DOI: 10.1111/joes.12517

ARTICLE

JOURNAL OF ECONOMIC SURVEYS WILEY

Policy experiments and financial regulation: Using laboratory methods to evaluate responses to the 2007–2009 financial crisis

Douglas D. Davis | **Oleg Korenok**

School of Business, Virginia Commonwealth University, Richmond, Virginia, USA

Correspondence
Douglas D. Davis, School of Business, Virginia Commonwealth University, Richmond, VA, USA.
Email: dddavis@vcu.edu

Funding information
National Science Foundation, Grant/Award Number: SES 1949112

Abstract

This paper reviews an emerging experimental literature that uses laboratory methods to both identify causes of the 2007–2009 financial crisis, and to assess the effectiveness of policies implemented in response. Papers reviewed include experiments conducted to evaluate central bank and Treasury responses to the crisis, experiments that study the consequences of interconnectedness between financial firms on financial system stability, and experiments conducted to evaluate policies intended to more effectively regulate specific types of financial institutions. Laboratory methods are ideally suited to investigating the consequences of untested policies in new environmental circumstances – just the situation provoked by the crisis. The continually evolving structure of the financial system suggests an expanded future role for laboratory methods in this area.

KEYWORDS
bank regulation, financial market stability, laboratory experiments

JEL CLASSIFICATION
C92, D82, G28

1 INTRODUCTION

A complex set of interrelated factors led to the 2007–2009 financial crisis, including, among other things, a global savings imbalance, the growth of an unregulated "shadow banking" sector that supported long term investments with short-term obligations, and a large real estate bubble instigated by the development and securitization of risky "subprime" mortgages, many of which were sustainable only as long as real estate prices continued to rise. The havoc to the financial sector caused by the collapse of this increasingly fragile structure extended to the real economy, as housing prices plummeted, interbank markets froze, and banks sharply restricted lending activity.[1]

Governments responded to the crisis with a host of interventions both on impact, to mitigate the damage, and subsequently to prevent similar situations from developing in the future.[2] Many of the interventions were novel and were predicated on untested assumptions about the underlying financial environment. In exploring the potential effectiveness of policy options, experimental methods have frequently been used as an aid to theory.[3] As observed by Falk and Heckman (2009), for assessing the effects of new policies experimental methods offer two key advantages: data can be easily and inexpensively collected, and policy variations can be exogenously modified in a controlled manner thereby allowing causal inference.

The use of experiments to explore policy responses to the financial crisis was no exception to the growing application of laboratory methods to study policy questions. In the last 8 years, an experimental literature studying policy responses to the crisis has emerged that offers useful insights about and qualifications to new financial regulations. This paper reviews this literature and uses the lessons learned to argue for an expanded use of laboratory methods to study the rapidly evolving nature of financial markets and the potential effectiveness of the associated necessary regulations.

Prior to proceeding, we make two preliminary comments. First, our intention in offering this review is to serve as a bridge between financial and experimental economists. For that reason, we discuss the experiment designs and the surrounding policy context in somewhat fuller detail than is typical in a review paper. We hope that careful discussion of the experimental designs will help financial economists better understand the experiments, while discussion of the policy context will provide insight for experimental economists into the factors regarded as causes of the crisis and the considerations that must be balanced in developing effective policy solutions. The second comment delimits the scope of this review. Specifically, we largely omit discussion of experiments evaluating the workhorse model of financial fragility by Diamond and Dybvig (1983), because of our focus on regulatory responses to 2007–2009 financial crisis and because it has been reviewed previously (see Dufwenberg, 2015 and Kiss et al., 2021).[4] For similar reasons, we do not review asset market experiments stemming from the environment developed by Smith et al. (1988). (See recent reviews by Palan, 2013 and Nuzzo & Morone, 2017)

Our review consists of three main sections. The Section 2 below reviews experiments conducted to evaluate implementation issues in direct actions of a central bank or Treasury. The subsequent Section 3 turns attention to experiments that evaluate policies that affect the fragility of the financial system as a whole. Section 4 reviews experiments that consider the effects of policies, both proposed and implemented, on specific financial institutions. These sections are divided into subsections. Each subsection starts with a review of events during the 2007–2009 crisis that triggered a need for policy actions. Next, we describe the new policies. After reviewing the policy changes, we describe an experiment design that evaluates the policy, or an implementation issue associated with the policy. We finish each subsection with a summary of experimental findings. Following the main sections, the paper concludes with a short summary discussion in Section 5.

2 | DIRECT ACTIONS OF THE TREASURY AND CENTRAL BANK

In the immediate aftermath of the crisis, governments were primarily concerned about restoring a stability and confidence in the financial system sufficient to stave off collapse. Pertinent policy issues included getting solvent but liquidity-stressed banks to access the discount window, finding ways to remove the “toxic” assets from bank balance sheets that were limiting interbank lending activity, and developing a better understanding of banks’ capacities to respond to stress. Questions about the appropriate responses to these issues motivated the experiments reviewed in this section.

2.1 | Discount window stigma

In the early stages of the crisis, liquidity stressed banks were reluctant to borrow from the central bank via the discount window, thus undermining the Federal Reserve System’s role as a lender of last resort. Many commentators argue that “stigma” drove this reticence. A bank’s need to access the discount window indicates liquidity stress at a minimum and may indicate insolvency.[5] Banks may be reluctant to use the discount window facility out of a fear that private investors will regard discount window access as a signal of insolvency.

