightened already in 1978 and 1979, before Paul Volcker became chairman.

interpretation of the loss function underlying the FOMC's policy framework during this period.¹⁹ The literature in the late 1980s and early 1990s was already examining the benefits of transparency. This is one instance when a regular conference on the framework during this period might have raised the issue of opportunism more clearly and fostered more discussion about transparency.

In the more recent period covered in the exercise, which spans from the 2000s to the onset of the Great Recession, the long-run response to unemployment has increased again in absolute value, with inflation and unemployment deviations carrying about the same weight in the reaction function. It is possible that the anchoring of inflation expectations and a "flat Phillips curve" have played a role in such a development. Here, we note that an optimal policy exercise with a credible inflation target would be consistent with a larger weight given to activity stabilization in a policy reaction function such as equation 3, when the slope of the Phillips curve becomes flatter (see, for example, Iakova 2007; Erceg and others 2018).


The time-varying nature of the reaction function given in equation 3 makes it complicated to talk about rules versus discretion, in that a changing unemployment response relative to inflation, or changes to the interest rate smoothing coefficient, could be interpreted as an exercise in discretion. Nevertheless, even with this flexible setup, it is possible to identify other important changes in the conduct of policy. In particular, figure 9 depicts the estimated error term v in the policy function. The dotted part of the line encompasses a period of high volatility in the early part of Paul Volcker's tenure associated with the operating procedure for nonborrowed reserves. Overall, it is apparent that the predictability of the rule has increased noticeably since the mid-1980s.

Another notable feature of the current policy environment is that changes in the policy rule are key to explaining the conduct of monetary policy after the liftoff from the zero lower bound. Figure 10 plots the predicted federal funds rate using the coefficients in equation 3 as estimated in 2008:Q4, vis-à-vis the actual, from 2015:Q4 to the present.²⁰ The simulation is static, in that it uses the actual lagged federal funds rate. Despite such a feature, it

19. This reverse engineering exercise posited a loss function in the absolute value of unemployment and the squared deviation of inflation from a short-run inflation target. This loss function induces a region of inactivity for sufficiently small inflation deviations. In these circumstances, the central bank optimally waits for a shock that moves inflation toward the long-run goal, pocketing gains along the way without deliberately altering the output gap.


20. Because the forecasts from the board staff are not yet publicly available, we use instead the FOMC's "Summary of Economic Projections."

Figure 9. Estimated Error Term in the Policy Rule Given in Equation 3, 1973:Q4–2008:Q4^a



Sources: Authors' calculations; Federal Reserve Board.
a. Noise in the policy rule, as captured by v .

Figure 10. Actual and Predicted Federal Funds Rate Targets, 2015:Q4–2018:Q2



Sources: Authors' calculations; Federal Reserve Board.
a. Predicted from 2008:Q4 with a modified constant to account for long-run values estimated by the "Summary of Economic Projections."

is apparent that the reaction function given in equation 3 with the 2008:Q4 estimated coefficients is a poor predictor of the FOMC's behavior in the most recent period. Modifying the intercept in equation 3 to account for changes in the assessment of the equilibrium value of the federal funds rate and the unemployment rate—as reported in the “Summary of Economic Projections” (SEP) over the simulation period—reduces but does not eliminate the difference. In other words, a decline in the estimate of the equilibrium federal funds rate has played an important role, but other factors have also been at play. These factors could be related to risk management considerations, and/or to shifts in the weights assigned in the rule to unemployment and inflation deviations. The potential for such shifts would point again to a flexible interpretation of the loss function underlying the FOMC's policy framework.


II.C. The Fed's Transparency and Credibility

The changes discussed so far to the policy framework in terms of the inflation goal and the systematic component of policy are also related to other elements of the framework, most notably transparency and the efficacy of Fed actions. In this regard, Refet Gürkaynak, Brian Sack, and Eric Swanson (2005) find evidence of excess sensitivity of longer-dated forward rates to economic news, which they argue is indicative of the public having to learn about the monetary authority's inflation target.²¹ More broadly, changing long-term inflation expectations could result from policymakers' lack of transparency or lack of credibility. Here, we revisit the relationship between a short-run spot interest rate and forward rates in the days when the Consumer Price Index (CPI) or the Producer Price Index (PPI) were released, and compare this reaction with nonrelease dates:

$$(4) \quad \Delta i_t^{F,j} = \alpha_0 + \alpha_1 \Delta i_t d_i (\text{release day} = 1) + \alpha_2 \Delta i_t (1 - d_i (\text{release day} = 1)) + \varepsilon_t.$$

The dependent variable $\Delta i_t^{F,j}$ is the daily change in the Treasury forward rate j years ahead, whereas the explanatory variable Δi_t is the daily change in the spot three-month Treasury bill yield. The dummy variable d_i takes the value of 1 on days when there was a CPI release or a PPI release,

21. Kozicki and Tinsley (2001), though not focusing specifically on economic news, also argue that movements in forward rates at the longer end of the maturity spectrum have been related to shifts in market perceptions of the policy target for inflation.

Figure 11. Estimated Coefficients for Equation 4, 1970–96 and 1997–2007^a

Sources: Authors' calculations; Federal Reserve Board.

a. The dashed lines are plus or minus 2 standard errors.

and a value of zero on the other days. The specification assumes that on release dates, the change in the spot three-month Treasury bill captures the “news” effect of the CPI or PPI release, and that the effect of the release on the forward rates can be assessed from its impact on the spot rate.²² We consider instantaneous forward rates spanning the maturities from 2 to 15 years ahead.²³ Near-term forward rates will be affected by cyclical variables, including expectations about monetary policy actions. Longer-term forward rates are determined by more persistent factors, including expectations about policymakers’ target for inflation.

Figure 11 reports estimation results for the coefficients α_1 and α_2 in equation 4 over two subsamples. The periods we consider are 1970–96

22. This assumption allows us to circumvent the issue of not being able to measure the news effect of the release using market survey data in the period before the 1990s. We focus on inflation release dates because these should capture potential shifts at the longer end of the maturity spectrum that are motivated by shifts in perceptions about the long-run inflation objective.


23. More detail about the data and estimation is provided in the appendix.

and 1997–2007.²⁴ The sample split is informed by our previous inference on the FOMC’s inflation objective. For the longer-dated forward rates, the reaction to inflation news is stronger than on nonrelease dates—that is, the heavier line is above the lighter line, or $\alpha_1 > \alpha_2$ —in the 1970–96 period. We take this finding as consistent with the view that since the 1970s and up until the late 1990s, financial markets had changing perceptions about the FOMC’s inflation goal, with those perceptions being influenced by news about inflation. The exercise cannot assess whether the way the public was revising expectations about the FOMC’s inflation objective was consistent with the FOMC’s changing target π_i^* as depicted in figure 6. Still, a lack of transparency about the inflation goal could have affected movements at the longer end of the expectations curve, an indication of the markets’ lack of confidence in the Fed’s commitment to bring inflation to a specific, well-understood inflation target. The more recent period, with no significant response of longer-dated forward yields to changes in the short-term Treasury bill both on inflation release dates and on nonrelease dates, is consistent with the public perceiving the policymaker as having a credible and stable inflation target.

Needless to say, this exercise provides at best partial answers to the evolution of the Fed’s transparency over time. And the findings for the 1970s and 1980s could have different interpretations. It is possible, for example, that the excessive reaction to inflation news at the longer end of the term structure was also a symptom of a lack of credibility in the 1970s. In the 1980s and early 1990s, the reasons for such a result could be different, and hinge not on the Fed’s credibility in its stance about inflation but on the FOMC’s vagueness about its long-run inflation goal. Nevertheless, the results suggest that the steps taken to increase credibility and transparency over time took long to manifest themselves in the form of the long end of the term structure becoming unresponsive to short-run inflation news. One potential reason for this finding is that the move toward increased credibility and transparency was incremental.

It is possible to identify a number of steps in this incremental process. The Humphrey-Hawkins Act of 1978 certainly increased transparency and solidified the importance of the dual mandate. This act clarified the goals of monetary policy, increased accountability to Congress, and provided an opportunity for a more transparent discussion of monetary policy actions.

24. We use daily data from the 1970s to the end of 2007, and exclude the more recent period because of the complications associated with the conduct of monetary policy at the zero lower bound.

Figure 12. Mentions of the Federal Reserve's Credibility, 1960–2012^a

Source: Federal Reserve Board.

a. Four-meeting moving-average term counts as a percentage of total words in FOMC transcripts, memorandums of discussions, and historical minutes.

And the detailed account given by Marvin Goodfriend and Robert King (2005) of the Volcker disinflation highlights Volcker's understanding of the importance of credible monetary policy actions vis-à-vis financial markets, and in particular the role of credibility in informing markets' expectations about inflation in the medium and longer runs. Figure 12's word counts show an increasing number of discussions at the FOMC table about credibility. Over time, this development led to important changes in the conduct of monetary policy. As discussed above, monetary policy actions became more predictable and more clearly anchored to the dual mandate goals.


Throughout the 1990s, several changes occurred regarding the transparency of monetary policy. The first tentative step toward greater transparency occurred in 1994, when the Federal Reserve began to include the intended change in the federal funds target in its statement. Although most financial market participants had been aware of the focus on the federal funds rate since 1987 or earlier, the Fed simply did not announce its new funds target before 1994. The adoption of a target of roughly 2 percent was discussed and agreed upon internally by the FOMC in 1996. From a transparency standpoint, it is interesting that such a target was not initially

made explicit to the public. It is possible that greater transparency was perceived as potentially carrying a credibility cost if that 2 percent target were subject to change at a future date. Another move toward greater transparency and predictability occurred during the slow recovery from the 2001 recession. By the late summer of 2003, the Fed had hit what it assumed was its effective lower bound. Given the low rate of inflation at the time, the FOMC indulged in rudimentary “forward guidance,” which was meant to provide the markets with its view of *future* policy actions.

Other important improvements to transparency have occurred during the most recent period, which is not covered in our empirical exercise. Forward guidance became a crucial element in the conduct of monetary policy when the federal funds rate was at the effective zero lower bound. In addition, the Federal Reserve began to hold quarterly press conferences to explain its policy actions to the public, and recently the Fed announced that these press conferences will be conducted after each meeting. Perhaps most important, starting in 2012 the FOMC provided a document explicitly outlining its monetary policy framework, including an explicit 2 percent inflation target. This framework document is explicitly voted on at the January meeting, and it certainly provides an annual opportunity to communicate changes in the monetary policy framework. However, more extensive, comprehensive, and public discussions of the policy framework, such as those conducted by the Bank of Canada, have not yet emerged through this process.

II.D. Other Strands of the Framework

The empirical exercises so far have highlighted some, but not all, of the relevant changes to the monetary policy framework. One important element that has not been discussed so far is the role of financial stability in monetary policy. For an empirical examination of how financial stability has affected the conduct of monetary policy over time in the context of a reaction function similar to the one considered in equation 3, we refer the reader to the research of Joe Peek, Eric Rosengren, and Geoffrey Tootell (2015). An enduring effect of the late-1990s “productivity revolution,” which was used as justification for the significant boom in equity prices relative to earnings, was an increase in the attention paid to asset prices and (more generally) financial stability in the policy discussion, as shown in figure 13. This focus subsided somewhat after the 2001 recession, which was caused in part by a significant reduction in prices of Internet-related stocks, but it was a precursor to the renewed focus on financial stability issues following the 2008 financial crisis. It is reasonable to view increased

Figure 13. Mentions of Financial Stability in the Policy Discussion, 1973–2009^a

Sources: Federal Reserve Board; Peek, Rosengren, and Tootell (2015).

a. Four-meeting moving-average term counts as a percentage of total words in FOMC transcripts, memorandums of discussions, and historical minutes.

attention to financial stability issues as a change in framework, tantamount in its extreme version to the adoption of a “ternary mandate.”

The other aspect of the evolution of the framework is more technical, and pertains to the instrument or set of instruments used for the conduct of monetary policy. A full discussion of these tools is beyond the scope of this analysis. Ben Bernanke (2006) provides a historical perspective on the use of monetary aggregates as a guide for monetary policy, and their eventual demise in favor of the federal funds rate as the primary tool of monetary policy. Nevertheless, it is important to highlight the use of the balance sheet as a policy tool during the Great Recession and the ensuing recovery, when the federal funds rate was at the effective zero lower bound. The way one judges the efficacy of those asset purchases has consequences for the urgency with which one perceives that the current framework needs to be changed.

II.E. Why Have a Formal Framework Review?

Overall, it seems reasonably clear that the monetary policy framework has changed along several important dimensions since the late 1960s.

Some changes occurred rapidly when necessitated by acute economic conditions, such as the Volcker disinflation. Other framework changes were considered and implemented gradually under relatively benign economic conditions as a response to long-standing problems with the framework, such as increases in transparency.

Concerning the merits of a formal framework review, several observations can be drawn from our historical assessment of framework changes. First, some of the past shortcomings in the conduct of monetary policy can be ascribed to missing elements of the framework laid out in section I. The lack of clarity on an inflation goal in the late 1960s and early 1970s contributed to a mix of fiscal and monetary policies in which output stabilization became the primary focus, and in which the responsibility of the central bank for price stability was more ambiguous. This subordination may also have been a consequence of the lack of independence and credibility at the time in pursuing monetary policy actions to reduce inflation. These issues were later exacerbated by bad luck in the form of a number of adverse supply shocks, with the stabilization of inflation again taking the backseat, in part because the sacrifice ratio was perceived as being too high.

Other elements of the framework came into being in different ways. With the Humphrey-Hawkins Act, Congress specified the dual goals of monetary policy. But it was Fed chairman Paul Volcker who ultimately reclaimed the role of monetary policy in controlling inflation. Although we have shown that the implicit inflation target under Volcker was subject to change, the direction of the change was unmistakably toward lower inflation, even if the disinflation was a bit “opportunistic.” Another element of the framework that emerged under Volcker (after the nonborrowed reserves operating procedure) was greater predictability of Fed actions. These elements in the conduct of policy were largely cemented under Greenspan’s chairmanship. Others came into focus later, and largely had to do with a better appreciation of the role of expectations and the associated importance of transparency in communicating long-run goals and policy intentions. In sum, many key elements of the framework that have been put in place have had an evolutionary aspect. The evolution of economic thinking and the lessons gleaned from historical experience have been instrumental in providing a better understanding of the ways to improve central bank design (Blinder 1998; Reis 2013). This evolution of our understanding of monetary policy exemplifies one way in which regular conferences might be helpful.

Second, it is interesting to note that the significant change in approach to policy taken under Volcker occurred with an essentially unchanged FOMC. Thus, this episode provides a stark example of the FOMC chair's crucial role in policymaking. It is notable that accounts of the history of U.S. monetary policy often identify regimes or frameworks with the FOMC chair at the time (for example, Romer and Romer 2013). Although our account of relevant framework changes suggests that such a view can be an oversimplification, only recently, under Chairman Ben Bernanke, has the FOMC taken steps to codify the policy framework and make it less chair-dependent. The 2012 introduction of the "Statement on Longer-Run Goals and Monetary Policy Strategy" is the clearest example, but there are many others. The increased transparency and accountability gained from the publication of the FOMC's projections forces a more focused discussion of the policy issues at stake in FOMC meetings. And the regular assessments of the risks to financial stability are a step toward better incorporating financial developments into the policy decision process. A key goal of a regular framework review of the type outlined in this paper is to continue the progression toward a more stable framework.

Third, the FOMC's large and diverse composition should in principle provide for a better representation of different viewpoints and a more continuous stress-testing of the framework in place. Although there is truth to this notion, the observation made just above about the importance of the Fed chair in the decisionmaking process also makes it clear that "the FOMC is not a simple democracy, but a consensus-driven organization, with the agenda set by the chair" (Bernanke 2016, 6). This approach to decisionmaking has benefits, but it might well be improved by devising processes to ensure that the status quo is regularly challenged and that necessary changes are adopted in a timely fashion. Again, a regular conference would help serve this function.

In sum, we view a formal framework review as a natural step following those already undertaken to strengthen the Federal Reserve as an institution. Though good policy will always benefit from a good chair, it will also benefit from a resilient framework. Undertaking a regular assessment of the framework that solicits input from varied sources increases accountability and transparency. It also helps to ensure that the framework in place is followed if it remains appropriate, or is changed if merited by economic circumstances or a new understanding. Thus, a regular formal framework review also improves the Fed's accountability, because a regular review

forces policymakers to better articulate to both the markets and the public at large the rationale for their actions in the context of the framework.

III. A Suggested Approach for How the Federal Reserve Should Regularly Evaluate the Monetary Policy Framework

Currently, the Federal Reserve reviews the framework document once a year. At issue is the depth at which it reassesses the document, the openness of the process, and the inputs from which it draws in reviewing. Instituting a less frequent and more thorough process for evaluating the Fed's monetary policy framework, such as that used by the Bank of Canada, sounds straightforward. But in practice, such a process would require decisions on a number of key features of the review. Here, we briefly outline the trade-offs involved with each of the key features, and we recommend one approach to the review process that we believe nicely balances these trade-offs.

III.A. When Should a Review Be Conducted?

A key factor for deciding when to conduct a review is whether the review's timing should be regular or state-dependent. Ideally, a central bank should be able to make effective changes to its operating framework whenever the need arises. The history recounted in section II suggests that the state of the economy has not always provided sufficient inducement to trigger a framework change. Section II implicitly provides possible guidelines for conditions that could prompt a state-dependent reevaluation:

—A significant deterioration in economic performance that is not readily linked to nonmonetary policy factors, perhaps along the lines of the loss function estimates (squared deviations of inflation from the target and of unemployment from the estimate of the natural rate) presented in figure 1;

—A significant change in the behavior of long-run inflation expectations, and other financial market signals that could imply a loss of efficacy and credibility, for example, along the lines of the results presented in figure 11;

—Or, on a brighter note, compelling evidence from new empirical research in the field that a superior framework exists.

In practice, such indicators and others are routinely examined by the Federal Reserve System's staff. The circumstantial evidence over the Federal Reserve's 105-year history, a portion of which has been examined in the previous section, suggests that changes to the framework have often occurred too slowly at key junctures, most notably during the Great Depression, when the persistent adherence to the gold standard

critically constrained the Fed's ability to respond to the crisis. Another widely studied example is the 1970s, when it took more than a decade to successfully address the significant rise in inflation. Given the nonsystematic way in which framework changes have occurred historically, and the mixed history of the timeliness and effectiveness of such changes, we suggest that the FOMC regularly reassess key elements of its framework at a fixed interval, perhaps more formally than the current annual signoff on the "Statement on Longer-Run Goals and Monetary Policy Strategy."

A choice to use a fixed frequency should not, however, be overly rigid. It would be foolish to assume that policymakers can anticipate all the circumstances that might require a change in the framework. Thus, even within a regular frequency review, it might be wise to allow for an escape clause that makes it possible to reassess the situation off the regular schedule.

III.B. Who Sets the Agenda, and Who Provides Input for the Review?

If a formal review process is undertaken, there are several options for how to structure such a review, most notably who sets the agenda and who participates in the review discussion. These decisions can be more important than it might seem. On one hand, outside political influence in setting the review's agenda could be viewed as eroding the Fed's independence. On the other hand, including outside voices in the review discussion (and making the discussion public) could go a long way toward building public accountability for the Fed's framework decisions.

Because the goal is to use the review as an input into the FOMC's decisions about its framework, we argue that agenda setting should be done primarily by the FOMC, whose members would be required to vote on changes, so there will be FOMC support for any changes that are considered.²⁵ Moreover, the FOMC members should know more than anyone else about the key issues with which they have been grappling. The Federal Reserve staff is constantly reassessing the framework; as a result, over the course of time there should be a fairly large inventory of topics from which the FOMC can choose.

Although we suggest that the FOMC and Federal Reserve staff should have primary responsibility for setting the review agenda, one cannot rule out the fact that consensus-building pressures at the FOMC might lead it to overlook dissenting views. To ensure that dissenting views are presented,

25. Any changes to the existing framework would likely be voted on at the following January organizational meeting, when the current framework is approved and other FOMC organizational changes occur.

it may be helpful to also have non-Fed economists and academics suggest possible topics for discussion at the review. Of course, the FOMC would ultimately decide which topics will be considered.

Although we argue that the FOMC should have primary responsibility for setting the agenda, it seems important for the review itself to include participants with many viewpoints, both internal and external. The details of how to include those with diverse viewpoints could vary. At one end of the spectrum, the review process could be only internal. In this case, the process could still draw on outside expertise, perhaps by surveying relevant research, but possibly also by soliciting external analyses. With respect to the meeting's outcomes, the FOMC would choose how much to disclose and when. At the other end of the spectrum, the analysis supporting the framework review could come solely from external contributors. Such a process might be akin to the processes at some central banks that have engaged an expert panel to provide an independent review of their monetary policy performance in recent years.²⁶ A key drawback of this option is that external reviewers might be unaware of all the internal work done by Fed staff to evaluate the framework. And one could argue that external evaluations have already taken place to some extent, as reflected in the volume of academic research and conferences devoted to this topic. Thus, on balance, we would argue for a synthesis of these two approaches, incorporating both internal and external inputs into the review.

III.C. What Should Be the Content of a Review?

One feature that should be common to all reviews is an evaluation of the current framework relative to agreed-upon criteria—for example, estimates of economic loss in recent years from a variety of loss functions, deviations from estimated policy rules, comparisons with optimal policy exercises, and deviations between SEP and market expectations (adjusted for other substantive and methodological differences). These elements of the review should provide a starting point for a discussion about potential changes to the framework. Much of this assessment could be compiled by the staff, but it might be augmented by conference participants' independent performance

26. Most often, these external reviews focus on monetary policy performance broadly defined, rather than on more specific aspects of the policy framework. Examples include the Norges Bank's "Norges Bank Watch," an annual report written by an independent committee of economists to evaluate the bank's monetary policy performance; Ingimundur Fridriksson's 2010 report on the Norges Bank's monetary policy process; the 2000 "Independent Review of the Operation of Monetary Policy," for the Reserve Bank of New Zealand; and a 2010 review of the Reserve Bank of Zimbabwe's monetary policy.

assessments. Candidates for altering the framework would then be considered, drawing on analyses from both staff and external participants.

The issues related to the elements of the framework outlined in section I are all potential candidates for a review, although as suggested above, we would likely shy away from issues that require alterations to the Federal Reserve Act. In section IV, we describe two challenges facing the current monetary policy framework that could be the subjects of a review. Needless to say, there are other issues worth considering. The historical review of framework changes in section II highlighted shifts in the conduct of monetary policy that were arguably related in part to changing views about the appropriate loss function to minimize. In the current circumstances, a discussion about the symmetry of losses related to unemployment—apart from inflationary consequences, does low unemployment imply losses as large as high unemployment?—could be an ideal focus for a framework review. Changes in economic structure and their influence on the conduct of monetary policy would also be candidates for discussion.

III.D. Who Decides If the Framework Needs to Change?

The element of the review process focusing on the question of who decides if the framework needs to change should be less controversial. The recommendations that emerge from the review will depend on who sets its agenda and how it is structured. Whether the review is internal, external, or a hybrid of the two will affect who provides the recommendations about potential framework changes. But in the final analysis, only the FOMC has the responsibility for making decisions about the framework, because the FOMC is the one body that is accountable to Congress and the public for monetary policy performance.

III.E. A Proposal for a Framework Review

A framework review could comprise a number of possible combinations along the dimensions we have described here. One candidate would be a purely internal review, an augmentation of the internal processes already in place. As suggested above, Federal Reserve staff members and FOMC members already devote considerable effort to evaluating the policy framework. But in our view, the internal processes suffer from shortcomings: The time allotted for discussion of the review of the annual framework document has been modest; the scope of questions discussed is normally relatively narrow; and to date, an alternative framework has not been discussed in depth, including a motion with an up-or-down vote on moving in a new direction. Given the historical record, it can be argued that changes

in the framework, when they have occurred, have not always been timely, and have often been heavily dependent on the chair's leadership. Thus, if one were to choose an internal process, it would be valuable to set aside significant time for the FOMC to discuss recent performance, to identify shortcomings in the framework and its implementation, and to consider changes to the framework that might address any shortcomings identified.

Although the point of a framework review is indeed to augment the FOMC's internal decisionmaking process, it is not clear that keeping the process entirely internal would achieve the desired results. After all, the FOMC has had the option to use its internal processes to alter its framework since its inception. But the historical record outlined in section II suggests that delays and/or ineffective changes to the monetary policy framework were in part due to issues with the FOMC's internal decision-making processes. The literature on monetary policy decisionmaking by committee is growing, but still small.²⁷ And the extent to which the vast literature on group behavior from social psychology can be readily applied to a committee such as the FOMC is not clear. Nevertheless, one cannot rule out the reality that a committee such as the FOMC may at times be subject to some of the same issues pertaining to performance, coordination, and polarization as those highlighted in the social psychology literature. Broadly speaking, this literature notes that accountability, transparency, and outside examination of the group decisionmaking process are potential ways of mitigating pitfalls associated with group behavior (Sibert 2006).

For these reasons, our preferred approach to a framework review is one that provides roles for the FOMC, for the Federal Reserve staff, and for outside specialists. An open, FOMC-designed evaluation with both internal and external input would increase transparency and accountability, and would broaden perspectives without ignoring the work done internally in the Federal Reserve System. There are issues to consider when opening up the review process, most notably the risk of politicizing the framework review. But such concerns may be mitigated by a review that employs evidence-based argumentation, an important tool to falsify claims driven only by political motives. And it is important to note that an opaque, internal process also has political risks. Making important changes in the monetary policy framework without clearly explaining the process and rationale to the public invites political backlash.

27. Notable contributions are the papers by Blinder and Morgan (2005), Gerlach-Kristen (2004), Riboni and Ruge-Murcia (2010), and Sibert (2006).

Specifically, we envision a review that occurs mostly at regular frequency. The Bank of Canada's five-year horizon seems a reasonable starting point. A potential option is to adapt this timing to take account of the term of the Federal Reserve chair—allowing for one framework reevaluation for each Fed chair's four-year term. As noted above, we favor the inclusion of an escape clause that allows for an off-schedule reassessment when necessary.

The FOMC should take the lead in setting the agenda, although external input could also be taken into consideration. At the very least, the review should include supporting work by staff explaining the issues and why they were selected. Once the agenda is set, a call for papers on the selected topics would allow interested researchers from academia, other central banks, think tanks, and the private sector to submit their ideas for consideration. Again, the FOMC would take primary responsibility for selecting contributions from among those submitted, perhaps in consultation with external experts. As discussed earlier, the review should include an evaluation of the current framework, which could include both staff-generated and external evaluations of monetary policy performance. The results of the research presented at the conference could be summarized by staff, detailing the findings and what they may imply for framework changes. With the results of this public conference, as well as additional internal work, the FOMC would be well positioned to take formal action on changes it judges to be appropriate.

III.F. The Potential Costs of a Regular Review

Although we believe our recommendation could modestly improve the Federal Reserve's performance over time, it is wise to consider the potential costs of undertaking such a review. In particular, the Fed should consider the effect that such a process might have on expectations and credibility. Just the existence of such a process might imply to markets that, say, the inflation goal was somewhat more subject to change than it is at present, which might in turn increase the uncertainty about long-run expectations of inflation. The consideration of a specific change in the lead-up to a formal evaluation, if it became public, could similarly increase uncertainty about the Fed's actions in coming years. Suppose, for example, that it became known that, like the Bank of Canada earlier, the FOMC was considering the merits of price-level targeting. Knowledge of this fact should shift some probability weight toward its adoption, and could imply a different trajectory for the funds rate and for inflation over the medium horizon.

It is not obvious how to mitigate such effects, apart from clear communication about the scope of the review and a gradual buildup of experience with routine framework evaluations. But it is important to recognize that such effects may be at play, and to work to minimize their impact on economic outcomes. It is also important to note that framework changes may be perceived as improvements, and thus help reduce any economic stress—just as the cost of unemployment fluctuations around the natural rate declined when inflation expectations became well anchored.

III.G. How a Formal Review Differs from the Current Process

As noted above, the staff members of the Federal Reserve System are continuously evaluating elements of the framework. How would the proposed review differ from the ongoing process? There are a number of dimensions in which this framework evaluation would deviate from the ongoing internal process. First, from the FOMC's perspective, the current annual process is more concerned with instituting minor changes than with introducing a major evolution in how the FOMC conducts monetary policy. The process that we are advocating would require a more significant amount of FOMC time to focus on the performance of the current and prospective framework.

And second, the current process does not typically include a performance evaluation of the current framework, in particular: (1) Is the current framework showing signs of stress, or is it expected to in the near future? (2) Has the FOMC deviated significantly from the current framework; and if so, for what reason? (3) Has the FOMC deviated significantly from its "normal" behavior (that is, its estimated policy rule); and if so, for what reason? (4) Have economic losses been larger than usual in recent history? Are some of these losses attributable to monetary policy?

The evaluation would entertain much more input from outsiders. It may not be that outsiders possess unique knowledge about how to improve the framework, but they would bring somewhat different perspectives, and they could reduce any tendency for institutional inertia or group think.

