American Economic Journal: Macroeconomics 2022, 14(4): 136–173
https://doi.org/10.1257/mac.20190337

Age Structure and the Impact of Monetary Policy[†]

By John V. Leahy and Aditi Thapar*

We exploit cross-sectional variation in the response of US states to an identified monetary policy shock to study how the impact of monetary policy varies with the age structure of the population. We find that the economy's response is weaker the greater the share of population under 35 years of age and stronger the greater the share between 40 and 65. We find that all age groups become more responsive to monetary policy shocks when the proportion of the middle-aged increases. We provide evidence consistent with middle-aged entrepreneurs starting and expanding businesses in response to an expansionary monetary shock. (*JEL* E23, E24, E32, E43, E52, J11, R23)

The world is undergoing a period of demographic change. Populations are aging in much of the developed world, especially Japan and continental Europe. At the other extreme, over 50 percent of the population of India is below the age of 25.[1] In China, the on-again, off-again one-child policy is likely to bring large fluctuations in the age distribution. Given the life cycle patterns in many economic activities such as saving, home purchases, education, and retirement, it is of interest to understand how changes in the age distribution affect the performance of the tools of economic policy. In this paper we focus on the relationship between the age structure of the population and monetary policy.

To understand how demographics affect monetary policy, we need a sample of economies with differing demographic structures, and we need a measure of monetary policy. We use the US states as our laboratory. The states differ in their age structures, and data are available on a consistent basis for a wide range of economic and demographic variables. Since the states share a single monetary policy, we do not have to control for policy differences across states or worry that policy might be responding to the demographic structure of a given state. This allows us to focus on how changes in the age structure influence the response to this common policy. Of course, states differ in other ways, some of which may be correlated with

*Leahy: University of Michigan and NBER (email: jvleahy@umich.edu); Thapar: University of Michigan (email: athapar@umich.edu). Simon Gilchrist was coeditor for this article. We would like to thank Viral Acharya, V.V. Chari, Jordi Gali, Christopher House, Amartya Lahiri, Pablo Ottonello, Lasse Pedersen, Morton Ravn, Nick Sly, Silvana Tenreyro, Chris Waller, Jonathan Willis, and Arlene Wong for helpful comments and discussions, and Jessica Meadows and Amanda West for valuable research assistance. Portions of this paper were completed while Thapar was visiting the Federal Reserve Bank of Kansas City as part of CSWEP's Summer Economic Fellows Program. We are grateful to the National Science Foundation for their support under grant SES-1919362. Earlier versions of this paper circulated under the title "Demographic Effects on the Impact of Monetary Policy."

[†]Go to https://doi.org/10.1257/mac.20190337 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

[1] Census of India (2011).

age structure, and we will attempt to control for these differences in what follows. Heterogeneity, however, is present (more or less) in nations as well.

For monetary policy, we follow the high-frequency approach to the identification of monetary policy shocks and consider changes in the federal funds rate over a short window around policy announcements.[2] We ask how the age distribution in a state affects the response of state income and employment to a monetary policy shock. We find that monetary policy is more effective—in the sense that income and employment respond more to a given change in the interest rate—the greater the share of the middle-aged population, defined as those aged 40 to 65, and less effective the greater the share of the young, defined as those aged 20 to 40. The share of the population over 65 has no discernible effect on the effectiveness of monetary policy.

This basic pattern in which the middle-aged amplify and the young dampen the effects of monetary policy appears robust. It holds for a variety of definitions of the monetary policy shocks and for a variety of empirical specifications. We consider shocks estimated from a structural VAR and shocks constructed from the minutes of the Federal Open Market Committee's meetings and the Federal Reserve Board's Greenbook forecasts (Romer and Romer 2004). We investigate altering the start and end dates of the sample, the removal of small states, and aggregation to census regions.

The effects of demography also appear sizable and comparable in size to the effects of a monetary policy shock on aggregate income. In the baseline analysis, a one-standard-deviation increase in the fraction of 30–34-year-olds reduces the impact of a monetary policy shock by approximately 40 percent, whereas an increase in the fraction of 50–54-year-olds increases the impact by a similar magnitude.

While we identify shocks to monetary policy, variation in the demographic structures of states is potentially endogenous. Absent instruments for the age distribution, we do the next best thing and control for state characteristics that may be correlated with demographics. We do this in two ways. First we interact the state fixed effect with the monetary shock. This controls for time-invariant state characteristics that may affect monetary policy transmission but leaves open the question of whether there might be some time-varying state characteristics that are correlated with the age structure. To address the latter concern, we consider specific state characteristics that might affect monetary policy transmission. We investigate three possibilities in detail: state industrial composition, state income, and state housing market characteristics. If, for example, a large middle-aged population were correlated with a large manufacturing sector, the greater response of manufacturing to monetary policy shocks could explain the greater sensitivity to interest rates. Alternatively, Cravino, Lan, and Levchenko (2018) have argued that richer individuals consume a greater share of sticky-price goods. If a large middle-aged population were correlated with high income, stickier prices could then explain the greater response of middle-aged states to monetary shocks. Finally, several authors have found relationships between housing markets and the efficacy of monetary policy. Di Maggio et al. (2017) find that the fraction of adjustable-rate mortgages

[2] See Gürkaynak et al. (2005); Gertler and Karadi (2015); and Nakamura and Steinsson (2018) for prominent examples of this approach.

affects the transmission of monetary policy, while Beraja et al. (2019) find that the recent behavior of house prices affects the transmission of monetary policy. A correlation of age structure with either characteristic could explain our results. We find no such correlations and therefore no evidence that industrial composition, income, or housing markets are driving our results. We also consider a number of other potential correlates including sex, race, and education. The basic pattern survives the inclusion of all of these variables.

Given the absence of an obvious confounding variable, we entertain the possibility that our effects are causal: that an exogenous increase in the proportion of the middle-aged population enhances monetary policy while an exogenous increase in the proportion of the young dampens its effects. We ask what mechanism might generate such an effect. Before considering potential mechanisms, we take a look at what other variables besides income and employment respond differently to monetary policy when the demographics change. First we consider population and unemployment. The idea is to see whether the response of employment to a monetary policy shock is due to hiring within the state or migration between states. If the response were due mainly to migration, then this would suggest caution in applying the results to sovereign nations, as migration between states is likely easier than migration between nations. We find that the response of unemployment mirrors the response of employment, suggesting an important role for internal adjustment. We look at the response of state populations and find some evidence that the age structure of the state affects the response of migration to a monetary policy shock. The response of migration to the age structure, however, appears to be largely orthogonal to the responses of income and employment. Tight monetary policy appears to increase outflows in states with more 35–40-year-olds and inflows in states with more elderly residents. Both are age groups that do not greatly affect the response of income or employment.

Next we ask what types of jobs and incomes are affected by the demographic structure of the economy. While all sectors and income types share the basic pattern, we find that the responses of construction employment and proprietors' income are both an order of magnitude larger than the rest, suggesting that they play a potential role in any mechanism.

A greater proportion of middle-aged residents could amplify the effects of monetary policy in a number of ways. On one hand, the income and employment of middle-aged residents could respond more to monetary policy shocks. Having more middle-aged residents would then explain the greater aggregate response. On the other hand, having more middle-aged residents could increase the responsiveness of the income and employment of all age groups. We find evidence for the second channel. We look at the response of income and employment broken down by age. We find that the incomes and employment of all age groups follow the same basic pattern: dampened if there are more young in the state and amplified if there are more middle-aged residents. It is not so much being middle-aged that matters as being surrounded by middle-aged people.

In the final section, we consider potential mechanisms by which a greater share of middle-aged residents might affect the transmission of monetary policy. It seems natural that any potential mechanism would have something to do with changes in economic behavior over the life cycle. We discuss three such mechanisms: the

young tend to be liquidity constrained, whereas home buyers and entrepreneurs tend to be middle-aged. This list is certainly not exhaustive. Each of these mechanisms can explain aspects of the data and each has its shortcomings. Theories based on liquidity constraints tend to predict that the young should be more responsive to monetary policy shocks, and, while middle-aged home buyers can help explain the sensitivity of construction to age structure, home buyers in their thirties look very much like home buyers in their forties, suggesting that both groups should have a similar impact on monetary policy.

Our preferred channel works through entrepreneurship. Selgado (2018) documents that entrepreneurs tend to be older than workers and that the share of entrepreneurs in the 35–60 age group is larger than the share of entrepreneurs in the 22–34 age group. Azoulay et al. (2019) argue that successful entrepreneurs are middle-aged, not young. If a larger share of middle-aged residents implies a larger number of potential entrepreneurs, then the elasticity of business formation with respect to a decline in interest rates is likely larger in states with a greater share of the middle-aged population. We find that business formation does in fact respond more strongly the greater the share of the middle-aged population. We find evidence for a multiplier by which the initial expansion increases business income, furthering the incentive to invest. We also find evidence of financial frictions. A greater fraction of middle-aged residents increases the response of small business employment to a monetary policy shock. The hypothesis is then that a monetary expansion increases the returns to starting a small business both by lowering the cost of borrowing and by increasing the profitability of a business once it is started. The greater the proportion of the middle-aged, the greater the response to these incentives. All income groups then benefit from this expansion.

In Section I we describe recent background literature that is relevant to our investigation. In Section II we introduce our empirical strategy. We present our main results in Section III. We discuss interpretations of our results in Sections IV, V, and VI. Section VII concludes.

I. Related Literature

This paper contributes to the small but growing literature on the relationship between monetary policy and the demographic structure of the economy.[3] Several recent papers study how monetary policy shocks impact the consumption of different age groups. Sterk and Tenreyro (2018) find that monetary policy shocks increase the consumption of those of working age relative to the old. In their Appendix, they argue that this result is driven by younger workers and durable consumption. Wong (2019) finds that the consumption of the young is more sensitive to monetary policy shocks and argues that young homeowners, who either obtain a new loan or refinance, are the primary drivers of this response. These studies focus on household

[3]There is a related literature considering the distributional effects of monetary policy. Prominent examples include Doepke and Schneider (2006); Coibion et al. (2017); Kaplan, Moll, and Violante (2018); Cravino, Lan, and Levchenko (2018); and Auclert (2019). Most of this literature focuses on how monetary policy affects income inequality.

responses. Using aggregated data, Cloyne, Ferreira, and Surico (2020) argue that consumption of households with mortgages and those who are renters rises after expansionary shocks, while outright owners do not respond. They argue that the first group (renters and households with mortgages) are typically younger, while the second group (outright homeowners) are typically older. One study that finds contrasting results is by Berg et al. (2019), who argue instead that older households as a group have a higher consumption response than younger households.

We find something very different—namely, that monetary policy is more effective the greater the number of the middle-aged and less effective the greater the number of the young. One potential reason for this difference is that we look at a different outcome variable. Sterk and Tenreyro and Wong look at the consumption of households and Cloyne, Ferreira, and Surico look at the consumption of households aggregated by housing tenure. In our state-level analysis, we look at the effects on income and employment, since quarterly US state consumption data are only available after 2005. It is very likely that consumption and income respond in different ways to monetary policy. The relationship between an individual's consumption and income depends to a large extent on the individual's position in the life cycle. Liquidity constraints are more likely to affect the young, and the middle-aged are generally in a better position to smooth consumption (Gourinchas and Parker 2002). It is therefore likely that the consumption of the young responds more to monetary policy shocks, while income and employment depend on the proportion of the middle-aged.

Another potential reason that our results differ from the previous literature is that our question is somewhat different. Rather than ask how monetary policy affects different age groups, we ask how changes in the age structure of the economy alter the impact of monetary policy shocks on aggregate economic performance. If aggregate responses are merely the sum of individual responses, then the answers to these two questions would be the same. One of our principle findings, however, is that the effects of age composition appear to work through general equilibrium effects, and that these general equilibrium affects hit all age groups in a similar manner. What matters for the response of the income of the young in Idaho is the overall age distribution in Idaho and not the fact that the group is young. For both of these reasons, we see our findings as consistent with and complementary to the existing literature.

One paper that considers how the age structure of the economy alters the impact of monetary policy is Imam (2015). Imam assesses the impact of monetary policy on unemployment and inflation as economies age. He first estimates a time-varying coefficient Bayesian VAR with stochastic volatility for a panel of the large OECD economies: Canadian, German, Japanese, UK, and US monetary policy shocks are identified by a Cholesky decomposition with money ordered last. Using the estimated impact of monetary policy on unemployment and inflation, he then proceeds to estimate the impact of aging on the variables of interest. He uses the old-age dependency ratio, the ratio of the number of people older than 65 years to those 15–64 years old, as his demographic variable. He therefore has only two age groups. He finds that policy has become less potent as the fraction older than 65 has risen. He does not consider how the middle-aged might differ from the young.

Several papers, including Fujiwara and Teranishi (2006); Kantur (2013); Carvahlo, Fererro, and Nechio (2016); Kara and von Thadden (2016); and Sterk and Tenreyro (2018) calibrate dynamic general equilibrium models to study the effects of demographic change on monetary policy. Most of these papers build on the model of Gertler (1999) in which individuals progress through two life stages: the young work and the old are retired. Our results suggest the need for a richer demographic structure. At a minimum one needs to distinguish between young and middle-aged workers, in addition to the retired.[4]

Our paper contributes to a literature that uses a panel of US states to answer aggregate questions. Most of these papers study fiscal policy. Shoag (2010), Nakamura and Steinsson (2014), and Chodorow-Reich (2019) analyze fiscal policy multipliers at the state level. Recently, Basso and Rachedi (2018) have studied the effects of fiscal shocks interacted with the proportion of the state population in three age bins and find that the fiscal multiplier is larger the greater the number of young people in the state. As concerns monetary policy, there are a number of recent papers that find regional disparities in the effects of interest rate reductions during the Great Recession. Di Maggio et al. (2017) find a greater pass-through from interest rates to consumption in zip codes with a greater fraction of adjustable-rate mortgages. Beraja et al. (2019) find that homeowners in MSAs that had experienced larger price declines were less able to refinance and take advantage of low interest rates.