Bank behavior in the fall of 2007 was certainly consistent with stigma. Despite massive liquidity problems, banks accessed the discount window only sparingly. As a first effort to stimulate discount window access, the Fed reduced the discount rate by 50 basis points in August of 2007, and increased the maturity of discount window loans from overnight to as long as 30 days. These changes were met with little success. In December 2007 the Fed tried a second approach, by opening a temporary liquidity program, the Term Auction Facility that was specifically designed to mitigate stigma. The Term Auction Facility differed from the discount window in a number of important ways. Among other things, rather than paying an announced rate, banks submitted two-part bids consisting of an interest rate and a loan amount. To reduce the chance of a bank being singled out for using Fed funds, the bid size was capped at 10% of available funds, in this way assuring at least ten simultaneous auction winners. Finally, to mitigate the perception that auction participation was motivated by pressing needs, auction winners could obtain funds only after a 3-day delay. The Term Auction Facility was an immediate success, despite the fact that the auction funds rate routinely exceeded the discount rate (Armantier et al., 2015).[6]

The Term Auction Facility was an emergency response to a serious if not catastrophic crisis when the normal regulatory mechanisms failed. Going forward, it would be well to correct the operational structure of the discount window so that it better serves its intended purpose. Key to understanding banks’ reluctance to access the discount window facility is an understanding of the mechanisms of and remedies for stigma. Armentier and Holt (2020a) explore this issue with an experiment that investigates the causes of and remedies for stigma. The experiment consisted of a series of twelve-player sessions. At the start of each session, half the participants were assigned the role of banks while the rest were investors. The sessions then proceeded as a series of two-period rounds. As shown in Figure 1 prior to the start of a period 1, each bank is randomly paired with an investor, and then privately informed about its solvency status (solvent or insolvent) and

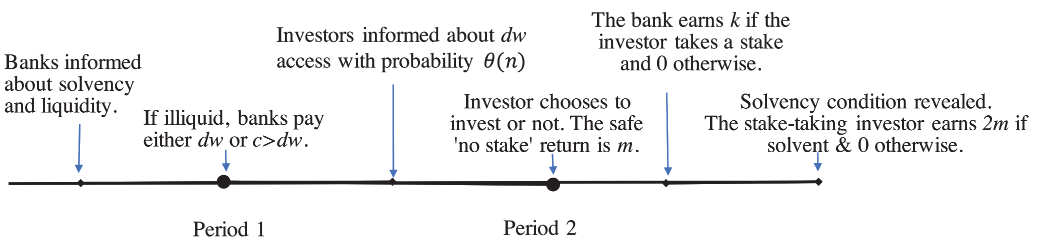

FIGURE 1 Sequence of round actions in the baseline treatment of the discount window stigma experiment [Colour figure can be viewed at wileyonlinelibrary.com]
Source: Armentier and Holt (2020a).

liquidity condition (liquid or illiquid) for the round. Although all insolvent banks are illiquid, a solvent bank may also have a temporary liquidity deficiency.[7] The bank's objective is for its investor to take a stake in the bank, in which case the bank earns a fixed return k, regardless of solvency. In period 1 illiquid banks address their liquidity deficiencies by choosing to pay either a low discount window fee dw that the investor may potentially observe, or a more costly but unobservable outside option $c > dw$. The bank cares about the observability of a discount window payment, because the payment reveals that the bank has a short term cash deficiency, which the investor may interpret as a signal of insolvency and thus choose to not take a stake in the bank.

Following the bank's period 1 decision, the investor is probabilistically informed about whether or not the bank accessed the discount window according to a detection probability. In period 2 the investor decides whether or not to take a stake in the bank. If the investor does not take a stake, the investor realizes a certain return m. If the investor does take a stake, the investor's earnings depend on the bank's solvency. In the event the bank is solvent, the investor's earnings double to $2m$. If the bank is insolvent the investor earns nothing.

A critical aspect of the game is the discount window detection probability, which is inversely related to the number of banks who chose to access the discount window. If only one bank pays dw, the probability of detection is very high (75% in the control treatments), while the probability of detection falls as more banks pay dw (to 50% if two banks pay dw, and 25% if more than two banks pay dw). Given the parameters Armentier and Holt used, two pure strategy equilibria exist for this game, a "no stigma" equilibrium where investors always choose to take a stake in their bank and so all illiquid banks access the discount window regardless of solvency, and a "stigma" equilibrium where investors never fund banks that were detected accessing the discount window, and, to avoid detection, no illiquid banks pay the lower dw fee.

Experimental results of a baseline treatment, summarized in the top row of Table 1 indicate a strong propensity for participants to coordinate on the stigma equilibrium. Investors chose to fund banks detected accessing discount window in only 4% of instances, compared to an 84% funding rate for banks not detected accessing the discount window. In turn, when illiquid, only 22% of solvent banks, and 18% of insolvent banks chose to pay dw rather than c.

Although reliably generating a stigma equilibrium outcome is in itself interesting, the most important feature of the baseline result is that it allows the evaluation of policies intended to mitigate stigma. Armentier and Holt studied three such policies. The first and simplest policy would be to reduce the cost of discount window access. The authors explored this option in a "low cost" treatment where the cost to banks of accessing the discount window is cut by half. As shown in the second row of Table 1, the access cost reduction improved discount window access only marginally (and insignificantly). Investors elected to take a stake in banks detected accessing the

TABLE 1 Average investment and bank discount window access decisions

Treatments	Investor behavior		(Illiquid) Bank decisions	
	***dw* access not detected (%)**	***dw* access detected (%)**	**Access *dw* if solvent (%)**	**Access *dw* if insolvent (%)**
Control	84	4	22	18
Low cost	83	14	32	33
Low detection	84	16	28	27
Random borrowing	84	45**	78***	71***

Note: Illustrated results represent average outcomes per treatment for the last 10 of the 25 rounds in each treatment session. *Key*: **,*** reject the null that the observed average in the random borrowing treatment does not differ from the observed average in any of the other treatments, $p < .05$, $p < .01$, respectively. No other differences across treatments were significant at $p < .10$.
Source: Armentier and Holt (2020a).

discount window in 14% of instances compared to a funding rate of 83% for undetected banks. In response, the number of illiquid banks that chose to pay *dw* to cover liquidity needs rose from less than 23% in the control treatment to roughly 33% in the low-cost treatment, an increase that was also insignificant. Results thus suggest that reducing the discount window costs alone is unlikely to overcome stigma.