IV. Is Now a Time When We Should Be Rethinking the Monetary Policy Framework?

Could the current framework be improved? For example, is it at risk of failure when the next downturn occurs? Is there a recognition among current FOMC members that a change should be considered now, perhaps consistent with other times when regimes were changed? The past

10 years have been marked by a record-sized recession and a financial crisis, the use of alternative tools to reduce the effects of the disruption, and a disappointingly long recovery back to full employment, despite the efforts undertaken during the crisis and in its aftermath. It would be difficult to say that economic performance during the recovery—specifically, the rate at which we reattained full employment—was completely satisfactory. Hence, the monetary policy framework is far from perfected; for a variety of reasons, more needs to be done. Despite the very significant changes over the past decade, the changes to date in its framework document, as detailed above, have been relatively minor. This may be one reason for having a more regularized schedule to discuss framework changes.

It is relevant to note that the two largest episodes of subpar economic performance in the Fed's postwar history have been the Great Inflation of the 1970s and the Great Recession and recovery that began at the end of 2007. In both these cases, one key failing has arguably been the fact that the Fed did not adequately address an emerging problem: whether or how to offset the rising inflation and inflation expectations in the first case, and how to overcome the lack of potent tools to offset recession in a low-inflation, low-real-rate environment in the second case.

In an important sense, these observations provide the strongest motivation for our recommendation for a regular review of the performance of the monetary policy framework. The economic environment is constantly changing, as shown by the two examples given above, when the economy was buffeted by large supply shocks and a declining real rate of interest. The framework must be flexible enough to adapt not only quickly but also effectively. The hope is that a regular review would ensure that the Fed would be ready to make the correct adjustments as soon as possible when they are required.

Here, we consider two high-level challenges currently facing the Fed's monetary policy framework, both of which might be viewed as requiring a change in framework: the potentially increased likelihood of protracted periods at the effective lower bound on interest rates; and the limited ability to stabilize the economy, including a chronic pattern of significantly overshooting full employment—a risk to which nonzero probability attaches in this cycle.

IV.A. The Effective Lower Bound on Interest Rates

We have been in a low-inflation regime for the better part of two decades. More recently, we have appeared to be in a regime of low real interest rates. These two imply that equilibrium nominal interest rates will, for some time,

be quite low by historical standards.²⁸ This in turn implies that the amount of policy buffer for conventional, short-term interest rate policy—the amount by which the central bank can lower its policy rate in response to an economic downturn—will likely be limited for some time. Thus, one motivation for considering alternative policy frameworks might be a desire to find one that would provide the central bank with a larger policy buffer.


During the Great Recession, we also learned about the efficacy of some key alternative monetary policy instruments. Most notably, the Fed's forays into quantitative easing (QE) and forward guidance provided an opportunity for researchers to estimate the effects of such policies on longer-term interest rates, on other asset prices, on inflation expectations, and on real economic outcomes (for example, Christensen and Rudebusch 2012; Gagnon and others 2011; D'Amico and King 2013; Hamilton and Wu 2012; and Swanson 2017). These findings bear on the confidence with which the Fed might use such tools in the future, which should in turn influence its comfort with a reduced policy buffer for its conventional instrument. If one accepts the median estimates of QE and forward guidance efficacy, and if one takes into account the difficulty experienced in returning the economy to full employment and target inflation following the Great Recession, one cannot assume that the current framework for monetary policy will necessarily provide enough potency to satisfactorily offset a modest to large-sized economic downturn, even combining the effects of conventional and unconventional policies. Thus, the prospect of a continued low-inflation, low-real-rate environment might well prompt consideration of monetary policy framework alternatives.

IV.B. Stabilizing the Economy Is Easier in Theory Than in Practice

When thinking about alternative policy frameworks in the form, for example, of adopting a price-level target, it is important to consider the record of monetary policy in stabilizing the economy. The top panel of figure 14 shows the four-quarter change in the unemployment rate, with recession shading, from 1949 to the present. The recurrent feature here is that whenever the unemployment rate increases by more than 0.5 percentage point, the economy always falls into a recession. The figure's bottom panel displays the unemployment rate gap over the same period. Whether using the latest-vintage estimates of the natural rate (the unemployment gap line) or real-time estimates (the Greenbook-based unemployment gap line), the figure

28. Kiley and Roberts (2017) assess the probability of becoming stuck at the effective lower bound from the perspective of two large econometric models.

Figure 14. The Record of Monetary Policy in Stabilizing the Economy, 1949:Q1–2018:Q3



Sources: Bureau of Labor Statistics; Congressional Budget Office; Federal Reserve Board; Haver Analytics; National Bureau of Economic Research.

shows a pronounced tendency for the unemployment rate to dip significantly and persistently *below* these estimates of the natural rate at the end of expansions. In every case, this overshooting is followed by a recession. The depth of the overshooting varies, and the magnitude of the ensuing recession varies, but the pattern is nearly perfect for postwar U.S. economic history.

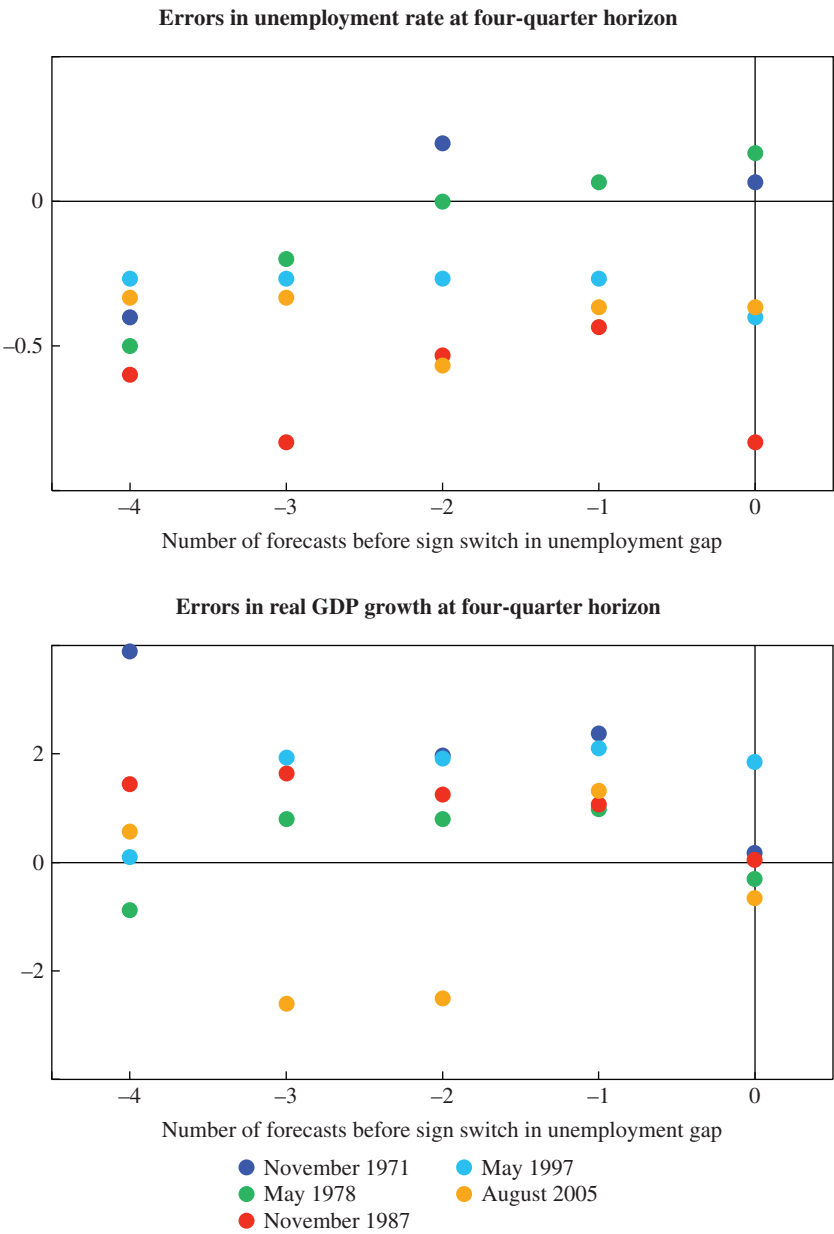
This limited ability to stabilize the economy could be due to events that are beyond the control of monetary policy. In some circumstances, the central bank has intentionally caused a recession, most notably in the case of the Volcker disinflation. It is relevant to note that for the three most recent downturns, staff forecasts as reported in the GB/TB have featured only a mild increase in the unemployment rate, to a level roughly consistent with the real-time estimate of the natural rate. In other words, the recurrent pattern has been one where the tightening of monetary policy has been expected to slow the economy down gently from above capacity to full employment. Ex-post, one might judge that monetary policy exacerbated the unexpected recession, but this is not what the Federal Reserve Board's staff was envisioning ex-ante. The limited ability to predict a recession is well known, and not just a feature of Fed forecasts. Here, we highlight the reality that once the unemployment rate starts to rise by a relatively modest amount, dynamics take hold that tend to push the economy into a recession in ways that standard linear models do not adequately capture.

There could be many reasons for the economy's tendency to overshoot full employment (or, correspondingly, to undershoot the natural rate of unemployment). The Fed (and other forecasters) could be surprised by the vigor of private growth late in the expansion, or by a late-recovery fiscal expansion (as occurred in the late 1960s, during the intensification of the Vietnam War). Still, one never sees an undershoot, by this definition—a landing “above the runway,” in which unemployment plateaus above the estimated natural rate before slipping into recession. The pattern is consistent enough that it should prompt thought about the role that monetary policy may play in this recurrent pattern.

The next figures provide some evidence on the Fed's forecasting errors, to see if they might help explain the economy's tendency to systematically overshoot full employment. Figure 15 examines the GB/TB forecasting errors for real GDP growth and unemployment at the four-quarter horizon.²⁹

29. Errors are computed using real-time, actual data from the Philadelphia Federal Reserve Bank's database, using the vintage of data eight quarters after the forecast was made. Data for longer-horizon forecasts are not as reliably available for the GB/TB data set, although the pattern for the available six- and eight-quarter-ahead unemployment forecasts is quite similar.

Figure 15. Greenbook/Tealbook Forecasting Errors at Turning Points in the Unemployment Gap



Sources: Bureau of Economic Analysis; Bureau of Labor Statistics; Federal Reserve Board.

The periods chosen are the forecast dates leading up to the time that the unemployment gap (as displayed above) changes sign from positive to negative. Because the data for the GB/TB are collected beginning in 1965, the first such episode that we can examine is late 1971. After that, there are four other episodes—1978, 1987, 1997, and 2005—along with a potentially late-breaking episode right now, with the quarterly average unemployment rate at 3.9 percent for 2018:Q2 and the CBO's latest estimate of the natural rate at 4.6 percent. Thus, unemployment has been below the current natural rate estimate since March 2017, or about 17 months—not a very long time by historical standards.

The top panel of figure 15 shows that at about the time when the unemployment rate begins to undershoot the natural rate (and correspondingly employment overshoots full employment), the GB/TB systematically overpredicts the unemployment rate four quarters ahead (errors are defined as actual minus forecast, so negative numbers indicate an overprediction). The same is true for eight-quarter forecasts, which are not shown.³⁰ As indicated in the bottom panel, at the same time, the GB/TB forecasts tend to systematically underpredict real GDP four quarters hence—which is consistent with the kind of Okun's law relationship that appears to be embedded in the GB/TB forecasting process.

Overall, errors forecasted around peaks and troughs could be just a reflection of the fact that the economy is hard to predict, especially at turning points.³¹ But it is also possible that these systematic errors and patterns at key junctures of the business cycle are indicative of more fundamental challenges that make it difficult to fine-tune the economy. Importantly for this paper, this inherent difficulty could be relevant when evaluating possible changes to the framework, such as price-level targeting. In such a regime, the need can arise to keep interest rates low for quite some time to offset the effects of the recession on the price-level gap, and subsequently to restrain the economy for some time, engineering a growth recession that brings inflation back down to target and employment back to full employment. The empirical record of policymakers' ability to engineer a growth recession that nicely lands the economy at full employment without morphing into a full-blown recession is not comforting. Similarly, a soft landing from an overheated economy—whether unexpected or not—to

30. Unfortunately, there are far fewer eight-quarter-ahead than four-quarter-ahead forecasts recorded in the GB/TB historical data set.

31. Recall that an optimal forecast will generally be less variable than the series being forecasted.

full employment has been a recurrent feature of past forecasts but not of actual outcomes.

IV.C. What Are the Alternatives?

If this is indeed an appropriate time to be considering the effectiveness of our current framework, what are the alternative approaches we should be considering? Key alternatives should probably include

- Inflation targeting with a different (higher) target rate;

- Adopting an inflation target range rather than a point target. The target range could vary with significant perceived and persistent changes in the equilibrium real rate of interest, per Rosengren (2018);

- Price-level targeting, including (1) conventional price-level targeting; and (2) opportunistic or asymmetric price-level targeting (in the wake of a large recession), thus making up for price-level misses on one side of the notional price path, when the policy rate hits the effective lower bound;³²

- Nominal GDP targeting, meaning a fixed combination of price-level and real-GDP-gap targeting;³³

- What should the loss function look like? Is the workhorse function the right one?

Other authors have reviewed the merits of these alternative frameworks, some in more detail than others. Those who have discussed the potential benefits (and costs) of raising the inflation goal to 3 or 4 percent have included Olivier Blanchard, Giovanni Dell’Ariccia, and Paolo Mauro (2010); Olivier Coibion, Yuriy Gorodnichenko, and Johannes Wieland (2012); Laurence Ball (2014); Paul Krugman (2014); John Williams (2016); and Lawrence Summers (2016). Estimates of the potential gains to a higher inflation goal are provided by José Dorich and others (2018) for Canada and by Michael Kiley and John Roberts (2017) for the United States. These authors find that in some circumstances, raising the inflation goal can provide substantial gains to macroeconomic stabilization.

The second option is discussed in a speech by Rosengren (2018), and is motivated by the observation that current estimates of equilibrium real short rates are quite low, implying (with a goal of 2 percent inflation) a low-equilibrium federal funds rate. Thus, the rationale is the same as those

32. See Bernanke (2017).

33. Nominal GDP targeting may be seen as a special case of a dual mandate policy that pursues price-level and output gap targeting, as it imposes weights of one on these two components of the GDP gap.

for most of the authors cited above. To date, no one has provided a numerical estimate of the benefits of an inflation goal that moves up and down with estimates of the equilibrium real rate.

The bodies of literature on price-level targeting and nominal GDP targeting are voluminous. Aspects of price-level targeting are discussed by Lars Svensson (1999) and by Coibion, Gorodnichenko, and Wieland (2012). Nominal GDP targeting, which implicitly imposes fixed weights of unity on both the real GDP gap and the price-level gap, has the advantage of ignoring the split between real and nominal activity, and thus in a sense automatically adjusts the policy rate when productivity growth (or other determinants of potential GDP growth) slows or speeds up. Robert Hall and N. Gregory Mankiw (1994), Stephen Cecchetti (1996), and Michael Woodford (2012) discuss the relative merits of nominal GDP targeting.

Although it is beyond the scope of this paper, in order to intelligently consider any of the alternatives presented above, we need to consider how we should evaluate the performance of historical and prospective monetary policy frameworks. Relatedly, the Fed (and any evaluation partners) would need to agree on how to assess the counterfactual of whether and/or how much an alternative framework might improve on the current one.

IV.D. Limits to the Alternative Monetary Policy Strategies

Most of the solutions mentioned above involve either temporary or permanent increases in the inflation goal. However, one must be realistic about how much comfort one should take in an indefinite increase of, say, 1 percentage point in the inflation goal, or a temporary increase of 2 percentage points in the goal. Starting from a steady state with a commensurately higher nominal interest rate would afford more latitude to lower interest rates—sometimes referred to as the amount of policy cushion—and would no doubt decrease the severity of a recession. But one must be realistic about the amount of relief such a framework would offer. Some of the studies cited above provide evidence bearing on the benefits of additional policy cushion (Kiley and Roberts 2017; Dorich and others 2018). But to simplify—using estimates from FRB/US, vector autoregressions, and the Federal Reserve Bank of Boston’s econometric models—every additional percentage point of federal funds rate decrease would yield about 1 percentage point more in real output, and an unemployment rate that is 0.5 percentage point lower. Thus, a regime with a 4 percent inflation goal would offset roughly 1 more percentage point of unemployment than

would a 2 percent regime. Though helpful, one should not expect such a framework to provide a complete solution to the types of monetary policy constraints faced, for example, during the Great Recession.³⁴

In addition, the apparently shallow slope of the Phillips curve makes the implementation of these policies more complicated. First, moving the economy to a significantly higher inflation rate today would entail a rather protracted period of subnatural rate unemployment. Second and related, on an ongoing basis, recessions that lowered inflation would similarly require protracted periods of low interest rates that, working through the Phillips curve, would move inflation back up to target. These periods of “low for long” would become a regular feature of macroeconomic policy under all these policies, and likely entails some risk of inducing either financial or macroeconomic instability.

The implication of these observations is that we should probably not rely on monetary policy alone, even with the best-designed framework, to take sole responsibility for economic stabilization. There are practical limits on the amount of stimulus that monetary policy could provide in the face of significant economic downturns. This observation implies that one should also consider whether there is an important role for fiscal policy in managing short-run fluctuations. And this of course is a topic for another paper.

V. Conclusions

We review some facts about monetary policy frameworks. First, they have changed quite a bit over time, with a frequency that is measured in years, but not decades. Second, they have changed for a variety of reasons. In some cases, such as the appointment of a new Fed chair in 1979, it was clear that economic performance had deteriorated, and a change was required.³⁵ In others, the economics profession’s understanding of monetary policy frameworks had evolved, and the Fed (often gradually) adapted

34. The same logic applies to the use of balance sheet policies (“quantitative easing” or “large-scale asset purchases” in Fed parlance) to stimulate the economy. The effects of these policies on interest rates to date have been of the same order of magnitude, and thus cannot be expected to offer more stimulus than policies that increase the amount by which short rates can be reduced.

35. One can of course debate whether the 1979 changes constitute a change in framework, or the correction of a misperception regarding the inflation/unemployment trade-off, or a recognition that discretionary deviations from the extant framework had been detrimental, and required a forceful return to the same framework. For the purposes of this paper, we take this to be a change in monetary policy framework.

to this change, as with the adoption of an explicit numerical inflation objective. In still others, key aspects of economic structure necessitated a change in framework, as in the failure of monetary aggregates to provide reliable indications of nominal GDP growth or inflation. Third, it seems best to characterize most changes in the framework as evolutions, rather than overnight revolutions. Recognition of framework deficiencies, recognition of key changes in economic structure, improvements in the profession's understanding of monetary economics—all these take time, and adoption normally lags recognition. Fourth, the distinction between a change in framework and a discretionary departure from a perfectly sound framework is subtle, but perhaps important. It matters because in some episodes, it may not have been the monetary policy framework but the lack of adherence to that framework that caused problems, and that necessitated a change in monetary policy implementation. Whether this change constituted the adoption of a new framework or better adherence to an old framework remains an open question.

Given this characterization of monetary policy frameworks, we believe the process that ensures adherence to a framework as well as the process for making needed changes to the framework can be improved. In particular, it is important that the Fed should consider a regular assessment of its monetary policy framework at a fixed interval and that this assessment provide a transparent evaluation of the current framework and how it could be improved or possibly changed. We hope that such a review process—in part, with the aid of outside contributors—would help the Fed more consistently adhere to its framework when it can continue to work well, and to make timely changes when it cannot. Although changes have regularly been made to the framework, an improved process would institutionalize the process of change, making the Fed less reliant on extraordinary leadership. As a transparent process, it would also help to hold the Fed accountable for adhering to the framework it announces, and to provide public and transparent justifications for changes to its framework. One can overstate the likely impact of such a process, but our judgment is that, over the long span of time, it could well help improve the economic outcomes delivered by the U.S. central bank.

Appendix

In this appendix, we provide information about data and estimation methods for equations 1, 3, and 4 reported in the text. We start with the inflation forecast equations, which we rewrite here for ease of exposition:

$$(1) \quad E_t \pi_{t+i} = \beta_{0,t} + \beta_{1,t} E_t \pi_{t+i-1}^4 + \beta_{2,t} E_t u_{t+i-1} + v_{i,t}, \quad i = 1, 2, 3.$$

In this system of equations, π_{t+i} denotes the annualized rate of inflation in quarter $t + i$, π_{t+i-1}^4 is the average of inflation prevailing over the four quarters from $t + i - 1$ to $t + i - 4$, and u_{t+i-1} is the level of the unemployment rate at $t + i - 1$. The operator E_t denotes a forecast made in quarter t . We consider forecasts of inflation one, two, and three quarters out, as indexed by i . We exclude the “nowcast” $E_t \pi_t$, because such a forecast is likely to be influenced by short-term factors that would not be adequately captured by equation 1. At each of the three forecasted horizons, the relationship is augmented by an error term, $v_{i,t}$. These errors are assumed to be persistent. In particular, we posit that $v_{1,t}$ evolves as an $MA(4)$ process. In each quarter t , we then have that $v_{2,t} = E_t v_{1,t+1} + \varepsilon_{2,t}$, and $v_{3,t} = E_t v_{1,t+2} + \varepsilon_{3,t}$, where the innovations $\varepsilon_{2,t}$ and $\varepsilon_{3,t}$ are such that $Cov(\varepsilon_{2,t}, \varepsilon_{2,t-j}) = Cov(\varepsilon_{3,t}, \varepsilon_{3,t-j}) = 0$ for any $j \geq 1$, but we allow $Cov(\varepsilon_{2,t}, \varepsilon_{3,t})$ to be different from zero.

The β coefficients in equation 1 are assumed to evolve as random walks, with uncorrelated innovations across coefficients. The coefficients remain the same at the three forecast horizons at which the relationship in equation 1 is estimated over, as only the timing of the variables is changing in accordance with i . This multiple-horizon aspect of the forecasts is especially useful for our purposes in that, under the plausible assumption that the same model is being used to forecast inflation at different horizons, it increases the degrees of freedom at the estimation stage, possibly allowing for a better identification of the coefficients of interest. The specification we use to model the inflation forecasts is admittedly simple, but it captures a fraction of the variation in the inflation forecasts, which, absent time variation in the estimated β coefficients, is already above 90 percent. The behavior of the inflation forecasts in equation 1 at the three different horizons is estimated jointly via maximum likelihood using the Kalman filter, over the period 1966:Q4–2017:Q4.

The Federal Reserve Board staff forecasts reported in the GB/TB are produced at every scheduled FOMC meeting, and the meetings have occurred at varying frequency but always more than once per quarter. To avoid estimation issues associated with uneven frequencies, we only consider one GB/TB per quarter, usually the one that coincides with the quarter’s middle month. When this is not possible, we consider the last GB/TB forecast made in any given quarter. Given that the staff’s forecasts are made public with a five-year lag, for the period from 2013 to the present we use for our analysis the FOMC’s economic projections. Specifically, for each SEP forecast that we consider, we take the middle point of the

published “central tendency” range. Unlike the staff’s forecasts, where the outlook is described at a quarterly frequency, the FOMC forecasts are less granular and follow a yearly frequency. We therefore interpolate these yearly forecasts to convert them to quarterly frequency. Such a procedure obviously injects additional noise into the exercise, but our findings are not affected qualitatively by the use of the publicly available FOMC projections for the most recent period. The inflation forecast is for the GDP deflator until 1985:Q4, for core CPI from 1986:Q1 to 2005:Q4, and for the core personal consumption expenditures deflator from 2006:Q1 to 2017:Q4. Given the forecasted horizons that we consider, the variable $E_t\pi_{t+i-1}^4$ includes a mix of forecasted and realized inflation. Whenever actual inflation is involved, we consider real-time realizations.

As concerns the estimated policy rule, it takes the form

$$(3) \quad \bar{ff}_t = \gamma_{0,t} + \gamma_{\bar{ff},t} \bar{ff}_{t-1} + \gamma_{\pi,t} E_t \pi_{t+3} + \gamma_{u,t} E_t u_{t+3} + v_t.$$

The only variable we have not already defined is the federal funds rate, \bar{ff} , which is given by the average value prevailing in the week after the FOMC meeting. The frequency is quarterly, with the same GB/TB selection criterion for the forecast variables $E_t \pi_{t+3}$ and $E_t u_{t+3}$ described above in the context of inflation equation 1. The error term, v_t , is assumed to follow an $MA(1)$ process. We allow for the variance in the error to exhibit breaks in 1979:Q4, 1986:Q1, and 1997:Q1. Accounting for the possible presence of heteroscedasticity in v_t is potentially important in order to correctly apportion time variation to the estimated coefficients. We posit that the γ coefficients evolve as random walks, with uncorrelated innovations across coefficients. The policy rule is estimated via maximum likelihood using the Kalman filter, over the period 1966:Q4–2017:Q4. We do not consider the post-2008 period because of the binding floor on the federal funds rate, whereas the starting date is dictated by the availability of the GB/TB forecast at the chosen horizon. For the exercise reported in figure 10, which uses forecasts of inflation and the unemployment rate from 2015:Q4 to the present, we consider the SEP projections as described above. In the context of univariate equation 3, maximum-likelihood estimates via the Kalman filter of the variance of the innovations in the random-walk processes underlying the time-varying coefficients γ will be biased toward zero. We use the median unbiased estimation procedure given by James Stock and Mark Watson (1998) to first estimate the variance in these innovations. Given such estimates, we then apply the Kalman filter to estimate the remaining parameters in equation 3.

Finally, in the daily-frequency regression

$$(4) \quad \Delta i_t^{F,j} = \alpha_0 + \alpha_1 \Delta i_t d_t (\text{release day} = 1) \\ + \alpha_2 \Delta i_t (1 - d_t (\text{release day} = 1)) + \varepsilon_t,$$

we use data on U.S. Treasury forward rates $i_t^{F,j}$ maintained by the Federal Reserve Board. These data are computed from U.S. Treasury yields, and the details of the computations can be found in the paper by Gürkaynak, Sack, and Jonathan Wright (2007). The dummy variable d_t takes the value of 1 on days when there is a CPI release or a PPI release, and a value of zero on the other days. The equation is estimated via ordinary least squares for each of the forward rates with maturity j going from 2 to 15 years. The estimates reported in figure 11 over the two subsamples that we consider, 1970–96 and 1997–2007, feature heteroscedasticity-consistent confidence bands. Estimation results in the earlier sample are sensitive to outliers, and for this reason we have excluded from the estimation observations featuring a daily change in the three-month Treasury bill rate Δi_t in excess of 50 basis points in absolute terms. Once outliers have been excluded, it is of interest to note that considering the 1970s, 1980s, and early 1990s separately produces estimates that are qualitatively similar to the ones reported for the entire 1970–96 period.

ACKNOWLEDGMENTS We thank our editors, Janice Eberly and James Stock, and our discussants, Jean Boivin and Stanley Fischer, for helpful comments. David Brown, Emily Freeman, and Peggy Gilligan provided excellent research assistance.

References

- Ball, Laurence M. 2014. "The Case for a Long-Run Inflation Target of Four Percent." International Monetary Fund Working Paper 14-92. Washington: International Monetary Fund.
- Bernanke, Ben S. 2006. "Monetary Aggregates and Monetary Policy at the Federal Reserve: A Historical Perspective." Paper presented at Fourth European Central Bank Central Banking Conference, Frankfurt, November 10.
- . 2016. "Federal Reserve Economic Projections: What Are they Good For?" Brookings blog, November 28. <https://www.brookings.edu/blog/ben-bernanke/2016/11/28/federal-reserve-economic-projections/>.
- . 2017. "Monetary Policy in a New Era." Paper presented at conference on Rethinking Macroeconomic Policy, Peterson Institute, Washington, October 12–13.
- Binder, Sarah A., and Mark Spindel. 2016. "Independence and Accountability: Congress and the Fed in a Polarized Era." Brookings Institution, Center for Effective Public Management. <https://www.brookings.edu/wp-content/uploads/2016/07/Monetary-politics.pdf>.
- Blanchard, Olivier J., Giovanni Dell'Ariccia, and Paolo Mauro. 2010. "Rethinking Macroeconomic Policy." *Journal of Money, Credit and Banking* 42 (supplement, 1, September): 199–215.
- Blinder, Alan S. 1998. *Central Banking in Theory and Practice*. MIT Press.
- Blinder, Alan S., and John Morgan. 2005. "Are Two Heads Better Than One? Monetary Policy by Committee." *Journal of Money, Credit and Banking* 37: 789–811.
- Boivin, Jean. 2006. "Has U.S. Monetary Policy Changed? Evidence from Drifting Coefficients and Real-Time Data." *Journal of Money, Credit and Banking* 38: 1149–74.
- Cecchetti, Stephen G. 1996. "Inflation Indicators and Inflation Policy." In *NBER Macroeconomics Annual 1995, Volume 10*, edited by Ben S. Bernanke and Julio J. Rotemberg. MIT Press.
- Christensen, Jens H., and Glenn D. Rudebusch. 2012. "The Response of Interest Rates to U.S. and U.K. Quantitative Easing." *Economic Journal* 122, no. 564: F385–F414.
- Clarida, Richard, Jordi Galí, and Mark Gertler. 2000. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory." *Quarterly Journal of Economics* 115, no. 1: 147–80.
- Coibion, Olivier, Yuriy Gorodnichenko, and Johannes Wieland. 2012. "The Optimal Inflation Rate in New Keynesian Models: Should Central Banks Raise Their Inflation Targets in Light of the Zero Lower Bound?" *Review of Economic Studies* 79, no. 4: 1371–1406.
- Conti-Brown, Peter. 2016. *The Power and Independence of the Federal Reserve*. Princeton University Press.