We are not the first to argue that the share of the middle-aged might have important consequences for economic activity. Feyrer (2007) finds that an increase in the proportion of workers between 40 and 49 increases productivity growth in a cross-section of OECD countries. Azoulay et al. (2019) show that founders of successful businesses tend to be middle-aged. Selgado (2018) argues that entrepreneurs tend to be middle-aged. We argue below that the reaction of middle-aged entrepreneurs to monetary policy shocks is a promising explanation for our results.

II. Empirical Implementation

We treat the US states as a panel of economies. There are several advantages to using the states as a laboratory. The primary advantage is that the United States is a monetary union, so all states are subject to the same monetary policy. This means that we can remove the first-order effect of monetary policy using time fixed effects and focus on the interaction of monetary policy shocks with local demography. A common monetary policy also removes many sources of heterogeneity. We do not have to worry about heterogeneity in monetary policy rules across states. We do not have to worry about monetary policy responding to the demographic makeup of each state.[5] We do not have to worry about differences that might arise from differences in the identification of monetary policy shocks across different nations. The states also have

[4] A couple of recent papers—Gagnon, Johanssen, and Lopez-Salido (2016) and Eggertsson, Mehrotra, and Robbins (2017)—use richer demographic structures to understand the decline in output growth and the real interest rate. These papers do not study the economy's response to monetary policy.

[5] Aksoy et al. (2019) argue that monetary policy responds to the age structure of the economy. In a cross-section of countries, they find that a larger working-age population tends to cause greater inflation, whereas a larger population of young dependents tends to reduce inflation.

a homogeneous legal and institutional structure. The problems that remain—such as immigration, trade, and heterogeneity of industrial structure—are present, more or less, in nations as well, and we will attempt to control for some of these influences in what follows.

In the remainder of this section we discuss the empirical implementation of our approach. We begin by describing the empirical specification, followed by a discussion of how we identify monetary policy shocks and a description of the data used in this paper.

A. *Specification*

We use variation across US states and across time in the share of population in different age groups to estimate how the population structure affects the response of income and employment to monetary policy shocks. We consider the following specification, which is a panel version of a Jordà projection (Jordà 2005):

$$\Delta \log\left[X_{s,t+i}\right] = \alpha^{i,a}\epsilon_t^m p_{s,t}^a + \phi^a p_{s,t}^a + \sum_{j=0}^{J}\beta^j Z_{s,t-j} + \gamma_s + \delta_t + u_{s,t}. \tag{1}$$

Here s indexes the US states, a indexes the age bins, and t indexes time. In our baseline specification we consider all 50 US states plus the District of Columbia. The age bins run every five years from age 20 to age 69, and then we group the population over 70 into an additional bin. By considering five-year age bins, we do not place any a priori restrictions on which age groups matter for monetary policy and how these age groups matter. The unit of time is a quarter.

The dependent variable is the log change in X_s between period $t-1$ and period $t+i$. Our main results take X to be either nominal personal income or private employment. The term ϵ_t^m is an exogenous shock to US monetary policy, $p_{s,t}^a$ is the percentage of the population of state s that belongs to age group a at time t, and $Z_{s,t-j}$ is a vector of additional control variables. We include the four-quarter growth rate in the population in the state in most specifications to control for trends in population growth. The state effects γ_s capture persistent differences in ages or industrial structures across states, and time fixed effects δ_t account for factors that are common to all states such as interest rates, the exchange rate, and aggregate inflation. The variable $u_{s,t}$ is the error term. Note that since we include time effects, we do not include the money shock ϵ_t^m separately in the regression.

All percentages, such as income growth, the monetary shock, and the population shares, are expressed in basis points.

We run this panel regression separately for each horizon i and each age bin a. The coefficients of interest are $\alpha^{i,a}$, which capture how the effect of a shock to monetary policy in period t on X_{t+i} varies with the proportion of the population in age group a. Since our monetary policy shocks represent shocks to interest rates, positive shocks reflect contractionary monetary policy. Positive values of $\alpha^{i,a}$ therefore imply that an increase in the proportion of the population in age group a reduces the effect of a shock to monetary policy on X. Negative values of $\alpha^{i,a}$, on the other hand, imply that as the population in age group a increases, the contractionary effects of an increase in interest rates become stronger.

For any given age group $\bar{a}$, estimating equation (1) for each $i = 1, 2, \ldots, H$ indicates how an increase in the proportion of the population in group $\bar{a}$ affects the response of X to a monetary policy shock at different horizons i. Alternatively, for any given horizon $\bar{\imath}$, estimating equation (1) indicates how an increase in the proportion of the population in different age groups a affects the response of monetary policy on X at the given horizon $\bar{\imath}$.

B. *Monetary Policy Shocks*

Why focus on monetary policy shocks? One might argue that interest rates are set at the national level and hence exogenous to state income and employment, so one can simply use changes in the interest rate as the measure of monetary policy. This argument, however, ignores the fact that monetary policy responds to other shocks—such as oil shocks, government spending shocks, or productivity shocks—which themselves may have differential effects at the state level that are correlated with the age composition of the state. We focus on monetary policy shocks to isolate the impact of monetary policy.

In our baseline specification, we follow the high-frequency approach to the identification of monetary policy shocks employed by Gürkaynak et al. (2005), Gertler and Karadi (2015), and Nakamura and Steinsson (2018).[6] As highlighted by Nakamura and Steinsson (2018), there are some concerns regarding the interpretation of these high-frequency shocks. We therefore consider the robustness of our results to other shocks in Section IIIB. We provide a detailed explanation of how we construct our shocks in Section 1 of the online Appendix. Here we provide a brief overview.

There are two steps in the construction. First, we obtain raw shocks from the movement in the federal funds rate implied by the current-month federal funds futures contract in a relatively short window of time around the FOMC announcements. We consider the change in the futures rate from 15 minutes before the announcement to 45 minutes after the announcement. By considering a narrow window around the FOMC announcement, one can be reasonably certain that no other news caused the change in the futures rates.

Second, following Ottonello and Winberry (2020), we aggregate the high-frequency shocks to a quarterly frequency using a weighted moving average of the shocks in the current quarter and the previous quarter. These weights take into consideration the fact that employment and personal income are the result of outcomes and decisions that accumulate over the entire quarter and that the monetary policy shock can only affect the subset of these decisions that are made subsequent to the shock. The weights are chosen so that an announcement at the end of quarter $t-1$ is essentially the same as an announcement at the beginning of quarter t.[7]

[6] Interest rate futures have been used to identify monetary policy shocks by Rudebusch (1998); Kuttner (2001); Soderstrom (2001); and Thapar (2008).

[7] Note that this second step introduces some serial correlation in the shocks, which we will have to control for in the estimation.

The shocks are illustrated in Figure A1 of the online Appendix. The identified shocks are small but not insignificant. The average high-frequency shock over the 1990–2008 sample is close to zero (-5 basis points), with a standard deviation of 11 basis points. Approximately 40 percent of shocks were contractionary while 60 percent are expansionary, and there are 70 percent more shocks than actual changes in the federal funds rate. Very often the absence of a change in interest rates is itself a shock.

Since our shocks only represent the unanticipated component of monetary policy at the time of the meetings, they may reflect only a fraction of exogenous monetary policy changes. This makes it difficult to judge the quantitative significance of the coefficients $\alpha^{i,a}$. To establish a benchmark we look at the effects of monetary shocks on the aggregate economy. We follow Ramey (2016) and estimate

$$\log[X_{t+i}] = \alpha^i \epsilon_t^m + \sum_{j=1}^{J} \beta^j \epsilon_{t-j}^m + \sum_{j=1}^{J} \gamma^j ffr_{t-j} + \gamma_s + \sum_{j=1}^{J} \delta^j Z_{t-j} + u_t, \tag{2}$$

where $i = 1, 2, \ldots, H$ represents the horizon, X is aggregate US real GDP, ϵ_t^m is our monetary policy shock, *ffr* represents the level of the federal funds rate, and Z (following Ramey) includes lags of X, current and lagged CPI inflation, the Commodity Research Bureau's spot commodity price index, and the unemployment rate. We also include current and lagged real GDP to calculate the federal funds rate impulse response function.[8] To account for serial correlation in the residuals, we estimate Newey-West standard errors. Note that all of these controls would be subsumed into the time effects in our baseline specification (1).

Figure 1 depicts the response of US real GDP, real personal income, private employment, and the federal funds rate to a 1 percentage point shock to monetary policy over a horizon of four years. We normalized the shock to 1 percentage point because the units are easy to interpret and relate to the discussion that follows. It should be noted that this scaling greatly magnifies the effects of monetary policy shocks relative to the data. The standard deviation of shocks in our sample is an order of magnitude smaller at 0.11.

The impulse responses in Figure 1 look like the effects of a monetary policy shock. As anticipated, this shock leads all three real variables to decline while the federal funds rate increases for four quarters and then declines. The decline in real GDP and real personal income become statistically significant five quarters after the shocks. The reduction in GDP reaches a maximum two to three years after the shock. We will compare the $\alpha^{i,a}$ in (1) to the β^j in (2) as a means of judging the quantitative significance of the interaction between demographics and monetary policy.

[8] We obtain the aggregate US data from the FRED data portal at the Federal Reserve Bank of St. Louis. We accessed the CRB index from Haver Analytics. Ottonello and Winberry kindly shared with us their high-frequency monetary policy shocks.

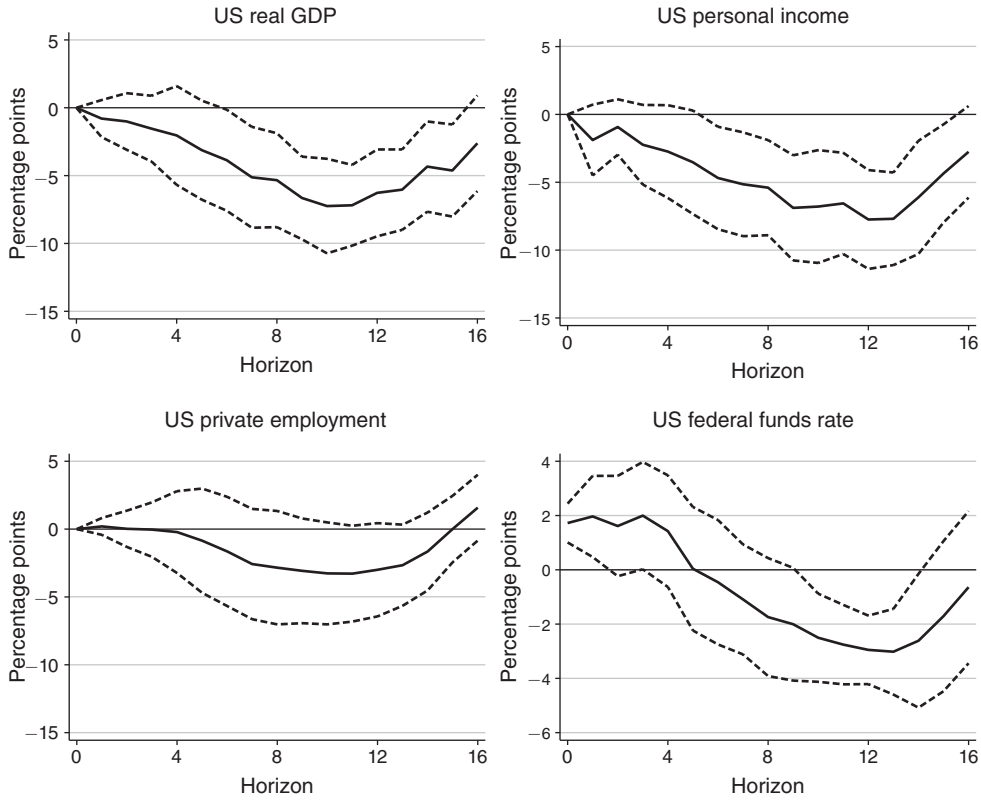

FIGURE 1. RESPONSE TO A MONETARY POLICY SHOCK

Notes: The solid line represents the response to a positive 1 percentage point increase in the high-frequency federal funds rate shock. The dashed line represents the 95 percent confidence intervals, given Newey-West standard errors.

C. *Data*

A detailed description of the data is in Section 5 of the online Appendix. Here we provide a brief overview. The federal funds futures market began in late 1988. Given that there was likely a learning period, we begin our sample in 1990:II. To avoid the zero lower bound on interest rates, we end in 2008:IV. We consider alternative start and end dates in the robustness section. The high-frequency monetary policy shocks were provided to us by Ottonello and Winberry.

We require state-level data on various macroeconomic variables and population estimates by state. We describe each of these in turn below. Further details are available in Section 5 of the online Appendix. In an ideal world, we would have quarterly state-level GDP and consumption data. Unfortunately, these data are available only from 2005 on. We use the next best alternative, namely, state personal income. Total state personal income, wages and salaries, dividends interest, and rental income and proprietors' income are obtained from the Regional Economic Accounts program at the Bureau of Economic Analysis. We use nominal personal income, since there are

no state-level price indices available for our sample period.[9] Deflating the nominal personal income of each state by a national price index would not alter our results, as the aggregate price index would simply be soaked up by the time fixed effects in the estimation.

Our primary source for state-level employment data is the Current Employment Statistics program conducted by the Bureau of Labor Statistics (BLS). State-level employment data are available beginning in 1990. The data are published at a monthly frequency, which we average to a quarterly frequency. Annually, the monthly estimates are revised based on a complete count of jobs derived from tax filings of employers. We use data on private employment as well as employment in three main industries: manufacturing, construction, and financial activities. For employment in the construction industry, data are available for only 44 states.[10] The baseline results for state personal income and employment have 3,825 observations: 18.75 years, 4 quarters, and 50 states and the District of Columbia.