A second policy would be to invest resources in increasing the confidentiality of discount window access. Armentier and Holt explore the effectiveness of this alternative in a "low detection" treatment, where the probability that an investor learns that a bank paid *dw* was reduced from 75% to 50% if only one bank paid *dw*, and from 50% to 25% if two or more banks paid *dw*. Results, summarized the third row of Table 1, are virtually the same as results in the "low cost" treatment, indicating that improved confidentiality is also not a promising remedy.

A third policy weakens the link between discount window access and potential insolvency by implementing a variant of a measure proposed for the Bank of England's discount window facility (Winters, 2012). The idea is to require banks to occasionally access the discount window in the normal course of events so as to reduce the visibility of any crisis access. Armentier and Holt investigate this approach with a "random borrowing" treatment in which one of the six banks is randomly selected each round and required to pay the *dw* fee, even if not illiquid. This forced discount window access is common knowledge to all participants, so investors knew that at least one of the banks accessing the discount window was doing so independent of any liquidity need. As can be seen in the bottom row of Table 1, the random borrowing condition substantially improved the discount window access rates. The rate of investors choosing to take a stake in banks detected as accessing the discount window rose from the 4% observed in the control treatment to 45%, and the discount window access rate of liquidity stressed banks rose from roughly 20% in the control treatment to over 70% in the random borrowing treatment. Armentier and Holt further find that the need to impose a random borrowing requirement may only be temporary. In a subsequent robustness treatment that still included both stigma and no stigma equilibria they found that discount window access rates remained high even after the random-borrowing requirement was lifted.

The contribution of this experiment merits some discussion. Of course, the laboratory environment is far simpler than the decision context faced by bank managers. Nevertheless, Armentier and Holt were able to isolate the stigma that is of concern to policymakers, and exploiting the control allowed by the laboratory, examine in isolation the effectiveness of policies intended to mitigate stigma. Experimental results appealingly parallel both the ineffectiveness of the Fed's discount rate reduction in the fall of 2007, and the success of the Term Auction Facility. Further, results of the random borrowing treatment, which emulates a proposal intended to reduce discount window stigma at the Bank of England, support the idea that required, regularized use of the discount window may facilitate its intended purpose during the next crisis.

2.2 | TARP auction design

Another issue that arose during the crisis was the problem of "toxic" assets. Uncertainty regarding the quality of mortgage backed securities offered as collateral by liquidity deficient banks seeking loans caused interbank lending activity to plummet and loan rates to skyrocket. To restore the normal functioning of the banking system, the Department of the Treasury considered using the $700 billion in emergency funding granted by Congress to acquire toxic mortgage backed securities from banks. To remove the largest volume of toxic assets at the lowest possible cost to taxpayers, the Department of Treasury proposed holding an auction.[8] Such an auction was novel, and Treasury officials were uncertain as to how to proceed. Given imperfect information about the quality of the assets, a standard low-price wins "grand" auction was not a desirable option because it would result in the sale of only the lowest quality assets, making the acquisition extraordinarily costly for the government. After consulting with several teams of academics, Treasury officials decided to focus on a "reference price" auction structure proposed by three experimental economists, Jacob Goeree, Charles Holt and Charles Plott, in consultation with the Fed's auction expert, Oliver Armentier. In the reference price auction bank bids are adjusted for the asset's expected quality.

To understand how a reference price auction works, consider a simple procurement auction with two bidders and two securities. Bidder 1 holds one unit of a low-quality Security A, which has a value of $5, and bidder 2 has two units of a high-quality Security B each of which have a value of $10. The auctioneer (e.g., Treasury) has a budget of $18, and seeks to remove as many securities from the market as possible, at the lowest cost. For simplicity we constrain bids to integers and break ties at random.

In a uniform-price grand auction, the lowest priced bid wins and is paid a price equal to the first rejected bid. In an equilibrium, each agent bids its value, yielding an auction price of $10, and the sale of only security A (since after the acquisition of security A, the remaining budget of $8 is insufficient to purchase a second security).[9] The government's purchase efficiency is 50% (5/10). Efficiency is so low because the government pays a high price to remove only the security with the lowest values from the market.

When information about the values of the Securities A and B are known, adjusting bids for reference prices can sizably improve purchase efficiency. The auctioneer assigns a reference price of 1 to security A, a reference price of 2 to security B, and then normalizes each bid by dividing it by its reference price. Bids with the lowest normalized price are accepted first until the auctioneer's budget is exhausted. In the equilibrium, bidder A bids value while bidder B, who has some market power because she has two securities, bids $12 for each of her securities, for a normalized bid of $12/2 = $6 each. The normalized equilibrium price is $6, and the government's budget of $18 is exhausted by acquiring Security A ($6 times reference price of 1) and one unit of Security B ($6 times reference price of 2). In this case, the purchase efficiency rises to 88.3% (15/18).

Despite the improved purchase efficiency suggested by the above example, Treasury officials were concerned about two aspects of the reference price auction. First, while the reference price auction is demonstrably more efficient than a grand auction when the values of assets are known, it is unclear how it might perform in the more pertinent circumstance where asset's value is uncertain to both banks and the government. Second, Treasury officials were worried about the possibility that bidders could use their informational advantage to manipulate auction outcomes. To mitigate bidders' exploitation of the auctioneer's reference price errors Treasury officials considered hiding reference prices from bidders until after the auction was complete.