- D'Amico, Stefania, and Thomas B. King. 2013. "Flow and Stock Effects of Large-Scale Treasury Purchases: Evidence on the Importance of Local Supply." *Journal of Financial Economics* 108, no. 2: 424–48.
- Dorich, José, Nicholas Labelle St-Pierre, Vadym Lepetyuk, and Rhys R. Mendes. 2018. "Could a Higher Inflation Target Enhance Macroeconomic Stability?" BIS Working Paper 720. Bank for International Settlements.
- Erceg, Christopher, James Hebden, Michael Kiley, David López-Salido, and Robert Tetlow. 2018. "Some Implications of Uncertainty and Misperception for Monetary Policy." Finance and Economics Discussion Series 2018-059. Washington: Board of Governors of the Federal Reserve System.
- Gagnon, Joseph, Matthew Raskin, Julie Remache, and Brian Sack. 2011. "The Financial Market Effects of the Federal Reserve's Large-Scale Asset Purchases." *International Journal of Central Banking* 7, no. 1: 3–43.
- Gerlach-Kristen, Petra. 2004. "Is the MPC's Voting Record Informative about Future U.K. Monetary Policy?" *Scandinavian Journal of Economics* 106, no. 2: 299–313.
- Goodfriend, Marvin, and Robert G. King. 2005. "The Incredible Volcker Disinflation." *Journal of Monetary Economics* 52, no. 5: 981–1015.
- Greenspan, Alan. 2004. "Risk and Uncertainty in Monetary Policy." Remarks at Meetings of American Economic Association, San Diego, January 3. <https://www.federalreserve.gov/boarddocs/speeches/2004/20040103/default.htm>.
- . 2005. "Reflections on Central Banking." Remarks at symposium sponsored by Federal Reserve Bank of Kansas City, Jackson Hole, Wyo., August 26. <https://www.federalreserve.gov/boarddocs/speeches/2005/20050826/default.htm>.
- Gürkaynak, Refet S., Brian Sack, and Eric Swanson. 2005. "The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models." *American Economic Review* 95, no. 1: 425–36.
- Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright. 2007. "The U.S. Treasury Yield Curve: 1961 to the Present." *Journal of Monetary Economics* 54, no. 8: 2291–2304.
- Hall, Robert E., and N. Gregory Mankiw. 1994. "Nominal Income Targeting." In *Monetary Policy*, edited by N. Gregory Mankiw. University of Chicago Press.
- Hamilton, James D., and Jing Cynthia Wu. 2012. "The Effectiveness of Alternative Monetary Policy Tools in a Zero Lower Bound Environment." *Journal of Money, Credit and Banking* 44 (supplement 1, February): 3–46.
- Iakova, Dora M. 2007. "Flattening of the Phillips Curve: Implications for Monetary Policy." IMF Working Paper 07/76. Washington: International Monetary Fund.
- Ireland, Peter N. 2007. "Changes in the Federal Reserve's Inflation Target: Causes and Consequences." *Journal of Money, Credit and Banking* 39, no. 8: 1851–82.
- Kiley, Michael T., and John M. Roberts. 2017. "Monetary Policy in a Low Interest Rate World." *Brookings Papers on Economic Activity*, Spring, 317–72.
- King, Robert G., Yang K. Lu, and Ernesto S. Pastén. 2008. "Managing Expectations." *Journal of Money, Credit and Banking* 40, no. 8: 1625–66.

- Kozicki, Sharon, and Peter A. Tinsley. 2001. "Shifting Endpoints in the Term Structure of Interest Rates," *Journal of Monetary Economics* 47, no. 3: 613–52.
- Krugman, Paul. 2014. "Inflation Targets Reconsidered." Paper prepared for Forum on Central Banking sponsored by European Central Bank, Sintra, Portugal, May 26.
- Meltzer, Allan H. 2002. *A History of the Federal Reserve, Volume 1: 1913–1951*. University of Chicago Press.
- . 2010. *A History of the Federal Reserve, Volume 2, Book 1: 1951–1969*. University of Chicago Press.
- . 2014. *A History of the Federal Reserve, Volume 2, Book 2: 1970–1986*. University of Chicago Press.
- Murray, John. 2018. "Why the Bank of Canada Sticks with 2 Percent Inflation Target." Paper presented at conference on Rethinking the Fed's 2 Percent Inflation Target, Hutchins Center on Fiscal and Monetary Policy, Brookings Institution, Washington, January 8.
- Orphanides, Athanasios. 2003. "Historical Monetary Policy Analysis and the Taylor Rule." *Journal of Monetary Economics* 50, no. 5: 983–1022.
- Orphanides, Athanasios, and David W. Wilcox. 2002. "The Opportunistic Approach to Disinflation." *International Finance* 5: 47–71.
- Peek, Joe, Eric S. Rosengren, and Geoffrey M. B. Tootell. 2015. "Should U.S. Monetary Policy Have a Ternary Mandate?" Paper presented at Federal Reserve Bank of Boston's Annual Economic Conference, October. <https://www.bostonfed.org/macropprudential2015/papers/Rosengren-Peek-Tootell.pdf>.
- Reis, Ricardo. 2013. "Central Bank Design." *Journal of Economic Perspectives* 27, no. 4: 17–44.
- Riboni, Alessandro, and Francisco J. Ruge-Murcia. 2010. "Monetary Policy by Committee: Consensus, Chairman Dominance, or Simple Majority?" *Quarterly Journal of Economics* 125, no. 1: 363–416.
- Romer, Christina D., and David H. Romer. 2002. "The Evolution of Economic Understanding and Postwar Stabilization Policy." In *Proceedings of the Economic Policy Symposium at Jackson Hole, Federal Reserve Bank of Kansas City*, 11–78.
- . 2013. "The Most Dangerous Idea in Federal Reserve History: Monetary Policy Doesn't Matter." *American Economic Review* 103, no. 3: 55–60.
- Rosengren, Eric S. 2018. "Considering Alternative Monetary Policy Frameworks: An Inflation Range with an Adjustable Inflation Target." Paper presented at Money, Models, and Digital Innovation Conference, Global Interdependence Center, San Diego, January 12.
- Rotemberg, Julio J., and Michael Woodford. 1997. "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy." In *NBER Macroeconomic Annual 1997, Volume 12*, edited by Ben S. Bernanke and Julio J. Rotemberg. MIT Press.
- Sargent, Thomas J. 2001. *The Conquest of American Inflation*. Princeton University Press.

- Sibert, Anne. 2006. "Central Banking by Committee." *International Finance* 9: 145–68.
- Stock, James H., and Mark W. Watson. 1998. "Median Unbiased Estimation of Coefficient Variance in a Time-Varying Parameter Model." *Journal of the American Statistical Association* 93, no. 441: 349–58.
- Summers, Lawrence H. 2016. "A Thought-Provoking Essay from Fed President Williams." <http://larrysummers.com/2016/08/18/6937>.
- Svensson, Lars E. O. 1999. "Price-Level Targeting versus Inflation Targeting: A Free Lunch?" *Journal of Money, Credit and Banking* 31, no. 3, part 1: 277–95.
- Swanson, Eric T. 2017. "Measuring the Effects of Federal Reserve Forward Guidance and Asset Purchases." NBER Working Paper 23311. Cambridge, Mass.: National Bureau of Economic Research.
- Williams, John C. 2016. "Monetary Policy in a Low R-Star World." FRBSF Economic Letter 2016–23, August 15.
- Woodford, Michael. 2002. "Inflation Stabilization and Welfare." *Contributions to Macroeconomics* 2, no. 1, article 1.
- . 2012. "Methods of Policy Accommodation at the Interest-Rate Lower Bound." In *Proceedings of the Economic Policy Symposium at Jackson Hole, Federal Reserve Bank of Kansas City*, 185–288.

Comments and Discussion

COMMENT BY

JEAN BOIVIN A full decade after the global financial crisis, there is little agreement about the lessons of the crisis for monetary policy *frameworks*—let alone any implementation of major changes. The role of financial stability considerations in the setting of monetary policy, I would argue, remains one of the most important unresolved questions. But that is not the only one. The risk of hitting the zero lower bound poses other questions: Should the level of the inflation target be higher? Would price-level targeting—or its close cousin, nominal GDP targeting—be a superior framework for dealing with drawn-out periods of disinflation? That these fundamental questions remain unresolved suggests the potential for significant improvements in the Federal Reserve’s monetary policy framework.

How should the Fed ensure that desirable framework changes are implemented in a timely manner? This paper by Jeff Fuhrer, Giovanni Olivei, Eric Rosengren, and Geoffrey Tootell tackles this important question. It argues for a Bank of Canada–inspired five-year formal review of the policy framework. The authors document that the Fed has made periodic framework changes anyway, but they have been ad hoc and delayed. Ensuring timely adoption while also avoiding undesirable changes requires a more systematic review with proper triggers. This would also ensure greater accountability and adherence to the Fed’s legally mandated objectives.

Conflict of Interest Disclosure: The author, who is the managing director and global head of research at BlackRock, did not receive financial support from any firm or person for this paper or from any firm or person with a financial or political interest in this paper. No outside party had the right to review this paper before circulation.

I agree with Fuhrer and his colleagues that a more robust and systematic framework review process is needed. As they make clear, framework changes did occur. But without a formal process in place, the changes were often reactive, responding to problems or emergency situations. Looking ahead, preventing or addressing future crises may require proactive change. As the authors imply, changes with such foresight seem unlikely to happen without a more formal review and assessment process.

Recommending a formal review process may sound reasonable, but such a process creates important challenges that need to be addressed before implementation.

The only challenge that Fuhrer and his colleagues highlight is the communication of framework changes to financial markets. This challenge is overstated: A framework review process does not create greater communication problems than the status quo. Framework changes occurred regularly, the authors argue. Financial markets must contend with these odds whether or not a formal review process exists. One could even argue that a more deliberate review process would also lead to more deliberate communication.

In my view, the elephant in the room is the greater challenge: the political economy of monetary policy. Fuhrer and his colleagues approach the framework review process largely from a technocratic perspective. They appear to assume that the Fed will be firmly in control of the process. The reality would be quite different, I believe. Any meaningful framework change will require democratic legitimacy. Importantly, ignoring the political economy considerations could put the Fed's operational independence at risk. In addition, the necessary political engagement in the process could create more inertia.

I believe the experience of the Bank of Canada—with its five-year inflation targeting renewal—illustrates both the importance of having a formal review process and the significant challenges.

RISKS TO DEMOCRATIC LEGITIMACY OR OPERATIONAL INDEPENDENCE The first key challenge: Many of the framework changes that Fuhrer and his colleagues document, or are being contemplated, require democratic legitimacy—this can only be achieved through government involvement. This includes any changes related to the fundamental objectives of monetary policy or any changes that may have significant distributional consequences. Choosing the appropriate level of inflation is one example. The government should be accountable for establishing the socially desirable level of inflation—a tax with distributional consequences—while taking into account the implications for the efficacy of monetary policy.

Addressing financial stability concerns also has the potential for important sectoral and distributional implications that would require agreement with the government.

This is why—and it is important to stress this—Canada’s inflation-targeting framework and its renewal constitute a *joint* agreement between the Bank of Canada and the Government of Canada. This is key. The renewal process is first and foremost an exercise in democratic accountability. It provides an occasion to formally evaluate whether the Bank of Canada is meeting its commitments to Canada’s citizens. In recent years, there has also been a deliberate objective to make the renewal process more transparent and to engage Canadians more directly. One manifestation of this evolution: The 2021 Inflation-Targeting Renewal Conference had its first webcast in the fall of 2017 (Bank of Canada 2017). This marked a major shift away from what had been an invitation-only, off-the-record event for policymakers and academics.

In contrast, other types of changes should be insulated from the political process. To achieve its mandated objectives, a central bank is responsible for decisions based upon the bank’s best understanding of the economy. As this understanding evolves, the central bank needs to have the flexibility to adapt its view on which model best represents the economy at that time. A concrete example mentioned by Fuhrer and his colleagues is the level of the neutral policy rate. This is not something that a central bank chooses or directly influences, but that conceptually plays a role similar to a north star for navigators (albeit on a cloudy night). Whether the neutral rate moves up or down is not a matter of democratic legitimacy. It is something a central bank needs to adapt to. The central bank needs to explain the evolution of the neutral rate to the public and needs to defend why it is adjusting policy in response. Political considerations should not interfere with a central bank’s decisions on which models it uses, or its assessment of the economy.

These distinctions are crucial—and more complicated—when it comes to unconventional policies, such as quantitative easing and a negative interest rate policy. Consider two important aspects. First, these issues are mostly operational, for the purpose of ensuring that a central bank has the necessary tool kit to achieve its legally mandated objectives. For this reason, these decisions should not become overly politicized. In this regard, it is worth noting that the Bank of Canada’s introduction of a framework for unconventional policy in 2009 occurred outside the five-year review process. It was part of the normal monetary policy process. Second, some tools can have direct distributional implications and blur the

distinction between fiscal and monetary policy. The introduction of such tools should then explicitly involve the government.

One of my main concerns with Fuhrer and his colleagues' proposal is that it fails to draw these important distinctions. Setting a review process without a clear understanding of when and how the government should be involved could be risky. Ignoring the potential societal trade-offs of these policy changes can lead to political backlash. On the flip side, a formal review process can invite greater attention and politicization to what should remain operational aspects of monetary policy, and under the purview of the central bank. A formal review process that does not clarify these boundaries risks undermining both the legitimacy and the operational independence of the Fed.

MORE TIMELY ADOPTION OR INERTIA? According to Fuhrer and his colleagues, an important reason for a formal review process is to ensure periodic and timely adoption of framework changes. This is not obvious to me.

Since Canada's renewal process began in 1991, there have been five renewals—most recently, in October 2016. This process has been effective in providing democratic accountability for the Bank of Canada's policies, and it has also allowed ongoing research on potential improvements to the inflation-targeting framework.

But over the past 25 years, how many actual framework changes have occurred through this process? None. Why? A little detour into the episode leading up to the 2011 inflation-targeting renewal is instructive. There are other instructive episodes, but I happened to have been directly involved in this one.

At the beginning of this review process, the Bank of Canada put two questions on the table: (1) Should it target a lower level of inflation (yes, "lower"!)? And (2) should it instead target a path for the level of prices?

However, the global financial crisis shed a different light on these questions—maybe the question is not whether the level of the inflation target should be lower, but in fact higher. The crisis also prompted a third question: To what extent should monetary policy take account of financial stability considerations?

This new question became *the* most pressing issue at the time. The 2011 renewal was the first opportunity to reassess the Bank of Canada's policy framework since the financial crisis. We wanted to ensure that we drew the right lessons from the crisis. We also needed to explicitly address the fact that price stability does not guarantee financial or economic stability. Part of the answer to this problem is that the first lines of defense were

elsewhere: building a resilient financial system, with improved regulations and micro- and macroprudential policies. But at the same time, we realized that these lines of defense could be insufficient. Perhaps monetary policy might have a more direct role to play—particularly if monetary policy itself can contribute to excessive risk-taking?

Did the importance of this question lead to a change in the policy framework? We had an intense debate within the Bank of Canada. One view was that no framework changes were needed because the inflation-targeting regime in Canada already provided the flexibility to respond to financial imbalances by simply allowing the Bank to bring inflation back to target over a longer time horizon. Another view was that a credible framework should specify how this flexibility with the inflation-targeting regime would be used. The Bank had not made completely clear whether it intended to use this flexibility to lean against a buildup of financial imbalances—for example, raising rates higher than was called for by the near-term inflation outlook. If the Bank ever used this flexibility, it might come as a surprise. According to this latter view, such a clarification needed to be explained as a framework change.

This was not a theoretical debate. The linkages between high housing prices and elevated household debt levels in Canada were causes for concern, particularly when policy rates were expected to be low for longer. Would the Bank need to tighten policy with the explicit intent of mitigating these financial stability concerns? Would this be understood by Canadians as consistent with the policy framework? At the time, these were live, actively considered questions about the policy stance and its outlook.

This is how the debate unfolded within the Bank of Canada. However, the joint agreement with the government meant the political implications needed to be considered. No matter how important this was from a public policy perspective, concerns about financial stability were just not a high priority for most Canadians—in fact, it was almost the opposite. Had Canada not just proven the robustness of its financial system during a crisis? Why take the risk of turning a nonissue into an issue? “If it ain’t broke, why fix it?” To be clear, this discussion was not about the monetary policy implications—whether the Bank should lean against financial imbalances—but rather whether the potential to “lean against the wind” should be communicated to the public as a framework change or simply a reaffirmation of the status quo.

The point is this: The political economy considerations can create understandable inertia and a bias toward the status quo. I am not a political

scientist, but I believe this is a manifestation of the literature on “path dependence” in political science—why institutions are sticky and political actors are incentivized to protect the status quo.

In the end, the outcome was somewhere in between. The renewal background document and the joint agreement put greater emphasis than ever before on the role that financial stability should play in monetary policy and how the Bank of Canada might react. But the Bank stopped short of characterizing this as a framework change. The summary statement in the joint agreement characterized the outcome as a reaffirmation of the existing framework:

These lessons reinforce the value of Canada’s flexible inflation-targeting framework, including its ability to respond to external shocks and its occasional role in supporting financial stability. (Bank of Canada 2011, 3; emphasis added)

The document also clarified explicitly for the first time that this flexibility could involve leaning against financial imbalances, with a technical box providing a concrete illustration:

Second, through a longer targeting horizon, monetary policy can also promote adjustments to financial excesses or credit crunches. . . . On the flip side, a tighter monetary policy that allows inflation to run below target for a longer period than usual could help to counteract pre-emptively excessive leverage and a broader buildup of financial imbalances. (Bank of Canada 2011, 33)

This framework clarification was soon put to work. Between April 2012 and September 2012, the Bank of Canada’s four press releases all concluded with this sentence on monetary policy stimulus:

The timing and degree of any such withdrawal will be weighed carefully against domestic and global economic developments. (Carney 2012a)

But then, in October 2012, this was changed to:

The timing and degree of any such withdrawal will be weighed carefully against global and domestic developments, *including the evolution of imbalances in the household sector.* (Carney 2012b; emphasis added)

For the first time, a policy statement indicated the possibility that the policy setting could be adjusted in response to financial imbalances. But in the October 2015 Monetary Policy Report opening statement, the Bank of Canada made this statement:

Our risk-management approach implied that, in the absence of any additional macroprudential measures, our actions would affect the balance of risks in opposite directions. Lowering interest rates could worsen vulnerabilities related to household debt at the margin, but it would also lessen the chances of the oil

price shock triggering financial stability risks. In the current context, getting the economy back to full capacity with inflation on target is central to promoting financial stability over the longer term. (Poloz 2015)

Like the previous statements, this acknowledged that financial stability played a role in the setting of monetary policy. But in this case, the Bank of Canada was *not* making a statement about its ability to lean against rising financial imbalances by tightening policy. It was signaling that it might keep policy *easier* to avoid worsening financial stability concerns. This is not about any inconsistency with Canada's flexible inflation-targeting framework: When financial vulnerabilities are high, the Bank's inflation-targeting framework allows it to respond in both directions with *easier or tighter* policy. But the framework provides little guidance as to when one is more appropriate than the other.

This episode shows that the review process allowed the Bank of Canada to clarify that it can lean against financial imbalances. And in some ways, it has already done so. This is important progress. But what principles govern how the Bank will respond in practice to the buildup of imbalances? The policy framework does not yet provide a clear answer. Why is this question not fully resolved? I think part of the explanation is that politically, it is easier to incorporate these considerations through a reaffirmation of the flexibility of the regime—defending the status quo—rather than making the case for a framework change.

CONCLUSION To conclude, I agree with Fuhrer and his colleagues that there is strong case to formalize a regular and systematic review process to consider the Federal Reserve's policy framework. I believe this is needed to guarantee the efficiency of the policy framework and its democratic legitimacy. But this process will inevitably have to involve the government, and this needs to be explicitly acknowledged. The process needs to delineate what type of framework changes should be jointly decided and what operational and implementation flexibility should be insulated from political influence. Otherwise, the process risks undermining the Fed's operational independence and could interfere with the timely adoption of necessary operational improvements.

REFERENCES FOR THE BOIVIN COMMENT

- Bank of Canada. 2011. "Renewal of the Inflation-Control Target." https://www.bankofcanada.ca/wp-content/uploads/2011/11/background_nov11.pdf.
- . 2017. "Monetary Policy Framework Issues: Toward the 2021 Inflation-Target Renewal." Workshop summary. <https://www.bankofcanada.ca/2017/09/monetary-policy-framework-issues-toward-2021-inflation-target-renewal/>.

Carney, Mark. 2012a. "Opening Statement before the House of Commons Standing Committee on Finance." <https://www.bankofcanada.ca/2012/04/opening-statement-24-april-2012/>.

———. 2012b. "Opening Statement before the House of Commons Standing Committee on Finance." <https://www.bankofcanada.ca/2012/10/opening-statement-30-october-2012/>.

Poloz, Stephen S. 2015. "Integrating Financial Stability into Monetary Policy." Paper presented at annual membership meeting of Institute of International Finance. <https://www.bis.org/review/r151012a.pdf>.

COMMENT BY

STANLEY FISCHER As its title states, this paper by Jeff Fuhrer, Giovanni Olivei, Eric Rosengren, and Geoffrey Tootell addresses the question of whether the Federal Reserve should regularly evaluate the need to modify the monetary policy framework. That is done in the first part of the paper, consisting of the introduction and the question "How do we define a monetary policy framework?" The second and longer part of the paper starts with a historical review of monetary policy frameworks since the 1960s, and then examines issues relating to the process of regularly evaluating the framework. The paper concludes with another question, asking whether this is a time when we should be rethinking the monetary policy regime.

Let me start by saying that I believe a central bank should have enough analytic and executive capacity to make monetary policy on its regular schedule, and at the same time should be reevaluating its monetary policy framework and considering potential changes. In this context, I find the concepts of "regularly evaluating" and "monetary policy framework" as set out in the paper not entirely clear; my view is that an institution of the importance of a central bank—and the Fed is certainly sufficiently important—should be able both to make monetary policy and to be reevaluating its policies and its policy framework as a regular part of its daily work. The extent of the changes that it wants to make at any moment, and the question of when to make them, are matters of its internal organization and its capacity for taking action.

A major central bank or a well-organized smaller central bank should be able to change its policy framework rapidly. And if it cannot, it should come up with a plan for making changes on the fastest possible schedule, if that is essential to the situation in which it finds itself. A change that could be made immediately if there had been enough time to plan it can be made on a slower, more gradual schedule if that is the best that can be done at a

moment in time.¹ Further, a well-run central bank should have made plans to deal with emergency situations, and carried out dry runs or table-top exercises of those plans.

Fuhrer and his colleagues discuss the Federal Open Market Committee's (FOMC's) "Statement on Longer-Run Goals and Monetary Policy Strategy," which has been published each January since 2012. As vice chair of the Fed, I was chairman of the FOMC's Subcommittee on Communications, and thus was responsible—together with the subcommittee members, Loretta Mester, Jerome Powell, and John Williams—for advising on revisions to be presented to the FOMC for the annual "Statement on Longer-Run Goals and Monetary Policy Strategy."

It was not easy going. The subcommittee's main goal during that period was to make the meaning of the statement on the symmetry of the inflation target clearer than it had been. After considerable discussion, the FOMC inserted the following words, namely, that it "would be concerned if inflation were running persistently above or below this objective."² This sounds trivial; but so far as I remember, no one regarded the discussion as illogical, and each variant on the statement had its own advantages and difficulties.

Although the statement has been evaluated regularly, if only for a few years, Fuhrer and his colleagues imply that they do not believe that the statement meets the definition of a Monetary Policy Framework (MPF). They are right.

Why? Because the authors' definition of an MPF "Statement on Longer-Run Goals and Monetary Policy Strategy" is appropriately much broader than the material presented in the Fed's annual statement. Their definition is that such an MPF statement comprises "the set of tools and processes by which the central bank attempts to define and attain its high-level economic goals"—and that is much more than is contained in the statement.

Fuhrer and his colleagues provide a definition of an MPF that includes:

- The governance structure of the central bank.
- A set of ultimate goals for the central bank.
- A loss function.
- The instruments that the central bank directly controls in order to achieve its key goals.

1. These comments are also probably relevant to institutions other than central banks.

2. The "this" in this statement refers to the inflation target. Federal Open Market Committee, "Statement on Longer-Run Goals and Monetary Policy Strategy, Adopted Effective January 24, 2012; as Amended Effective January 26, 2016, https://www.federalreserve.gov/monetarypolicy/files/FOMC_LongerRunGoals_20160126.pdf.

—*Transparency*, including the predictability of monetary policy and “managing expectations” and forward guidance. The transparency category is very broad, and its empirical implementation is difficult to summarize. But the concept—perhaps under another name—is clearly extremely important.

—Rules *and* discretion: the systematic component of monetary policy As Fuhrer and his colleagues argue, “To the extent that a central bank’s behavior can be well described by a policy rule, . . . policy predictability will be enhanced, and the transmission channel will be more effective.” “Discretion” is included because “truly optimal policy may deviate noticeably from simple rules under certain conditions.”

—The central bank’s depiction of the economy, that is, the model, “broadly speaking.”

By this point, the definition of the MPF has become blurred. Is the model the MPF? Clearly not; for the goals of policy—the governance of the central bank, and particularly its policy process—are parts of the MPF. But does a shift from an assumption about the consumption function, or expectations generation, imply a change in the MPF? That is not clear.

Presumably, if we follow the approach of Fuhrer and his colleagues, we need to use the word “significant” somewhere—for example, by stating that the MPF changes when a revision implies significant changes in monetary policy for a given state of the economy. Clearly, the word “significant” is a weasel word, which would allow many an argument about whether the MPF will be or has been changed by a specific change or set of changes in the actions of the central bank in a given state of the economy.

And this raises the question of whether we need the concept of “evaluation” of the MPF. Decisionmakers and researchers should be evaluating what the central bank is doing whenever they are making decisions on its behalf. Someone who wants to change some aspects of an organization’s behavior should have to discuss their recommendation with colleagues and managers, who should evaluate the proposed change and decide whether to implement it. What will be evaluated is a suggestion for changing one or more of the ways in which the firm or institution carries out its tasks.³

Fuhrer and his colleagues cite a 2004 paper by Alan Greenspan that discusses the centrality of risk management to monetary policy decisions. But

3. Fuhrer and his colleagues also note, later in the paper, “As Fed insiders, we wish to emphasize that the framework is to some extent always under discussion and debate.”

they conclude that it is beyond the scope of their paper to provide a precise analytical framework for the Fed's response to risks. I will return to this issue toward the end of my comment.

A REVIEW OF MPFs SINCE THE 1960s Consistent with the conventional view, the authors identify the breakdown of the post–World War II MPF as having occurred in the middle to late 1960s. They add that the next major regime began with the appointment of Chairman Arthur Burns, and they identify this regime as continuing the upward trend in inflation that began in the second half of the 1960s.

Continuing the chronology, they talk about the inflation of the 1970s, which they characterize as changing views about inflation dynamics that resulted in the inflation process appearing to have become accelerationist—a change that they attribute to Milton Friedman's presidential address to the American Economic Association. In light of oil price increases, they mention cost-push inflation and Chairman Burns's 1970 view that cost-push inflation could not be dealt with successfully from the monetary side, and that it would be a great mistake to try doing so. They point out, in addition, that estimates of longer-run inflation increased steadily over the course of the 1970s—but that long-run responses to inflation and unemployment have been relatively stable over time.

They note that the emphasis on output stabilization and the sacrifice ratio “raises the much-discussed issue of Fed independence” from the political cycle. They mention favorably the Humphrey–Hawkins Act of 1978, which provided the basis for the FOMC's dual (perhaps triple) mandate of “maximum employment, stable prices, and moderate long-term interest rates.”

Then the good news: “The target-less regime began to crumble by the end of the 1970s.” This crumbling was brought about by the appointment of Paul Volcker as chairman of the Fed. The authors say that primitive discussions about Fed independence, which were often couched in terms of “credibility,” then began. If the discussions of Fed independence at that time were primitive, that was partly because the discussion in the United States did not take into account some of the foreign literature—including, for example, the works of Richard Sayers, who was the historian of the Bank of England and a professor at the London School of Economics and Political Science.⁴

4. See, for instance, Richard Sidney Sayers, *The Bank of England 1891–1944, Volume 1* (Cambridge: Cambridge University Press Archive, 1976).

The authors' summary of the period through the end of the Volcker era states that "it seems reasonably clear that the monetary policy framework has changed along several important dimensions since the late 1960s." And so it has.

THE PROCESS OF REGULARLY EVALUATING THE FRAMEWORK Fuhrer and his colleagues are enthusiastic about the possibility of regularly evaluating the framework. They are impressed, as I have been in the past, by the Canadian law requiring the chair of the central bank and the finance minister to reach a new agreement on the goals of monetary policy every 5 years. I have heard rumors that some Canadian monetary policymakers are now contemplating turning the 5 years into 10 years.

The authors mention 12 elements that will need to be dealt with in putting a regular evaluation of monetary policy in place. Some of these elements give cause for concern. Let me focus first on two of them.

First: "Would a more open and regular evaluation of the monetary policy framework improve policy in the United States?" In this regard:

—One of the most important changes in the behavior of the economy in recent years in the U.S. and in other countries is that there is now considerable global agreement on the inflation target, namely, 2 percent. This is a significant achievement, for it means that there is now close to a consensus about the long-term value of money in the United States and Canada, the euro zone, and in Australasia—something that did not exist at all under the gold standard, or in the days of the effectiveness of the Bretton Woods framework.