When we need payroll or employment broken down by age group, by demographic characteristics, or by industry, we use data from the Census Bureau's Longitudinal Employer-Household Dynamics (LEHD) Program. The LEHD provides data on the number of jobs and payroll. We construct earnings per worker as the ratio of payroll to jobs. The LEHD provides data for employees broken down by the following predefined age bins: 14–18 years, 19–21 years, 22–24 years, and then ten-year bins from 25–34 years to 55–64 years, followed by one age bin for 65–99 years. The LEHD also provides data broken down by sex, educational attainment, race, and industry at the four-digit NAICS level. We use the latter to construct payroll for tradable and nontradable industries following Mian and Sufi (2014). It should be noted that LEHD data are not available for all states at all dates. It is an unbalanced panel, with states entering the panel at different times. Massachusetts, for example, enters in 2010 and is therefore not included in our sample. The LEHD samples have approximately 2,300 observations, with different variables having slightly more or slightly less.

State population data are obtained from the Census Bureau. Over our sample period, the data are available on an annual frequency.[11] We do not interpolate the data and assume that population remains constant within the calendar year. Note, $p^a_{s,t}$ is defined as the number of individuals in the age bin a in state s at date t, divided by the total population in the state at that time, multiplied by 100.

Establishment births and deaths data are published by the BLS's Business Employment Dynamics program and are based on microdata collected from the Quarterly Census of Employment and Wages (QCEW). The data include all establishments subject to state unemployment insurance laws. The BLS estimates that approximately 98 percent of all employment is covered in the QCEW.

For housing prices, we use the Federal Housing Finance Agency's seasonally adjusted purchase-only house price index. The index is estimated using sales price

[9] Since 2008 the BEA has begun publishing annual regional price parities, which compare prices in states relative to overall prices for the United States.

[10] Construction industry employment is not included for Delaware, the District of Columbia, Hawaii, Maryland, Nebraska, South Dakota, or Tennessee.

[11] In 2015 the BEA began publishing quarterly state population estimates, which are based on unpublished beginning-of-the-month state population estimates from the Census Bureau.

data for single-family houses whose mortgages have been purchased or securitized by Fannie Mae or Freddie Mac. The state-level home ownership rate is published by the Census Bureau. We use data on the fraction of mortgages that are adjustable-rate mortgages from the Federal Housing Finance Agency's Monthly Interest Rate Survey. The state-level home ownership rate and the fraction of ARM's are available at an annual frequency.[12] We do not interpolate the data and assume that each remains constant within the calendar year.

III. Results

A. *Baseline Results*

Our baseline specification takes X to be either state personal income or private employment and Z to be the change in state population between t and $t-4$. The monetary policy shocks are the high-frequency shocks discussed in Section IIB.

Figure 2 presents estimates of $\alpha^{i,a}$ for $i = 6$ in equation (1) when X is state personal income and state private employment, respectively. Coefficients and p-values may be found in Appendix Table 2. The various five-year bins for a are on the horizontal axis. The vertical axis represents $\alpha^{i,6}$ in percentage points. The solid line represents the (smoothed) point estimates. For a 1 percentage point unexpected increase in the federal funds rate, the point estimate $(\alpha^{6,a})$ represents the effect of an increase in the population share of group a by 1 percentage point on the growth rate of the variable of interest. The vertical bars depict the 95 percent confidence intervals for the estimated coefficients. These confidence intervals are estimated using the approach in Driscoll and Kraay (1998), which is robust to general forms of spatial and temporal correlation in the error terms.

The left panel presents results for personal income. The coefficients are positive for the young, aged 20–39, and statistically significant at the 1 percent level for the 25–29 and 30–34 age groups. The coefficients are negative for the middle-aged, 40–64-year-olds, and statistically significant at the 1 percent level for all but the 40–44 and 60–64 age groups. The coefficients are small and insignificant for the old, aged 65 and above. The positive coefficients for the young imply that if the fraction of the population in this age group increases, then an unanticipated increase in interest rates raises personal income after six quarters relative to the effect that monetary policy has on the US economy as a whole. Note that this does not mean that personal income rises, as the total effect of the money shock on state income is a combination of the first-order effect subsumed into the time effect and the interaction effect considered here. Since we typically expect an increase in interest rates to reduce personal income, a positive coefficient implies a weaker negative effect of monetary tightening. Similarly, the negative coefficients for the middle-aged imply a stronger impact of monetary policy on personal income the larger this age group is.

The absolute size of the $\alpha^{6,a}$ have very little meaning. The magnitude of the coefficients depends on the size of our estimated monetary policy shocks as well as the

[12] Since 2005 quarterly home ownership data for states have been published by the Census Bureau.

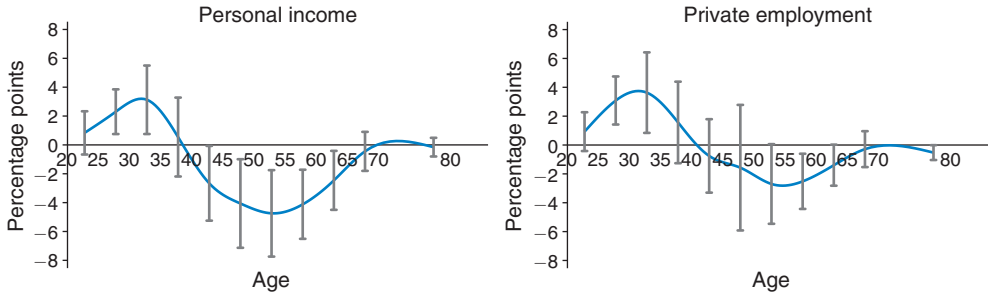

FIGURE 2. EFFECTS BY AGE BIN (SIX QUARTERS AFTER MONETARY SHOCK)

Notes: The dependent variable is indicated by the title of each graph. The horizontal axes represent the age groups a. The vertical axes represent the effect on the dependent variable of a positive 1 percentage point shock to the federal funds rate interacted with a 1 percentage point increase in the population share of group a. The solid line represents the point estimates. The vertical bars represent the 95 percent confidence intervals, given Driscoll-Kraay standard errors.

relationship between these shocks and the overall path for interest rates. We have normalized the size of the monetary policy shock to 1 percent for convenience. A 1 percentage point shock is not only very large relative to the shocks that we observe in our sample, but could be associated with a very different change in interest rates, especially after six quarters. While we cannot interpret the magnitude of the coefficients, we can get some sense of how large the effects are by comparing the $\alpha^{6,a}$ to the effects of monetary policy shocks on the US economy that we estimated in Section IIB. This comparison keeps the magnitude of the monetary shocks fixed. In Figure 1, we see that a 1 percent shock to the federal funds rate reduces personal income by almost 5 percentage points after six quarters. Now consider a one-standard-deviation change in the percentage of the population that is 30–34. This is an increase in p^{30-34} of 0.64 percentage points.[13] Given $\alpha^{6,30-34} = 3.13$, a one-standard-deviation increase in the share of 30–34-year-olds reduces the impact of a 1 percentage point increase in the federal funds rate by approximately 2 percentage points or 40 percent of the aggregate effect. Similarly, a one-standard-deviation increase in the share of 50–54-year-olds equal to 0.39 percentage points together with $\alpha^{6,50-54} = -4.74$ implies that a one-standard-deviation increase in the share of 50–54-year-olds increases the negative effect of a 1 percentage point increase in the federal funds rate by 1.85 percentage points. Again, this is approximately 40 percent of the aggregate effect. Hence, the variation in the effects of a monetary policy shock induced by demography appear to be of a magnitude similar to the direct effects of monetary policy on the economy as a whole.

It appears that the sum of the coefficients in Figure 2 is approximately zero. Given that an increase in one age group must show up as a decrease in another age group, it is tempting to argue that the coefficients must sum to zero. This is not necessary, however, for three reasons. First, we have omitted the 0–19 year old age group. Second, theoretically, it is the variance-weighted sum of coefficients

[13] This is the standard deviation of p^{30-34} after removing time effects. This is the relevant variation, as discussed in Section IV.

that sums to zero.[14] Finally, since we run the regressions separately for each age group, variation in the coefficients on the other variables may affect this homogeneity property.

The right panel of Figure 2 illustrates the effects on private employment. The pattern closely mimics the results for personal income but with some small differences in the turning points and a more moderate response for the middle-aged groups. The effects of monetary policy on private employment are dampened if there are more 20–34-year-olds in the economy, the effects of monetary policy are strengthened if there are more 40–59-year-olds in the economy, and they are largely neutral if there is a change in the population over 60 years of age.

The pattern that we see in Figure 2—with positive coefficients for the young, negative coefficients for the middle-aged, and small and statistically insignificant coefficients for the old—is a pattern that we will see often in what follows. Generally, estimates will be stronger and statistically more significant for income than employment, reflecting the fact that both wages and employment respond in similar ways to the shock, so that the response of the product of the two components is stronger than the response of each component separately. Results for the young and the middle-aged are remarkably stable across alternative specifications. Results for the old are generally statistically insignificant but a bit less stable. There are a few specifications in which a greater share of the old reduces the impact of monetary policy and a few specifications in which a greater share of the old increases the impact of monetary policy.

We presented the results for $i = 6$ because this is the horizon for which the effects of age appear strongest. We get qualitatively similar effects at other horizons over the first three to four years following a monetary shock. Figure 3 presents the $\alpha^{i,a}$ for personal income as functions of the horizon i for each age bin a. The effects of age tend to grow over the first six quarters: positive for the young and negative for the middle-aged. The maximal impact is at about six quarters. After six quarters the coefficients remain elevated for some age groups, especially the young, and die out for others, primarily those between 40 and 54. The effects are always small and statistically insignificant for the old.

B. *Robustness*

Monetary Shocks.—One obvious question is whether our results are driven by our choice of monetary policy shocks. Nakamura and Steinsson (2018) have raised questions regarding the interpretation of shocks identified though the high-frequency movements of asset markets around policy announcements. These shocks may reflect unanticipated moments in policy that are orthogonal to current economic conditions, or they may reflect the communication of the Federal Reserve's information regarding current economic conditions.

We consider two alternative identification schemes. First, we consider the shocks that Romer and Romer (2004) construct from the minutes of Federal Open Market

[14] If all of the age groups are included in a single regression, then the OLS coefficient is $(X'X)^{-1}(X'Y)$. It is $X'Y$ that must sum to zero.

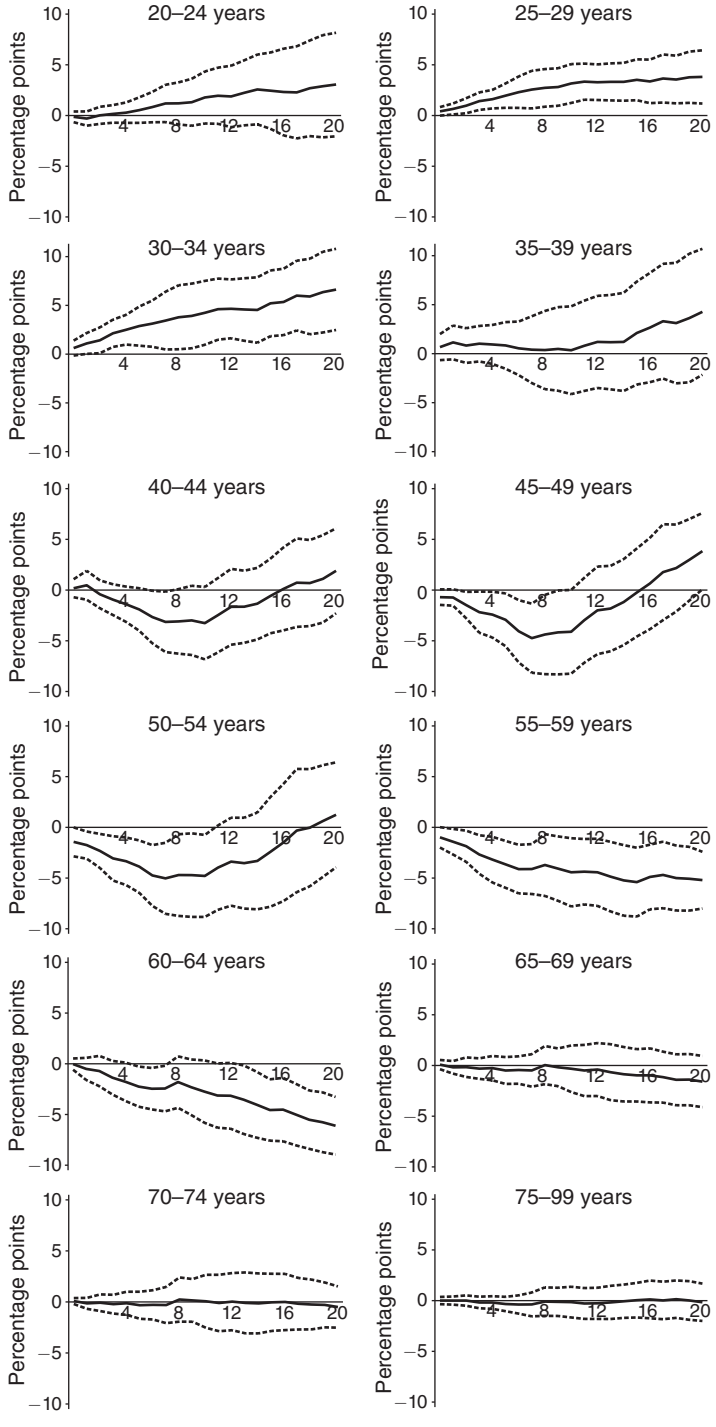

Figure 3. Impulse Response Functions of Personal Income

Notes: Each figure depicts the effects over time of a 1 percentage point increase in the share of the population in a specific age group on the responsiveness of the growth rate of state personal income to a positive 1 percentage point shock to the federal funds rate. The age group is indicated by the title of the graph. The *x*-axis represents quarters after the monetary policy shock. Solid lines represent point estimates. Dashed lines represent 95 percent confidence intervals based on Driscoll-Kraay standard errors.

Committee's meetings and the Federal Reserve Board's Greenbook forecasts. These are extended through the end of 2007 by Wieland and Yang (2019). Second, we consider shocks estimated from a structural VAR that includes US real GDP, the GDP deflator, the CRB spot commodity price index, and the quarterly average federal funds rate. The VAR is estimated over the 1960–2008 sample period and monetary policy shocks are identified via a Cholesky decomposition with the interest rate ordered last.