To demonstrate the efficacy of the reference price mechanism and to evaluate Treasury concerns, Armantier et al. (2013) conducted a laboratory experiment in which groups of six participants, each heterogeneously endowed with differently valued assets, participated in a series of auctions. In addition to treatments featuring a grand auction and a reference price auction with known asset values, these investigators evaluated two additional treatments. In one treatment, asset values were only known stochastically. Traders each received a private signal drawn from a uniform distribution centered at the asset's true value. The auctioneer was even less well informed than traders, receiving a signal drawn from a uniform distribution with a range twice that of the traders. As a consequence, reference prices were highly inaccurate. (In one auction, for

TABLE 2 Reference price auction efficiencies

Treatment, reference prices	Mean purchase efficiencies
Grand auction (none)	76.6
Secret, noisy	85.4**
Announced, noisy	87.6
Announced, accurate	93.6**

Note: Each mean purchase efficiency is the average of the eight sessions in each treatment. ** indicates rejection of the null that the mean efficiency in one row does not differ from that in the row above it at $p < .01$ (Wilcoxon signed rank test). Source Armantier et al. (2013).

example, the reference price of a security was set at one-ninth of its true value.) The final treatment replicated procedures in the third treatment, with the difference that reference prices set by the auctioneer were hidden from bidders.

As can be seen from the purchase efficiencies in Table 2, announced accurate reference prices significantly improved purchase efficiency, from 76.6% in the grand auction baseline to 93.6% in the announced accurate reference price auction. More interesting still, in the announced noisy reference price treatment mean purchase efficiency fell only marginally from 93.6% to 87.6%. Moreover, keeping reference prices secret from bidders did not improve purchase efficiencies. Purchase efficiency in the secret noisy reference price auction at 85.4% differed insignificantly from the announced noisy reference price auction.

A subsequent look "inside the box" at individual bidding behavior reveals the reason why secret reference prices failed to improve purchase efficiency. As Treasury officials expected, bidders respond strategically to reference prices, raising bids in response to an overly high reference price. However, when the reference price was very low, bidders who owned that security recognized that they were at a disadvantage and needed to compensate with lower bids to remain competitive. Adjustments to overly high and low reference prices were offsetting, and yielded net results that are statistically indistinguishable from sessions where reference prices were hidden.

Ultimately, the Department of Treasury changed course, and rather than conduct the TARP auctions, decided to use the funds allocated by Congress to boost banks' capital by taking equity positions in the banks. Nevertheless, the laboratory results show that a reference price auction has the potential to improve efficiency compared to a grand auction, even given substantial uncertainty about asset's values and inaccurate reference prices. The results also indicate that hiding reference prices neither hurts nor improves auction efficiency. These results demonstrate the usefulness laboratory experiments in evaluating competing designs of public policies when polices are to be implemented in complex environments under time pressure.[10]

2.3 | Optimal disclosure of bank stress test results

To help restore confidence in the financial system in the aftermath of the crisis, the Federal Reserve Bank conducted the Supervisory Capital Assessment Program, which consisted of "stress tests" that evaluated institutions' capacities to survive conditions of extreme financial duress. Subsequently, the Dodd Frank Act (2010) mandated continued periodic repetitions of these stress tests.

An intensely debated aspect of stress tests regards the specificity with which regulators should disclose test results. A natural position, consistent with current Federal Reserve policy, would be to disclose results for individual banks; disclosure allows for more informed bids which in turn generates market prices that more fully reflect an asset's value. Disclosure, however, is not without costs. The theoretical results in Hirshleifer (1971) suggests that information disclosure would critically limit the opportunity of banks who suffer cash deficiencies to borrow funds from banks with excess cash. An optimal disclosure policy involves balancing of the informational benefits of disclosure against the potential costs. Goldstein and Leitner (2018) show that one way to balance benefits and costs of disclosure is to pool groups of banks into disclosure classes.

The following example illustrates how variations in disclosure conditions can improve banks' ability to maintain long term investments. Consider the situation of a bank that is endowed with an asset that yields a return equal to the sum of an *intrinsic value* and a *random dividend*. The intrinsic value represents the bank's type and takes on one of three equally likely outcomes, $6, $9 or $13. The random dividend represents the bank's idiosyncratic risk, which is realized only after transactions decisions are completed. The random dividend is drawn from a uniform distribution that is centered on the bank's type realization and has a range of $10. Three competitive buyers offer bids for the asset based on information about the asset supplied by a regulator. In light of the winning bid, the bank decides to either sell the asset or hold it to maturity, under the condition that if the bank sells at a market price of $10 or more, or holds the asset to maturity and the final asset value exceeds $10, the bank's return increases by a $15 bonus. This bonus represents the bank's success at maintaining its long-term investment projects.[11]

The discontinuity in bank payoffs creates room for information pooling. If a regulator, interested in maximizing bonus realizations (e.g., mature long-term investment projects), discloses no information, buyers would bid the average of the three possible intrinsic values, or $9.33, and banks would refuse to sell, since any sale would preclude a bank from realizing the bonus.[12] In this case bonuses would be realized only when the mature asset value exceeds $10, which in expectation will occur in 43% of instances.[13] Alternatively, the regulator might fully disclose intrinsic value information. In this case, banks would agree to sell only the asset with a $13 intrinsic value, since a bank's expected return from $6 and $9 realizations, both of which include a probability of realizing the $15 bonus, exceeds buyers' bids. The certain sale of the asset with an intrinsic value of $13 raises slightly the expected bonus incidence for a full disclosure to 50%.[14] As a third option, the regulator could adopt a partial disclosure policy consisting of a pooling of the $9 and $13 intrinsic values along with a separate disclosure of the $6 intrinsic value. Such a partial pooling would substantially increase the incidence of bonuses. Although banks would forego sales for the $6 realizations, the mean value of the two highest intrinsic values, $11, would result in sales for both $9 and $13 value realizations, because at a price of $11, the bank would realize a certain return of $26 (including the $15 bonus), which strictly exceeds the bank's expected return from holding either asset to maturity.[15] Thus, under this partial disclosure condition the expected bonus incidence rate rises to 70%.[16]