—The assessment should not be tied to the term of the chair, for the Fed is not a political body with the holder of the chair as its leader; rather, we should see the Fed as making rules and implementing monetary policy as a technical matter of what is good for the economy, preferably for the long run.

—For essentially the same reason, *mutatis mutandis*, the term of the regular evaluations should have nothing to do with the federal political cycle.

—In brief, we would be worsening, not improving, monetary policy if we tied its execution to a single individual, or to the political cycle.

Second: "The consideration of a specific change in the lead-up to a formal evaluation, if it became public, could similarly increase uncertainty about the Fed's actions in coming years." Canadian policymakers with whom I have spoken do not seem to worry about this potential problem. I do. But we should also note that the Canadians have undertaken an evaluation of the potential need for changes in the MPF six times, and did not change the target inflation rate after the first evaluation.

Two more of the 12 questions asked by Fuhrer and his colleagues are also important.

First: Could a more formal, and open, review process improve how well [the Fed] adhere[s] to [its] current framework? Could transitions to a new framework be made more effectively? A central bank not burdened by the need to invoke meetings of the large number of people that would be involved in the kind of framework for an evaluation that the authors seem to be considering, could move much more rapidly in changing whatever features of the MPF are deemed necessary to change.

And second, in this section, the authors ask on whose inputs the Fed should draw. They present an impressive list of candidates, namely:

- other central banks;
- academics whose research focuses on central bank issues;
- participants in financial markets, who are part of the transmission mechanism of Fed actions to other markets;
- other policy institutions and think tanks that focus on central banking issues;
- representatives of the federal government; and
- members of the financial press.

Fuhrer and his colleagues present a quite long discussion of whether to include representatives of the federal government among those whose opinions should be consulted. This is a delicate issue, for representatives of the federal government would probably ask whether they should not be the first consultees, rather than any of the four groups above them in the list—that is, other central banks, academics, financial market participants, and other policy institutions and think tanks. Any process for evaluating revisions of the operation of the central bank will need to wrestle with the fact that “independence” for the central bank has a considerable element of independence from the political system in its meaning. Also, the authors’ discussion is not sufficiently sensitive to this difficulty, which could mean that an attempt to improve the MPF would instead contribute to politicizing the revision process.

In this regard, it is important that we take into account the difference between (1) parliamentary political systems, in which the government typically has an effective majority in the parliament, and in which the central bank can make agreements with the government by discussions with only a few political policymakers; and (2) the congressional system, in which an agreement on changing the monetary policy framework could well require reaching an agreement with many political policymakers. Anyone thinking about this process could come to understand the difference between

operating in a parliamentary system like that of Britain and in countries whose political systems are more similar to the American congressional system by comparing the process by which the Bank of England was made independent in 1997–98 with the creation of the third U.S. central bank—the Federal Reserve—that is described in Roger Lowenstein’s book *America’s Bank* (2015), which details the political maneuvering that led to the creation of the Fed in 1913.

IS NOW A TIME TO RETHINK MPFs? This penultimate section examines three issues that suggest to Fuhrer and his colleagues that we should be rethinking the monetary policy regime. The first issue is the effective lower bound on the interest rate. This problem has been analyzed in depth, and there has been some progress in dealing with it, including the possibilities of negative interest rates and of quantitative easing. Neither of these possibilities fills policymakers or the average citizen with joy, but they do suggest ways of responding to an interest rate that goes negative or threatens to do so.

The second issue is that of the chronic overshooting of the unemployment target in most of the recoveries in the U.S. economy since 1949. Fuhrer and his colleagues suggest that we have never seen an “undershoot,” in which the unemployment rate plateaus above the estimated natural rate before the economy slips into a recession. In this context, one might want to think of the supply-side recessions of the 1970s and 1980s as examples of an undershoot of this type.⁵ Nonetheless, more prominently, there has been a repeated and chronic overshooting of the unemployment rate target, in which the policymakers have in each cycle since 1949 “shown a pronounced tendency” to drive the unemployment rate below their estimates of the natural rate at the end of the cycle, as is the case at this moment—except that we do not know yet when this cycle will end. Further, they note that in every case the overshooting has been followed by a recession.

Having voted for a path that undershoots the target unemployment rate in the current cycle, I believe the tendency arises from the fact that although our utility function is quadratic around a higher unemployment rate than the present one, we continue to believe that less unemployment with inflation very close to the inflation target is a good place to be. Or, in other words, the main reason for our arguing that the natural rate of unemployment is higher than the current rate of unemployment is our fear of the inflation that is likely to follow if we push the unemployment

5. In the discussion here, the words “undershoot” and “overshoot” sometimes change their meanings for the direction of monetary policy changes.

rate down too far, rather than our studies of the greater efficiency of the labor markets at a higher unemployment rate than the current one.

The third instance is our increasing reliance on unobservable or estimated variables. At Jackson Hole, we heard from Chair Powell of the exploding stars— r^* , u^* , and π^* (Powell 2018). It is indeed a problem that we need to estimate these variables, but there is a firm analytic definition of each of these stars that makes them worth studying.

In discussing the alternatives to our current procedures, Fuhrer and his colleagues mention the possibility of operating with an inflation target range rather than a point target. This is what is contained in the revised Bank of Israel law, passed in 2010. The inflation target in Israel is 2 percent, with a target range of 1 to 3 percent. This approach works, and in my opinion the United States would be better off adopting a target range than a single number—but that will not solve the problem of the zero lower bound. Rather, it was included in the Israeli law to solve the problem of an economy with a large trade sector where movements in the exchange rate have frequent, large, and rapid effects on the inflation rate.

At the end, Fuhrer and his colleagues suggest that one—or perhaps many—should consider whether there could be an important role in stabilization policy for fiscal policy. Having raised the question, they move on rapidly by stating that that is a topic for another paper, and, they implicitly add, for another day. But it is clear that dealing better with crises will require better coordination between monetary and fiscal policy than has been typical in the United States.

CONCLUSION I would like to conclude these remarks by noting the importance of two topics that will require more discussion than they have received in the already-far-reaching paper by Fuhrer and his colleagues. The first is the treatment of risks and uncertainty in the policy process. We need to better prepare the U.S. monetary system for unexpected events than was evident in 2008 and 2009. I speak with the understanding that we will not be able to prepare precisely for future crises, for each crisis is somewhat different—and, as was noted in the discussion of Ben Bernanke's paper (2018), developments in each crisis will depend also on changes in the approaches to each crisis by monetary policymakers. But what policymakers in institutions for whom fighting crises is an important part of their mission certainly realize is that preparations made in exercises to deal with future crises themselves generate improvements in the institutions' capacity to contend with crises. To put it clearly, there is typically no manual that a central bank can simply pull off the shelf to deal effectively with a crisis—but thinking about and carrying

out desk exercises for dealing with a crisis will reveal changes that need to be made to improve the management of future crises.

Second, it is clear that problems in bank supervision and regulation played an important role in the processes that led up to the Great Financial Crisis. We do not yet have an efficient system for coordinating the responses of different agencies for dealing with future financial crises. And it may well take another round of reform in the supervisory/regulatory system to improve the treatment of the problems that are now visible in this system, including questions about the operation of the lender of last resort.

In summary, Fuhrer and his colleagues have written an ambitious and interesting paper, with a high level of productivity, even on a per author basis. However, the paper does not make a persuasive case for monetary policy reforms to be undertaken on a fixed schedule, rather than continuously, as the profession and the central bank make progress in their research, and as the world throws up new challenges and facts that will require changes in current and future economic policies. This paper will likely be seen as an important early contributor to that discussion as it progresses.

REFERENCES FOR THE FISCHER COMMENT

- Bernanke, Ben S. 2018. "The Real Effects of Disrupted Credit: Evidence from the Global Financial Crisis." In this issue of *Brookings Papers on Economic Activity*.
- Lowenstein, Roger. 2015. *America's Bank: The Epic Struggle to Create the Federal Reserve*. New York: Penguin Press.
- Powell, Jerome H. 2018. *Monetary Policy in a Changing Economy*. Speech given at "Changing Market Structure and Implications for Monetary Policy," symposium sponsored by Federal Reserve Bank of Kansas City, Jackson Hole, Wyo., August 24. <https://www.federalreserve.gov/newsevents/speech/powell20180824a.htm>.

GENERAL DISCUSSION Alan Blinder remarked that in some cases, institutional inertia at the Federal Reserve can be an advantage rather than a hindrance. For example, both the Supreme Court and the Federal Reserve are less subject to politics and, as a consequence, are less subject to the more volatile policy whims of Congress. Blinder also remarked that the continuous optimization of policy taught in economics courses should probably not apply to the Fed. Regarding politics, Blinder remarked that there are clear differences between changes in the Fed's framework and changes that would actually require alterations to the Federal Reserve Act.

For example, changes to the dual mandate would require a revision of the law by Congress. However, changing the inflation target or estimates of the neutral rate of interest would not require any change in the law, but are nonetheless very important to policy. For example, revisions to estimates of the neutral rate of about 1 percentage point would dramatically change monetary policy. Blinder argued that changes at the Fed that require legislation should be very rare, and only occur when the Fed is consistently failing to meet its mandate. Changes to the Fed's framework that do not require Congress's approval should be done continuously rather than as some discrete event that invites political participation.

Alice Rivlin agreed with Blinder that the Federal Reserve should constantly be revising its policy framework; she believes it currently operates as such. In fact, she noted that the Fed's focus on monetary policy is proportionally greater than its focus on other policies, in part because the Federal Open Market Committee (FOMC) is a receptive audience. The framework for fiscal policy does not generate the same degree of attention because there is no equivalent audience. She noted that one of the dangers of formalizing a process around the Fed's framework is to overemphasize monetary policy and deemphasize things like fiscal policy that deserve more attention.

Frederic Mishkin noted his conflicting views vis-à-vis a formal review of the monetary framework. On the affirmative side, he referenced the successful push for adopting inflation targeting as part of the Fed's framework and its emphasis on the role of institutions. Good institutional governance and less reliance on individuals was a crucial argument for the inflation-targeting framework. He noted that there were several episodes in which the Fed did not formally consider its framework and instead relied on individuals, which led to major policy mistakes. For example, he noted the policies of former Fed chair Arthur Burns. The Burns Fed overemphasized employment as part of its mandate, and as a result facilitated the Great Inflation of the 1970s. The lesson of these policies, Mishkin argued, is to not have overconfidence about the natural rate of unemployment but to instead focus on inflation and what it might tell policymakers about the natural rate of unemployment. A second example was then-Fed chairman Alan Greenspan's opposition to institutional transparency at the Fed. Mishkin described a past FOMC meeting at which inflation targeting was discussed. Then-San Francisco Fed president Robert Parry, who was considered more conservative, and then-Fed vice chair Janet Yellen, who was considered more liberal, agreed on the need for an inflation target. However, Chair Greenspan shut down the conversation and forcefully opposed an explicit

2 percent inflation target. Both these examples illustrate why a formal review process could improve governance and policy at the Fed.

However, Mishkin also agreed with commenters Stanley Fischer and Jean Boivin about the importance of political economy. Though a formal review process involving outside actors might improve the Fed's democratic accountability, there are some questions about how such a process would be implemented. For example, Mishkin cited congressional proposals to "audit the Fed." These proposals, he argued, are really about ending the Fed entirely, because they would allow Congress to counteract specific interest rate decisions by the central bank, hindering its independence. Former senator Ron Paul, one of this policy's main proponents, has explicitly written that he wants to end the Fed. Such a proposal shows the risks of a potential "audit" of the Fed's framework. Mishkin noted that certain countries with parliamentary systems, such as Sweden, have successfully used such a process by focusing on the framework rather than the central bank's day-to-day decisions. Mishkin observed that other countries that have implemented a successful review of the framework, such as Canada, have made the process more technocratic. These countries have avoided "big picture" issues and therefore sidestepped politics. He pointed out that a regular review process, conducted perhaps every three years, would make sense, but that implementing it effectively without criticizing current policy or interfering with the Fed's independence would require subtlety and a focus on the details.

Philipp Hartmann praised the authors for their paper and overall contributions to important debates over the years on central banks. Although the European Central Bank (ECB) is currently not conducting a review of its monetary policy strategy, he thought he would bring some of the euro area practices into the debate, so participants could make comparisons with the United States. First, he noted that the ECB undertakes regular and relatively frequent internal reviews of specific monetary policy instruments, programs, or frameworks. Second, he noted a major review of the ECB's monetary policy strategy in 2003, which is described in detail in the Hartmann and Smets paper presented earlier at the panel. This review led to important clarifications and changes with respect to the ECB's quantitative definition of price stability, the role of the "monetary pillar" in its strategy, and the structure of the ECB president's introductory statement at monetary policy press conferences. *Ex ante*, these two types of reviews do not actively involve parties outside the Eurosystem, except when a consultant is hired from academia or another central bank. Nevertheless, apart from undertaking its own analyses, the ECB collects large volumes of materials from academia and takes the views of outside

commentators and stakeholders into account. Ex post, the results of such reviews may then be publicly announced.

In what concerns the possibility of actively involving non-central bank stakeholders or the general public in reviews of the ECB's monetary policy framework, Hartmann referred to specificities in Europe compared with mature countries with federal systems, such as the United States. First, it would be very complicated to involve stakeholder groups from the 19 different euro area countries. Moreover, it is very challenging to communicate on such complicated matters with the public and the media of all these countries. The existence of different languages, policy cultures, and traditions often leads to the same message being interpreted differently across the euro area countries.

Regarding the metrics to be used in such a review of the monetary policy framework, Hartmann noted that a standard central bank loss function, as widely used in the academic literature and also in the paper by Fuhrer and his colleagues, could not be a primary one for the ECB's case. This is due to the ECB's single mandate with a primary objective of price stability, which allows it to pursue other objectives such as full employment only in lexicographic order.¹ Hartmann discussed the issue with the Harvard University economist Benjamin Friedman, but ultimately felt that the ECB could use such loss functions at best as a secondary input. Hartmann also noted that for central banks in general, it now seems particularly important to discuss the role of financial stability in their monetary policy frameworks. For a review involving outside stakeholders, however, this might significantly enlarge the number of parties.

Finally, Hartmann referred to an effect of a public review that had not yet been mentioned in the discussion (but perhaps hinted at by the last remark of Frederic Mishkin). The fact that aspects of the existing monetary policy framework are being questioned could have implications for the transmission of current policy and, ultimately, policy effectiveness. This transitional risk would need to be actively managed through communication until all uncertainties are resolved again, probably without being able to entirely exclude the possibility of unintended adverse effects.

Olivier Blanchard recommended the International Monetary Fund's internal review process as a potentially useful model. This process is conducted internally but by independent reviewers, it relies on academics, and

1. A similar point is made by John Driffill and Zeno Rotondi, "Monetary Policy and Lexicographic Preference Ordering," CEPR Discussion Paper 4247 (Washington: Center for Economic Policy Research, 2004).

it focuses on specific aspects of the IMF's framework rather than on the overall mandate. Blanchard suggested that this process works reasonably well and sidesteps political controversy.

Eric Rosengren responded to the commenters and other discussants by noting that their paper was not focused on changes to the monetary policy framework that would involve changes to the Federal Reserve Act. Many changes in the framework had occurred in the past without changes to this law, and these types of changes were their focus in the paper. Rosengren noted that the framework can be broader than just focusing on goals. For example, the statement on longer-run goals and policy strategy released by the Fed every January highlights the loss function, operating instruments, and transparency. He also emphasized that changing the framework without involving outside parties or being transparent does not eliminate the political economy problems raised by the commenters and discussants. He therefore argued that public discussion does have some merit as part of a review process. For example, many of the Fed's transparency initiatives, including those highlighted by Frederic Mishkin, were considered inconceivable in the past. In fact, the Fed would have been better off if it had moved earlier toward greater transparency. How the review process itself should be structured would be up for debate, but making it regular and structured to ensure that changes to the framework do indeed occur seems like the vital point.

Jeff Fuhrer noted that the convening process that would be a part of the review of the monetary policy framework that he and his colleagues propose in their paper is not supposed to definitively set the Fed's framework. Instead, it is supposed to be used to obtain input as the FOMC considers alternatives. The Fed would be responsible for determining the agenda and inviting attendees, and it could be selective in doing so. The paper also does not argue for having the process occur around elections, and the authors do not believe the review process should be politicized. Finally, Fuhrer noted that the review process conducted by the Bank of Canada could not be directly transplanted and used in the United States, due to the different relationship between the government and the central bank in the two countries. In Canada, the government is more directly involved in the process. The authors suggest in their paper that the process should include input from government experts, while being careful to maintain the Fed's independence in deciding on its framework's technical aspects.

*Symposium on Monetary Policy
at the Effective Lower Bound*

Monetary Policy at the Effective Lower Bound: Less Potent? More International? More Sticky?

ABSTRACT This paper discusses whether monetary policy at the effective lower bound (ELB) is less effective, generates greater international spillovers, or is “stickier” than conventional monetary policy. It argues that monetary policy at the ELB can be potent and that there has thus far been no convincing evidence that it has greater international spillovers through capital flows and exchange rates than comparable adjustments in interest rates. It may be more challenging to raise rates off the ELB than to raise rates from higher levels—possibly due to counterbalancing effects through the exchange rate—although there are only anecdotes to support this stickiness rather than any formal, empirical evidence.

During the 2008 global financial crisis, many advanced economies lowered their policy interest rates to their effective lower bounds (ELBs). In some countries, these interest rates are still there. In the future, there is a good chance that many central banks will operate at the ELB more often, especially given the fall in the global neutral interest rate (r^*) and the high probability that the next slowdown will come before interest rates are raised to levels from which they could be lowered enough to provide a substantial stimulus. Understanding how monetary policy at the ELB is different from “conventional” monetary policy is therefore critical for thinking about monetary policy in the future.

This paper explores three ways in which monetary policy at the ELB may differ from more “conventional” monetary policy—defined as primarily

Conflict of Interest Disclosure: The author did not receive any financial support from any firm or person for this paper or from any firm or person with a financial or political interest in this paper. She is currently not an officer, director, or board member of any organization with an interest in this paper.

consisting of changes in the central bank's main policy rate. First, it asks whether monetary policy at the ELB is less effective, making it difficult for the central bank to meaningfully support the economy. Second, it asks if monetary policy at the ELB has larger international spillovers—through larger effects on the volume and volatility of capital flows or on exchange rates. Third and finally, it discusses whether the ELB is “sticky,” in the sense that adjustments in monetary policy around the ELB generate disproportionate feedback effects that make it harder to tighten this policy.

Each of these questions addresses concerns that have been raised about monetary policy at the ELB—concerns that could provide reasons to adjust monetary frameworks in order to reduce the probability of reaching the ELB in the future. I do not venture into this broader debate, but simply focus on whether these arguments for concern about the ELB are valid. My attempts to answer these questions are far from definitive; if anything, the discussion suggests the need for more careful analysis of these important questions.

The preliminary evidence, however, suggests that these concerns about the ELB may be overstated. Monetary policy made with “unconventional tools” can be effective at the ELB, assuming there are no political constraints on using these tools. There is also little convincing evidence to date that monetary policy at the ELB has greater international effects than would occur through comparable adjustments in interest rates on the volume or volatility of capital flows, or on exchange rates. Whether raising interest rates after being at the ELB is more challenging than raising rates from more normal levels is an open question—and one that has been even harder to answer, given the small number of countries that have thus far successfully exited the ELB. In fact, all these questions are difficult to answer because any changes in the effectiveness and channels of monetary policy since the 2008 crisis could reflect changes related to operating at the ELB—or the many other structural changes in the global economy that have occurred over this period. On a more positive note, if the current improvement in global growth and inflation continues, there should be more examples of countries exiting the ELB and therefore more evidence to help answer these questions.

I. Is Monetary Policy at the ELB Less Potent?

One of the concerns most frequently cited about central banks operating at the ELB is that they will not have sufficient ammunition to provide a stimulus in response to the next slowdown. In the decades before the 2008 crisis,

adjustments to interest rates were the primary tool used by central banks to stimulate the economy. For example, in the United Kingdom, there were eight business cycle slowdowns from 1980 through 2010, during each of which the Bank of England reduced interest rates by an average of 3.75 percentage points.¹ In the United States, there were seven business cycle slowdowns over the same period, and the U.S. Federal Reserve reduced interest rates by an average of 4.59 percentage points. If interest rates are at the ELB, then these types of reductions will not be possible. If central banks cannot provide stimulus through other mechanisms, and if fiscal policy is constrained due to high deficits or political constraints, countries could face periods of slower and more volatile growth. This is a key concern behind arguments to adjust inflation targets and reduce the probability of being at the ELB.

One challenge to this line of reasoning, however, is that reductions in interest rates are not the only channel by which central banks can provide stimulus. The global financial crisis of 2008 and the subsequent prolonged recovery, combined with monetary policy at the ELB in many advanced economies, prodded many central banks to experiment with other forms of stimulus. Some were more potent than others, and the effectiveness of many is still widely debated.² Some policies that seemed to be effective at the time may have worked due to specific characteristics of the crisis period (such as poor market liquidity), so that they would be less effective in stimulating the economy during less stressed periods.

Nonetheless, my experience on the Monetary Policy Committee (MPC) at the Bank of England convinced me that these unconventional tools can be effective, even outside crisis periods. In fact, central banks can stimulate the economy in a number of ways when at the ELB—even if most central bankers (myself included) would prefer to return to an era when adjustments in monetary policy were made primarily through adjustments in interest rates.

More specifically, before I joined the MPC in 2014, the Bank of England had embarked on several rounds of quantitative easing from 2009 to 2012.³ Most studies of this experience suggest that this provided a meaningful stimulus to the U.K. economy. For example, Martin Weale and Tomasz

1. For the details of these calculations, see Forbes (2015).

2. For a summary of the evidence, and more skeptical view of the effectiveness of asset purchases in the United States, see Greenlaw and others (2018).

3. For information on these programs and different estimates of their effects, see Joyce, Tong, and Woods (2011) and Haldane and others (2016).

Wieladek (2016) estimate that, on average over this period, asset purchases worth 1 percent of GDP boosted U.K. GDP by about 0.25 percent. This estimated impact of asset purchases worth 1 percent of GDP is roughly equivalent to the impact of a reduction of 25 basis points in the Bank Rate (the policy interest rate set by the Bank of England) on U.K. GDP—according to very rough rules of thumb. Total asset purchases as of 2012 were £375 billion, equivalent to about 20 percent of the U.K.’s GDP at the time, which would imply a boost to GDP of about 5 percent—the equivalent of reducing the Bank Rate by 5 percent. Of course, these are only rough estimates and do not incorporate the many other factors that were affecting the economy at this time; but even if they are off by half, they still suggest that monetary policy at the ELB was able to provide a meaningful stimulus.

I admit, however, that I was always skeptical of these types of estimates, especially given that some of the large estimated benefits from quantitative easing (QE) over this period likely arose from its ability to improve the liquidity and functioning of stressed financial markets. Would QE provide a similar stimulus when markets were functioning well? This was a critical question for the MPC in 2016, after the U.K. voted to leave the European Union (the “Brexit” vote), and most surveys suggested that economic growth would slow sharply. The policy interest rate was near what was then believed to be the ELB, and the majority of the MPC’s members wanted to provide more support for the economy than could be achieved by lowering interest rates to the ELB. Were there other monetary policy tools that could provide a meaningful stimulus at this time?

In August 2016, the majority of the MPC’s members voted for a four-pronged easing program: to reduce Bank Rate by 25 basis points; to purchase an additional £60 billion in government bonds; to purchase £10 billion in corporate bonds; and to start a Term Funding Scheme (TFS) that would provide contingent and targeted funding for banks to encourage them to pass on the reduction in Bank Rate to borrowing costs for businesses and households. The Bank of England’s staff simulated the effects of this four-pronged package under model assumptions that the asset purchase programs would provide some stimulus, but less than the average effects from earlier rounds of QE. The reduction in Bank Rate was expected to account for less than one quarter of the total stimulus from the package—with most of the stimulus resulting from the additional purchases of government bonds. If interest rates were not at the ELB, the MPC would have had to lower Bank Rate by roughly 100 basis points to get the same estimated aggregate effect on GDP growth and inflation.

Table 1. U.K. Financial Market Indicators after the August 2016 Stimulus

<i>Indicator</i>	<i>Cumulative change between August 3, 2016, and:</i>	
	<i>August 4, 2016</i>	<i>September 30, 2016</i>
U.K. 10-year gilt yield (percent)	–17	–11
Sterling investment-grade corporate bond spreads (basis points)	–10	–17
Sterling high-yield corporate bond spreads (basis points)	–8	–20
FTSE All-Share (index)	1.5	4.2
U.K.-focused companies' equity prices (index: August 3, 2016 = 100)	0.9	2.2
Sterling Exchange Rate Index (January 2005 = 100)	–1.3	–2.9

Source: Bank of England (2016, box on 2–3).

Although it is impossible to estimate the exact effects of this program, and especially the effects of its individual components because their joint announcement may have amplified their impact, the available evidence suggests that the asset purchase programs and TFS provided a meaningful amount of stimulus to the economy. In fact, they appear to have provided an even larger boost than expected. For example, the Bank of England analyzed financial market data in the period after the package was announced and concluded that “if anything, the impact was slightly greater than had been anticipated.”⁴ Although the reduction in interest rates had largely been priced in before the announcement of the four-pronged package, table 1 shows that other market prices (which primarily reflect the impact of the “unconventional” components of the package) adjusted in ways that would support the economy. The sterling Exchange Rate Index depreciated, and the spread on 10-year gilt yields and various corporate bonds fell. The *Financial Times* Stock Exchange (known as FTSE) All-Share Index and equity prices for U.K.-focused companies increased. Funding costs of U.K. banks also decreased (likely supported by the TFS). All these price adjustments are in the same direction that traditionally follows an unexpected easing in monetary policy, suggesting that the unexpected and unconventional components of the four-pronged package also acted to ease financial conditions.

4. See Bank of England (2016, box on 2–3).

Although it is difficult to directly connect these developments to changes in the real economy, these movements in financial indicators are key channels by which a monetary stimulus traditionally supports economic growth and inflation. Data for the subsequent year also suggest that the package supported the economy in ways that would normally occur from easing monetary policy when not constrained by the ELB. For example, retail interest rates for households and businesses fell.⁵ Lena Boneva, Calebe de Roure, and Ben Morley (2018) estimate that the corporate bond purchase program reduced the spreads of eligible bonds by 13 to 14 basis points (compared with foreign bonds issued by the same set of firms), and boosted values for other U.K. assets that were not eligible for the purchase program.

All these estimates are imprecise; it is impossible to know the counterfactual, and different monetary tools will undoubtedly have different effects in different economies (as well as different effects at different times in the same economy). Nonetheless, they suggest that central banks do have tools available to stimulate the economy other than lowering interest rates. As a result, central banks are not necessarily “out of ammunition” just because they are at their ELB. Of course, there are also constraints on these types of unconventional policies. For example, asset purchases will be limited by the size of the relevant asset market, and political constraints could limit the ability of some countries to use these types of unconventional tools (such as in the United States). Nonetheless, the fact these tools are available, and that they can be potent even when markets are functioning well, should alleviate some of the concerns about the potency of monetary policy at the ELB.

II. Is Monetary Policy at the ELB More International?

A second common concern about monetary policy at the ELB is that it works through different channels than traditional monetary policy. There are a range of ways this could occur. For example, if monetary policy at the ELB is done more through forward guidance, then it could have larger effects on the longer end of the yield curve (relative to those on short-term rates) than occurs with adjustments in policy rates. Or, if monetary policy at the ELB is adjusted more through asset purchases, it could have greater effects on specific asset prices and therefore have different distributional implications. Here, however, I focus on two ways in which monetary

5. Ibid.

policy at the ELB could have greater effects through international channels, and thereby generate larger global spillovers.⁶ More specifically, do interest rates at the ELB in advanced economies stimulate excessive volumes or volatility in capital flows to other countries? And when countries are at the ELB, do adjustments in monetary policy have greater effects on the exchange rate?

II.A. Capital Flows around the ELB


Prominent policymakers in emerging markets have complained that QE and near-zero interest rates in major advanced economies stimulate excessive capital flows to emerging markets—which have been described in colorful terms as “currency wars” by Guido Mantega (Brazil’s former finance minister) and as a “monetary tsunami” by Dilma Rousseff (Brazil’s former president). They argue that these “surges” of capital flows can lead to challenges, such as elevated asset prices and currency appreciation, and also increase vulnerabilities from the inevitable “sudden stop” when the abundant capital inflows reverse. There is no doubt that volatile capital flows create challenges for emerging markets—especially those with weaker institutions and financial systems. There is also evidence that monetary policy in advanced economies is an important driver of global capital flows, although most research suggests that it is only one of a number of factors driving capital movements (with other variables, such as global risk, often being more important).⁷ The key questions, however, are if interest rates near the ELB in advanced economies tend to aggravate the surges in capital flows to emerging markets, and if they exacerbate excess volatility in capital flows.

It is difficult to test these hypotheses formally, partly due to the limited episodes at which interest rates in major economies have been near the ELB, and partly because there is no clear benchmark for determining the optimal level of capital flow volumes or volatility. Nonetheless, as an informal test, it is useful to look at recent patterns in capital flows to assess if they appear to have been elevated or more volatile during the last decade, when interest rates in advanced economies have often been at the ELB.