Figure 4 compares the effects on personal income and employment for the various shock specifications. In the top three rows, the sample is the same as in our baseline specification (1990–2008). Given that different shocks have different scales, we normalize by the standard deviation of the shocks. The vertical axis is $\alpha^{i,a} \times sd(\epsilon_t^m)$. For ease of comparison, all the figures are on the same scale. The first row reproduces the results from our baseline specification. Regardless of shocks, the effectiveness of monetary policy is dampened for the younger age groups, is strengthened for the middle-aged groups, and is about the same as the aggregate effect for the older age groups. While the basic shape and magnitude of the response is similar across the shock series, there are differences. The most responsive groups tend to be a bit younger for the Romer and Romer shocks and a bit older for the structural VAR shocks. Similarly, the transition from positive to negative coefficients also occurs at an earlier age for the Romer and Romer shocks and later for the structural VAR shocks.

With these alternative shock series, we can expand the sample period for state personal income.[15] The last two rows in Figure 4 show the results for personal income for the 1980–2008 sample. The results for the VAR shocks become more significant and look much more like the baseline results. The results for the Romer and Romer shocks are largely unaffected by the sample period.

Other Robustness Checks.—We tried a number of alternative specifications. These results are discussed in detail in Section 2 of the online Appendix. Here we provide a brief list of the robustness checks that we ran. The bottom line is that the basic pattern seen in Figure 2, in which the young dampen the effects of monetary policy and the middle-aged amplify the effects of monetary policy, survives all of these modifications.

- In the baseline specification we run separate regressions for each age group. We tried running a single regression with all of the age groups. To avoid a proliferation of regressors, we aggregated the age bins into three age groups: 20–39, 40–59, and 60+, which broadly correspond to the three age groups identified in our baseline results.[16]
- We aggregated population and personal income by BEA regions.

[15] Since quarterly state-level private employment data begin in 1990, we do not extend the sample for this variable.

[16] We also ran a single regression for all of the age groups using ten-year bins. The estimate for 30–39 is positive and significant, and the estimate for 40–49 is negative and significant. The estimates for the other age groups are small and insignificant. These results appear in Section 2 of the online Appendix.

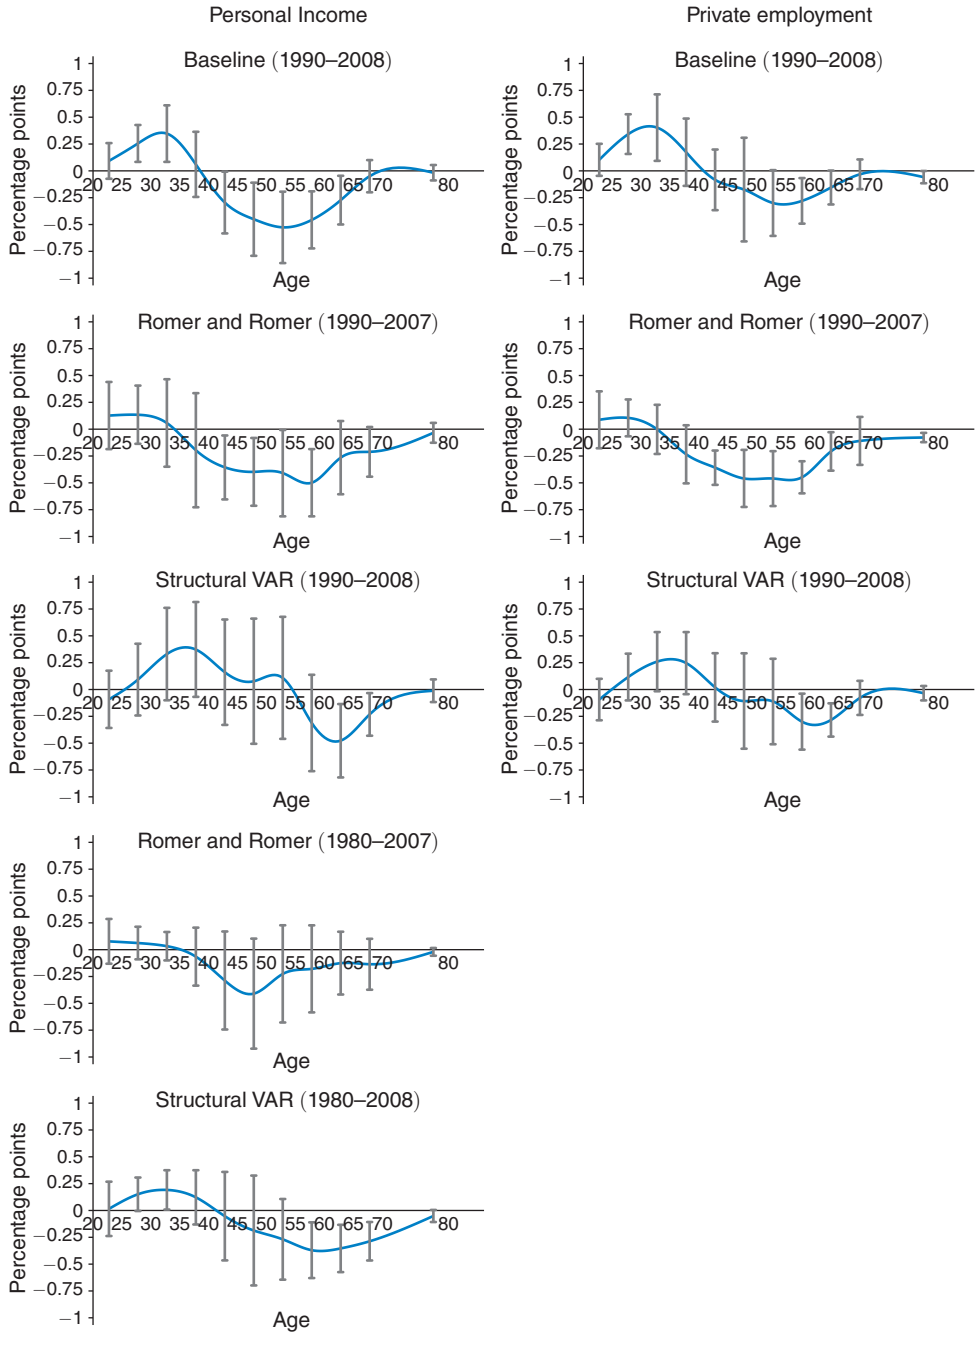

FIGURE 4. ALTERNATIVE MONETARY SHOCKS

Notes: The horizontal axes represent the age groups a. The vertical axes represent the effect on the dependent variable of a positive 1 percentage point shock to the federal funds rate interacted with a 1 percentage point increase in the population share of group a. The solid line represents the point estimates. The vertical bars represent the 95 percent confidence intervals, given Driscoll-Kraay standard errors. The baseline and VAR specifications are estimated through 2008; Romer and Romer data end in 2007.

- We dropped the five smallest states by population (Alaska, North Dakota, Vermont, Washington DC, and Wyoming).
- We do not have a long enough sample to properly analyze subsample stability. We do, however, investigate whether our results are sensitive to altering the beginning and the end of the sample. First, we truncated all data at the end of 2006 to avoid any contamination with the housing market crash and the Great Recession. Second, we began our sample in 1995, after the FOMC began issuing press releases following every FOMC meeting. If anything, beginning in 1995 strengthened the economic and statistical significance of the estimates.

IV. Are the Effects Causal?

While we identify plausibly exogenous shocks to monetary policy, variation in the demographic structure of states is potentially endogenous. Absent instruments for the young, middle-aged, and old, we do the next best thing and control for state characteristics that may be correlated with the state's demographic structure.

Note that the time and state effects in the baseline specification (1) absorb common trends and some state differences, but not all. Since the money shock is common across states, the time effects control for trends in population growth that are common across states. To see this, consider adding a constant c_t to $p^a_{s,t}$ for all s at a given date t. Then the interaction term becomes $\left(p^a_{s,t} + c_t\right)\varepsilon^m_t$, and $c_t\varepsilon^m_t$ is absorbed by the time effect. Given that the money shock varies over time, however, the state fixed effects do not absorb average population differences across states. To see this, consider adding a constant c_s to $p^a_{s,t}$ for a given state s for all t. The interaction term becomes $\left(p^a_{s,t} + c_s\right)\varepsilon^m_t$, and $c_s\varepsilon^m_t$ is not absorbed by either the time or state fixed effect. The relevant variation in (1) is therefore the population shares in a given age bin across states and time net of common time trends.

To illustrate this variation, Figure 5 represents the demographic composition of 50–54-year-olds across states in 1990 and 2000. To construct the figure we first removed time fixed effects from the population share of this age group. We then sorted the residuals into five quintiles. Lighter colors represent states with higher residuals. For example, the population share of 50–54-year-olds in Hawaii rose from 0.3 percentage points below average in 1990 to 0.3 percentage points above average in 2000. In Indiana, on the other hand, this population share fell from 0.2 percentage points above average in 1990 to 0.1 percentage points below average in 2000. Although there is substantial variation in the rankings across the two time periods, given that the state effects do not fully control for state characteristics, it is of interest to investigate whether our results are driven by some state characteristics that are correlated with the population shares.

Our first experiment is to include the interaction between the state fixed effect and the monetary policy shock as an additional regressor. This will soak up $c_s\varepsilon^m_t$ and control for level differences in population across states that might be correlated with fixed state characteristics. For example, southern states may have more old and young residents but less industry. The results are presented in Section 2 of the online Appendix, particularly Figure A5. The basic pattern is the same as in our baseline specification. There are a few notable differences. The response of personal

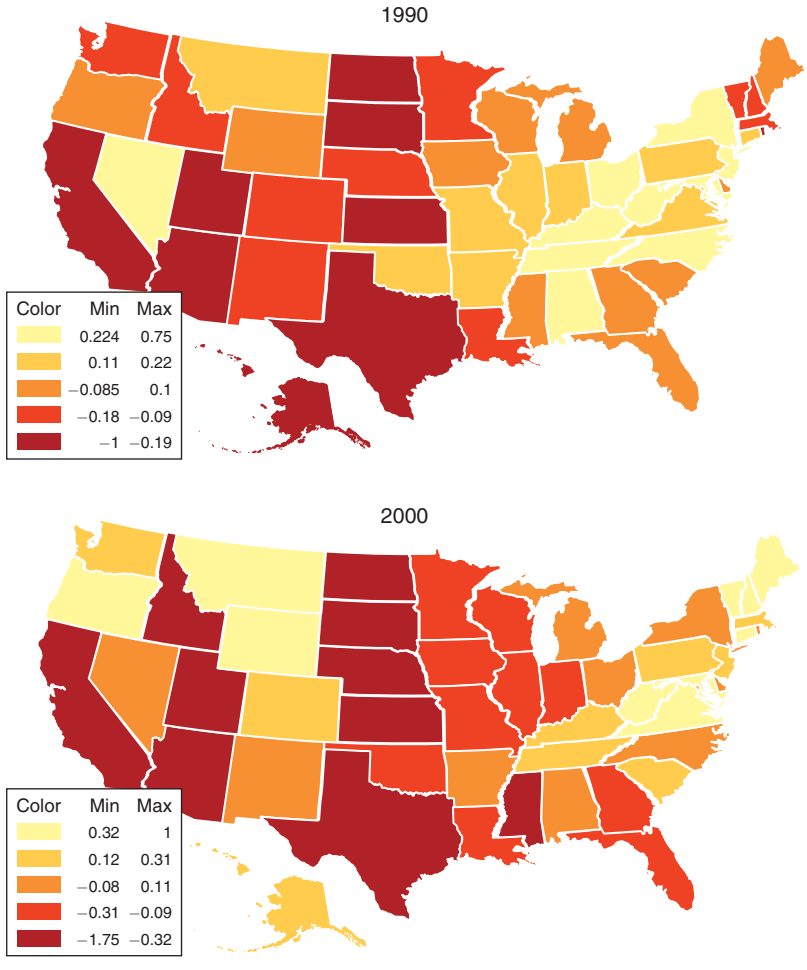

FIGURE 5. CHANGING DEMOGRAPHIC MAKEUP OF THE UNITED STATES: AGES 50–54

Note: The figures represent the residuals of the share of the population aged 50–54 years after removing time fixed effects.

income is amplified. For both personal income and employment, statistical significance improves for the 25–34 and 40–49 age groups while worsening for those over 50. There is also weak evidence that a greater fraction of older people may reduce the effects of monetary policy on employment. Given that including these interaction terms does not affect the general pattern of the baseline results, in what follows we will report the results for our baseline specification, a specification that does not include the shock interacted with the state fixed effects. We will comment whenever including these interaction terms qualitatively alters the results.

Interacting the shock with state fixed effects controls for a correlation between average population in the state and average characteristics of the state. The possibility remains that some characteristics may change over time in ways correlated with the change in demographic structure. To investigate this issue we consider whether the population structure of a state is correlated with certain characteristics. We begin

with the industrial composition of the state. We consider three industries that have been tied to monetary policy: manufacturing, financial services, and construction. We consider manufacturing and construction because these industries tend to be more cyclical than other industries. A state with a large manufacturing or construction share would be expected to respond more strongly to a monetary policy shock. We consider financial services because the effects of monetary policy, through banks' lending, could explain a state's greater response.

We include the fraction of employment in manufacturing, construction, and financial services both as additional regressors and interacted with the monetary policy shock. The addition of these additional regressors has almost no effect on the estimates of $\alpha^{6,a}$. We therefore present the results in Figure A6 of the online Appendix. To understand why there is no effect, we consider the correlation between the share of employment in these industries and the proportion of the population between 40 and 64 years (both net of time effects). The correlation is -0.04 for manufacturing, -0.08 for financial activities, and -0.14 for construction.[17] All three correlations are small and negative. None of these sectors vary positively with the share of the middle-aged, as would be expected if these sectors were driving the results.

Industrial structure is not the only state characteristic that may confound our results. Cravino, Lan, and Levchenko (2018) argue that households with higher incomes tend to purchase goods with stickier prices. Since middle-aged households tend to have high incomes, and since stickier prices tend to be correlated with the response to monetary policy shocks, our results could reflect this mechanism. We therefore add state income interacted with the monetary policy shock as an additional regressor. Regression results appear in Figure A7 in the online Appendix. Again, we find no evidence that this mechanism is driving the results. The correlation between state personal income and the proportion of the population between 40 and 64 years of age (both net of time effects) is -0.11. Again, the correlation is small and negative.