To behaviorally evaluate the effects of disclosure rule variations on bonus incidences Cox et al. (2021) conduct an experiment that used the parameters in the example just discussed.[17] A first treatment consisted of repeated interactions between a bank and three buyers, as specified above. As the bonus incidences shown in the top row of Table 3 illustrate, in this case Cox et al. observe no beneficial consequences of a partial disclosure regime.[18] Bonus incidences in the partial disclosure regime did not differ significantly from those observed in the no and full disclosure regimes, and in fact were numerically a bit lower. The main driver of this outcome was persistent buyer overbidding, a sort of "winner's curse". In particular, winning bids for the $9.33 expected intrinsic value in the no disclosure regime and for the $9 intrinsic value in the full disclosure regime consistently exceeded $10. As a consequence, bonus incidences in the no disclosure and full disclosure

TABLE 3 Average bonus incidences

Treatments	Investor information		
	Low (%)	Partial (%)	Full (%)
Real buyers	68††	60††	62††
Simulated buyers	35***††	61††	40***††
Predicted	43	70	50

Note: Illustrated results represent average outcomes per treatment for the last eight of the 15 rounds in each treatment sequence. *Key*: *** reject the null that the entry does not differ significantly from the value in the partial information condition for the treatment, $p < .01$. †† reject the null that the entry does not deviate from the predicted bonus incidence, $p < .05$. No other differences across treatments were significant at $p < .10$. Source: Cox et al. (2021).

regimes exceeded predicted levels. Bank sales decisions also deviated from optimizing behavior, with banks frequently ignoring the likelihood of the $15 bonus when making sales decisions.

To examine the persistence of nonoptimal bank sales decisions and the impact of these decisions on bonus incidences, Cox et al. conducted a second treatment which followed procedures identical to those in the first treatment, except that buyer bids were simulated to equal the asset's intrinsic expected value in light of the information disclosed. As summarized in the second row of Table 3 in this case bonus incidences under partial disclosure exceed those in no disclosure and full disclosure regimes by 26 and 21 percentage points, respectively, with both differences significant at $p < .01$. Despite the improvement in bonus incidences, Cox et al. observe that banks still deviate substantially from optimal decisions. As can be seen from the crosses in the second row of the table, even with simulated bidders realized bonus incidences were significantly below the equilibrium prediction in each disclosure regime. The lower than predicted bonus incidences were driven by a propensity for banks to focus on an asset's expected (intrinsic) value exclusive of the bonus when making a sales decision. Banks accepted bids that were below their expected payoff from not selling in about one third of instances and rejected bids above $10 that would guarantee the bonus nearly half of the time.

In sum, Cox et al. interpret their results as suggesting that variations in disclosure conditions can be expected to have little effect on lending behavior if bids generically exceed assets' expected values. Such a circumstance might arise in good economic times. For example, in the years following the crisis, the economy expanded rapidly and the banking sector expanded along with it. In such an environment, the Federal Reserve's current practice of publishing disaggregated results is unlikely to adversely affect risk sharing activity between banks.[19] To the extent that winning bids approximate assets' expected values, however, as might occur in times of economic stress, a partial disclosure policy that balances benefits and costs of information disclosure can importantly improve risk pooling behavior, and this result is robust to substantial deviations of bank managers from the behavior predicted in the model.

3 | POLICIES AFFECTING THE FINANCIAL SYSTEM AS A WHOLE

One of the primary lessons of 2007–2009 financial crisis for regulatory authorities was that to ensure the stability of the financial system, an exclusive focus on the well-being of individual financial institutions was insufficient. Rather, effective prudential regulation requires a more holistic approach. As a result, the recommendations emerging from the Basel III Agreement, many of which were institutionalized in the United States through the Dodd Frank Act (2010),

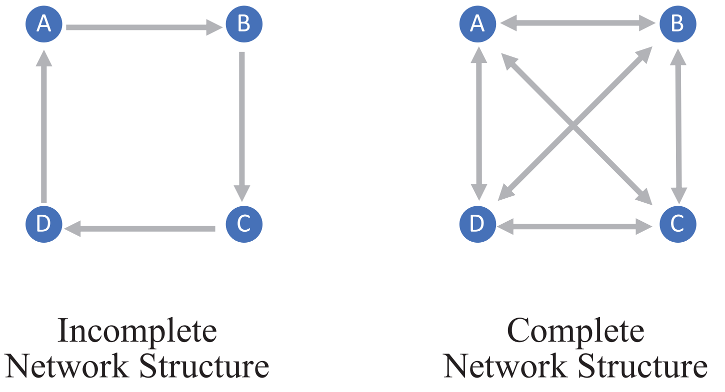

FIGURE 2 Symmetric banking network structures [Colour figure can be viewed at wileyonlinelibrary.com]
Source: Duffy et al. (2019).

resulted in a series of regulations intended to improve the stability of the banking system in part by ensuring the solvency of systemically important banks. This section reviews three experiments that respectively illustrate the importance of interconnectedness on financial stability, justify a focus of attention on systemically important banks, and illustrate an unintended consequence of new regulations caused by bank interconnectedness.

3.1 | Bank interconnectedness and financial contagion

Financial firms' insolvencies both triggered the 2007–2009 financial crisis and through the impacts of those insolvencies on broader financial system, caused its spread. For example, Lehman Brother's collapse was associated with a $423 billion contraction in the U.S interbank lending market Gorton (2010), which in turn forced other banks to the brink, requiring government bailouts (Morgan Stanley) or to be sold off (e.g., Merrill Lynch). The unanticipated sensitivity of other financial institutions to the Lehman Brothers collapse made it clear to regulatory authorities that effective prudential policy required a better understanding of the interrelationships between financial institutions.