6. Monetary policy at the ELB could also generate international spillovers by affecting foreign market prices. For analyses of whether these spillovers differ when monetary policy is conducted through QE or adjustments in interest rates, see Curcuru and others (2018). Most research finds no consistent differences in the spillovers from conventional and unconventional monetary policy.

7. See Forbes and Warnock (2012) and Rey (2013).

Figure 1. Capital Inflows to Emerging Markets as a Percentage of Emerging Market GDP, 2000–2016^a



Sources: For capital flow data, the Emerging Market Capital Flows database of the Institute of International Finance, May 2018; for GDP and interest rate data, the International Monetary Fund's *World Economic Outlook* database, April 2018.

a. Capital inflows are nonresident capital flows (changes in liabilities) to emerging markets as a percentage of emerging market GDP. These include foreign direct investment, portfolio investment, and other investment. The lighter shaded area excludes foreign direct investment. The average interest rate is the average of the policy rate for the U.S. and U.K. and discount rate for the euro area and Japan in each year.

Figure 1 shows a first piece of evidence: gross global capital inflows to emerging markets as a percentage of emerging market GDP from 2000 through 2017.⁸ The figure also shows the average interest rate set by four major central banks (the U.S. Federal Reserve Board, European Central Bank, Bank of England, and Bank of Japan) over this period. During the last decade, when interest rates have been around the ELB in these major economies, it is hard to make the argument that capital inflows to emerging markets have been “excessive”—at least compared with precrisis patterns.

8. Capital inflows are annual nonresident capital flows (changes in liabilities) to emerging markets, based on the Emerging Market Capital Flows database of the Institute of International Finance, May 2018. GDP and interest rate data are from the IMF's *World Economic Outlook* database, April 2018. The interest rate is the annual average of the policy rate for the U.S. and U.K. and the discount rate for the euro area and Japan.

More specifically, gross capital inflows to emerging markets averaged 4.0 percent of emerging market GDP from 2010 through 2017, below the five-year average before the crisis (of 5.2 percent from 2003 to 2007). Even in 2010, when capital flows to emerging markets rebounded as many economies experienced rapid recoveries, capital inflows never reached their peak of 2006. These patterns even continue to hold for the more volatile capital flows that are more tightly linked to monetary policy (shown in the lighter shading in figure 1).⁹ These more volatile capital flows only averaged 1.9 percent of emerging market GDP from 2010 to 2017, as compared with 2.6 percent from 2003 to 2007. Granted, the volume of capital inflows to emerging markets may still be elevated relative to optimal levels, and may be large enough to create challenges for many countries, but the period of very low interest rates in major economies does not appear to have accelerated these flows relative to when interest rates were higher.

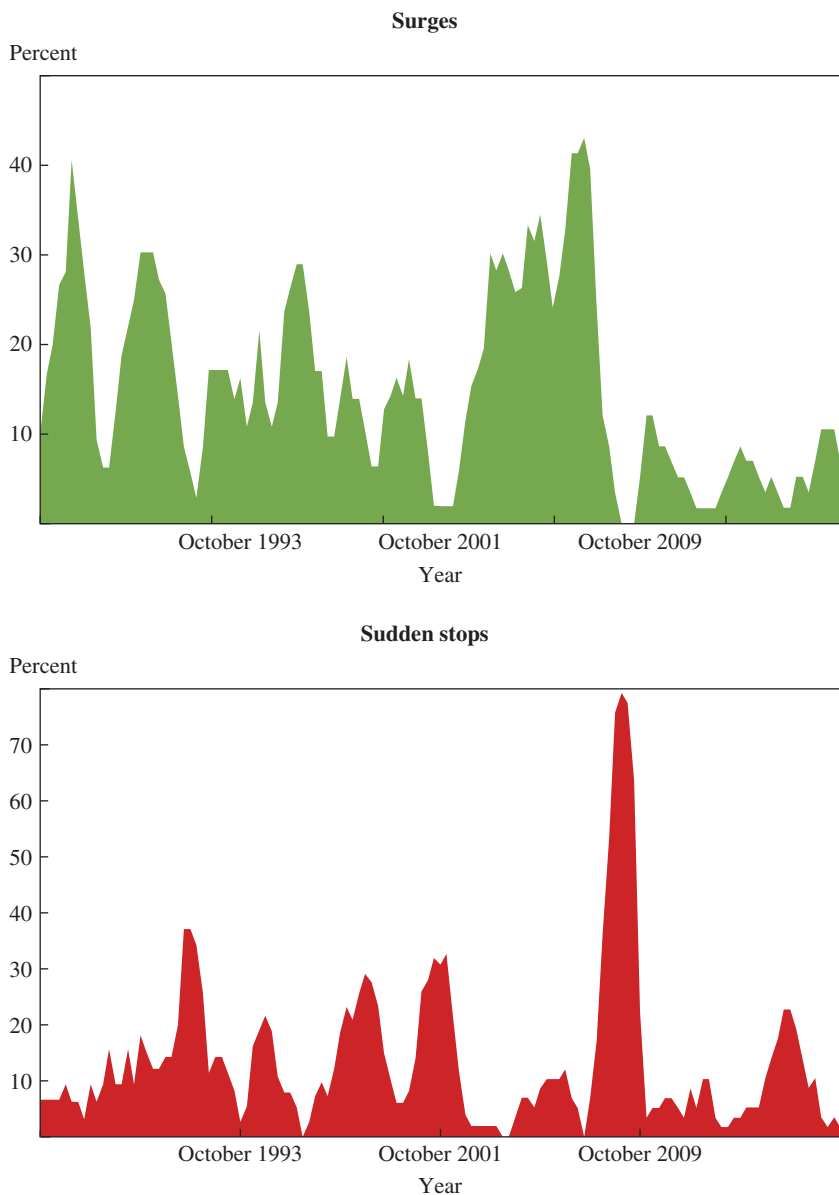
Many emerging markets, however, are more concerned about the volatility in capital inflows than about the volumes, and especially the occurrence of “sudden stops” and “surges” of capital inflows. Therefore, to assess whether capital flows to emerging markets are more volatile around the period of interest rates at the ELB in advanced economies, I use the technique developed by Forbes and Francis Warnock (2012) to calculate the occurrence of surges and sudden stops in capital flows from abroad, based on whether there are unusually large increases or decreases in foreign capital flows relative to historic country-specific trends. More specifically, this methodology uses gross quarterly capital inflow data and defines a “surge” as a period that includes an increase in year-over-year changes in four-quarter gross capital inflows that is more than 2 standard deviations above the historic average for at least one quarter. A “sudden stop” is defined symmetrically, requiring a decrease in gross capital inflows that is more than 2 standard deviations below the historic average.¹⁰

Figure 2 shows the share of the sample that experienced surges and stops from 1985 through 2017, using updated data and a slightly different sample from when this methodology was introduced by Forbes and

9. More volatile capital flows are defined as portfolio flows and “other” investment flows, the latter of which are largely bank flows. They exclude foreign direct investment.

10. Each surge and stop episode is defined as lasting for all consecutive quarters for which the year-over-year change in annual gross capital flows is more than 1 standard deviation above or below the historical average. The length of each episode is required to be greater than one quarter. Data are primarily from the IMF’s *International Financial Statistics*, supplemented with country sources. See Forbes and Warnock (2012) for details.

Figure 2. The Incidence of Surges and Stops of Capital Flows from Abroad, October 1985–October 2017^a



Sources: The underlying data are primarily quarterly capital flow data from the International Monetary Fund's *International Financial Statistics*, supplemented with country sources.

a. This figure shows the percentage of countries in the sample experiencing a surge or sudden stop in capital inflows from abroad each quarter. The methodology is discussed in the text and is described in more detail by Forbes and Warnock (2012).

Warnock (2012). The figure's top panel does not suggest any increase in the share of countries experiencing surges of capital inflows during the period of interest rates near the ELB in advanced economies. In fact, capital flow surges are even less frequent since 2009 than during the 1990s, and much less frequent than during the period of relatively high interest rates preceding the 2008 crisis. The bottom panel also suggests that there was not an unusual number of sudden stops. Although the incidence of sudden stops increased around the "taper tantrum" in 2013–14 (to peak at about 20 percent of the sample), this was not unusual when compared with the cycles experienced over the 20 years before the 2008 crisis—a period when interest rates in major economies were not near the ELB.

II.B. Exchange Rate Sensitivity around the ELB

Even if interest rates around the ELB do not seem to have generated an unusually large volume or increase in the volatility of capital flows, adjustments in monetary policy around the ELB could still be generating unusual international spillovers through their effects on exchange rates. This is another angle of the concerns about "currency wars"; unconventional monetary policy could have greater effects on the exchange rate than a comparable stimulus provided through changes in policy interest rates. (In fact, a larger effect on the exchange rate could mute the subsequent adjustments in capital flows.) These concerns were serious enough that they were the topic of a Group of Seven meeting in 2013 and were discussed at the group's resulting special statement establishing ground rules to address the potential effects on exchange rates of different monetary policy tools.¹¹ The research of Christopher Neely (2015) is frequently cited as evidence supporting these concerns; it finds that the U.S. Federal Reserve's announcements of QE had larger effects on the dollar than non-QE announcements. This analysis, however, does not control for the fact that the average stimulus provided by the QE announcements was larger than that by the non-QE announcements.

Nonetheless, there are reasons why unconventional monetary policy could have larger effects on exchange rates than a comparable stimulus provided by adjusting interest rates. Unconventional monetary policy appears to work more through the term premium (and therefore long-term securities), whereas conventional monetary policy works more through short-term rates (and therefore money market rates). Unconventional monetary

11. See Group of Seven (2013).

policy may be interpreted as a longer-term commitment to a path of monetary policy over a longer period, whether in the form of a commitment to asset purchases over an indefinite period or state-contingent forward guidance. Any of these channels could cause a monetary stimulus at the ELB to have a larger effect on the exchange rate than more conventional changes in policy interest rates. This could, in turn, generate greater spillovers and challenges for any emerging markets that subsequently experienced sharp currency appreciations.¹²

Whether monetary policy at the ELB has a larger effect on exchange rates is an important question—but one that is extremely difficult to identify and test. Several papers (such as the one by Glick and Leduc 2015) have tried to assess one piece of the puzzle: if exchange rates respond differently to changes in short-term than long-term rates. These papers generally find no significant difference, although identification is a challenge, given that movements in short-term rates tend to correspond to movements in long-term rates. Several studies (Glick and Leduc 2015; Curcuru 2017; Ferrari, Kearns, and Schrimpf 2017) have also found that the responsiveness of the dollar to U.S. monetary policy announcements or U.S. monetary policy surprises rose after the 2008 crisis. This could have resulted from structural changes not directly related to the form of monetary policy, however, which may have made the dollar more responsive to all forms of monetary policy over the last decade.

Stephanie Curcuru and others (2018) and Jan Hatzius and others (2017) take a different approach—and find somewhat different (albeit not contradictory) results. Curcuru and others (2018) tackle the identification challenge by assuming that asset purchases affect the term premium (and therefore longer-term bond rates), whereas conventional monetary policy only affects short-term rates. Based on this assumption, it finds that QE does not generate significantly larger spillovers (in terms of dollar movements, as well as other financial market measures) than conventional monetary policy. Instead, it finds evidence of the opposite: that a given increase in expected interest rates has more than double the effect on the dollar than the same increase in the term premium (which is assumed to be accomplished through asset purchases). Hatzius and others (2017) reach similar conclusions in an analysis that regresses exchange rates on components of the yield curve and also assumes a larger effect of asset purchases on the

12. Brainard (2017) has an excellent discussion of these issues. It models the different spillovers from adjusting interest rates versus asset purchases, and shows how the spillovers will vary based on the country's exchange rate regime and output gap.

term premium. Two challenges to these studies, however, are the restrictiveness of the identification assumptions and the lack of a broader understanding of what has been causing movements in the term premium over the last decade.


All in all, whether unconventional monetary policy used at the ELB has a larger effect on exchange rates than a comparable adjustment in monetary policy made through interest rates is still an open question—and a prime target for future research. Although there are valid arguments why monetary policy at the ELB could have larger international effects through exchange rates, as well as through the volume and volatility of capital flows, there is little convincing evidence to date that this has occurred.

III. Is Monetary Policy at the ELB More Sticky?

A closely related issue is whether exchange rate adjustments at the ELB make it more difficult to raise interest rates and exit the ELB. More specifically, does the first increase in the policy interest rate from the ELB—or even providing guidance on the intent to do so—cause a larger exchange rate appreciation than would occur for a comparable increase in interest rates from a higher starting point? Because appreciations tend to reduce import price inflation and headline inflation (especially when the appreciation corresponds to a monetary policy shock, as shown in Forbes and others 2018), the subsequent drag on inflation could make it more difficult to justify an increase in interest rates. If the appreciation caused by forward guidance of a forthcoming exit from the ELB were large enough, it could even prevent the exit from the ELB. Or, if a large appreciation were caused by the first increase in interest rates off the ELB, it could make it more difficult to raise interest rates again—leading to an unusually slow tightening cycle. In other words, does excessive exchange rate sensitivity around the ELB make interest rates more “sticky”?

Although there has been no empirical work assessing these effects (to the best of my knowledge), my experience at the Bank of England suggests that the ELB may in fact be “sticky.” More specifically, when I started on the MPC in July 2014, the MPC had recently provided guidance that raised expectations that Bank Rate would soon be increased—the first increase in the policy interest rate since 2009. The top panel of figure 3 shows market expectations for U.K., U.S., and euro area interest rates about that time, indicating that investors expected this increase in U.K.

Figure 3. Expected Interest Rates and the Sterling Exchange Rate in 2014^a



Source: Based on data from the Bank of England (2014).

a. Curves of market expectations for interest rates are estimated using instantaneous forward overnight index swap rates in the 15 working days to August 6, 2014. The Sterling Exchange Rate Index is based on January 2, 2007 = 100.

rates to occur within the next six months.¹³ Sterling had also been appreciating sharply (the bottom panel of figure 3)—with the exchange rate index already up about 12 percent by October 2014 (from its recent low in March 2013). This appreciation would continue over the next few months (peaking at over 15 percent) and have a number of effects on the economy. For example, it contributed to tighter financial conditions and slower growth in net exports—both of which would be a drag on GDP growth and therefore inflation in the future. The currency appreciation was also expected to reduce import prices and Consumer Price Index (CPI) inflation. Using the Bank of England's rough rule of thumb at the time, a 12 percent appreciation would be expected to reduce the level of import prices by about 11 percent and the CPI by over 3 percent over the next few years—very large effects.¹⁴


Moreover, these effects of sterling's appreciation on inflation were expected to have first-order importance for the appropriate path for monetary policy. Figure 4 shows the results of a simulation I did at that time (in Forbes 2014), using the more complicated dynamic stochastic general equilibrium model used by the Bank of England to capture the full effects of the appreciation, combined with data on the economy that existed in October 2014.¹⁵ The MPC inflation forecast (the red line) incorporated the effects of sterling's sharp appreciation to date. This forecast suggested that inflation would remain below the 2 percent inflation target over the next year, implying that interest rates would not need to be tightened as much or as quickly as suggested by the market curve. In contrast, the simulated path of inflation (the black line) assumes that sterling did not appreciate and instead remained at its 2013:Q1 level. The simulation predicts that inflation would have been well over the 2 percent target for the next few years.

13. Market expectations are measured by instantaneous forward overnight index swap rates from Haver and the Bank of England.

14. The rule of thumb at the time was that the pass-through from movements in the sterling Exchange Rate Index was 90 percent to import prices and then 30 percent to headline CPI (so that a 10 percent depreciation corresponds to a 9 percent increase in the level of import prices and 3 percent increase in the level of the CPI). This rule of thumb was subsequently adjusted so that the pass-through to import prices was reduced to 60 percent (and there was no change in the second stage of the pass-through).

15. This simulation compares the path of CPI inflation predicted in the latest *Inflation Report* relative to a situation in which the exchange rate had remained at its 2013:Q1 level and there had been no other shocks or changes in policy. The shift in the exchange rate is assumed to result from an exogenous exchange rate shock, and the shaded bands around the black line capture the range of outcomes based on different assumptions for the persistence of the appreciation.

Figure 4. U.K. Inflation Forecast and Simulated Inflation Path Assuming No Sterling Appreciation, 2013:Q1–2016:Q1



Source: Forbes (2014).

Note: CPI = Consumer Price Index. The shaded area is COMPASS's predictions of CPI inflation if the exchange rate remained at its 2013:Q1 level, under assumptions of different degrees of persistence of the appreciation. The appreciation is assumed to be exogenous, with no other changes in policy and no other shocks.

Although it is impossible to know what the MPC would have decided in this counterfactual situation, it is likely that interest rates would have been lifted off the ELB sooner if the exchange rate had not appreciated so sharply and substantially dampened the expected path of inflation. Instead, exit from the ELB was delayed for an extended period—and the next move in U.K. interest rates was actually down (after the Brexit vote) instead of up. U.K. interest rates were only lifted above 0.5 percent in August 2018—four years after this period of serious consideration of exiting from the ELB. Granted, much of this delay was due to other subsequent shocks (such as the sharp decline in commodity prices in 2015 and uncertainty about the Brexit vote), but the initial move off the ELB would likely have occurred before these additional shocks if sterling had not appreciated so sharply when interest rates were at the ELB.

Of course, sterling would still have appreciated if the expected 2014 increase in interest rates occurred at a level of interest rates above the ELB. The key question is whether the appreciation during this episode was larger

than it would have been if rates were not at the ELB. This is a more difficult question to answer, but a comparison with historic episodes suggests that sterling was more sensitive than would normally be expected. More specifically, an increase of 25 basis points in interest rates is usually assumed to correspond to a sterling appreciation of about 0.25 to 1 percent.¹⁶ This band reflects historic averages as well as model estimates, and suggests that the exchange rate movement in 2014 and early 2015 was meaningfully larger than would be expected based on expected changes in monetary policy.

There are several reasons why exchange rates could be more sensitive to changes in monetary policy as countries attempt to move away from the ELB. First, the initial movement away from the ELB is likely to occur through forward guidance about the near term, especially because central banks tend to be even more cautious than usual and not to want to create surprises when raising interest rates for the first time in an extended period. Forward guidance—especially if focused on imminent changes in policy—would likely have a large effect on short-term interest rates, which may be more closely linked to exchange rate movements. Second, raising interest rates off the ELB after an extended period of monetary stimulus may be seen as signaling a major shift in policy, which will affect not only short-term rates but also the whole market curve, and in a stronger way than normally occurs. Similarly, it could be interpreted as showing a shift in confidence about the economic outlook, similar to the “Delphic effect” found by Emi Nakamura and Jón Steinsson (2018). Finally, the relative size of the change in interest rates when starting at such a low level may matter; for example, raising interest rates by 25 basis points is a doubling of interest rates if moving from an ELB of 0.25 percent, but only about a 10 percent increase if moving from a level of 2.0 percent. The relative increase in carry costs or other prices related to the increase of 25 basis points in interest rates could cause disproportionate effects on currency trading and other pricing.

If there is a “stickiness” to raising interest rates from the ELB, assessing the magnitude of this effect is challenging. Not only are there limited examples to assess, but any such effects will also undoubtedly differ across countries and over time. Factors that would determine the magnitude of any such stickiness include whether other countries are also tightening monetary policy at the same time; the sensitivity of the currency to interest

16. The lower estimate reflects the rule of thumb from the Bank of England’s COMPASS model under a set of standard assumptions, described by Burgess and others (2013). The higher number is estimated by Forbes and others (2018).

rates; and the sensitivity of inflation, financial conditions, and exports to exchange rate movements.

Given these challenges, it is not surprising there has not yet been a formal study of any of these channels that could make adjusting interest rates at the ELB sticky. There are, however, numerous anecdotes from countries other than the U.K., which would support the hypothesis that it has been harder to exit from the ELB than expected. For example, as of June 2017, despite seven years of solid global economic growth above 3 percent, no advanced economy (other than Hong Kong and the United States) had been able to maintain an increase in interest rates since 2011. In fact, at that time, nine countries that had tried to “lift off” and raise interest rates after 2009 had then reversed the rate increase (see Forbes 2017). Even the U.S., the advanced economy able to raise interest rates the most from its ELB, was only able to do so after a very slow start; it was a full year between the date when the U.S. Federal Reserve first raised interest rates above the ELB and its next rate increase. Granted, the simultaneous challenges for so many countries in exiting the ELB may also reflect common global developments—such as a decline in the global equilibrium interest rate. Nonetheless, it also may reflect additional challenges and a stickiness in raising interest rates from the ELB.

IV. Conclusions

Research on monetary policy at the ELB is only in its infancy. An empirical analysis of whether monetary policy functions differently at the ELB is complicated by the fact that the last decade when many advanced economies were at the ELB coincided with many other structural economic changes—changes that would also affect the functioning of monetary policy. Nonetheless, here I have drawn on what we know to date, including my experience setting monetary policy in the United Kingdom, to assess the validity of three different concerns about monetary policy at the ELB. Is monetary policy less potent at the ELB? Does it generate greater international spillovers (through capital flows and exchange rates)? And does the ELB make monetary policy stickier and make it harder to raise rates when appropriate? The last set of concerns is more speculative, but the first two have been raised as reasons to avoid the ELB when possible—potentially justifying changes to monetary policy frameworks.

The discussion in this paper, however, suggests that monetary policy at the ELB can still be potent, and does not necessarily generate any greater

international spillovers through capital flows and exchange rates than comparable adjustments in interest rates. It may be more challenging to raise rates off the ELB than to raise rates from higher levels—possibly due to counterbalancing effects working through the exchange rate—although there are only anecdotes to support this stickiness rather than any formal empirical evidence. The debate on these issues will continue—albeit hopefully not as long as countries have been mired at the ELB.

References

- Bank of England. 2014. *Inflation Report*, August. <https://www.bankofengland.co.uk/inflation-report/2014/august-2014>.
- . 2016. *Inflation Report*, November. <https://www.bankofengland.co.uk/inflation-report/2016/november-2016>.
- Boneva, Lena, Calebe de Roure, and Ben Morley. 2018. “The Impact of the Bank of England’s Corporate Bond Purchase Scheme on Yield Spreads.” Bank of England Staff Working Paper No. 719. <https://www.bankofengland.co.uk/-/media/boe/files/working-paper/2018/the-impact-of-the-bank-of-englands-corporate-bond-purchase-scheme-on-yield-spreads.pdf?la=en&hash=6C4947C7C536F27A8ED86165C521F301C01EE9E6>.
- Brainard, Lael. 2017. “Cross-Border Spillovers of Balance Sheet Normalization.” Speech to Board of Governors of the Federal Reserve System, New York, July 11. <https://www.federalreserve.gov/newsevents/speech/brainard20170711a.htm>.
- Burgess, Stephen, Emilio Fernandez-Corugedo, Charlotta Groth, Richard Harrison, Francesca Monti, Konstantinos Theodoridis, and Matt Waldron. 2013. “Appendices to Working Paper No. 471: The Bank of England’s Forecasting Platform: COMPASS, MAPS, EASE, and the Suite of Models.” <https://www.bankofengland.co.uk/-/media/boe/files/working-paper/2013/the-boes-forecasting-platform-compass-maps-ease-and-the-suite-of-appendix>.
- Curcuru, Stephanie. 2017. “The Sensitivity of the U.S. Dollar Exchange Rate to Changes in Monetary Policy Expectations.” Federal Reserve Board International Finance and Discussion Papers, Notes 2017-09-22.
- Curcuru, Stephanie, Steven Kamin, Canlin Li, and Marius Rodriguez. 2018. “International Spillovers of Monetary Policy: Conventional Policy vs. Quantitative Easing.” International Finance Discussion Paper 1234, Board of Governors of the Federal Reserve System.
- Ferrari, Massimo, Jonathan Kearns, and Andreas Schrimpf. 2017. “Monetary Policy’s Rising FX Impact in the Era of Ultra-Low Rates.” Working Paper No. 626, Bank for International Settlements. <https://www.bis.org/publ/work626.htm>.
- Forbes, Kristin. 2014. “The Economic Impact of Sterling’s Recent Moves: More Than a Midsummer Night’s Dream.” Speech at Bank of England, London, October 1. <https://www.bankofengland.co.uk/speech/2014/the-economic-impact-of-sterlings-recent-moves-more-than-a-midsummer-nights-dream>.
- . 2015. “Low Interest Rates: King Midas’ Golden Touch?” Speech at Bank of England, London, February 24. <https://www.bankofengland.co.uk/speech/2015/low-interest-rates-king-midas-golden-touch>.
- . 2017. “Failure to Lift.” Speech at London Business School, London, June 22. <https://www.bankofengland.co.uk/speech/2017/failure-to-launch>.
- Forbes, Kristin, Ida Hjortsoe, and Tsvetelina Nenova. 2018. “The Shocks Matter: Improving Our Estimates of Exchange Rate Pass-Through.” *Journal of International Economics* 114 (September): 255–75.

- Forbes, Kristin, and Francis Warnock. (2012). "Capital Flow Waves: Surges, Stops, Flight and Retrenchment." *Journal of International Economics* 88, no. 2: 235–51.
- Glick, Reuven, and Sylvain Leduc. 2015. "Unconventional Monetary Policy and the Dollar: Conventional Signs, Unconventional Magnitudes." Working Paper No. 2015-18, Federal Reserve Bank of San Francisco.
- Greenlaw, David, James Hamilton, Ethan Harris, and Kenneth West. 2018. "A Skeptical View of the Impact of the Fed's Balance Sheet." U.S. Monetary Policy Forum Paper. <https://research.chicagobooth.edu/-/media/research/igm/docs/2018-usmpf-report.pdf?la=en&hash=D8BE7A0F78D72A6762918282D5A56A2E76349AED>.
- Group of Seven. 2013. "Statement by G7 Finance Ministers and Central Bank Governors." February 12. www.g8.utoronto.ca/finance/fm130212.htm.
- Haldane, Andrew, Matt Roberts-Sklar, Tomasz Wieladek, and Chris Young. 2016. "QE: The Story so Far." Bank of England Staff Working Paper 624. <https://www.bankofengland.co.uk/working-paper/2016/qethe-story-so-far>.
- Hatzius, Jan, Sven Stehn, Nicholas Fawcett, and Karen Reichgott. 2017. "Policy Rate vs. Balance Sheet Spillovers." Goldman Sachs Economics Research, Global Economics Analyst, July 28.
- Joyce, Michael, Matthew Tong, and Robert Woods. 2011. "The United Kingdom's Quantitative Easing Policy: Design, Operation and Impact. *Bank of England Quarterly Bulletin*, third quarter. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1933696.
- Nakamura, Emi, and Jón Steinsson. 2018. "High-Frequency Identification of Monetary Non-neutrality: The Information Effect." *Quarterly Journal of Economics* 133, no. 3: 1283–1330.
- Neely, Christopher. 2015. "Unconventional Monetary Policy had Large International Effects." *Journal of Banking and Finance* 52: 101–11.
- Rey, Hélène. 2013. "Dilemma Not Trilemma: The Global Financial Cycle and Monetary Policy Independence." Proceedings of the Jackson Hole Symposium, sponsored by Federal Reserve Bank of Kansas City. <https://www.kansascityfed.org/publicat/sympos/2013/2013Rey.pdf>.
- Weale, Martin, and Tomasz Wieladek. 2016. "What Are the Macroeconomic Effects of Asset Purchases?" *Journal of Monetary Economics* 79: 81–93.

The Efficacy of Large-Scale Asset Purchases When the Short-Term Interest Rate Is at Its Effective Lower Bound

ABSTRACT The Federal Reserve on net purchased almost \$4 trillion in additional securities between March 2009 and December 2014. Although the initial announcements of these policies were associated with dramatic market reactions, these responses were soon reversed. The overall market reaction to news surprises from the Federal Reserve over this period was increases, not decreases, in interest rates. It is hard to disentangle the effects of the purchases themselves from new information about economic fundamentals. My conclusion is that it is difficult to estimate accurately what large-scale asset purchases accomplished, but the magnitude of the effect is likely smaller than commonly believed.


The traditional instrument of monetary policy is the short-term interest rate, which was stuck near zero in a number of the world's largest economies over much of the last decade. Central banks in the United States, Europe, and Japan purchased many trillions of dollars in securities in an effort to provide the stimulus that their traditional policy instrument could not. The U.S. Federal Reserve increased its holdings of Treasury securities, mortgage-backed securities, and agency debt from under \$600 billion at the start of March 2009 to over \$4.4 trillion by the end of 2014 (see figure 1). What did these large-scale asset purchases (LSAPs) accomplish?

Many standard macroeconomic and finance models predict that LSAPs would not affect any nominal or real variable of interest if the traditional policy rate is at its effective lower bound (ELB). If being at the ELB means

Conflict of Interest Disclosure: The author did not receive any financial support from any firm or person for this paper or from any firm or person with a financial or political interest in this paper. He is currently not an officer, director, or board member of any organization with an interest in this paper.