Finally, we consider whether the age structure of a state might be correlated with some characteristic of local housing markets. Wong (2019) finds that young homeowners are more likely to refinance when interest rates fall and that this contributes to their greater consumption response. Di Maggio et al. (2017) and Beraja et al. (2019) find regional disparities in the impact of monetary policy that correlate with characteristics of the housing market such as the fraction of adjustable-rate loans and the change in house prices. To address these concerns we include variables such as home ownership, house price changes, and the fraction of adjustable-rate loans. We included each variable by itself as well as interacted with the monetary policy shock. These results are presented in Figure A8 of the online Appendix. Again, we find that the basic pattern holds. There is no evidence that any of these variables are driving the results.[18]

For completeness we considered other variables related to the demographic structure of the economy, such as the percentage of employees who have college degrees, who are White, or who are male. These results are presented in Figure A9

[17] Using the level of industry employment instead of the fraction of industry employment to total private employment gives similar correlations.

[18] The correlation between the share of adjustable-rate mortgages and the proportion of the population between 40 and 64 years of age (both net of time effects) is -0.05.

in the online Appendix. We included each variable as well as the variable interacted with the monetary policy shock. None of these variables had a qualitative effect on the estimates of $\alpha^{6,a}$.

While we cannot conclude with certainty that the results of Section IIIA are not due to the correlation of population structure with some other state characteristic, we find no evidence that this is the case.

V. Inspecting the Mechanism

Given the absence of an obvious covariate that explains the basic pattern, we entertain the notion that these effects are causal and ask why shifts in the age distribution might affect monetary policy. We begin by painting a picture of how the demographic structure affects the response of other variables to economic shocks.

First, we investigate whether these effects reflect migration between states or the expansion of opportunities within a state. Migration varies with age and is likely to be an important margin of adjustment for US states, where the borders are porous.

Next, we ask which industries and what types of incomes are most sensitive to variations in the demographic structure. Both occupation and type of income are likely to be age dependent.

Finally, we ask whether the greater share of middle-aged residents amplifies the effects of monetary policy because the middle-aged themselves respond more to the shock and a greater number of middle-aged residents increases the response, or because all age groups respond more to monetary policy shocks when the proportion of the middle-aged is larger. The former points to something particular about the middle-aged, the latter to some general equilibrium mechanism by which the actions of middle-aged individuals propagate and affect the entire economy.

In the next three sections, we discuss possible economic mechanisms consistent with these patterns and the potential role of middle-aged entrepreneurs, in particular.

A. *Migration*

One of the margins by which states adjust to shocks is through migration (Blanchard and Katz 1992). This raises the question of whether the responses in Figure 2 result from hiring within the state or from migration between states. It also raises the question of whether our results for US states can be applied to nations, since migration between states is likely easier than migration between nations.

To investigate this question we consider the responses of population and unemployment to monetary policy shocks interacted with the age bins. The idea is that migration should be reflected in the response of population, whereas hiring from within the state should show up in unemployment.

In the left panel of Figure 6, we take X to be the log change in population over six quarters, $\ln P_{s,t+6} - \ln P_{s,t-1}$, where $P_{s,t}$ is the population of state s at date t.[19] The response of population looks very different than the response of employment in

[19] For population, we also lag population growth by one period, since the dependent variable is population growth since $t-1$.

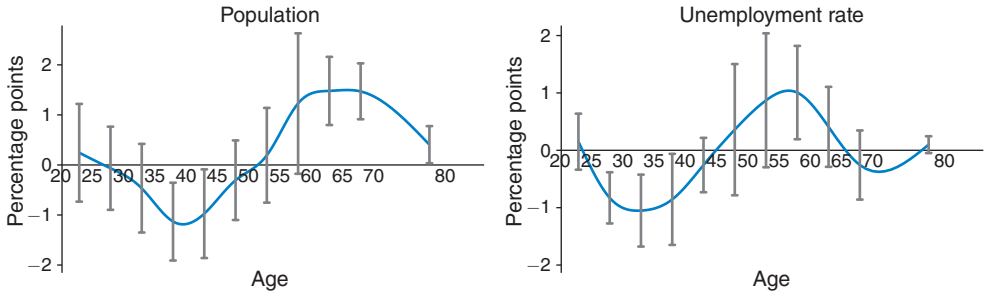

FIGURE 6. EFFECTS ON POPULATION AND UNEMPLOYMENT

Notes: The dependent variable is indicated by the title of each graph. The horizontal axes represent the age groups a. The vertical axes represent the effect on the dependent variable of a positive 1 percentage point shock to the federal funds rate interacted with a 1 percentage point increase in the population share of group a. The solid lines represent the point estimates. The vertical bars represent the 95 percent confidence intervals, given Driscoll-Kraay standard errors.

Figure 2. Whereas the response of employment and personal income partitioned the population into three three distinct age groups, the response of population partitions the population into only two. Population is more responsive in states with a greater share of the population younger than 50 and less responsive the greater the share over 50. This likely reflects the fact that the young are more mobile than the old.[20] The response of population is strongest for those 35–45 and those older than 60. The response of income and employment is generally the weakest for these age bins. Moreover, the response of population is less than half the response of employment. We therefore conclude that migration, while potentially important, is not the main factor behind our results.

In the right panel, we take the dependent variable to be the change in the unemployment rate, $U_{s,t+6} - U_{s,t-1}$. This figure looks very similar to Figure 2. The sign of the response is flipped, as an increase in unemployment tends to coincide with a decline in employment. Together these two figures point to the importance of internal adjustment mechanisms and increase our confidence that conclusions that we draw for the US states should apply to nation-states.

B. *Industries and Incomes*

Next we ask what types of jobs and incomes are affected by the demographic structure of the economy. Figure 7 presents $\alpha^{6,a}$ for various age bins a where we take X to be different types of employment. For ease of comparison, the axes are scaled to be the same across the figures. The response of manufacturing employment to a 1 percentage point increase in the federal funds rate looks very much like the response of total private employment in the baseline analysis. Employment in the financial activities sector, however, appears to be somewhat different from overall private employment. The groups that enhance monetary policy are slightly younger,

[20] As might be expected, the magnitude of the response is muted if we drop the five states with the smallest populations. The basic shape remains the same.

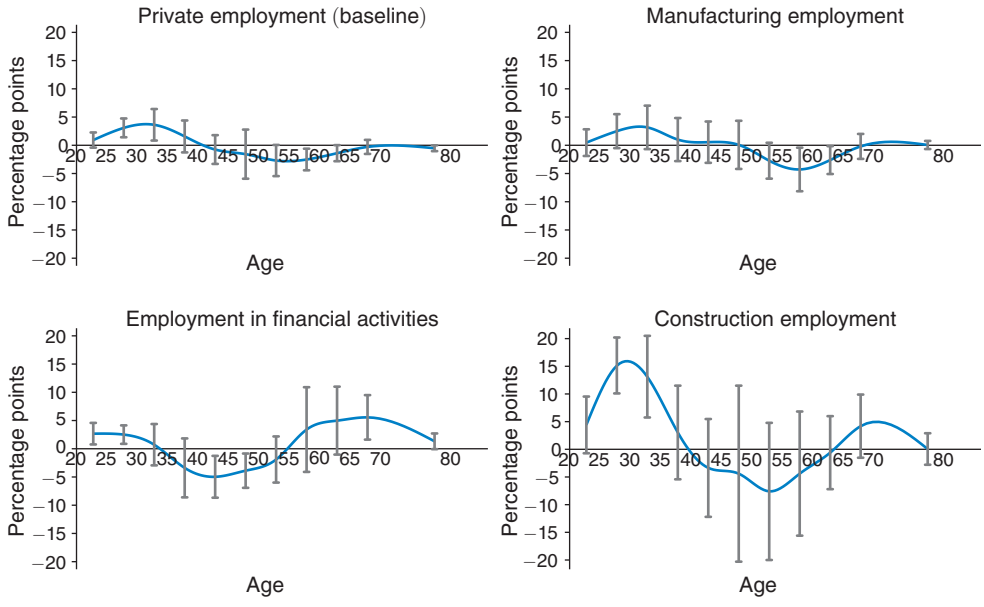

FIGURE 7. EFFECTS ON EMPLOYMENT BY SECTOR

Notes: The dependent variable is indicated by the title of each graph. The horizontal axes represent the age groups a. The vertical axes represent the effect on the dependent variable of a positive 1 percentage point shock to the federal funds rate interacted with a 1 percentage point increase in the population share of group a. The solid lines represent the point estimates. The vertical bars represent the 95 percent confidence intervals, given Driscoll-Kraay standard errors. Construction employment data are available in only 44 states.

35–50, and a greater share of the older age bins significantly reduces the effect of monetary policy shocks. This latter result is driven mostly by states like Alaska and Wyoming, which have a relatively small fraction of the population involved in financial sector activities. The response of construction employment is where most of the action is. It has a shape similar to our baseline analysis, but the magnitude of the response is much larger.

Figure 8 illustrates the responses of different types of income to our interacted shock. The figure presents $\alpha^{6,a}$ for various age bins a where we take X to be different types of income. The responses of wages and salaries and the responses of dividends, interest, and rent look very similar in shape and magnitude to our baseline specification. The main difference is in the response of proprietors' income. Here the negative response of the middle-aged is much more pronounced, although imprecisely estimated.[21]

[21] The response of proprietors' income in Figure 8 is large but imprecisely measured. If we include the monetary shock interacted with state fixed effects, the response of proprietors' income remains large and is statistically significant. Results available from the authors upon request.

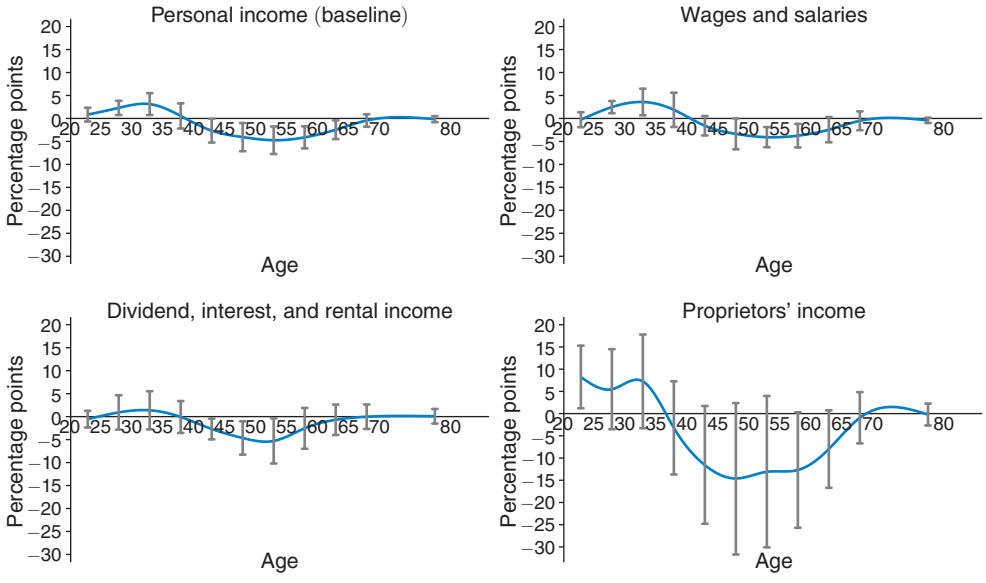

FIGURE 8. EFFECTS ON DIFFERENT TYPES OF INCOME

Notes: The dependent variable is indicated by the title of each graph. The horizontal axes represent the age groups a. The vertical axes represent the effect on the dependent variable of a positive 1 percentage point shock to the federal funds rate interacted with a 1 percentage point increase in the population share of group a. The solid lines represent the point estimates. The vertical bars represent the 95 percent confidence intervals, given Driscoll-Kraay standard errors.

C. *Aggregate or Individual Channels*

Our next experiment is to replace X with $X^{\bar{a}}$, the income of age group $\bar{a}$, to investigate whether the mechanism works at the level of the individual or in the aggregate. The LEHD provides the number of jobs and payroll for employees. We construct earnings per worker as the ratio of these two variables. The publicly available data are available for the following predefined age bins: 19–21 years, 22–24 years, ten-year bins from 25–34 years to 55–64 years, and one age bin for 65–99 years. We construct the income age groups $\bar{a}$ to match the LEHD age bins.

Let $\alpha^{i,a}(\bar{a})$ denote the response of $X^{\bar{a}}$ to $\varepsilon^m p^a$. If $\alpha^{i,a}(\bar{a})$ is the same for all a, then the response of each age group $\bar{a}$ is independent of the age distribution in the state. In this case, the mechanism likely operates at the individual level, and the aggregate response reflects the sum of individual responses. If, on the other hand, $\alpha^{i,a}(\bar{a})$ is the same for all $\bar{a}$, then all age groups $\bar{a}$ respond similarly to an increase in the proportion of the population in bin a. In this case the mechanism likely operates at the aggregate level.