The workhorse model for analyzing factors affecting financial fragility of banks is the Diamond and Dybvig (1983) banking model. Although the model and its variations illustrate the stabilizing effects of policies that affect individual banks, such as deposit insurance and liquidity suspension, it does not address interrelationships between financial institutions. The pioneering work of Allen and Gale (2000) extends the Diamond and Dybvig framework to a banking network, and demonstrates the importance of the network structure of a banking system on financial stability.

Duffy et al. (2019) reports an experiment based on an implementation of the Allen and Gale (2000) model. The experiment was conducted to evaluate the effects of variations in the interconnectedness of banks on financial stability, but also lends insight into the effects of policies that reduce liquidation costs on bank contagion. Figure 2 illustrates the bank networks examined in the experimental environment. Each bank has four depositors: two "impatient" depositors who must withdraw early due to short term liquidity needs, and two "patient" depositors who, provided that the bank remains solvent, would enjoy higher utility by maintaining funds in the bank until asset maturity, but may withdraw early if they fear that the bank may become insolvent. In the incomplete network structure, shown in the left panel of Figure 2, the banks are linked unidirectionally, Bank A deposits with Bank B; Bank B deposits with Bank C, and so on. In the complete network structure shown in the right panel of the Figure, each bank places deposits in and accepts

TABLE 4 Estimated withdrawal odds ratios (relative to the complete network structure with $r = 0.20$)

Network structure	Liquidation value	
	$r = 0.20$	$r = 0.40$
Incomplete	3.498**	0.293*
Complete	1	0.201***

Key: *,**, *** reject the null that the probability of a depositor withdrawal in a network structure/liquidation value treatment cells do not differ from that in the complete network $r = 0.2$ cell, $p < .10,.05, .01$ respectively. Source: Duffy et al. (2019).

deposits from all three other banks. Allen and Gale show that in both network configurations the first-best outcome can be achieved. The network structure, however, critically affects the banking systems' stability in response to even a small additional shock. In the incomplete network, a bank impacted by an unanticipated additional withdrawal suffers bankruptcy, and as a consequence induces insolvency in the bank connected to it, setting off a chain of further insolvencies through the network. In the complete network, on the other hand, insolvencies can be confined to the bank facing an additional early withdrawal. The impact on the other banks is muted since the shocked bank divides its obligations among each of the other banks in the network. In addition to varying network structure, Duffy et al. consider a two asset liquidation values, a low value $r = 0.20$ and a higher value $r = 0.40$. Increased liquidation values make individual banks less susceptible to fragility and for that reason may serve as a proxy for a regulation that requires higher capital surcharges.[20]

Experimental results, summarized by the odds ratios in Table 4 clearly indicate that network structures matter.[21] Given a low liquidation value, depositor withdrawals were 3.498 times more likely with an incomplete network than with a complete network. On the other hand, the liquidation value also importantly affects withdrawals. Given a complete network structure withdrawal rates are 4.978 (e.g., 1/0.201) times lower when the liquidation value is 40% rather than 20%. Duffy et al. further report that, given an incomplete network structure, the original shock spreads to an average of 2.5 other banks when the liquidation value is 20% while it spreads to 0.5 other banks when the liquidation value is 40%. Finally, increases in the liquidation value powerfully mitigate the effects of an incomplete network structure. While the difference is only marginally significant, early withdrawal rates given a high liquidation value in an incomplete network are lower than early withdrawal rates given a complete network structure and a low liquidation value.

The stabilizing effects of increased liquidation values suggests that regulations which reduce liquidation costs or make liquidation unnecessary, such as risk-based capital surcharges, may ease the effects of network-induced contagion even in relatively unstable network settings.

3.2 | Interconnectedness and systemically important banks

One important dimension of network structures that the Duffy et al. (2019) experiment design could not capture regards the importance of "core" or systemically important banks. Banking reforms agreed to by the G-20 in the Basel III accord (that in the U.S. were codified in the Dodd-Frank legislation) focus primarily on systemically important banks and nonbank financial institutions that were identified by their size and interconnectedness.

Choi et al. (2017) reports an experiment that provides useful evidence illustrating the importance of core institutions on the stability of a financial system. In a core-periphery network highly connected banks interact with each other as well as with poorly connected peripheral banks.

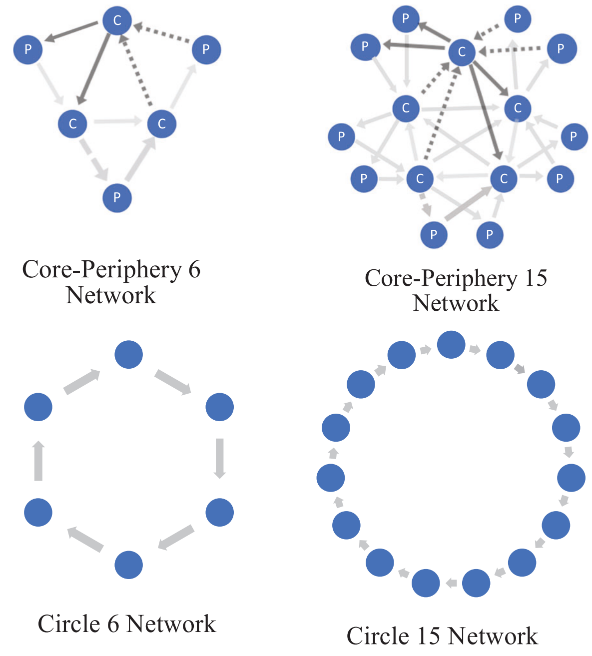

FIGURE 3 Circular and core-periphery networks with six and 15 banks [Colour figure can be viewed at wileyonlinelibrary.com] (*Source*: Choi et al., 2017).