Figure 1. Federal Reserve Holdings of Securities, November 19, 2008, to December 27, 2017^a


Billions of dollars



Source: Federal Reserve Statistical Release H.4.1.

a. QE = quantitative easing. Sum of Federal Reserve holdings of Treasury securities, mortgage-backed securities, and agency debt, plus unamortized premiums less unamortized discounts, Wednesday values, in billions of dollars.

that further increases in the monetary base would yield essentially zero marginal liquidity benefits to a holder of the monetary base, purchasing any asset with the newly created base should not change the price of any state-contingent claims, and so should have zero effect on asset prices or spending decisions in many models (Eggertsson and Woodford 2003). Richer models allow for the possibility of some effects. For example, buying long-term assets may commit the fiscal or monetary authority to a different state-contingent path for distortionary taxes or inflation (Hamilton and Wu 2012; Eggertsson and Proulx 2016). Or if some assets confer unique benefits on certain institutions—for example, as collateral for repurchase agreements or to satisfy capital requirements—there could also be real effects from altering the supply of these special assets (Woodford 2012; Caballero and Farhi 2017). Real effects can also arise in models where some individuals are unable to hold certain assets (Cúrdia and Woodford 2011; Gertler and Karadi 2011; Chen, Cúrdia, and Ferrero 2012; Greenwood and Vayanos 2014). Granting the potential relevance of such mechanisms, the

Figure 2. The Interest Rate on a 10-Year U.S. Treasury Security on March 18, 2009^a


Source: Cboe DataShop (n.d.).

a. Calculated as 10 times the price of the ^TNX futures contract based on the 10-year Treasury constant-maturity rate each minute during March 18, 2009.

magnitude of the effect that can be achieved by LSAPs is very much an empirical question.

Figure 2 plots some dramatic evidence that might seem to settle this issue. The graph shows the interest rate on a 10-year U.S. Treasury security each minute of the day on March 18, 2009. At 2:15 p.m. EDT on that day, the Federal Open Market Committee (FOMC) issued a statement announcing its intention to purchase up to an additional \$1.15 trillion in mortgage-backed securities, long-term Treasury securities, and agency debt beyond the purchases announced previously. Within minutes of this announcement, the long-term Treasury rate fell by 50 basis points. It would be impossible to argue that the cause of this decline was something other than the Fed's announcement. When one adds this together with a few other dramatic moves—such as the 20-basis-point drop on November 25, 2008, when the Fed announced its initial intention to purchase up to \$600 billion in mortgage-backed securities and agency debt—it seems one can make a strong case that the first quantitative easing (QE1), as the first phase of LSAPs came to be called, may have lowered long-term yields by 100 basis points or more.

However, it is then interesting to look at figure 3, which shows what happened after the Fed's subsequent meeting on April 29, 2009. The

Figure 3. The Interest Rate on a 10-Year U.S. Treasury Security on April 29, 2009^a

Source: Cboe DataShop (n.d.).

a. Calculated as 10 times the price of the ^TNX futures contract based on the 10-year Treasury constant-maturity rate each minute during April 29, 2009.

Fed did not announce any change in plans for LSAPs in this statement, and indeed confirmed its intention to continue conducting the purchases announced on March 18. Yet just as we are forced to conclude that something the Fed said on March 18 caused the 10-year yield to fall, it is equally clear that something the Fed said (or did not say) on April 29 caused the yield to jump up by almost 10 basis points. What was it?

Here was the assessment of William Sullivan, chief economist at JVB Financial Group, as quoted in the Reuters bond market wrap-up for that day (Reuters 2018):

Treasuries prices fell because the Fed's statement has been adjusted to confirm its observation that some "green shoots" of stability and potential improvement in the economic environment are evident. . . . Also, some observers perhaps thought that the Fed would be able to increase the amount of Treasury and mortgage-backed securities purchases over and above the amount they delineated at the March policy meeting. So it doesn't look as if they will increase the size of those purchase programs.

To the extent that Sullivan's second explanation is accurate—that the market was surprised not to see additional purchases beyond those that had been announced on March 18—it raises the possibility that the initial 50-basis-point drop on March 18 should not be interpreted as the effect

Table 1. Comparison of Federal Open Market Committee Statements on March 18 and April 29, 2009

<i>March 18 statement</i>	<i>April 29 statement</i>
Information received since the Federal Open Market Committee met in January indicates that the economy continues to contract.	Information received since the Federal Open Market Committee met in March indicates that the economy has continued to contract, though the pace of contraction appears to be somewhat slower.
Job losses, declining equity and housing wealth, and tight credit conditions have weighed on consumer sentiment and spending.	Household spending has shown signs of stabilizing but remains constrained by ongoing job losses, lower housing wealth, and tight credit.
Weaker sales prospects and difficulties in obtaining credit have led businesses to cut back on inventories and fixed investment. U.S. exports have slumped as a number of major trading partners have also fallen into recession.	Weak sales prospects and difficulties in obtaining credit have led businesses to cut back on inventories, fixed investment, and staffing.
Although the near-term economic outlook is weak, the Committee anticipates that policy actions to stabilize financial markets and institutions, together with fiscal and monetary stimulus, will contribute to a gradual resumption of sustainable economic growth.	Although the economic outlook has improved modestly since the March meeting, partly reflecting some easing of financial market conditions, economic activity is likely to remain weak for a time. Nonetheless, the Committee continues to anticipate that policy actions to stabilize financial markets and institutions, fiscal and monetary stimulus, and market forces will contribute to a gradual resumption of sustainable economic growth in a context of price stability.

Source: Statements of the Federal Open Market Committee.

of the policy the Fed actually implemented, but rather as a potential effect of some policy that markets thought the Fed might implement, though in practice it did not actually do so. Sullivan's first interpretation—that the market was responding to the Fed's more optimistic assessment of economic fundamentals—was the primary factor cited in the rest of the Reuters news account. The April 29 statement made significant changes in the words that the Fed used to describe the economy. The Fed sounded considerably less pessimistic on April 29 than it had on March 18 (see table 1).

A market response to these improvements in the Fed's outlook could be interpreted in two different ways. One view maintains that the Fed's information about the economy is a strict subset of the market's. According to this view, the market knew the true condition of the economy, and it had

a guess about but did not know the Fed's assessment for sure. On learning that the Fed was more optimistic than anticipated, the market participants would revise their expectations of future monetary policy, now expecting perhaps fewer LSAPs or an earlier liftoff from the ELB. The second view is that the market and the Fed each has some information about the economy that the other does not have. According to this view, the release of the Fed's more optimistic assessment rationally leads to an upward revision of the market's forecast of economic fundamentals, and could lead to higher interest rates per this mechanism.


Direct comparisons of private forecasts with those of the Federal Reserve's Greenbook have demonstrated that the Fed has some information that is useful for forecasting output and inflation beyond what is known to the private sector (Romer and Romer 2000; Faust and Wright 2009). If the Fed knows some things that private analysts do not know, and private analysts know some things that the Fed does not know, the rational response of a private actor to the revelation of the Fed's economic assessment is to revise his or her own assessment (Melosi, forthcoming; Miranda-Agrippino and Ricco 2018). Much research has convincingly shown that this channel is an important component of the typical market response to Fed statements and actions. Jeffrey Campbell and others (2012) found that from 1994 to 2007, when the Fed announced an interest rate that was higher than the market anticipated, this was associated with a move to *lower* forecasts of unemployment and *higher* forecasts of inflation in the Blue Chip consensus forecast, exactly the opposite of what is predicted by the first view (the Fed is going to be more contractionary than anticipated) and exactly what is predicted by the second view (the economy is in better shape than people thought). Emi Nakamura and Jón Steinsson (2018) confirmed this finding in a careful analysis of high-frequency data through 2014. Additional evidence in support of this view has been provided by Silvia Miranda-Agrippino and Giovanni Ricco (2018) and by Aemil Lakdawala and Matthew Schaffer (2018).

If information that the economy was in better shape than many private analysts had previously concluded was indeed one factor driving rates up on April 29, 2009, we also need to allow the possibility that the Fed's negative economic assessment, and not just the LSAPs, were factors driving rates down on March 18. To the extent that is the case, it would mean that the 50-basis-point drop observed on March 18 is an overestimate of the effect of LSAPs themselves on the long-term rate.

It is even more telling to note the scale of the vertical axis in figures 2 and 3. The 10-year rate began March 18 at 2.97 percent and began April 29

Figure 4. Cumulative Change in the 10-Year Yield of U.S. Treasury Securities on Fed Days, November 20, 2008–December 1, 2017^a

Basis points



Source: Adapted by the author from the data set given by Greenlaw and others (2018).

a. Shaded regions denote periods of bond purchases under QE1 (January 1, 2009–March 31, 2010), QE2 (November 3, 2010–June 30, 2011), and QE3 (October 1, 2012–October 29, 2014).


at 3.00 percent. Thus some sort of news arriving after the March 18 meeting and before the April 29 meeting led to a complete reversal of the dramatic drop of 50 basis points on March 18. And by the end of April 29, the rate was significantly higher than it had started out before the March 18 announcement. Was this information arriving between March 19 and April 28 news about what the Fed was going to do, or news about other fundamentals that matter for bond prices?

A recent paper by David Greenlaw and others (2018) used two approaches to try to answer this question. Their first approach was to note the date of every single FOMC meeting, release of minutes, or speech by the Fed chair about the economy or monetary policy. They called these “Fed Days.” Figure 4 plots the cumulative change in the 10-year rate coming only on Fed Days from November 20, 2008, to December 1, 2017.¹ After some dramatic initial drops, the overall movement of the market on Fed Days subsequent to March 18, 2009, was up for the remainder of the bond

1. This figure is adapted from exhibit 4.2 in Greenlaw and others (2018). The latter begins November 1, 2018, whereas figure 4 begins November 20, 2018, just before the first announcement of QE1 on November 25. Note that November 25 is not included in the definition of “Fed Days” because it was not the date of an FOMC meeting, minutes release, or Fed speech, but rather took the form of an unscheduled Fed announcement.

Figure 5. Cumulative Change in the 10-Year Yield of U.S. Treasury Securities on Reuters Fed News Days, November 20, 2008–December 1, 2017^a

Basis points



Source: Adapted by the author from the data set given by Greenlaw and others (2018).


a. Shaded regions denote periods of bond purchases under QE1 (January 1, 2009–March 31, 2010), QE2 (November 3, 2010–June 30, 2011), and QE3 (October 1, 2012–October 29, 2014).

purchases of QE1, a period during which the Fed was intending that its LSAP would help hold rates down. The overall market move on Fed Days during both QE2 and QE3 was also unquestionably up, not down.

The second approach taken by Greenlaw and others (2018) was to look at every day when the 10-year yield changed by more than 1 standard deviation and study the Reuters bond market wrap-up for that day. If Reuters described news about the Fed as the primary driver of bond prices on that day, it was designated a “Reuters Fed News Day.” If Reuters listed the Fed as one of two contributing factors, the day was given a weight of $\frac{1}{2}$. Figure 5 plots the cumulative change in the 10-year rate on Reuters Fed News Days. By including a larger set of days than considered in figure 4 on which there was information released to the market about Fed policy, these suggest a bigger role for Fed announcements in bringing rates down in the fall of 2008. But the conclusion remains that the overall effect of news from the Fed after March 18 and throughout QE2 and QE3 was to drive interest rates higher.

Another event that many people consider convincing evidence of the importance of LSAPs came on May 22, 2013, when Fed chair Ben Bernanke suggested in congressional testimony that the Fed might slow the rate of monthly net bond purchases within the next three FOMC meetings.

Figure 6. Cumulative Change in the 10-Year Yield (in Basis Points) of U.S. Treasury Securities, January 1, 2013–December 31, 2013



Source: Greenlaw and others (2018, exhibit 5.3).

Note: The vertical line is at May 21, the day before Bernanke's warning.

The 10-year yield rose 11 basis points that day, a development that subsequently came to be referred to as the “Taper Tantrum.” But this was the only change in May that either of the methodologies used by Greenlaw and others (2018) would associate with news from the Fed. Notwithstanding, the rate was up overall 45 basis points in May (figure 6). The key factors identified by Reuters as driving yields higher in May included a strong employment report on May 3 and favorable housing and consumer sentiment data released on May 28.

It is also worth noting the market's nonresponse to the Fed's more recent decision to begin reducing the size of its balance sheet. Greenlaw and others (2018) noted that both the Blue Chip consensus and the primary dealer survey in January 2017 were anticipating that the Fed would not begin reducing its balance sheet until June 2018. These surveys expected that total Fed assets would still be \$3.8 trillion to \$4.0 trillion by the end of 2019. The actual shrinkage began in October 2017, three quarters earlier than the market initially expected, and announced a target balance sheet for the end of 2019 of \$3.6 trillion. Significant information arrived during 2017 that the Fed was going to contract sooner and faster than many expected. But it is difficult to identify any significant market reaction to this. Greenlaw and others described this as the “Shrinkage Shrug.”

These observations raise doubts not just about the magnitude of the effects of LSAPs but also about the whole strategy of identifying the

effects of monetary policy using high-frequency event studies, which has become the dominant approach in empirical economic research. The Fed's announcements in November and December 2008 and March 2009 came at times when news of a deteriorating economy was arriving from multiple sources. Investors (and the Fed) were trying to sort out exactly what it all meant. Bond prices would be particularly sensitive to the Fed's assessment of economic fundamentals in this setting. Likewise, in April 2009 and May 2013, investors had already been seeing a number of more favorable indicators, and accordingly may have responded more strongly to optimistic assessments from the Fed.

The idea behind high-frequency identification is that one can measure the isolated contribution of each source of news by the market response within a narrow window of the first release of this news. Consider taking this view to its logical conclusion. Equity futures tumbled 5 percent within hours after Donald Trump was predicted to win the 2016 presidential election in the evening of November 8, only to regain it all by noon the next day. According to the strict event study methodology, the interpretation would have to be that Trump's election did indeed take 5 percent off the value of U.S. corporations, but that some other shock within hours added this amount back. A more natural interpretation is that there are limits to investors' ability to understand, within minutes, all the implications of untested and unclear policies (Wolfers and Zitzewitz 2018). Moreover, the Fed's announcements reveal not just actions that it is going to take but also its best assessment of economic fundamentals. The Fed's assessment can be important information for me for purposes of refining my own assessment of economic fundamentals. Separating the contributions of these two factors is challenging.

Let me emphasize what I am *not* concluding from these observations. I have been discussing only the effects of the Fed's LSAP programs and their huge expansion initiated in March 2009. This does not say anything about the efficacy of the Fed's emergency lending facilities implemented in the fall of 2008 (and mostly phased out by the end of 2009). Evidence from the idiosyncratic responses of different banks and money market funds to the lending facilities suggests that these programs may well have had beneficial effects (for example, Duygan-Bump and others 2013). Nor am I suggesting that LSAPs had no effects on bond prices. As noted above, figure 2 makes such a claim difficult to defend. But I do conclude that it is very hard to accurately estimate the magnitude of exactly what LSAPs accomplished, and that the magnitude of their true effects is likely to be smaller than many central banks believe.

References

- Caballero, Ricardo, and Emmanuel Farhi. 2017. "The Safety Trap." *Review of Economic Studies* 85, no. 1: 223–74.
- Campbell, Jeffrey R., Charles L. Evans, Jonas D. M. Fisher, and Alejandro Justiniano. 2012. "Macroeconomic Effects of Federal Reserve Forward Guidance." *Brookings Papers on Economic Activity*, no. 1: 1–80.
- Cboe DataShop. No date. "Historical Data." <https://datashop.cboe.com/equity-quotes>.
- Chen, Han, Vasco Cúrdia, and Andrea Ferrero. 2012. "The Macroeconomic Effects of Large-Scale Asset Purchase Programmes." *Economic Journal* 122, no. 564: F289–F315.
- Cúrdia, Vasco, and Michael Woodford. 2011. "The Central-Bank Balance Sheet as an Instrument of Monetary Policy." *Journal of Monetary Economics* 58, no. 1: 54–79.
- Duygan-Bump, Burcu, Patrick Parkinson, Eric Rosengren, Gustavo A. Suarez, and Paul Willen. 2013. "How Effective Were the Federal Reserve Emergency Liquidity Facilities? Evidence from the Asset-Backed Commercial Paper Money Market Mutual Fund Liquidity Facility." *Journal of Finance* 68, no. 2: 715–37.
- Eggertsson, Gauti B., and Kevin Proulx. 2016. "Bernanke's No-Arbitrage Argument Revisited: Can Open Market Operations in Real Assets Eliminate the Liquidity Trap?" NBER Working Paper 22243. Cambridge, Mass.: National Bureau of Economic Research.
- Eggertsson, Gauti B., and Michael Woodford. 2003. "Zero Bound on Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity*, no. 1: 139–233.
- Faust, Jon, and Jonathan H. Wright. 2009. "Comparing Greenbook and Reduced Form Forecasts Using a Large Realtime Dataset." *Journal of Business & Economic Statistics* 27, no. 4: 468–79.
- Gertler, Mark, and Peter Karadi. 2011. "A Model of Unconventional Monetary Policy." *Journal of Monetary Economics* 58, no. 1: 17–34.
- Greenlaw, David, James D. Hamilton, Ethan Harris, and Kenneth D. West. 2018. "A Skeptical View of the Impact of the Fed's Balance Sheet." NBER Working Paper 24687. Cambridge, Mass.: National Bureau of Economic Research.
- Greenwood, Robin, and Dimitri Vayanos. 2014. "Bond Supply and Excess Bond Returns." *Review of Financial Studies* 27, no. 3: 663–713.
- Hamilton, James D., and Jing Cynthia Wu. 2012. "The Effectiveness of Alternative Monetary Policy Tools in a Zero Lower Bound Environment." *Journal of Money, Credit and Banking* 44, no. s1: 3–46.
- Lakawala, Aemit, and Matthew Schaffer. 2018. "Federal Reserve Private Information and the Stock Market." Working paper, Michigan State University.
- Melosi, Leonardo. Forthcoming. "Signaling Effects of Monetary Policy." *Review of Economic Studies*.

- Miranda-Agrippino, Silvia, and Giovanni Ricco. 2018. "The Transmission of Monetary Policy Shocks." Working paper, Bank of England.
- Nakamura, Emi, and Jón Steinsson. 2018. "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect." *Quarterly Journal of Economics* 133, no. 3: 1283–1330.
- Reuters. 2018. "Treasuries: Bonds Fall as Fed Gives Hopeful Outlook." April 29. <https://www.reuters.com/article/markets-bonds/treasuries-bonds-fall-as-fed-gives-hopeful-outlook-idUSN2943617720090429>.
- Romer, Christina D., and David H. Romer. 2000. "Federal Reserve Information and the Behavior of Interest Rates." *American Economic Review* 90, no. 3: 429–57.
- Wolfers, Justin, and Eric Zitzewitz. 2018. "'The Standard Error' of Event Studies: Lessons from the 2016 Election." *AEA Papers and Proceedings* 108: 584–89.
- Woodford, Michael. 2012. "Methods of Policy Accommodation at the Interest-Rate Lower Bound." In *The Changing Policy Landscape* (Federal Reserve Bank of Kansas City, Jackson Hole, Wyoming), 185–288.

The Federal Reserve Is Not Very Constrained by the Lower Bound on Nominal Interest Rates

ABSTRACT I survey the literature on monetary policy at the zero lower bound (ZLB) and effective lower bound (ELB) to make three main points: First, the Federal Reserve's forward guidance and large-scale asset purchases are effective monetary policy tools at the ZLB/ELB. Second, during the 2008–15 U.S. ZLB period, the Fed was not very constrained in its ability to influence medium- and longer-term interest rates and the economy. And third, the risks of the Fed being significantly constrained by the ELB in the future are typically greatly overstated. I conclude that the Federal Reserve is not very constrained by the lower bound on nominal interest rates.

In December 2008, the Federal Reserve lowered the federal funds rate to essentially zero, where it remained until December 2015. Because U.S. currency carries an interest rate of zero, it is essentially impossible for the Fed to set the federal funds rate substantially below zero without triggering widespread conversion of deposits into currency. This constraint is commonly referred to as the zero lower bound (ZLB)—or as the effective lower bound (ELB), to acknowledge that the bound may be somewhat negative rather than literally zero.

The existence of the ZLB/ELB has led many researchers to conclude that it imposes a substantial constraint on the Fed's ability to conduct monetary policy in a low-interest-rate environment (Krugman 1998; Williams 2009; Kiley and Roberts 2017). In this paper, I survey the recent literature to demonstrate exactly the opposite: that the ZLB/ELB has not been, is not,

Conflict of Interest Disclosure: The author did not receive any financial support from any firm or person for this paper or from any firm or person with a financial or political interest in this paper. He is currently not an officer, director, or board member of any organization with an interest in this paper.

and almost certainly will not be a significant constraint on the Federal Reserve—in the past, the present, or the foreseeable future. This conclusion follows from three main observations. First, the Federal Reserve’s forward guidance and large-scale asset purchases are effective monetary policy tools at the ZLB/ELB; in fact, they are about as effective as the federal funds rate in normal times. Second, during the 2008–15 U.S. ZLB period, the Fed was not very constrained in its ability to affect medium- and longer-term interest rates and the economy. And third, the risks of the Fed being significantly constrained by the ELB in the future are typically greatly overstated.

1. The Federal Reserve Has Additional Monetary Policy Tools Available

The first main observation is that the Federal Reserve has other monetary policy tools available beyond just changes in the current federal funds rate. In particular, there is a large and growing body of literature on the effectiveness of forward guidance—that is, communication by the Federal Reserve about the likely future path of the federal funds rate over the next several quarters—and of large-scale asset purchases (LSAPs)—which are purchases by the Federal Reserve of hundreds of billions of dollars of longer-term U.S. Treasury securities and mortgage-backed securities.

Theoretically, financial markets and firms are forward-looking, so firms’ investment decisions depend not just on the current short-term interest rate but also on the path of expected future short-term interest rates over the next several years. A simple way to formalize this observation is with a standard New Keynesian investment/saving curve,

$$(1) \quad \hat{y}_t = E_t \hat{y}_{t+1} - \alpha \hat{r}_t + \varepsilon_t,$$

which can be solved forward, assuming $\lim_{j \rightarrow \infty} E_t \hat{y}_{t+j} = 0$, to get

$$(2) \quad \hat{y}_t = -\alpha E_t \sum_{j=0}^{\infty} \hat{r}_{t+j} + \varepsilon_t,$$

where t indexes periods, \hat{y}_t is the output gap, \hat{r}_t is the deviation of the one-period real interest rate from its steady state, E_t denotes the mathematical expectation conditional on information at time t , and ε_t is a mean-zero shock. The infinite sum in equation 2 illustrates how the Fed can affect

the current output gap by changing people's expectations about the future path of \hat{r}_{t+j} as well as the current value of \hat{r}_t itself. David Reifschneider and John Williams (2000) and Gauti Eggertsson and Michael Woodford (2003) use this fact to show that, even at the ZLB, the Fed can still stimulate the economy as long as it can credibly commit to a lower path of short-term interest rates in the future, when the ZLB is no longer binding.

Empirically, Refet Gürkaynak, Brian Sack, and Swanson (2005, henceforth GSS) showed that changes in the federal funds rate alone were not sufficient to explain financial market reactions to announcements by the Federal Reserve Federal Open Market Committee (FOMC), and that there was a second dimension of monetary policy that was being missed. GSS developed a measure of forward guidance based on high-frequency changes in a range of federal funds and eurodollar futures contracts around FOMC announcements, orthogonalized to the change in the current federal funds rate. They showed that forward guidance had highly statistically significant effects on financial markets and dramatically increased the explanatory power of FOMC announcements for financial market responses. GSS also showed that large movements in their measure of forward guidance were associated with FOMC statements about the future path of the federal funds rate. Their results demonstrated that forward guidance is effective, and have since been updated and confirmed for the U.S. and other countries by, among others, Claus Brand, Daniel Buncic, and Jarkko Turunen (2010); Stefania D'Amico and Mira Farka (2011); Jeffrey Campbell and others (2012); Matteo Leombroni and others (2017); and Swanson (2018).

An equally large and growing body of literature finds that the Fed's LSAPs have had economically and statistically significant effects on longer-term bond yields. These analyses range from high-frequency event studies of the United States (Gagnon and others 2011; Krishnamurthy and Vissing-Jorgensen 2011) and the United Kingdom (Joyce and others 2011) to historical studies of "Operation Twist" in 1961 (Swanson 2011) to lower-frequency, monthly studies of U.S. Treasury yields and spreads vis-à-vis the Treasury's supply from 1919 to 2008 (Krishnamurthy and Vissing-Jorgensen 2012) to monthly no-arbitrage term structure models with quantity effects (Greenwood and Vayanos 2012; Hamilton and Wu 2012). Swanson (2011) and Williams (2013) survey these estimates, but a common benchmark is that \$600 billion in LSAPs causes the 10-year Treasury note yield to fall about 15 basis points.

Swanson (2018) estimates the effects of both forward guidance and LSAP announcements on financial markets and finds that they are roughly as effective as changes in the federal funds rate in normal times. Table 1

Table 1. Estimated Effects of Changes in the Federal Funds Rate, Forward Guidance, and Large-Scale Asset Purchases on Financial Markets, 1991–2015^a

Type of change	Treasury yields			Baa yield	Standard & Poor's 500	Dollar/yen exchange rate
	2-year	5-year	10-year			
Change in federal funds rate [<i>t</i> statistic]	3.68*** [13.91]	2.04*** [8.56]	0.95*** [3.74]	0.56 [1.51]	−0.30*** [−6.98]	−0.13*** [−3.04]
Change in forward guidance [<i>t</i> statistic]	4.85*** [5.87]	5.09*** [5.49]	3.92*** [5.04]	1.99*** [3.96]	−0.19*** [−3.92]	−0.18*** [−3.71]
Change in LSAPs [<i>t</i> statistic]	−0.32 [−1.01]	−3.71*** [−6.49]	−5.68*** [−7.17]	−4.22*** [−5.59]	0.04 [0.76]	0.25*** [3.97]

Source: Swanson (2018).

a. Coefficients β are from regressions $\Delta y_t = a + \tilde{F}\beta + \varepsilon_t$, where t indexes FOMC announcements, y denotes a given bond yield or log asset price, \tilde{F} are the monetary policy factors estimated by Swanson (2018), and Δ is the change in a 30-minute window bracketing each FOMC announcement (1-day window for Baa yield). Coefficients are in units of basis points (for bond yields) or percentage points (for stock prices and exchange rates) per the standard deviation change in each monetary policy instrument. Bootstrapped t statistics are in square brackets; *** denotes statistical significance at the 1 percent level. See the text and Swanson (2018) for details.

summarizes the results.¹ The table's first four columns report the effects of changes in the federal funds rate, forward guidance, and LSAPs on 2-, 5-, and 10-year Treasury yields and Moody's index of Baa-rated corporate bond yields, in units of basis points per standard deviation change in the policy instrument.² Thus, the effect of an increase of 1 standard deviation in the federal funds rate is about 3.7 basis points on the 2-year Treasury yield; for forward guidance, the effect on the 2-year yield is bigger, about 4.9 basis points per standard deviation change; and for LSAPs, the effect is smaller, about -0.3 basis points, and not statistically significant. LSAPs primarily affect longer maturities and the federal funds rate affects shorter maturities, but overall the three policies have effects on yields that are broadly comparable in magnitude. This is further supported by the last two columns of table 1, which report the effects on the Standard & Poor's 500 and the dollar/yen exchange rate, in units of percentage points per standard deviation change in each policy instrument. The effects of all three policies have the signs one would expect—higher interest rates imply lower stock prices and dollar appreciation—and are roughly comparable in magnitude. These results all suggest that forward guidance and LSAPs are effective monetary policy tools; in fact, they are about as effective as changes in the federal funds rate in normal times.

Looking beyond asset prices, some researchers have used detailed bank-level data to show that LSAPs have significant effects on bank lending. Alexander Rodnyansky and Olivier Darmouni (2017) show, via a differences-in-differences analysis of quarterly U.S. bank-level data, that banks that owned more LSAP-eligible mortgage-backed securities increased business lending in response to the Fed's LSAPs. Marco Di Maggio, Amir Kermani, and Christopher Palmer (2016) apply a similar differences-in-differences analysis to monthly loan-level U.S. mortgage originations to show that conforming (eligible for purchase by Fannie Mae and Freddie Mac)

1. Some researchers, such as Campbell and others (2012), distinguish between two types of forward guidance announcements by the Federal Reserve—those that convey information about the economy versus those that only convey information about monetary policy. Swanson (2018) does not try to separately identify these two types of forward guidance announcements, so the estimates given in table 1 represent an average forward guidance announcement effect.


2. The standard deviation of surprise changes in the federal funds rate is 8.8 basis points, measured over the period from 1991 to 2008; the standard deviation of surprise changes in forward guidance is 6 basis points in the one-year-ahead expected federal funds rate, as measured by eurodollar futures from 1991 to 2015; and the standard deviation of a surprise LSAP announcement is about \$250 billion in long-term bond purchases, measured over the period 2009–15. See Swanson (2018) for details.

mortgage originations increased in response to the Fed's LSAPs.³ Michael Koetter, Natalia Podlich, and Michael Wedow (2017) analyze quarterly German bank-level, security-by-security data to show, via differences-in-differences, that German banks that held more eligible securities for the European Central Bank's Securities Markets Programme increased their lending in response to the program. Thus, the effects of LSAPs extend beyond just a high-frequency change in financial market prices.

Some researchers have argued that, even though LSAPs have had a significant effect on financial markets on impact, those effects have tended to die out over time (Greenlaw and others 2018). The monthly and quarterly bank-lending studies described above provide evidence against this view—after all, if the financial market effects rapidly died out, why would banks increase their lending over subsequent months and quarters? Swanson (2018) also studies the persistence of financial market responses to LSAP announcements and finds that they were very persistent, with the exception of the very large and perhaps special “QE1” (first quantitative easing) announcement on March 18, 2009. On that date, bond yields fell dramatically in response to the FOMC announcement, but then began to rise over subsequent weeks, as the Dow Jones Industrial Average gained over 13 percent and the Fed's bank stress tests, released May 7, 2009, turned out better than markets expected; thus, the markets' behavior in those weeks may not be representative of the longer-run effects of LSAPs more generally. Figure 1 reproduces two graphs from Swanson (2018) that show a tendency for the effects of LSAPs to die out when the March 18, 2009, observation is included (left panel), but not when that one very influential announcement is excluded (right panel).