We find strong evidence for aggregate mechanisms. Figure 9 depicts $\alpha^{6,a}(\bar{a})$ for different choices of $\bar{a}$ and for different choices of X. Note, $X^{\bar{a}}$ is either average earnings per worker, the number of jobs, or total payroll in bin $\bar{a}$. The overall picture

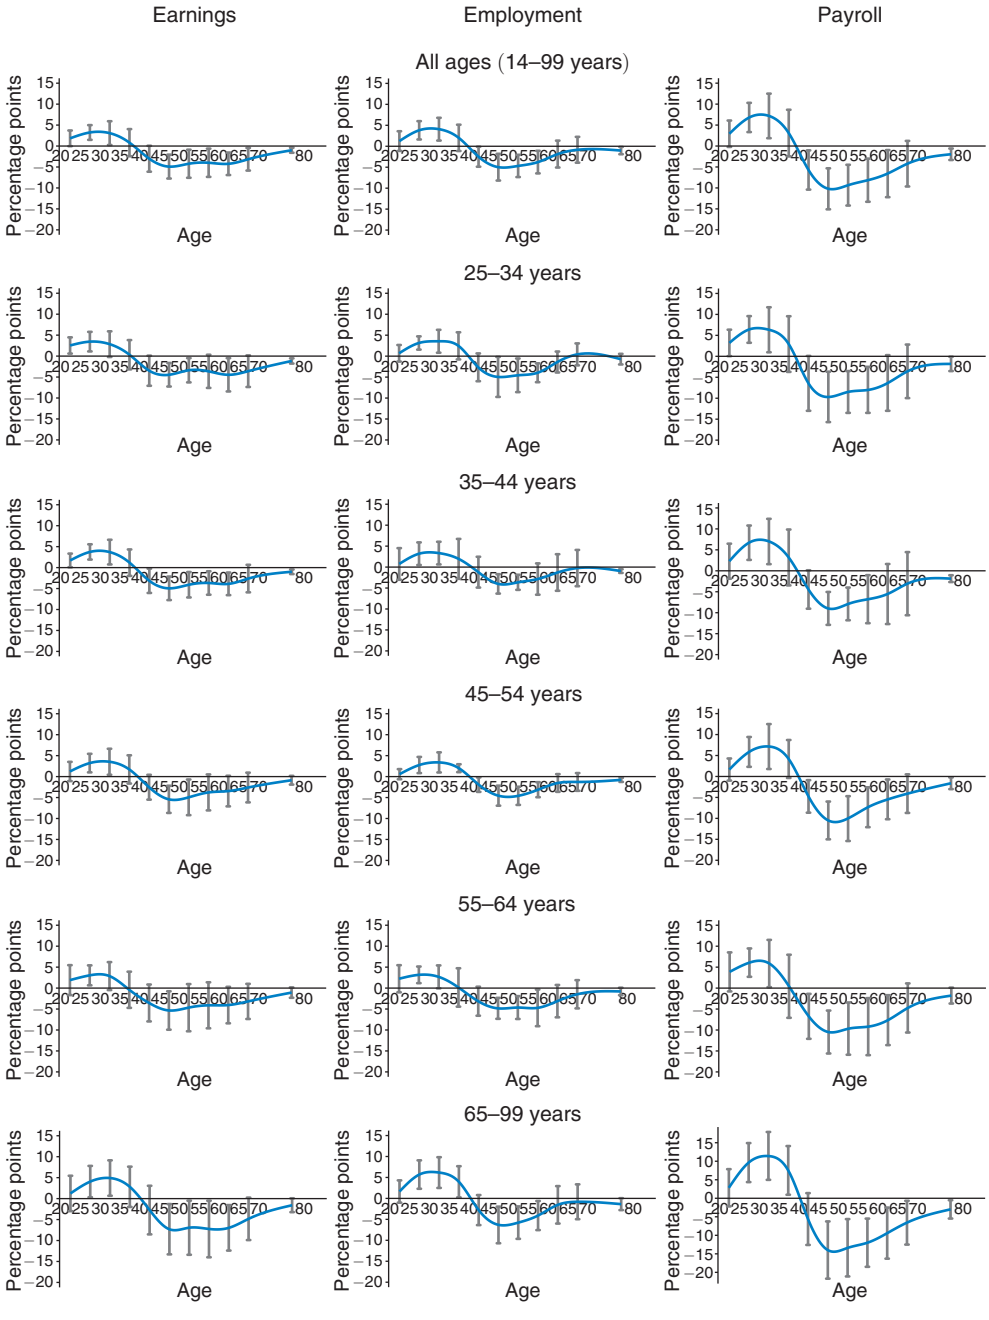

FIGURE 9. EFFECTS ON EARNINGS AND EMPLOYMENT

Notes: Each graph represents the response of earnings, employment, or payroll of all individuals in a specific age group (listed in the title of each column) to a positive 1 percentage point shock to the federal funds rate interacted with a 1 percentage point increase in the population share of group *a* (listed along the horizontal axis of each figure). The vertical bars represent the 95 percent confidence intervals, given Driscoll-Kraay standard errors.

remains similar to Section IIIA above.[22] The response of earnings per worker, jobs, and payroll is independent of one's own age group and depends primarily on the age structure of the state in which one lives. It is not so much being middle-aged as living in a state with lots of middle-aged residents that explains the response of income to a monetary policy shock.[23]

As further evidence that the mechanism operates through an aggregate channel, Figure 9 shows that the mechanism appears to work both through the extensive margin (the number of jobs) and the intensive margin (earnings per worker). The combined effect on total payroll is the product of the two. This combined effect is stronger and statistically more significant.

VI. Potential Mechanisms

In this section we consider potential mechanisms by which a greater share of the middle-aged might affect the transmission of monetary policy. It seems natural that any potential mechanism would have something to do with changes in economic behavior over the life cycle. We discuss three possibilities: the middle-aged tend to save, they tend to purchase homes, and they tend to start businesses.

A. *Life Cycle Saving*

Recent work on heterogeneous agent models has shed light on the importance of liquidity constraints and income effects in the transmission of monetary policy (Kaplin, Violante, and Moll 2018). The consumption of households with little in the way of liquid assets tends to track income. Loose monetary policy, which raises income, tends to increase consumption among these liquidity-constrained agents. At first glance, it would appear difficult to explain our results with this mechanism. After all, middle age is the time at which saving and asset accumulation tends to be highest (Gourinchas and Parker 2002). Even accounting for the accumulation of illiquid assets, Kaplin, Violante, and Weidner (2014) find that the fraction of hand-to-mouth households declines steadily over the life cycle from 50 percent at the age of 20 to 20 percent at the age of 80. One would therefore expect that a greater share of young households would enhance the effects of monetary policy. It is true that these authors also find that the wealthy hand-to-mouth tend to be middle-aged and that this might help explain the effect of middle-aged households on monetary policy. This mechanism, however, still has difficulty explaining why a greater fraction of the young should mitigate the effects of monetary policy, particularly relative to the old.

[22] Note that the response of payroll in the LEHD is slightly stronger and more statistically significant then that of personal income from the BEA. There are two reasons for this difference. First, the variables measure different things. Personal income includes not only wages and salaries but also Social Security and other government benefits, dividends and interest, and income from business ownership. Second, the LEHD sample is different, as states enter at different dates.

[23] When we include the interaction between the monetary shock and the state fixed effect as an additional control, Figure 9 is largely unchanged, except $\alpha^{i,a}$ is often positive and significant for the old (age 65 and over).

It is important to note that our results are not inconsistent with these models, nor are they inconsistent with the greater consumption response of young households to a monetary policy shock as found by Sterk and Tenreyro (2018), Wong (2019), and Cloyne et al. (2020). In fact, our results say nothing about the average consumption response of the young relative to the old. Lacking data on consumption at the state level prior to 2005, our focus is on income, not consumption. It is entirely possible, even likely, that the consumption of the young is more responsive to these changes in income than is consumption of the middle-aged or the old. More importantly, the average response of each group is netted out by the time fixed effects. Our results only say that an increase in the fraction of certain age groups has effects relative to this average response. Moreover, these effects appear to be common to all age groups. A greater share of the middle-aged increases the response of all groups' incomes to a monetary shock.

One possible story is that age structure affects the strength of the income effects that drive the consumption response in these heterogeneous agent models. A greater number of middle-aged residents makes these income effects stronger. A greater number of young dampens these income effects.

B. *Housing*

Home buyers tend to be middle-aged. According to the National Association of Realtors, the median age of first-time home buyers is in the low thirties and the median age of all home buyers is in the late forties.[24] Older individuals tend to purchase more expensive homes and are more likely to own multiple residences (NAR 2020). This suggests a channel by which loose monetary policy spurs the home purchases of the middle-aged. These purchases then feed construction, and the benefits spread to the economy as a whole. This channel is consistent with our earlier finding that the response of construction employment to monetary policy shocks is very sensitive to the age structure of the state.

The data that we have on home purchases by age are not kind to this home purchase channel. Such a channel would predict that increasing the fraction of 30–39-year-olds and 40–54-year-olds would have similar effects since these groups make up similar factions of home buyers and purchase homes of similar sizes and expense (NAR 2020). These groups, however, have opposing effects on the transmission of monetary policy. Of course, it is not the average behavior of these groups that we are interested in but their responses to interest rates. It is possible that, on average, the housing choices of the two groups are similar, but that each responds differently to monetary policy. The younger group, for example, may be comprised of more first-time buyers with limited resources and may react to a rise in housing prices that accompanies loose monetary policy by saving more to meet higher down payments.

It is also important to note that our sample period coincides with the rapid rise in home equity lines of credit (HELOCs) in the 1990s and the housing boom of the

[24] See https://www.nar.realtor/newsroom/real-estate-story-ideas/median-age-of-home-buyers-trending-upward.

2000s. The median age of a homeowner with a HELOC was 50 in 2000 (Cavanaugh 2007). It is possible that the homeowners in their forties and fifties accumulated more home equity during this period and therefore had greater financial capacity to respond to reductions in interest rates. These households may have responded to the reductions in interest rates by borrowing more against their homes in order to finance home improvements and consumption. These expenditures then spilled over to the economy as a whole, and, in particular, to the construction sector.

Each of these stories would appear to require older households to increase consumption and investment in response to lower interest rates; however, Wong (2019) and Cloyne et al. (2020) argue that younger households—in particular, younger households that adjust their mortgages—tend to increase their consumption in response to monetary policy shocks. The key to reconciling these conflicting views may lie in the details: exactly how households are categorized, exactly how each age group's consumption and investment respond, and exactly how these actions propagate through the economy. Properly testing either the housing demand or housing finance stories would require data on home purchases or debt by state and by age at a quarterly frequency over a sufficiently long sample period. Most publicly available datasets are lacking in one of these dimensions. We leave a definitive analysis for future research.

C. *Entrepreneurs*

Entrepreneurs, especially successful entrepreneurs, tend to be middle-aged (Selgado 2018, Azoulay et al. 2019). This suggests a channel by which a greater number of middle-aged residents implies a greater number of active and potential entrepreneurs. The story goes as follows: with more middle-aged residents and more potential entrepreneurs, business formation and expansion responds more aggressively to a reduction in interest rates. This investment raises both demand and supply. Income and employment follow. These first-round effects are followed by knock-on effects. The relatively larger expansion in states with more entrepreneurs tends to increase business income. The rise in business income further increases the incentive to invest and provides entrepreneurs with the resources to take advantage of investment opportunities and expand employment. The result of all this investment is a general rise in output and incomes that benefits all age groups, spurred by the entrepreneurial activities of the middle-aged. Moreover, the increased incentive to invest tends to cause the middle-aged to substitute out of consumption, helping to reconcile our results with the earlier literature that finds it is the consumption of the young that responds more to monetary shocks.

This story is consistent with our earlier finding that business profits respond more the greater the share of middle-aged residents. It would also appear consistent with the greater responsiveness of construction. Looking further, we find evidence for other aspects of the story. We find that establishment births respond more strongly to monetary policy the greater the share of middle-aged residents. (See Figure 10, left panel.) Establishment births slightly lead income and employment. They peak five quarters after the shock, one quarter prior to the peak in income. Interestingly, we find no consistent response of establishment deaths. The

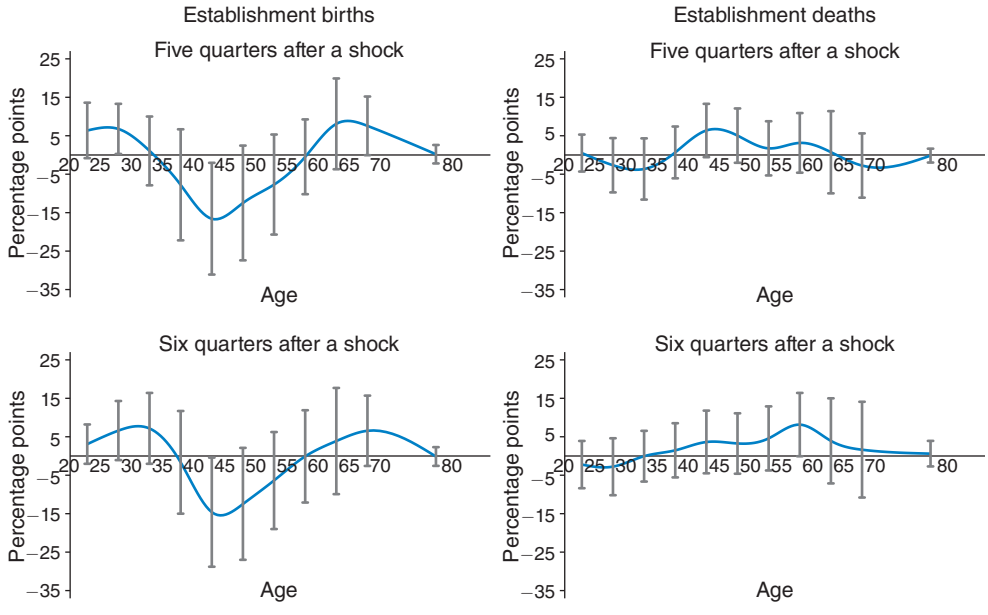

FIGURE 10. EFFECTS ON BUSINESS FORMATION

Notes: The left-hand panels show the effects on establishment births five and six quarters after the shock. The right-hand panels show the effects on establishment deaths. The horizontal axis represents the age groups *a*. The vertical axis represents the effect on the dependent variable of a positive 1 percentage point shock to the federal funds rate interacted with a 1 percentage point increase in the population share of group *a*. The solid line represents the point estimates. The vertical bars represent the 95 percent confidence intervals given Driscoll-Kraay standard errors.

response of deaths is generally insignificant and the estimates are unstable over time. (See Figure 10, right panel.)[25] In some sense this is not surprising. There are two effects of a monetary expansion. The first is to make existing firms more profitable. This effect tends to reduce establishment deaths. The second is to increase the entry of new firms run by entrepreneurs with marginal ability. This effect tends to increase the rate of business failure. The net effect on establishment deaths is therefore ambiguous.