As Choi et al. observe, empirical studies identify a variety of financial markets as having core-periphery structures. Included among their cited examples are the federal funds market (Bech & Atalay, 2010), interbank markets (e.g., Boss et al., 2004) and the US Federal Reserve Bank loans program (Battiston et al., 2012), so assessing the stability of core-periphery network structures relative to easier to analyze structures such as a circular network is an important task. The upper left panel of Figure 3 illustrates a core-periphery network with six banks. As indicated for a representative core bank by the dark gray lines, each core bank is linked to the other two core banks as well as to two peripheral banks. On the other hand, as the thickened light bars at the bottom of the panel illustrate for a representative peripheral bank, each peripheral bank is linked only to two core banks. Increasing the network size magnifies the disparity in the number of links between core and peripheral banks, as can be seen from the 15-bank core-periphery network shown in the upper right panel of Figure 3, where each core bank is linked to 8 banks, while each peripheral bank remains linked to two core banks. These complex linkages contrast sharply with the 6-bank or 15-bank versions of incomplete or circular network shown in the bottom panels of Figure 3, where every bank is linked to only two other banks.

Choi et al. behaviorally evaluates two factors that may affect contagion in response to a shock. First, they evaluate the relative importance of network structures on shock transmission (as discussed by Allen & Gale, 2000, and further developed by Acemoglu et al., 2015). Choi et al. examine contagion in the circular and core-periphery networks with both six and 15 banks shown in Figure 3. Second, to assess the contagion-inducing effects of informational uncertainty about the location of the shock (analyzed by Caballero & Simsek, 2013), the authors examine two information conditions, an informed condition, where the identity of the shocked bank is made public, and an uninformed condition, where the identity of the shocked bank remains private.

The mechanics of the Choi et al. design differ markedly from those just reviewed in Duffy et al. (2019). Rather than evaluating contagion emanating from depositors' withdrawal decisions, Choi et al. study the asset purchase and sales decisions of linked banks. The experiment consisted of a series of three period "rounds", $t = 0,1,2$. At the outset of each round, each bank is endowed with one unit of a long-term asset that matures in period 2, and a series of short-term loans and obligations to other banks, as shown by the series of inward and outward arrows to and from each node in Figure 3. To motivate trade, each bank is randomly assigned a unique asset value. Prior to the start of period 0, one bank is hit with a shock in the form of an obligation to pay a certain amount of money in period 1.[22] Period 0 consists of a 90 s double auction market in which the shocked bank as well as banks with low asset values sell assets to banks with high asset values. Following the close of trade, the computer clears all cross holdings of debt in period 1. Any bank unable to satisfy its obligations is bankrupt. Finally, in period 2, assets mature and payoffs are determined for participants who did not go bankrupt.

Choi et al. find that network structure importantly affects contagion, particularly in the thicker networks. The authors report that in the 15-bank core-periphery network banks broached a 60% contagion threshold (e.g., 60% of banks went bankrupt) in 43% of instances in the informed condition and in 50% of instances in the uninformed condition. In the comparable 15-bank circular network, banks broached the 60% contagion threshold in only 4% of instances for both informed and uninformed conditions. On the other hand, as these same comparisons suggest, making information about the location of the shock public had little effect on contagion. Evidently, the simple awareness of a shock, which indicates to traders an excess demand for liquidity, is far more important to contagion than knowing the source of the excess demand.

More importantly, Choi et al. also find that within the core-periphery networks the location of the shock prominently affects contagion. Specifically, bankruptcy rates are considerably higher when a core bank rather than a peripheral bank is hit with the shock. For example, given informed banks in the 15-bank core-periphery network, a shock to a core bank resulted in networks broaching the 60% contagion threshold in 65% of instances, compared to only 36% of instances following a shock to a peripheral bank. This result is particularly notable because the contagion-inducing effects of shocks to core vs. peripheral banks are unclear a priori. Although a shock to a core bank impacts a larger number of linked institutions than a shock to a peripheral bank, the effect of the shock to the core bank is more broadly distributed.

Results of this experiment provide evidence supporting the focus of Basel III, and the Dodd-Frank Act on systemically important "core" banks. Core-periphery networks are more susceptible to contagion than circular structures, particularly when a highly interconnected bank suffers a shock. In such networks, attending to the health of systemically important financial institutions can potentially prevent or mitigate the financial contagion.

3.3 | Liquidity regulation and the interbank market

Although results of the first two experiments reviewed in this section clearly support a regulatory focus on bank interconnectedness, not all of the policies enacted via the Basel III agreement fully attend to interconnectedness's consequences. One particularly controversial policy regards the recommendation by the Basel Committee on Banking Supervision to require banks to hold a sizable buffer of liquid assets.[23]

Financial economists have long been concerned that liquidity requirements may impede the functioning of the interbank market in times of stress. In the National Banking Era (1863-1913), for example, liquidity requirements were the primary prudential tool, and despite substantial required obligations, the era was characterized by multiple panics. Moreover, some experts from the era believed that liquidity requirements were at least partly responsible (Sprague, 1910). A

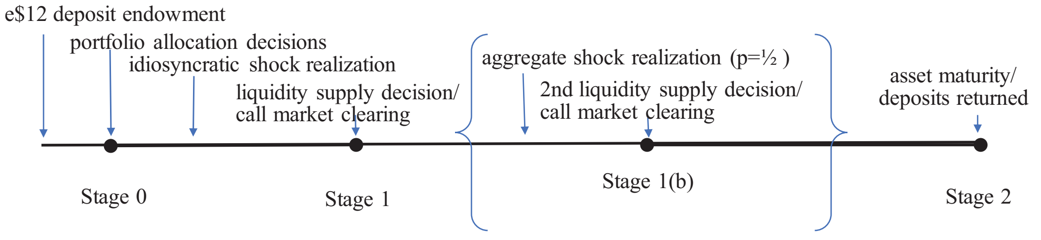

FIGURE 4 Sequence of moves in a liquidity regulation game [Colour figure can be viewed at wileyonlinelibrary.com]
(*Source*: Davis et al., 2019).