To sum up thus far, there is a great deal of evidence that both forward guidance and LSAPs are effective monetary policy tools—in fact, they are about as effective as changes in the federal funds rate in normal times. There is also very strong evidence that LSAPs have affected bank lending. Although some have argued that the effects of LSAPs are not persistent, this view seems to be driven by one very influential FOMC announcement on March 18, 2009, which may have been special for a number of reasons. Excluding this one announcement, the estimated effects of LSAPs on financial markets have been essentially completely persistent.

3. To be precise, Rodnyansky and Darmouni (2017) and Di Maggio and others (2016) find that the Fed's LSAP purchases of mortgage-backed securities had a significant effect on bank lending; the Fed's purchases of long-term Treasury securities during the QE2 program did not seem to have such an effect.

Figure 1. Estimated Effects of LSAPs on the 10-Year, Zero-Coupon Treasury Yield^a

Source: Swanson (2018).


a. LSAPs = large-scale asset purchases. Estimated effects of LSAPs on the 10-year, zero-coupon Treasury yield, for different horizons h from 1 to 120 business days, including and excluding the influential March 18, 2009, “QE1” announcement. Estimated coefficients $\hat{\gamma}_h$ (solid lines) and bootstrapped ± 1.96 -standard error bands (dashed lines) are from regressions $y_{t-1+h} - y_{t-1} = \gamma_h - F_t + \varepsilon^{(h)}$. Restricted coefficient estimates $\gamma_h = a e^{-b(h-1)}$ (dash-dotted lines) are from the same set of regressions estimated jointly via nonlinear least squares. See the text and Swanson (2018) for additional details.

II. The Federal Reserve Was Not Very Constrained by the ZLB from 2008 to 2015

The second main observation is that, during the 2008–15 ZLB period, the Federal Reserve was not very constrained in its ability to affect medium- and longer-term interest rates and the economy. A quick way to see this is shown in figure 2, which plots the federal funds rate and 1-, 2-, 5-, and 10-year, zero-coupon Treasury yield from 2007 to 2017.⁴ Although the federal funds rate was virtually zero and never changed from December 2008 to November 2015, the 2-year Treasury yield—which is

4. Zero-coupon yields are from the data set given by Gürkaynak, Sack, and Wright (2007), available from the Federal Reserve Board’s website at <https://www.federalreserve.gov/pubs/feds/2006/200628/200628abs.html>.

Figure 2. The Federal Funds Rate and 1-, 2-, 5-, and 10-Year Zero-Coupon Treasury Yields, 2007–18^a



Sources: Author's update of figure from Swanson and Williams (2014), using yield curve data from Gürkaynak, Sack, and Wright (2007).

a. See the text for details.

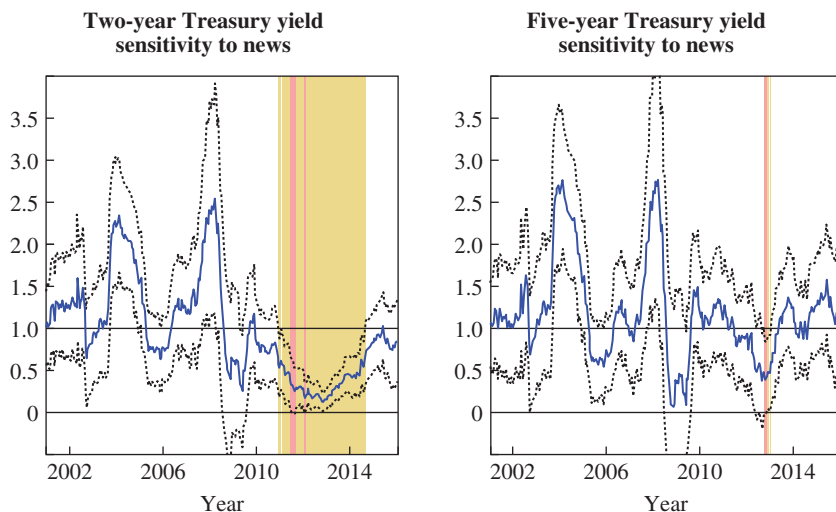
a better measure of the overall stance of monetary policy, as can be seen in equation 2—averaged about 55 basis points during this period and fluctuated substantially over time, ranging between 16 and 140 basis points and moving up or down every day in response to macroeconomic data releases, FOMC announcements, and other news.

Swanson and Williams (2014) formalize this point by estimating how responsive the 2-year and other Treasury yields are to major macroeconomic announcements, relative to a benchmark sample from 1990 to 2000, when the ZLB was not a constraint. That is, they run daily-frequency regressions of the form

$$(3) \quad \Delta y_t = \gamma^\tau + \delta^\tau \beta X_t + \varepsilon_t,$$

where t indexes business days, Δy_t is the one-day change in the 2-year Treasury yield (or other yield), X_t is an n -dimensional vector of major macroeconomic data releases that day (such as nonfarm payrolls, the Consumer Price Index, and GDP), β is an n -dimensional vector of parameters containing the normal responsiveness of the 2-year Treasury yield to each of those releases, and the parameters γ^τ and δ^τ are scalars that vary over time, with δ^τ normalized to a unit mean over the baseline sample from

Figure 3. Time-Varying Sensitivity Coefficients δ^τ from Regression Equation 3 Applied to Two-Year and Five-Year Treasury Yields^a



Source: Author's update of figures from Swanson and Williams (2014).

a. Dotted lines denote heteroskedasticity-consistent ± 2 -standard error bands, $\delta^\tau = 1$ corresponds to normal Treasury sensitivity to news, and $\delta^\tau = 0$ corresponds to complete insensitivity. Light shaded regions denote δ^τ significantly less than 1; darker shaded regions denote δ^τ significantly less than 1 and not significantly different from 0. See the text and Swanson and Williams (2014) for additional details.

1990 to 2000.⁵ Thus, the scalar δ^τ captures the overall sensitivity of the 2-year Treasury yield to major macroeconomic announcements around a given date, with $\delta^\tau = 1$ corresponding to normal sensitivity to news and $\delta^\tau = 0$ to complete insensitivity to news.

Figure 3 plots the time-varying sensitivity coefficients δ^τ from regression equation 3 for the 2- and 5-year Treasury yields from 2001 to 2015. The solid line in each panel plots the estimated value of δ^τ on each date τ , while the dotted lines depict heteroskedasticity-consistent ± 2 -standard error bands. The horizontal lines are drawn at 0 and 1 as benchmarks for

5. On most days, there is no news about a given macroeconomic statistic; thus, if the first column of X corresponds to nonfarm payrolls, then that column would be zero on every date t except once per month, when the nonfarm payrolls data are released. On each nonfarm payrolls announcement date, the first column of X_t contains the surprise component of the announcement—that is, the actual released value of nonfarm payrolls less the median market expectation from the day before. The scalar parameters γ^τ and δ^τ are estimated over rolling 250-business-day windows, while the vector β is fixed over the whole sample. See Swanson and Williams (2014) for details.

comparison, corresponding to the cases of complete insensitivity to news and normal sensitivity, respectively. Light shaded regions denote periods when the estimated value of δ^r is significantly less than unity; in addition, if the hypothesis that $\delta^r = 0$ cannot be rejected, then the region is shaded darker. Thus, darker shaded regions correspond to periods in which the Treasury yield was essentially insensitive to news, while lighter shaded regions correspond to periods when the yield was partially—but not completely—unresponsive to news.

The left panel of figure 3 shows that, from 2008 to 2011, the 2-year Treasury note yield's sensitivity to news was essentially never significantly less than normal. From 2011 to mid-2014, the 2-year yield's sensitivity did drop below normal, but was still greater than zero, except for two very brief periods near the end of 2011. Thus, despite the fact that the federal funds rate (and other short-term interest rates) were completely constrained by the ZLB throughout the period 2009–15, the 2-year Treasury yield continued to respond substantially to macroeconomic news. The 5-year Treasury yield, in the figure's right panel, was essentially never constrained by the ZLB during this period.

Carlos Carvalho, Eric Hsu, and Fernanda Nechio (2016) examine the effects of Federal Reserve communications on medium- and longer-term interest rates and come to the same conclusion. These researchers use textual analysis of newspaper articles in *Factiva* around each FOMC announcement to measure the perceived change in the Fed's "hawkishness" or "dovishness." They run regressions of medium- and longer-term Treasury yields on this measure of Fed communication, analogous to equation 3 above. Carvalho, Hsu, and Nechio (2016) show that their text-based measure of Fed communication had economically and statistically significant effects on 2-, 5-, and 10-year Treasury yields throughout the 2008–15 ZLB period, with results that are similar to those of Swanson and Williams (2014) and figure 3, above. Their results provide direct evidence that the Fed was never very constrained in its ability to move medium- and longer-term interest rates throughout the period 2008–15.

Additional indirect evidence supporting this conclusion is provided by the macroeconomic vector autoregression (VAR) studies done by Jing Cynthia Wu and Fan Dora Xia (2016) and by Davide Debortoli, Jordi Galí, and Luca Gambetti (2018). Wu and Xia (2016) use an affine term structure model to estimate a "shadow federal funds rate" during the ZLB period—that is, a hypothetical negative federal funds rate that summarizes the effects of the Fed's unconventional monetary policies on the yield

curve at each date. They estimate a VAR for output, inflation, and the shadow federal funds rate from 1960 to 2013, where the shadow federal funds rate is set equal to the federal funds rate in the pre-ZLB period, and find no evidence of a structural break in the VAR between the pre-ZLB and ZLB periods. They conclude that the Fed was able to affect the macroeconomy during the ZLB period in much the same way as it did before, albeit through unconventional rather than conventional monetary policy.

Debortoli, Galí, and Gambetti (2018) estimate VAR models with time-varying parameters and come to the same conclusion. They find no evidence of a change in the U.S. economy's responses to a technology shock or a demand shock in the pre-ZLB and ZLB periods. They also show that their methods would detect clear evidence of such a change if the economy followed a standard New Keynesian model and monetary policy was conducted by a Taylor-type interest rate rule that faced a ZLB constraint. They conclude that the Fed's unconventional monetary policies during the ZLB period were essentially a perfect substitute for changes in the federal funds rate.

Finally, Arsenios Skaperdas (2017) performs a multisector analysis of the U.S. economy from 1970 to 2012 and from 1988 to 2012. He ranks sectors by their interest rate sensitivity in the pre-2008 period, with sectors like construction, mining, and transportation being the most interest-sensitive and health care and services the least sensitive. If interest rates were kept artificially higher than normal by the ZLB in the period 2008–15, then the interest-rate-sensitive sectors of the economy should have performed relatively worse than they did after the previous 1990–91 and 2001–3 recessions. He shows that this was not the case; interest-rate-sensitive sectors performed about as well after the period 2007–9 as they did after previous recessions. Like Wu and Xia (2016) and Debortoli, Galí, and Gambetti (2018), he concludes that the Fed's forward guidance and LSAPs during the ZLB period were able to lower medium- and longer-term interest rates in much the same way as in previous recessions.

To sum up the results of this section, the Fed was not very constrained in its ability to affect medium- and longer-term interest rates throughout the ZLB period. Moreover, explicit tests for a structural break or change in macroeconomic behavior in 2009 fail to find any evidence that the economy behaved differently during the ZLB period than before, suggesting that the Fed's unconventional monetary policies during this period were a close substitute for changes in the federal funds rate.

III. The Risks of Being Constrained by the ELB in the Future Are Overstated

Finally, the risks of the Fed being significantly constrained by the ELB in the future are typically greatly overstated. There are three main reasons for this overstatement. First, the federal funds rate must be constrained by the ELB for several quarters, rather than just one quarter, to have a noticeable effect on the economy. Second, central banks in Europe have demonstrated that the ELB is substantially below zero; at least -0.75 percent, and probably a bit below -1 percent. And third, even in those rare cases when the federal funds rate is at the ELB for several quarters, the Fed has alternative monetary policy tools available to it, as discussed above.

Lawrence Christiano, Martin Eichenbaum, and Sergio Rebelo (2011) study the effects of fiscal policy in a standard, medium-scale New Keynesian model at the ZLB. They show that when the ZLB constrains the short-term interest rate for 8 or 12 quarters, the fiscal multiplier is substantially larger than normal because the standard monetary policy response to the fiscal shock is shut down. However, they also show that when the ZLB binds for only four quarters, then the fiscal multiplier is *not* any larger than normal (Christiano, Eichenbaum, and Rebelo 2011, n. 11). The intuition for this result is straightforward; according to equation 2 above, the output gap today is determined by the entire expected path of the federal funds rate, not just the federal funds rate today. If the federal funds rate is only constrained by the ZLB for a few periods, then the effect on the sum in equation 2 is relatively small, and the effect on the economy is correspondingly small. This helps to clarify that the ZLB is not a significant constraint on the economy unless it binds for several quarters (for example, eight or more).

In addition, several central banks in Europe have shown that the ELB is substantially less than zero. In December 2014, the Swiss National Bank lowered the target for its short-term policy rate to -0.25 percent, followed by an additional cut to -0.75 percent in January 2015, where it has since remained. In Sweden, the Riksbank lowered its short-term policy rate to -0.1 percent in February 2015, followed by several additional rate cuts that brought it down to -0.5 percent in February 2016, where it has since remained. For the euro area, the European Central Bank reduced the lower end of its policy rate corridor, the deposit facility rate, to -0.1 percent in June 2014, followed by several additional cuts that lowered it to -0.4 percent in March 2016, where it has since remained; importantly, money market interest rates have traded near the lower end of the European Central Bank's

corridor throughout this period. In Denmark, the Nationalbank lowered its deposit rate to -0.2 percent in July 2012 and eventually reduced it to -0.75 percent in February 2015, although it has since raised it to -0.65 percent. All these central banks have maintained negative policy rates for several years with no widespread conversion of deposits into currency.⁶ Evidently, the ELB in Europe is substantially below zero, at least -0.75 percent and probably a bit below -1 percent.

A concern that is sometimes raised regarding negative policy rates is that they might not pass through to other interest rates in the economy. For retail deposit rates, there is some evidence that this is the case (Eggertsson, Juelsrud, and Wold 2017). However, Rafael De Rezende (2017) finds no difference in the pass-through from changes in the Swedish policy rate to Swedish government bond yields during the negative policy rate regime relative to the period before. Rima Turk (2016) shows that policy rate cuts in Sweden and Denmark passed through to money market rates and bank loan interest rates in those countries to the same extent during the negative interest rate regime as before, and Turk (2016) and Daniel Gros and others (2016) report that banks increased a variety of fees on retail customers as a substitute for charging those depositors an explicitly negative interest rate. Selva Demiralp, Jens Eisenschmidt, and Thomas Vlassopoulos (2017) analyze quarterly, individual, euro area bank balance sheet data and find that banks increased lending in response to policy rate cuts in the negative policy rate regime by at least as much as before.⁷ Overall, the pass-through from negative monetary policy rates to other financial market rates does not seem to be inhibited by the policy rate being negative.⁸

6. This is especially remarkable given that the European Central Bank offers €500 denomination notes and the Swiss National Bank offers CHF1,000 notes.

7. To be precise, Demiralp and others (2017) compare more versus less retail-deposit-funded banks. They show that more retail-deposit-funded banks, which are hit harder by negative policy rates, were relatively more likely to increase lending, reduce reserves, and increase government securities holdings in response to interest rate cuts in the negative policy rate regime than before.

8. A second, related concern regarding negative policy rates is that they might depress bank profitability, which in turn might reduce bank lending or have other deleterious effects on the economy. However, because retail deposit rates are less than the policy rate, this argument applies to low *positive* interest rates just as much as it applies to negative rates; thus, if the Fed were willing to lower the federal funds rate from 0.75 percent to 0, it should be essentially just as willing to lower the funds rate into negative territory. In addition, Lopez, Rose, and Spiegel (2018), Turk (2016), and Gros and others (2016) find no decrease in bank profitability in the negative policy rate regimes in Europe and Japan, because banks' increased fee income and capital gains offset their retail deposit interest expenses.

Together, these two observations—that the ELB is significantly less than zero and must bind for eight quarters or more to have noticeable effects on the economy—imply that previous estimates of the risks of the Fed facing a significant ELB constraint in the future are typically greatly overstated. For example, Michael Kiley and John Roberts (2017) define the ELB to be 0 percent—a ZLB—and then simulate a structural model to count the number of quarters in which the federal funds rate is less than or equal to zero, even if this episode lasts for just one quarter. Reifschneider and Williams (2000) and Williams (2009) perform calculations very similar to those of Kiley and Roberts (2017), albeit with a less pessimistic shock distribution.⁹ Obviously, these calculations greatly overstate the number of times the short-term interest rate drops below a more realistic ELB of -0.75 percent for eight quarters or more, which is the economically relevant question.

Even in those rare cases when the nominal interest rate does fall that far for that long, the simple calculations done by Williams (2009) and Kiley and Roberts (2017) ignore the existence of unconventional monetary policies such as forward guidance and LSAPs. As shown in the previous two sections, there is extensive evidence that these policies are effective and provide a close substitute for changes in the federal funds rate.

IV. Caveats

Although the observations above are supported by a wide variety of papers, data sets, and methods, there are still a few caveats to keep in mind. First, in a very severe ELB scenario, in which the federal funds rate is expected to be at the lower bound for more than eight quarters, the effectiveness of forward guidance could become much lower than in the past. This almost happened in 2012 (see figure 3), when financial markets expected the ZLB to be a constraint for long enough that the two-year Treasury yield's sensitivity to news fell substantially. In the end, the two-year yield's sensitivity never fell to zero, but if such a severe ELB constraint arose in the future, the two-year Treasury yield could cease to be a viable instrument of forward guidance. In principle, the Fed could work around this constraint by extending its forward guidance to even longer horizons, but in practice the Federal Reserve's chair may have difficulty committing his or

9. Kiley and Roberts (2017) draw shocks from the empirical distribution of shocks to the U.S. economy from 1970 to 2015, which implies that the United States will face another financial crisis and Great Recession every 40 to 45 years, on average.

her successor to a given path for the federal funds rate. However, even in such a dire situation as this, the Fed still has the ability to conduct LSAPs and influence financial markets and the economy through that channel. As shown in table 1 above, LSAPs have effects that are similar in magnitude to those of the federal funds rate and forward guidance, but operate substantially farther out along the yield curve.

Second, there may be political constraints that make it difficult for the Fed to use LSAPs and negative interest rates. During the 2008–15 U.S. ZLB period, LSAPs seemed to be poorly understood by the public and in many cases evoked strong negative reactions, such as being called “almost treasonous” by Texas governor Rick Perry (Wearden 2011). And even though the Fed never used negative interest rates in 2008–15, the idea evokes similarly vehement opposition from many commercial and investment bankers, presumably due to fears about bank profitability; for example, Deutsche Bank chief executive John Cryan argued that they have “fatal consequences,” Allianz chief economic adviser Mohamed El-Erian called them an “insane experiment,” Janus Capital financial manager Bill Gross said that “capitalism . . . cannot function” with them, Barclays CEO Jes Staley stated that “they are not helpful,” and BlackRock chief executive Larry Fink told shareholders they bring “potentially dangerous financial and economic consequences” (Cox 2016; Flynn 2016; Alban 2016; Wenik 2016). Given this opposition, it may be more difficult for the Fed to pursue these policies than the literature surveyed above would suggest.

V. Conclusions

The Federal Reserve is not significantly constrained by the lower bound on nominal interest rates. This conclusion is supported by three main observations. First, the Fed’s forward guidance and LSAPs are effective monetary policy tools, about as effective as changes in the federal funds rate in normal times. Second, during the 2008–15 U.S. ZLB period, the Fed was not significantly constrained in its ability to affect medium- and longer-term interest rates and the economy. And third, the risks of the Fed being significantly constrained by the ELB in the future are typically greatly overstated. These observations are supported by dozens of papers analyzing a variety of countries and data sets and using a wide variety of methods, ranging from high-frequency financial market responses to no-arbitrage term structure models to macroeconomic VARs to quarterly bank-level lending data. Although there are a few caveats to keep in mind, the overall conclusion is robust to these concerns.

References

- Alban, Jennifer. 2016. "Bill Gross Steps Up Criticism of Fed, Central Banks for Zero Rates." Reuters Business News, August 31.
- Brand, Claus, Daniel Buncic, and Jarkko Turunen. 2010. "The Impact of ECB Monetary Policy Decisions and Communication on the Yield Curve." *Journal of the European Economic Association* 8: 1266–98.
- Campbell, Jeffrey, Charles Evans, Jonas Fisher, and Alejandro Justiniano. 2012. "Macroeconomic Effects of Federal Reserve Forward Guidance." *Brookings Papers on Economic Activity*, Spring, 1–53.
- Carvalho, Carlos, Eric Hsu, and Fernanda Nechio. 2016. "Measuring the Effect of the Zero Lower Bound on Monetary Policy." Federal Reserve Bank of San Francisco Working Paper 2016–06.
- Christiano, Lawrence, Martin Eichenbaum, and Sergio Rebelo. 2011. "When Is the Government Spending Multiplier Large?" *Journal of Political Economy* 119: 78–121.
- Cox, Jeff. 2016. "Deutsche CEO: Negative Rates Have 'Fatal Consequences.'" CNBC, August 25.
- D'Amico, Stefania, and Mira Farka. 2011. "The Fed and the Stock Market: An Identification Based on Intraday Futures Data." *Journal of Business and Economic Statistics* 29: 126–37.
- Debortoli, Davide, Jordi Galí, and Luca Gambetti. 2018. "On the Empirical (Ir) Relevance of the Zero Lower Bound Constraint." Barcelona GSE Working Paper 1013.
- Demiralp, Selva, Jens Eisenschmidt, and Thomas Vlassopoulos. 2017. "Negative Interest Rates, Excess Liquidity, and Bank Business Models: Banks' Reaction to Unconventional Monetary Policy in the Euro Area." Unpublished manuscript, European Central Bank.
- De Rezende, Rafael. 2017. "The Interest Rate Effects of Government Bond Purchases Away from the Lower Bound." *Journal of International Money & Finance* 74: 165–86.
- Di Maggio, Marco, Amir Kermani, and Christopher Palmer. 2016. "How Quantitative Easing Works: Evidence on the Refinancing Channel." NBER Working Paper 22638. Cambridge, Mass.: National Bureau of Economic Research.
- Eggertsson, Gauti, Ragnar Juelsrud, and Ella Getz Wold. 2017. "Are Negative Nominal Interest Rates Expansionary?" NBER Working Paper 24039. Cambridge, Mass.: National Bureau of Economic Research.
- Eggertsson, Gauti, and Michael Woodford. 2003. "The Zero Interest-Rate Bound and Optimal Monetary Policy." *Brookings Papers on Economic Activity*, Spring, 139–235.
- Flynn, Finbarr. 2016. "Negative Interest Rates Confront Growing Chorus of CEO Critics." Bloomberg, April 27.
- Gagnon, Joseph, Matthew Raskin, Julie Remache, and Brian Sack. 2011. "The Financial Market Effects of the Federal Reserve's Large-Scale Asset Purchases." *International Journal of Central Banking* 7: 3–43.

- Greenlaw, David, James Hamilton, Ethan Harris, and Kenneth West. 2018. "A Skeptical View of the Impact of the Fed's Balance Sheet." NBER Working Paper 24687. Cambridge, Mass.: National Bureau of Economic Research.
- Greenwood, Robin, and Dimitri Vayanos. 2012. "Bond Supply and Excess Bond Returns." *Review of Financial Studies* 27: 663–713.
- Gros, Daniel, Christophe Blot, Paul Hubert, Maria Demertzis, and Guntram Wolff. 2016. "How Do Low and Negative Interest Rates Affect Banks' Activity and Profitability in the Euro Area?" *European Parliament Monetary Dialogue* 2016–04, November.
- Gürkaynak, Refet, Brian Sack, and Eric Swanson. 2005. "Do Actions Speak Louder than Words? The Response of Asset Prices to Monetary Policy Actions and Statements." *International Journal of Central Banking* 1: 55–93.
- Gürkaynak, Refet, Brian Sack, and Jonathan Wright. 2007. "The U.S. Treasury Yield Curve: 1961 to the Present." *Journal of Monetary Economics* 54: 2291–2304.
- Hamilton, James, and Jing Cynthia Wu. 2012. "The Effectiveness of Alternative Monetary Policy Tools in a Zero Lower Bound Environment." *Journal of Money, Credit, and Banking* 44: 3–46.
- Joyce, Michael, Ana Lasasoa, Ibrahim Stevens, and Matthew Tong. 2011. "The Financial Market Impact of Quantitative Easing in the United Kingdom." *International Journal of Central Banking* 7: 113–61.
- Kiley, Michael, and John Roberts. 2017. "Monetary Policy in a Low Interest Rate World." *Brookings Papers on Economic Activity*, Spring 2017, 317–72.
- Koetter, Michael, Natalia Podlich, and Michael Wedow. 2017. "Inside Asset Purchase Programs: The Effects of Unconventional Policy on Bank Competition." European Central Bank Working Paper 2017.
- Krishnamurthy, Arvind, and Annette Vissing-Jorgensen. 2011. "The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy." *Brookings Papers on Economic Activity*, Fall, 215–65.
- . 2012. "The Aggregate Demand for Treasury Debt." *Journal of Political Economy* 120: 233–67.
- Krugman, Paul. 1998. "It's Baaack! Japan's Slump and the Return of the Liquidity Trap." *Brookings Papers on Economic Activity*, Fall, 137–187.
- Leombroni, Matteo, Andrea Vedolin, Gyuri Venter, and Paul Whelan. 2017. "Central Bank Communication and the Yield Curve." Unpublished manuscript, Boston University.
- Lopez, Jose, Andrew Rose, and Mark Spiegel. 2018. "Why Have Negative Interest Rates Had Such a Small Effect on Bank Performance? Cross Country Evidence." NBER Working Paper 25004. Cambridge, Mass.: National Bureau of Economic Research.
- Reifschneider, David, and John Williams. 2000. "Three Lessons for Monetary Policy in a Low Inflation Era." *Journal of Money, Credit, and Banking* 32: 936–66.
- Rodnyansky, Alexander, and Olivier Darmouni. 2017. "The Effects of Quantitative Easing on Bank Lending Behavior." *Review of Financial Studies* 30: 3858–87.

- Skaperdas, Arsenios. 2017. "How Effective Is Monetary Policy at the Zero Lower Bound? Identification through Industry Heterogeneity." Federal Reserve Board Finance and Economics Discussion Series 2017–073.
- Swanson, Eric. 2011. "Let's Twist Again: A High-Frequency Event-Study Analysis of Operation Twist and Its Implications for QE2." *Brookings Papers on Economic Activity*, Spring, 151–88.
- . 2018. "Measuring the Effects of Federal Reserve Forward Guidance and Asset Purchases on Financial Markets." NBER Working Paper 23311. Cambridge, Mass.: National Bureau of Economic Research.
- Swanson, Eric, and John Williams. 2014. "Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates." *American Economic Review* 104: 3154–85.
- Turk, Rima. 2016. "Negative Interest Rates: How Big a Challenge for Large Danish and Swedish Banks?" IMF Working Paper WP/16/198.
- Wearden, Graeme. 2011. "Rick Perry: 'Quantitative Easing Akin to Treason.'" *The Guardian*, August 11.
- Wenik, Ian. 2016. "Barclays (BCS) CEO Staley: 'Negative Interest Rates Are Not Helpful.'" *The Street.com*, September 12.
- Williams, John. 2009. "Heeding Daedalus: Optimal Inflation and the Zero Lower Bound." *Brookings Papers on Economic Activity*, Fall, 1–37.
- . 2013. "Lessons from the Financial Crisis for Unconventional Monetary Policy." Paper prepared for panel discussion at NBER Conference on Lessons from the Financial Crisis for Monetary Policy, October 18. <http://www.frbsf.org/our-district/press/presidents-speeches/williams-speeches/2013/october/research-unconventional-monetary-policy-financial-crisis>.
- Wu, Jing Cynthia, and Fan Dora Xia. 2016. "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound." *Journal of Money, Credit and Banking* 48: 253–91.

Comments on Monetary Policy at the Effective Lower Bound

ABSTRACT Constraints on the setting of short-term interest rates due to the effective lower bound are likely to bind more often in the future than in the past if the neutral real rate of interest remains in the neighborhood of 1 percent. This paper argues that the Federal Open Market Committee should commit to pursuing a “lower-for-longer” or “makeup” strategy for setting short-term rates when the zero bound binds. This strategy is consistent with the goal of targeting 2 percent inflation, on average, over the business cycle. A “lower-for-longer” approach would improve economic performance during zero-lower-bound episodes and avoid an erosion of inflation expectations.