Further evidence comes from looking at the response of earnings, employment, and payroll of firms broken down by firm age and size. Figures 11 and 12 in Appendix Section B show the response at six quarters. While most age and size groups show the familiar patter of positive coefficients for the young, negative coefficients for the middle-aged, and small and statistically insignificant coefficients for the old, this pattern tends to be amplified for young firms and for small firms. Table 1 presents the coefficients for firms 0–5 years of age and firms that are 11 years or older for the two age groups that have the strongest positive and negative effects on monetary policy in our baseline results, 30–34 and 45–49. The point estimates for the younger

[25] When we include the interaction between the monetary shock and the state fixed effect as an additional control, Figure 10 is largely unchanged except that in the case of establishment deaths, the coefficient for the 40–44 -year-old bin becomes positive and significant.

TABLE 1—EFFECTS ON YOUNG FIRMS

Age group	Firm age	Earnings	Employment	Payroll
30–34	0–5 years	5.40	6.68	12.39
		(2.80)	(2.77)	(5.21)
		[0.058]	[0.018]	[0.020]
	11+ years	2.57	2.90	5.87
		(1.37)	(1.33)	(2.57)
		[0.065]	[0.032]	[0.025]
45–49	0–5 years	−8.94	−7.24	−16.17
		(2.18)	(3.11)	(4.59)
		[0.000]	[0.023]	[0.001]
	11+ years	−4.33	−5.18	−10.04
		(1.68)	(1.34)	(2.40)
		[0.012]	[0.000]	[0.000]

Notes: The coefficients represent the effect on the dependent variable of a positive 1 percentage point shock to the federal funds rate interacted with a 1 percentage point increase in the population share of 30–34 and 45–49 year olds. Point estimates are listed first, followed by Driscoll-Kraay standard errors in parentheses and *p*-values in square brackets. Results are presented for earnings per worker, employment, and payroll for 0–5-year-old firms and firms older than 11 years.

firms are at least 50 percent larger than those of the older firms, and the difference is significant at the 10 percent level. The greater importance of young and small firms is reminiscent of the effects found in Gertler and Gilchrist (1994), in which monetary policy has a greater impact on small firms. Gertler and Gilchrist interpret this as evidence of financial frictions.

One might worry that the share of the middle-aged is simply proxying for firm size. It is true that the share of the middle-aged is correlated with the share of small business employment, as would be expected if the middle-aged tended to be the ones to start businesses. The correlation between the share of the population 40–64 years old and the share of employment in businesses with fewer than 500 employees is 0.32 (after removing time effects). It does not appear, however, that the middle-aged population is simply a proxy for small business. Online Appendix Figure A10 shows that the basic pattern seen in Figure 2 survives the inclusion of the small business share interacted with the monetary policy shock. The evidence is therefore consistent with a story in which a monetary expansion prompts entrepreneurs to start and expand small businesses, and this expansion leads to a rise in the incomes of all age groups.

Another concern might be that state-level effects might be more pronounced in nontradable goods industries and that these industries might have more small and young firms. In Appendix C we look at the response of tradable and nontradable goods industries following the categorization in Mian and Sufi (2014). Given the noisiness of the estimates of tradables, we cannot reject that the age pattern in the responses of both groups of industries are similar to each other and to our baseline response. As one might expect, however, the pattern of response for nontradables matches the baseline response more closely, so we cannot rule out the possibility that a correlation between young firms and nontradable industries is driving the relative response of young firms and small firms.

One challenge for the entrepreneur story is to explain the response of large firms. In this story, the comovement of large and small firms, as well as young and old firms and tradable and nontradable industries, is the result of spillovers from increased entrepreneurial activity. This is also what generates the common pattern in age-specific employment and earnings found in Section VC. A greater number of middle-aged residents leads to a greater response of small and young firms, and these pull the rest of the economy in their wake.

VII. Conclusion

We consider the effects of monetary policy shocks on income and employment in US states and how the effects of these shocks differ depending on the age structures of the states. We find that there are three distinct age groups in each state. The young, who are under 35 years of age, tend to dampen the effectiveness of monetary policy on personal income and employment. The middle-aged, who are between 40 and 64 years old, tend to amplify the impact of monetary policy. The old, who are over 65 years of age, do not affect the impact of monetary policy one way or another. We find that these effects work through the aggregate, not the individual: what matters for the income of a given age group is not the age group itself, but the age composition of the state. We find that these results cannot be explained by the industry composition of a state. We propose a mechanism by which a greater share of middle-aged residents increases the elasticity of business formation in response to monetary policy.

Our results are complementary to the findings of Sterk and Tenreyro (2018); Wong (2019); and Cloyne et al. (2020). These authors study the effect of monetary policy on the consumption of different age groups. We study how the effect of monetary policy on aggregate income varies with the age structure of the economy. An interesting direction for future research would be to construct a model consistent with both sets of facts. In the model, liquidity constraints would cause the consumption of the young to respond strongly to monetary policy shocks. The middle-aged would respond not by consuming more but by investing, starting business, and hiring, all activities that lead the economy to expand and increase the incomes of all groups.

Appendix A. Baseline Coefficient Estimates

Table 2 presents the coefficient estimates behind Figure 2.

Table 2—Baseline Results

	Personal income			Private employment		
	Driscoll-Kraay			Driscoll-Kraay		
	Coefficient	Standard error	*p*-value	Coefficient	Standard error	*p*-value
20–24 years	0.83	0.75	0.271	0.93	0.67	0.174
25–29 years	2.30	0.78	0.004	3.08	0.83	0.000
30–34 years	3.13	1.19	0.010	3.63	1.40	0.011
35–39 years	0.54	1.37	0.693	1.57	1.42	0.272
40–44 years	−2.66	1.30	0.044	−0.76	1.28	0.555
45–49 years	−4.06	1.54	0.010	−1.57	2.18	0.474
50–54 years	−4.74	1.50	0.002	−2.70	1.39	0.055
55–59 years	−4.12	1.20	0.001	−2.52	0.96	0.010
60–64 years	−2.46	1.03	0.019	−1.40	0.71	0.054
65–69 years	−0.45	0.68	0.513	−0.29	0.625	0.647
70–99 years	−0.15	0.32	0.633	−0.53	0.26	0.045

Notes: The first column contains the age groups, *a*. The coefficients represent the effect on the dependent variable of a positive 1 percentage point shock to the federal funds rate interacted with a 1 percentage point increase in the population share of group *a*. There are 3,825 observations for 50 states and the District of Columbia.

Appendix B. Firm size and Age

Figures 11 and 12 present estimates of (1) in which the dependent variable X is income and employment broken down by firm size and age. Firm size in Figure 11 is measured by the number of employees. To put these figures into perspective, Table 3 presents the fraction of employment and of payroll comprised by each age and size group. Entries are averages across the sample.

Firms of all types follow the pattern in which the impact of monetary policy is dampened the greater the share of the young and amplified the greater the share of the middle-aged. The effects are similar for employment, income per worker, and total payroll, although the effects on total payroll are stronger than the effects on the number of jobs or earnings per worker. There is also some weak evidence that the pattern is stronger for the extensive margin (the number of jobs) than the intensive margin (earnings per worker) for young and small firms. There is weak evidence that the effects are exaggerated for younger and smaller firms.

We investigate the relative responses of young and small firms further in Figure 13. Here we estimate (1) taking X to be the fraction of young or small firms. We do this exercise for both employment and payroll. The relative responses of the payroll of young firms and of the employment of small firms exhibit the familiar pattern, indicating that the responses of these variables for young and small firms are exaggerated. These effects are statistically significant for the fraction of small-firm employment. The responses of the fraction of young-firm employment and small-firm payroll are relatively noisy and only vaguely resemble the basic pattern.

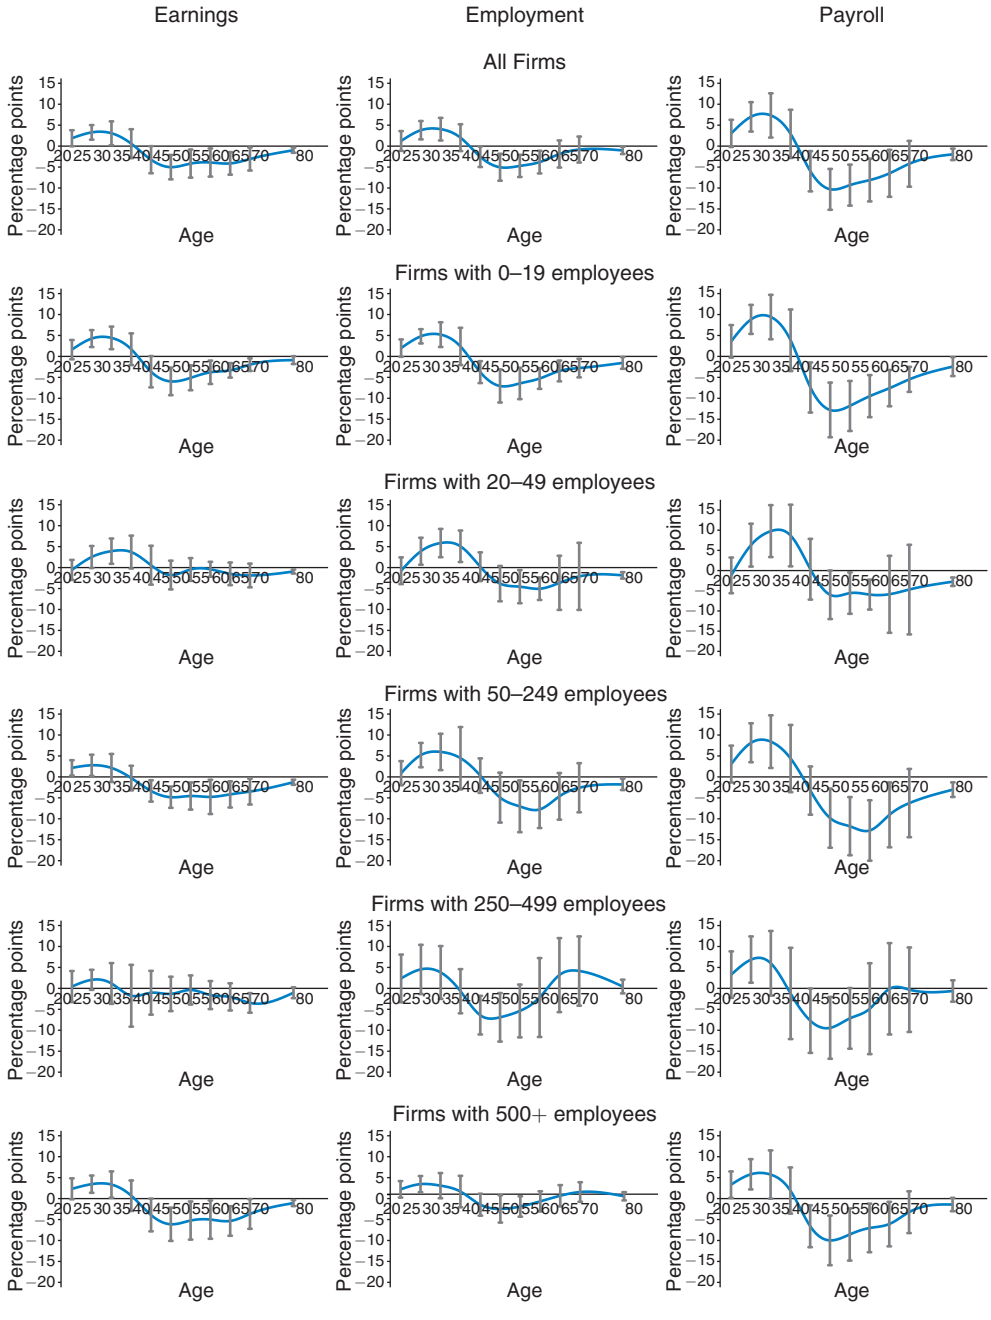

FIGURE 11. RESPONSES OF DIFFERENT FIRM SIZES

Notes: The horizontal axes represent the age groups a. The vertical axes represent the effect on the dependent variable of a positive 1 percentage point shock to the federal funds rate interacted with a 1 percentage point increase in the population share of group a. The solid lines represent the point estimates. The vertical bars represent the 95 percent confidence intervals, given Driscoll-Kraay standard errors.

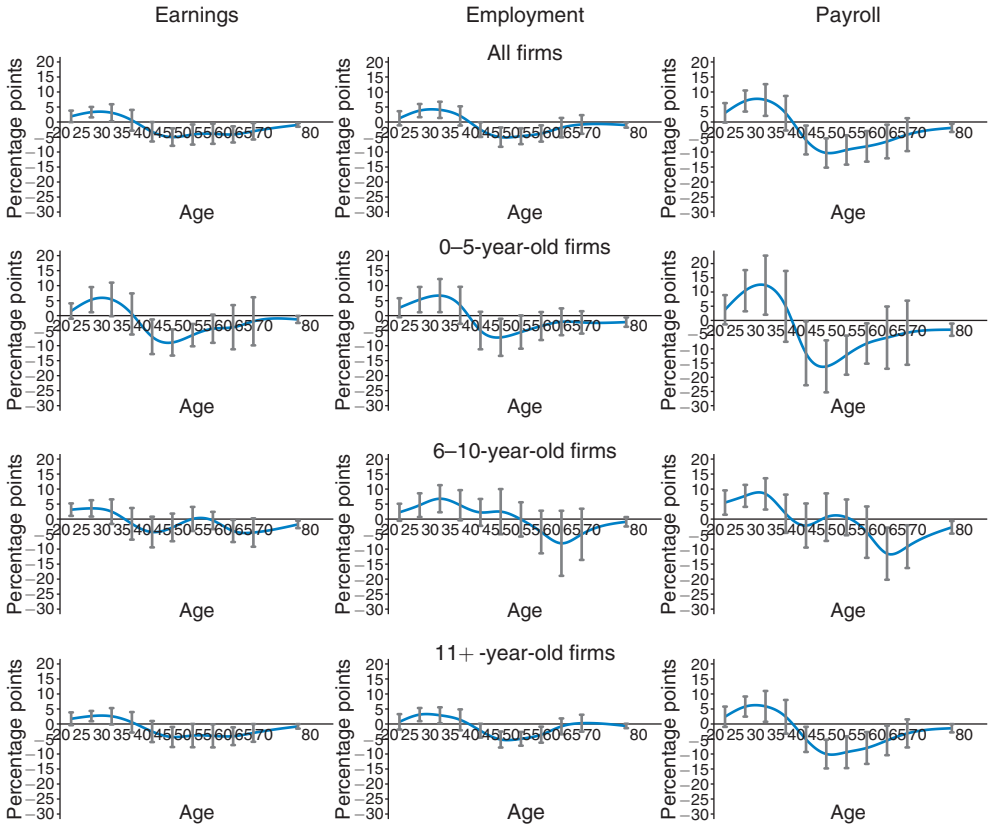

FIGURE 12. RESPONSES OF FIRMS BY AGE

Notes: The horizontal axes represent the age groups a. The vertical axes represent the effect on the dependent variable of a positive 1 percentage point shock to the federal funds rate interacted with a 1 percentage point increase in the population share of group a. The solid lines represent the point estimates. The vertical bars represent the 95 percent confidence intervals, given Driscoll-Kraay standard errors.