In the coming years, the Federal Reserve faces the significant issue of how to provide the accommodation the economy needs to recover from future downturns. This issue is important because constraints on the setting of short-term rates due to the zero (or effective) lower bound on interest rates may well bind more often in the future than they have in the past. Michael Kiley and John Roberts (2017) recently showed that the zero bound would constrain monetary policy 40 percent of the time if the neutral nominal short-term rate is 3 percent and the Federal Open Market Committee (FOMC) conducts policy by following a standard monetary policy rule, such as the Taylor rule. The consequence would be poor economic performance with significant shortfalls in output and employment during zero-lower-bound episodes. In addition, with inflation averaging about 2 percent when the zero lower bound does not bind and often declining to below 2 percent when it does, inflation, on average, will fall short of the

Conflict of Interest Disclosure: The author did not receive any financial support from any firm or person for this paper or from any firm or person with a financial or political interest in this paper. She is currently not an officer, director, or board member of any organization with an interest in this paper.

FOMC's 2 percent target. Such a persistent inflation shortfall could erode inflation expectations over time, compounding the zero-lower-bound constraint by lowering the normal nominal short-term rate consistent with any given neutral real rate.

The increased relevance of the zero lower bound reflects the fact that the neutral real rate of interest (r^*) looks to have declined considerably in recent decades while inflation expectations have become well anchored around the FOMC's 2 percent target. There is a good deal of uncertainty about both the current magnitude and future evolution of r^* . Empirical estimates are sensitive to methodology. However, updated estimates from Thomas Laubach and John Williams's (2003) model place r^* at currently just a bit under 1 percent. And FOMC participants estimate that the longer-run normal real federal funds rate is in this same range, with the median estimate of r^* between 0.75 and 1.00 percent. The causes of the decline in r^* , which is also evident in other advanced economies, are uncertain, but appear to be structural and persistent. They include low productivity growth; declining trend labor force growth, reflecting aging societies; and an increased preference for safe assets.

To improve economic performance, the FOMC could consider a number of approaches. Some involve the deployment of unconventional tools, such as longer-term asset purchases, interventions to directly target longer-term yields (similar to the Bank of Japan's yield curve control approach); and negative nominal interest rates. Other approaches, such as raising the inflation target or adopting price or nominal GDP targeting, entail a change in policy goals. I have argued that asset purchases have worked and thus should remain in the Fed's tool kit. But this tool will not likely fully alleviate the zero-lower-bound problem. The other approaches that I have mentioned deserve study and debate, but I see considerable disadvantages with each of them. Their shortcomings were recently summarized by my colleague, Ben Bernanke (2017). I agree with his assessments and, given the space constraints of this paper, I instead focus on an approach I consider promising—one that is evolutionary, practical, and has the potential to significantly mitigate the adverse effects of the zero lower bound.

I believe the FOMC should seriously consider pursuing a “lower-for-longer,” or “makeup,” strategy for setting short-term rates when the zero lower bound binds and, ideally, articulate its intention to do so before the next zero-lower-bound episode. The phrase “lower for longer” is due to David Reifschneider and John Williams (2000), who suggested that the Fed, during a zero-lower-bound period, could keep track of the cumulative deviations of short-term rates from the recommendations of a simple rule

(the Taylor rule) and then “work off” or “make up” these accommodation shortfalls over time by holding short-term rates lower for a longer time than the rule would recommend. Under this strategy, the average level of short rates over a period of, say, 10 years from the onset of a zero-lower-bound episode would be essentially unaffected by the zero-lower-bound constraint. If this strategy is understood and credible, it should cause long-term rates to decline when the zero lower bound begins to bind by about as much as would occur in the absence of any effective lower bound at all—that is, if the FOMC could set negative rates. Under the empirically reasonable assumption that what matters for aggregate spending is the entire expected path of short-term rates rather than just the current level, this strategy enables the Fed to provide substantial additional accommodation during zero-lower-bound episodes. The strategy also potentially supports aggregate demand by raising inflation expectations, thereby lowering real long-term rates relative to a Taylor rule-type baseline.

A key characteristic of lower-for-longer strategies is that they do not treat “bygones” as “bygones.” In determining the timing of exit from the zero lower bound and the subsequent path of short-term rates, the FOMC must consider not only the current state of the economy—the levels of the output gap and inflation, as in the Taylor rule—but also a measure of past performance shortfalls during the zero-lower-bound period—either cumulative shortfalls in monetary accommodation or cumulative shortfalls in output and inflation relative to the FOMC’s targets. In effect, the FOMC would augment the usual factors incorporated in standard rules with an adjustment reflecting the severity of the zero-lower-bound episode. One technique is to characterize the appropriate policy path by reference to a “shadow” rate of interest that cumulates accommodation shortfalls. Reifschneider and Williams (2000) and Kiley and Roberts (2017) have proposed policies along these lines. An alternative approach is to characterize the same shortfalls in metrics relating to economic outcomes. Bernanke recently proposed such an approach—“flexible temporary price level targeting.” He suggests that the FOMC hold rates at the zero lower bound at least until the cumulative shortfall in inflation from a 2 percent trend during the zero-lower-bound period has been eliminated and until unemployment has also at least declined to its natural rate.

This lower-for-longer strategy is attractive for several reasons. First, it is evolutionary; it builds on the existing flexible inflation targeting framework with an unchanged 2 percent longer-run inflation objective. This framework is well understood and, in my view, has contributed considerably to the attainment of good macroeconomic performance in the United

States. It has helped to anchor inflation expectations, which, in turn, has enabled the FOMC to all but ignore the inflationary implications of supply shocks—and to instead focus on stabilizing employment.

This approach would build on the forward guidance the FOMC offered after 2008, which was explicitly intended to bring down long-term rates. The FOMC adopted forward guidance pertaining to the path of short-term rates that increasingly shifted market expectations, in effect promising to hold them below rule-based recommendations for a substantial time into the recovery—the essence of the lower-for-longer approach. A full evaluation of the impact of this guidance is difficult, because it evolved over time, along with the market's understanding of the economic impact of the financial crisis, and was complemented by asset purchases. However, long rates fell 20 basis points when the FOMC announced in August 2011 that the funds rate would stay at zero at least through mid-2013. The yield curve moved down further as the date moved out and the FOMC, in 2012, adopted threshold-based guidance, pledging that it would not raise rates at least until unemployment declined below 6.5 percent if inflation was projected to run no higher than 2.5 percent. Also, professional forecasters significantly reduced their estimates of the unemployment rate that would prevail at the time of liftoff.

The FOMC subsequently provided guidance that further pushed out the likely date of liftoff, and it only began to raise short-term rates when the unemployment rate had declined to 5 percent, close to estimates of the longer-run normal rate of unemployment at the time. The FOMC's guidance commanded broad FOMC support and, in my view, a significant degree of commitment, although it did not meet the “Odyssean” standard. When the FOMC finally began to raise rates, it promised a “gradual” approach, holding the funds rate below Taylor-type rule recommendations, even with an adjustment for the estimated decline in r^* . In addition, the FOMC has more recently emphasized the symmetry of the 2 percent inflation objective, and recent projections envision an overshoot of the 2 percent target in 2020, although there has been no statement or indication that the FOMC intends or considers it appropriate to “make up” for the cumulative inflation shortfall. This general approach, including the adoption of explicit quantitative thresholds, was consistent with the spirit of lower-for-longer-type recommendations and provides a basis on which the FOMC could build.

For the lower-for-longer approach to work well in future zero-lower-bound episodes, the FOMC needs to make a credible statement endorsing such an approach, ideally before the next economic downturn. This could take the form of a revision of the FOMC's “Statement on Longer-Run Goals

and Monetary Policy Strategy,” or it could be couched as an addendum to this statement on “Committee Guidelines for Implementing Policy at the Zero Lower Bound.” Such a statement should enunciate the approach; show that the FOMC understands and embraces its implications; and, to the extent possible, provide quantitative guidance about how the policy would be implemented. It would be important for the FOMC to emphasize to both the public and Congress that the lower-for-longer approach is consistent with the Federal Reserve’s congressional dual mandate to pursue maximum employment and price stability.

The FOMC could explicitly endorse the approach that it will set short-term rates lower for longer than would be called for by standard monetary policy rules when the zero lower bound binds. It could elaborate that in setting the policy path, it will take previous performance shortfalls into account. It could establish quantitative thresholds consistent with such a strategy. For example, it could indicate that, after a period of very weak economic activity and inflation below 2 percent, it would generally be appropriate to wait to raise rates at least until the unemployment rate has declined to estimates of its normal longer-run level and inflation has stably returned to 2 percent—typical requirements of the policy path under any implementation.

As Bernanke has suggested, the FOMC could go further, stating that a condition for raising rates is that cumulative shortfalls of inflation from 2 percent have been erased. It could emphasize that once the Fed begins to raise short-term rates, it expects to close the gap with normal rule-based recommendations only gradually, in order to compensate for the shortfall in accommodation provided during the period when policy rates were constrained and as an appropriate risk management strategy. It could also emphasize that it anticipates that the additional stimulus provided by this approach will result in a period of exceptionally low unemployment, and that inflation would likely overshoot the FOMC’s symmetric 2 percent target for a time, emphasizing the desirability of compensating for a previous shortfall to avoid an erosion of inflation expectations. It could articulate that the FOMC’s objective is to achieve inflation near 2 percent, on average, over the business cycle.

The lower-for-longer strategy entails some costs and risks that need to be assessed and managed. By keeping interest rates unusually low after the zero lower bound no longer binds, the lower-for-longer approach promises, in effect, to allow the economy to boom after a zero-lower-bound episode. Unemployment will typically undershoot the longer-run normal rate of unemployment—a development that, in and of itself, I consider beneficial.

Inflation would also typically rise above the 2 percent inflation target for some period, albeit not on a permanent basis. One could argue, in this regard, that such an overshoot helps to keep inflation at 2 percent “on average” (as is explicit in Bernanke’s proposal), and that it is desirable because it makes the long-run level of prices more predictable—although this line of reasoning raises the question of whether the FOMC should more generally pursue a flexible price-level targeting strategy—an approach that would involve a very substantial alteration in the FOMC’s policy framework. A prolonged period of inflation above 2 percent could potentially unanchor inflation expectations; and prolonged boom conditions could undermine financial stability. These concerns may militate in favor of some “tempering” in the application of the lower-for-longer approach. They raise the issue, as well, of whether the FOMC could credibly commit to such a plan. Market participants could well question whether the FOMC would allow the economy to “overheat,” and they might see an incentive for the FOMC to renege. Although the FOMC can never bind future committees to a particular course of action, I think that incorporating a set of widely supported principles into the FOMC’s strategy statement would ameliorate this problem. Let me conclude by saying that I consider this approach worthy of consideration by the FOMC and of more general public debate.

References

- Bernanke, Ben S. 2017. "Monetary Policy in a New Era." Paper prepared for conference on Rethinking Macroeconomic Policy, Peterson Institute for International Economics, Washington, October 12–13.
- Kiley, Michael T., and John M. Roberts. 2017. "Monetary Policy in a Low Interest Rate World." *Brookings Papers on Economic Activity*, Spring, 317–72.
- Laubach, Thomas, and John C. Williams. 2003. "Measuring the Natural Rate of Interest." *Review of Economics and Statistics* 85, no. 4: 1063–70.
- Reifschneider, David, and John C. Williams. 2000. "Three Lessons for Monetary Policy in a Low-Inflation Era." *Journal of Money, Credit and Banking* 32, no. 4, pt. 2: 936–66.

GENERAL DISCUSSION James Hamilton responded to comments from Eric Swanson. He argued that Swanson's theory that unconventional monetary policy can have large effects at the effective lower bound (ELB) is more related to forward guidance than to large-scale asset purchases (LSAPs). Hamilton noted that his own choice in studying the Federal Reserve's policy announcement in March 2009 was because it was the most obvious, clear example of a "true" LSAP program, not because its effects seemed to die out. Hamilton further noted that Swanson counted the March 2009 FOMC policy announcement by the Fed as a "pure" policy shock. Hamilton instead argued that disentangling the actual effect of the LSAP program from the signaling effect was quite difficult. He further clarified that his argument is not that LSAPs had no effect, but that their effects are very easy to overestimate.

Janet Yellen agreed with Eric Swanson's conclusion that the bulk of evidence does suggest that LSAPs had a significant effect on interest rates and the economy. She did, however, agree with James Hamilton that it is difficult to disentangle the direct effect of asset purchases from changes in the public's expectations of future interest rates and the public's interpretation of the Fed's outlook for the economy when asset purchases are announced. She noted that although they would not be her preferred tool, she believed that LSAPs should remain a part of the Fed's tool kit. She noted that it was difficult to see the adverse side effects of LSAPs, despite the concerns by many that these effects would be large. In addition, she noted that forward guidance also seems quite important; hence, her focus on enhancing its role as part of the Fed's framework.

Eric Swanson shared some of James Hamilton's concerns about high-frequency event study estimates of the effects of the Fed's LSAPs. Estimates of LSAP effects should not rely only on these types of studies. However, he noted that substantial evidence from a variety of methods, data sets, and countries shows similar effects. These other studies thus corroborate the high-frequency event study estimates.

Kristin Forbes asked Eric Swanson about the practicality of implementing negative interest rates in the United States. She suggested that although a number of countries implemented negative interest rates after the Great Recession, the feasibility of such a policy varied across countries. For example, in the United Kingdom, the Bank of England had worried that setting interest rates below 0.5 percent would lead to problems for building societies and other segments of the financial sector. Would other countries, particularly those with similarly large financial sectors, face similar constraints—or other types of concerns?

Second, Forbes asked Yellen about her proposed “makeup” strategy for monetary policy after an ELB episode. Though Forbes understands the intuition behind “overshooting” and providing additional monetary policy accommodation after a period at the ELB, she wonders about the time horizon for such a policy. For example, if inflation expectations are well grounded, but inflation had been below target for almost a decade, should the makeup period also extend for as long as a decade? Forbes suggested a shorter period.

Janet Yellen responded to Forbes by noting that, if such a makeup policy were put in place before a recession, significant undershooting of the Fed’s inflation target might not actually occur in the first place, because the Fed would be promising to provide additional accommodation in the future, shifting the public’s expectations. If such a long period of undershooting did occur, then Yellen acknowledged the risk that a long makeup policy could unanchor inflation expectations by overshooting for too long. In such a scenario, the Fed might want to temper its approach.

Raghuram Rajan noted that, as an academic and former leader of a central bank, he was surprised at how many of the basic facts about monetary policy during the global financial crisis are still in dispute. He noted debate during discussion of Ben Bernanke’s paper in the panel’s previous session about whether Taylor Rule residuals were positive or negative before the crisis, and consequently whether policy was too accommodative. In the context of the current discussion, he noted ongoing debate about whether the effects of the LSAPs were actually positive, and whether unconventional monetary policy had a significant effect on exchange rates and capital flows. He viewed this debate as an indicator that further research is still needed.

Rajan focused on Kristin Forbes’s comments, and he noted his pleasure that someone was willing to publicly acknowledge that exchange rates matter for setting monetary policy, as many central bankers had refused to do in the past. Rajan asked Forbes whether exchange rates are an important channel for central bankers to consider when conducting unconventional monetary policy. For example, it seemed that one reason the European Central Bank (ECB) began asset purchases was because most other central banks were conducting asset purchases, even while the ECB was not, causing an unwanted appreciation of the euro. He asked whether exchange rates and the cross-border transmission of policy should be considered in future debates about unconventional monetary policy.

Joshua Hausman noted his surprise that the panel did not discuss the case of Japan's unconventional monetary policy, and he asked what lessons they would take from unconventional monetary policy conducted as part of Prime Minister Shinzo Abe's "Abenomics" policy program. Particularly, he noted the country's failure to get inflation to its 2 percent target, despite dramatic and unconventional policy actions, such as yield curve control (that is, setting the 10-year government bond yield at 0 percent).

Mark Gertler acknowledged James Hamilton's critique of event study analysis in evaluating the effects of quantitative easing (QE), but Gertler struck a more optimistic tone about the policy. Specifically, he noted that the Federal Reserve's purchases of agency mortgage-backed securities were quite effective. The mechanism through which the policy acted was to affect excess returns on long-term securities that were unusually high due to financial market frictions. Without these frictions, such a policy would not have been effective, because speculators would arbitrage away the excess returns. During the global financial crisis, however, these frictions were clearly present because excess returns persisted. For government bonds, the excess returns manifested through higher term premiums. Though term premiums are quite difficult to measure, the interest rate spreads on asset-backed securities over government bonds were quite elevated. After QE1 was first announced in December 2008 and then implemented in March 2009, these spreads compressed. The same occurred after the announcement and implementation of QE3. Gertler viewed this as evidence of QE's effectiveness. Hamilton asked a clarifying question as to whether Gertler believed these financial frictions were persistent through the end of 2014, after QE3 was implemented. Gertler expressed his confidence that frictions were present for a few years after the crisis, but that it was not clear exactly how long they persisted.

Janet Yellen commented on the difficult experience of Japan, and saw the country's experience as a warning for how important it is to get inflation up after a binding ELB period. Yellen noted that inflation expectations in Japan had likely fallen and that such a phenomenon is a very difficult process to reverse.

Eric Swanson noted that monetary policy in Japan in the early 2000s was not very good. For example, the Bank of Japan only conducted asset purchases in a very superficial way, buying very-short-term government bonds. As a result, the Bank of Japan's policies had relatively little effect and the Japanese economy seems to have fallen into a deflationary expectations equilibrium. Swanson pointed to research by S. Boragan Aruoba, Pablo Cuba-Borda, and Frank Schorfheide that shows the United States is

in a “normal” equilibrium and Japan is in a deflationary equilibrium.¹ Once Japan fell into a deflationary equilibrium, Swanson argued, it became much harder for it to extricate itself and return to the normal equilibrium.

Kristin Forbes noted that Japan could be a good case study of how unconventional monetary policy works through exchange rates, perhaps to a greater degree than conventional monetary policy, as suggested by Raghuram Rajan. Japan’s unconventional monetary policy resulted in large exchange rate movements that seemed to fuel much of the policy’s economic stimulus. Studying this dynamic also addresses the potential spillovers of unconventional monetary policy working through different channels. Forbes expressed her surprise, when she joined the Bank of England as a Monetary Policy Committee member, about how exchange rate movements were incorporated into inflation forecasts and monetary policy decisions. Exchange rate shocks were modeled at the Bank of England as exogenous and resulting from risk shocks; little thought was given to other reasons why exchange rates might move and how this could determine the effects. Recent research, however, shows that exchange rate movements resulting from monetary policy shocks can result in much larger pass-through effects to inflation than exchange rate movements caused by other shocks. However, none of this was discussed at the Bank of England at the time. Forbes speculated that a possible reason may be that economists have been so ingrained to think that they cannot explain exchange rate movements, deterring them from attempting to model them. She argued that exchange rates should be a key part of the conversation on monetary policy, especially in countries like the United Kingdom where exchange rate pass-through effects can be large. She noted, however, that while emerging markets would also like exchange rates to be a bigger part of the conversation about the spillovers from monetary policy in other countries, this would be difficult to implement for advanced economies’ central banks, whose mandates are usually politically constrained to focus on the domestic sphere.

Jeff Fuhrer made two points. He first suggested looking at different ways that monetary policy might have been constrained, or ways in which it could have done more during the Great Recession. For example, he suggested looking at central banks’ loss functions in the wake of the Great Recession to measure the degree of overall welfare loss incurred across

1. Aruoba, S. Boragan, Pablo Cuba-Borda, and Frank Schorfheide, “Macroeconomic Dynamics Near the ZLB: A Tale of Two Countries,” *Review of Economic Studies* 85, no. 1 (2018): 87–118, <https://doi.org/10.1093/restud/rdx027>.

different economies despite sizable monetary and fiscal actions. Such an analysis might suggest that monetary policy could have done more to right the economy. Of course, all these estimates of the loss function are dependent on the models used. Second, he noted the danger of monetary policy solutions that rely too much on expectations. For one, during a crisis, the public might not find policymakers' promises of action in the future to be credible compared with actual action taken at the time of the crisis. He noted that is striking how little economists know about how expectations are actually formed, given how much monetary policy depends on expectations. He pointed to research by the University of California, Berkeley, economist Yuriy Gorodnichenko and the University of Texas–Austin economist Olivier Coibion on expectations formations, as well as work by the Harvard economist Andrei Shleifer and his colleagues.

Athanasios Orphanides remarked on monetary policy at the ELB. He noted that it would be useful for central banks to cut interest rates faster as they approach the ELB, opting to reach the ELB quicker than conventional monetary policy rules would recommend. He and Volker Wieland recommended in a 2000 paper that the Bank of Japan implement this strategy, to no avail.² Likewise, they had difficulty convincing policymakers to implement a similar strategy in 2008. He emphasized that the idea of “saving ammunition,” or waiting to cut interest rates down to the ELB during a time of crisis, should be permanently discarded.

Orphanides also asked the panel about negative interest rates. He noted that although there are no limits to the size of QE, there could be political effects of QE that are quite large. He noted that, instead, the Fed might consider announcing ahead of time how low it would be willing to cut interest rates in the next recession. Would it be willing to go to -1 or -1.5 percent? Announcing this ahead of time would change expectations about how likely it would be that the ELB actually binds in the future, and therefore might decrease the possibility that more controversial policies like QE would be needed at all.

Philipp Hartmann explained the ECB's experience in implementing negative interest rates. The ECB cut its deposit facility rate to -0.4 percent in four small steps between June 2014 and March 2016. New studies are now coming out about the effect of these policies, and most of the research suggests that negative interest rates worked in the euro area.

2. Athanasios Orphanides and Volker Wieland, “Efficient Monetary Policy Design Near Price Stability,” *Journal of the Japanese and International Economies* 14, no. 4 (2000): 327–65, <https://doi.org/10.1006/jjie.2000.0452>.

However, negative interest rates might not be the most powerful instrument because there is a limit to how low they can go, and therefore to how much accommodation can be provided. For example, research by Markus Brunnermeier and Yann Kobe suggests that below a certain level the policy could become counterproductive.³ The effects of the ECB's policy worked through the interest rate and bank lending channel. First, there was a "twist" and a "shift" in the yield curve. The twist was a result of negative rates acting as a charge on cash hoarding and triggering portfolio shifts toward long-term bonds compressing the term premium. The shift was simply a result of the removal of the nonnegativity constraint on future expected short-term rates. The second, and perhaps more surprising, positive effect of negative interest rates was through the bank lending channel. Several studies by ECB economists suggest that negative interest rates increased lending. Florian Heider, Farzad Saidi, and Glenn Schepens find this effect for the syndicated loans of banks with a relatively large share of market-based funding relative to retail-deposit based funding (because wholesale funding rates can go negative but retail rates do not, and therefore retail banks do not benefit from funding relief through negative policy rates).⁴ Jens Eisenschmidt and Smets and Selva Demiralp, Eisenschmidt, and Thomas Vlassopoulos present evidence of positive lending effects for broader credit measures, including banks with large retail deposit bases.⁵ The former research also finds the pass-through to lending rates to remain unchanged. It should, however, be kept in mind that negative rates were introduced by the ECB in parallel with other unconventional monetary policy measures, notably targeted long-term refinancing operations and asset purchase programs. Therefore, for some of the studies not all the lending effects can be associated with negative rates alone. But they can be seen as "activating"

3. Markus Brunnermeier and Yann Kobe, "The Reversal Interest Rate," working paper, Princeton University, March 2018.

4. Florian Heider, Farzad Saidi, and Glenn Schepens, "Life Below Zero: Bank Lending under Negative Policy Rates," ECB Working Paper 2173, August 2018, forthcoming in *Review of Financial Studies*, <https://www.ecb.europa.eu/pub/pdf/scpwps/ecb.wp2173.en.pdf?de58f4c0f6cf96f0a7d99014d0ed2454>.

5. Jens Eisenschmidt and Frank Smets, "Negative Interest Rates: Lessons from the Euro Area," unpublished paper, ECB, March 2018, <https://cepr.org/sites/default/files/40021%20Eisenschmidt%20Smets%202017%20March%202018.pdf>; Selva Demiralp, Jens Eisenschmidt, and Thomas Vlassopoulos, "Negative Interest Rates, Excess Liquidity, and Retail Deposits: Banks' Reaction to Unconventional Monetary Policy in the Euro Area," ECB Working Paper 2283, May 2018, <https://www.ecb.europa.eu/pub/pdf/scpwps/ecb.wp2283~2ccc074964.en.pdf?fb6d4de645fdd3ea2f6b24834bfd82c>.

the excess reserves induced in the system through those other measures, which would otherwise have remained idle.

Hartmann further discussed the effect of the exchange rate and agreed that the effects are country dependent. In the ECB's case, the effects also seemed to be state dependent—sometimes, the exchange rate effects of monetary policy were larger; other times, not so much. Most of the available evidence, however, suggests that there are fewer international spillovers from ECB unconventional monetary policy, particularly to emerging markets, because the euro is less of a global currency than the dollar.

Hartmann suggested that the panel discuss the cost-benefit analysis of different unconventional monetary policies given the potential for unintended side effects. He referenced this year's economic conference sponsored by the Federal Reserve Bank of Boston, which had captured this perspective very nicely.⁶ Given the lexicographic ordering of objectives in the ECB's mandate, it is somewhat hard to say with precision which weight potential side effects should receive. He wondered how the Fed evaluated the cost-benefit trade-offs of unconventional monetary policy, for example, given their potential effects on financial stability, and how this analysis pairs with the macroprudential regulation and supervision of financial institutions.

Michael Kiley agreed with many of the discussants that the evidence supported the efficacy of QE; at the same time, he argued that this evidence primarily focuses on the effect of QE on financial markets, and not on the transmission of these effects to real activity or inflation. Although research in the latter areas is limited, he thought the evidence suggested that QE may have been less effective on these dimensions—which are the ultimate objective of monetary policymakers—than Eric Swanson suggested. He expressed his support for the types of makeup policies that Yellen laid out in her presentation but wondered if they would be fully credible and appreciated by the public. He suggested, instead, consideration of the simpler approach of raising the inflation target to about 3 percent. Though he acknowledged that such a move might be unpopular, he argued that another large QE program of \$2 trillion to \$3 trillion might be just as unpopular.

Kristin Forbes first addressed the comments by Athanasios Orphanides. She disagreed with his comment that there are no constraints on QE. She

6. Conference on "What Are the Consequences of Long Spells of Low Interest Rates," Federal Reserve Bank of Boston, Boston, September 7–8, 2018, <https://www.bostonfed.org/consequences2018/agenda/>.

noted that QE is constrained by the characteristics of the assets and the size of the asset pool from which a central bank can purchase assets. Specifically, if a central bank could buy corporate bonds or other assets, it probably should not purchase corporate assets of the financial institutions that they regulate, or companies with significant risks. These restrictions—plus the size of the overall pool—place limits on the size of QE. But there are also other options to provide an unconventional stimulus than purchasing assets. She referenced an effective program implemented by the Bank of England known as the “funding for lending scheme.” In this program, the Bank of England set up an incentive system where banks received a subsidized lending rate for increasing their total lending. The goal of the program was to more efficiently pass through the reduction in interest rates by the Bank of England to customers as interest rates approached the ELB. Forbes and other Bank of England officials were surprised at the level of participation in the program despite the slower growth of the U.K. economy. She noted that this type of program could merit further investigation by other countries.

Forbes also agreed with Philipp Hartmann that cost-benefit analysis is important in monetary policy. She noted that it is important not to disregard certain costs as outside the mandate of a given central bank, because these costs can accumulate over time, especially when interest rates are low.

Eric Swanson addressed the question from Jeff Fuhrer regarding whether the Federal Reserve could have provided additional monetary policy accommodation during the global financial crisis. Swanson argued that, in retrospect, the Fed definitely could have done more, particularly from 2009 to 2011. Stronger forward guidance could have brought the two-year Treasury yield down from 1 percent during this period to close to 0 percent. He contended that the only reason the Fed did not do this was because it was still figuring out how to communicate forward guidance to the markets more effectively. It took the Fed until August 2011, when it implemented date-based forward guidance, to figure out how to better lower interest rate expectations.

Regarding the idea for a higher inflation target raised by Michael Kiley, Swanson mentioned a symposium at the San Francisco Federal Reserve Bank that he attended at which every participant opposed raising the inflation target.⁷ The Boston College economist Peter Ireland particularly argued

7. “A New Target for Monetary Policy,” symposium at the Federal Reserve Bank of San Francisco, March 2, 2018.

at the symposium that, based on research in the 1990s by Martin Feldstein and others, inflation as low as 3 percent could still be quite costly.⁸

Janet Yellen addressed the question of negative interest rates. She noted that when the Federal Reserve was considering cutting interest rates down to 0 percent instead of the range of 0 to 0.25 percent to which it ultimately did cut rates, it was worried about the distortionary effects of very low rates. Particularly, it was concerned about the functioning of money markets and that banks would not or could not pass through very low or negative rates to retail depositors. Yellen noted her surprise that so many European countries and Japan were able to cut interest rates to as negative a level as they did. She noted that there is some research showing the effects of negative interest rates should be evaluated through the bank lending channel, and that there could be adverse side effects. She noted that this is a topic worth studying for the Fed in the future.

Regarding a higher inflation target, Yellen suggested that solving the problem of the ELB would actually require an inflation target higher than 3 percent. An inflation target that high would call into question whether the Fed would be meeting its price stability mandate, and she doubted that Congress would consider such a high inflation target as consistent with price stability. Therefore, such a policy change would be unpopular politically.

8. Martin Feldstein, "The Costs and Benefits of Going from Low Inflation to Price Stability," in *Reducing Inflation: Motivation and Strategy*, edited by Christina Romer and David Romer (University of Chicago Press, 1997).