TABLE 3—SHARES OF EMPLOYMENT AND PAYROLL BY AGE AND SIZE OF FIRMS

	Employment (percent)	Payroll (percent)
Age		
0–5 years	14.9	12.3
6–10 years	10.2	8.9
11+ years	74.9	78.8
Size		
0–19 employees	21.7	18.2
20–49 employees	10.9	9.8
50–249 employees	16.5	15.6
250–499 employees	5.7	5.5
500+ employees	45.3	50.9

Notes: The first column contains the age or size of the firms. The second and third columns contain the fraction of employment and payroll of each type of firm to all firms. Each entry is the average over the 1990–2008 sample period.

Figure 13. Effect on Employment and Payroll Share of Young and Small Firms Relative to All Firms

Notes: The horizontal axes represent the age groups a. The vertical axes represent the effect on the dependent variable of a positive 1 percentage point shock to the federal funds rate interacted with a 1 percentage point increase in the population share of group a. The solid lines represent the point estimates. The vertical bars represent 95 percent confidence intervals, given Driscoll-Kraay standard errors.

Appendix C. Tradables and Nontradables

To the extent that young and small firms are more concentrated in nontradable industries, one concern is that the estimated response of young and small firms might just be due to a larger response of nontradable industries. Figure 14 takes X to be total payroll in tradable and nontradable industries. We use two definitions of tradability following Mian and Sufi (2014). Both definitions use four-digit NAICS industry-level data. A list of industries for each definition by sector is available in the online Appendix to Mian and Sufi (2014).

The first column groups industries as tradable or nontradable based on the level of geographic concentration of the industries. Since nontradable industries are needed everywhere, they should be geographically dispersed. The production of tradable goods tends to be more concentrated geographically, since these goods require greater specialization and a larger scale of operation relative to nontradable goods. The authors use the share of industry employment to total employment in each county to create a geographic Herfindahl index. Higher index scores indicate greater geographic concentration. We follow Mian and Sufi (2014) and classify the top 25 percent of industries, per the Herfindahl index, as tradable and the bottom 25 percent of industries as nontradable.

The second column groups industries based on the extent of retail and world trade the industries are engaged in. Restaurants and the retail sector are categorized

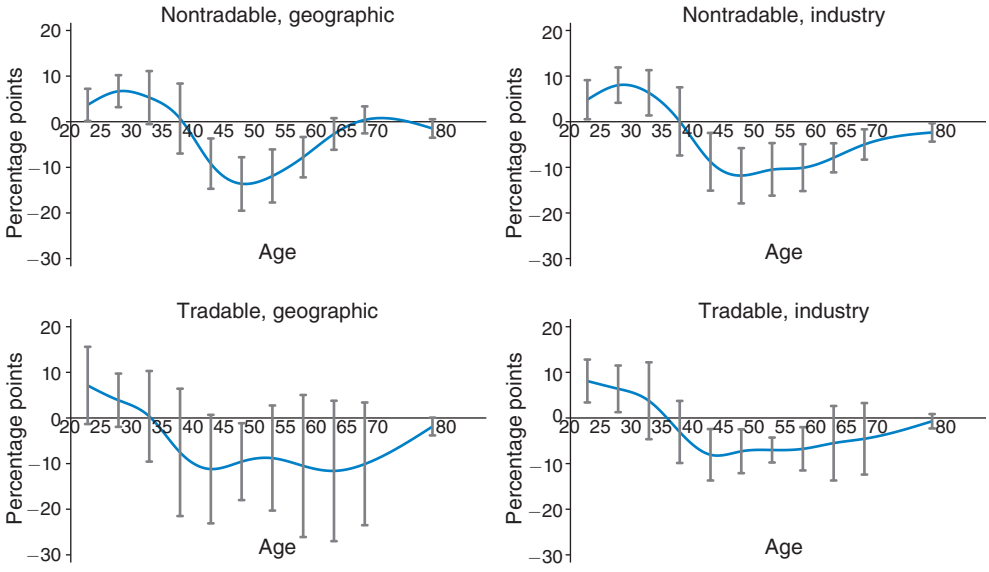

FIGURE 14. PAYROLL IN NONTRADABLE VERSUS TRADABLE INDUSTRIES

Notes: The horizontal axes represent the age groups *a*. The vertical axes represent the effect on the dependent variable (payroll by industry) of a positive 1 percentage point shock to the federal funds rate interacted with a 1 percentage point increase in the population share of group *a*. The solid lines represent the point estimates. The vertical bars represent the 95 percent confidence intervals, given Driscoll-Kraay standard errors.

as nontradable industries. An industry is classified as tradable if either imports plus exports for the industry exceed $500 million or imports plus exports equal at least $10,000 per worker.

The usual pattern of more young residents dampening the effects of monetary policy and more middle-aged residents exacerbating the effects of monetary policy holds for both tradable and nontradable industries. Our results imply that young and small firms are not simply proxying for tradable industries.

REFERENCES

Aksoy, Yunus, Henrique S. Basso, Ron P. Smith, and Tobias Grasl. 2019. "Demographic Structure and Macroeconomic Trends." *American Economic Journal: Macroeconomics* 11 (1): 193–222.

Auclert, Adrien. 2019. "Monetary Policy and the Redistribution Channel." *American Economic Review* 109 (6): 2333–67.

Azoulay, Pierre, Benjamin F. Jones, J. Daniel Kim, and Javier Miranda. 2019. "Age and High-Growth Entrepreneurship." *American Economic Review: Insights* 2 (1): 65–82.

Basso, Henrique, and Omar Rachedi. 2018. "The Young, the Old, and the Government: Demographics and Fiscal Multipliers," Bank of Spain Working Paper 1837.

Beraja, Martin, Andreas Fuster, Erik Hurst, and Joseph Vavra. 2019. "Regional Heterogeneity and the Refinancing Channel of Monetary Policy." *Quarterly Journal of Economics* 134 (1): 109–83.

Berg, Kimberly A., Chadwick C. Curtis, Steven Lugauer, and Nelson C. Mark. 2019. "Demographics and Monetary Policy Shocks." NBER Working Paper 25970.

Blanchard, Olivier Jean, and Lawrance F. Katz. 1992. "Regional Evolutions." *Brookings Papers on Economic Activity* 23 (1): 1–75.

Chodorow-Reich, Gabriel. 2019. "Geographic Cross-Sectional Fiscal Spending Multipliers: What Have We Learned?" *American Economic Journal: Economic Policy* 11 (2): 1–34.

Carvalho, Carlos, Andrea Ferrero, and Fernanda Nechio. 2016. "Demographics and Real Interest Rates: Inspecting the Mechanism." *European Economic Review* 88: 208–26.

Cavanaugh, Linda. 2007. "Home Equity Lines of Credit–Who Uses This Source of Credit?" Census 2000 Brief C2KBR-37.

Census of India. 2011. "C-14: Population in Five Year Age Group by Residence and Sex." https://censusindia.gov.in/nada/index.php/catalog/1541/download/4618/DDW-0000C-14.xls (accessed August 2, 2022).

Cloyne, James, Clodomiro Ferreira, and Paolo Surico. 2020. "Monetary Policy When Households Have Debt: New Evidence on the Transmission Mechanism." *Review of Economic Studies* 87: 102–29.

Cravino, Javier, Ting Lan, and Andrei A. Levchenko. 2018. "Price Stickiness along the Income Distribution and the Effects of Monetary Policy." *Journal of Monetary Economics* 110: 19–32.

Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia. 2017. "Innocent Bystanders? Monetary Policy and Inequality." *Journal of Monetary Economics* 88: 70–88.

Di Maggio, Marco, Amir Kermani, Benjamin J. Keys, Tomasz Piskorski, Rodney Ramcharan, Amit Seru, and Vincent Yao. 2017. "Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging." *American Economic Review* 107 (11): 3550–88.

Doepke, Matthias, and Martin Schneider. 2006. "Inflation and the Redistribution of Nominal Wealth." *Journal of Political Economy* 114 (6): 1069–97.

Driscoll, John C., and Aart C. Kraay. 1998. "Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data." *Review of Economics and Statistics* 80 (4): 549–60.

Eggertsson, Gauti, Neil Mehrotra, and Jacob Robbins. 2017. "A Model of Secular Stagnation: Theory and Quantitative Evidence." NBER Working Paper 23093.

Feyrer, James. 2007. "Demographics and Productivity." *Review of Economics and Statistics* 89 (1): 100–109.

Fujiwara, Ippei, and Yuki Teranishi. 2008. "A Dynamic New Keynesian Life-Cycle Model: Societal Aging, Demographics, and Monetary Policy." *Journal of Economic Dynamics and Control* 32 (8): 2398–2427.

Gagnon, Etienne, Benjamin Johannsen, and David Lopez-Salido. 2016. "Understanding the New Normal: The Role of Demographics." Finance and Economics Discussion Series 2016-080. Washington: Board of Governors of the Federal Reserve System.

Gertler, Mark. 1999. "Government Debt and Social Security in a Life-Cycle Economy." *Carnegie-Rochester Conference Series on Public Policy* 50: 61–110.

Gertler, Mark, and Simon Gilchrist. 1994. "Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms." *Quarterly Journal of Economics* 109 (2): 309–40.

Gertler, Mark, and Peter Karadi. 2015. "Monetary Policy Surprises, Credit Costs, and Economic Activity." *American Economic Journal: Macroeconomics* 7 (1): 44–76.

Gourinchas, Pierre-Olivier, and Jonathan A. Parker. 2002. "Consumption over the Life Cycle." *Econometrica* 70 (1): 47–89.

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 (1): 55–93.

Imam, Patrick A. 2015. "Shock from Graying: Is the Demographic Shift Weakening Monetary Policy Effectiveness." *International Journal of Finance and Economics* 20 (2): 138–54.

Jordà, Òscar. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 95 (1): 161–82.

Kantur, Zeynep. 2013. "Aging and Monetary Policy." Unpublished.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. 2018. "Monetary Policy According to HANK." *American Economic Review* 108 (3): 697–743.

Kara, Engin, and Leopold von Thadden. 2014. "Interest Rate Effects of Demographic Changes in a New Keynesian Life-Cycle Framework." *Macroeconomic Dynamics* 20 (1): 120–64.

Kuttner, Kenneth N. 2001. "Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market." *Journal of Monetary Economics* 47 (3): 523–44.

Leahy, John V., and Aditi Thapar. 2022. "Replication Data for: Age Structure and the Impact of Monetary Policy." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E130082V1.

Mian, Atif, and Amir Sufi. 2014. "What Explains the 2007–2009 Drop in Employment." *Econometrica* 82 (6): 2197–2223.

Nakamura, Emi, and Jón Steinsson. 2014. "Fiscal Stimulus in a Monetary Union: Evidence from U.S. Regions." *American Economic Review* 104 (3): 753–92.

Nakamura, Emi, and Jón Steinsson. 2018. "High-Frequency Identification of Monetary Non-neutrality: The Information Effect." *Quarterly Journal of Economics* 133 (3): 1283–1330.

National Association of Realtors. 2020. "2020 Home Buyers and Sellers Generational Trends Report." https://www.nar.realtor/sites/default/files/documents/2020-generationaltrends-report-03-05-2020.pdf.

Ottonello, Pablo, and Thomas Winberry. 2020. "Financial Heterogeneity and the Investment Channel of Monetary Policy." *Econometrica* 88 (6): 2473–2502.

Ramey, V.A. 2016. "Macroeconomic Shocks and their Propagation." In *Handbook of Macroeconomics*, Vol. 2, edited by John B. Taylor and Harald Uhlig, 71–162. Amsterdam: Elsevier.

Romer, Christina D., and David H. Romer. 2004. "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review* 94 (4): 1055–84.

Rudebusch, Glenn D. 1998. "Do Measures of Monetary Policy in a VAR Make Sense?" *International Economic Review* 39 (4): 907–31.

Selgado, Sergio. 2018. "Technological Change and Entrepreneurship." Unpublished.

Shoag, Daniel. 2010. "The Impact of Government Spending Shocks: Evidence on the Multiplier from State Pension Plan Returns." Unpublished.

Söderström, Ulf. 2001. "Predicting Monetary Policy with Federal Funds Futures Prices." *Journal of Futures Markets* 21 (4): 377–91.

Sterk, Vincent, and Silvana Tenreyro. 2018. "The Transmission of Monetary Policy through Redistributions and Durable Purchases." *Journal of Monetary Economics* 99: 124–37.

Thapar, Aditi. 2008. "Using Private Forecasts to Estimate the Effects of Monetary Policy." *Journal of Monetary Economics* 55 (4): 806–24.

Wieland, Johannes F., and Mu-Jeung Yang. 2019. "Financial Dampening." *Journal of Money Credit and Banking* 52 (1): 79–113.

Wong, Arlene. 2019. "Refinancing and the Transmission of Monetary Policy to Consumption." Unpublished.