primary concern at the time was that the required reserves failed to provide banks with a cushion against stress, because use of the reserves would violate the requirement.[24] Apart from this historical concern, contemporary research suggests that liquidity requirements may hamper interbank market performance in times of stress, not because stressed banks may not use their reserves, but because of the incapacity of unstressed banks to use their reserves to address the liquidity needs of stressed banks via the interbank market.[25]

An experiment by Davis et al. (2020) provides behavioral insight into the way that liquidity requirements may impede the functioning of interbank markets. The experiment, based on a model by Gale and Yorulmazer (2013), investigates combinations of changes in the shock type and the presence or absence of a liquidity requirement. Sessions consisted of a series of three-stage periods illustrated in Figure 4. Prior to stage 0 each of eight banks is given a $12 endowment of deposits. Then in stage 0 each bank divides its deposits in $1 increments between investments in long term assets that mature to a value of $2 in stage 2, and cash that may be used to satisfy withdrawal demands arising from a non-stochastic, or “idiosyncratic” shock in stage 1 and/or to acquire assets in the interbank market. Before stage 1, four banks experience an early withdrawal demand of $8. In stage 1, after exhausting cash, shocked banks sell assets to the unshocked banks in a one-sided call market. If, after selling its assets, a shocked bank still faces a cash deficiency, that bank becomes insolvent and suffers a $4 bank resolution penalty.

The most interesting treatments examine behavior in a “compound shock” environment.[26] Following asset trades in response to the idiosyncratic shock, two of the four initially unshocked banks receive an additional $8 withdrawal demand with probability ½. As indicated by the bracketed entry labelled 1(b) in Figure 3, in the event of a second shock realization, a second call market takes place to allow newly cash deficient banks to sell assets.

In a symmetric equilibrium for the baseline treatment each bank invests 6 assets, no bank suffers bankruptcy, and each of the initially invested assets mature for a total of 48 assets, as shown in the top row of Table 5.[27] Comparison of the maximum sustainable investment level for the baseline compound shock treatment with that for the autarkic “no trade” prediction, shown in the bottom row illustrates the enhanced liquidity transformation that interbank trade allows. Without an interbank market, each bank, acting as a narrow “unit bank”, could invest only $4 in the long term asset to ensure enough cash on hand for a possible $8 shock. Thus, in the autarkic, solution only 32 assets mature each period, 16 assets less than the number that interbank trade facilitates.

In a second liquidity-regulated treatment, each bank is required to maintain a $4 cash reserve. Banks may use their required reserves to satisfy their own cash demands if shocked, but may not use those reserves to buy assets on the interbank market to satisfy the liquidity needs of

TABLE 5 Liquidity regulation experiment results

Treatment	Investments (Assets)			Bankruptcy rate (Banks)	
	Max. sustainable	Mature	Deviations from unregulated sustainable[a]	Overall[b]	2nd stage[b]
Baseline	48	42.79	5.21	1.49	0.86
Liquidity-regulated	40	35.85	12.15	1.79	1.38**
Autarkic	32				

[a]All mature investment outcomes deviate significantly from unregulated sustainable reference levels.
[b]Bankruptcy rates include only periods where a second stage shock occurred. **,*** rejection of the null that the liquidity requirement does not affect bankruptcy rates within a shock type, $p < .05$, $p < .01$, respectively. Source Davis et al. (2019).

other banks. It is the incapacity of banks to help other banks through interbank market loans that impacts sustainable mature asset volumes. Given the liquidity requirement, the aggregate equilibrium investment volume in the liquidity regulated treatment falls from 48 to 40 units, as shown in the second row of Table 5, a 50% reduction in the gains from trade that interbank trading allows.[28,29]

Experimental results, summarized in the middle columns of Table 5 clearly illustrate the high costs of liquidity requirements. The mean mature investment rate of 42.79 assets in the baseline treatment is 5.21 units below the maximum sustainable level of 48 assets. In the counterpart liquidity-regulated treatment the mean mature investment rate falls below the same reference level by 12.51 units, or by more than twice as many assets, and leaves mature investment levels only marginally higher than that sustainable in the autarkic condition. Moreover, liquidity requirements fail to provide the desired cushion against bankruptcies. To the contrary, mean bankruptcy rates in the baseline compound shock treatment, at 1.49 banks, are actually lower than the 1.79 banks in the liquidity-regulated treatment. As entries in the rightmost column of Table 5 indicate, the increased bankruptcy rate in the liquidity regulated treatment is driven entirely by increased bankruptcies in response to the second shock – precisely the periods where the liquidity regulation is intended to help the most. Analysis of individual decisions indicates that the failure of liquidity requirements to mitigate bankruptcies in this environment is driven by the unintuitive nature of the equilibrium response to the liquidity requirement. Given that banks are already required to hold cash in reserve, presumably to deal with a shock, it makes little obvious sense to hold still more cash to address a second stage shock, as the efficient equilibrium requires.

In sum, results of Davis et al. (2020) indicate that liquidity requirements do not appear to improve the capacity of the interbank market to respond to unanticipated aggregate shocks, and in fact may actually increase the financial system's fragility during a crisis. The stylized nature of the environment that Davis et al. evaluate limits the policy inferences that can be drawn from their results. In particular, Davis et al. are careful to observe that the model on which their experiment is based excludes a central bank. Nevertheless, as observed in a recent policy analysis by Yankov (2020) the liquidity requirements imposed under Basel III quite clearly have the effect of reducing the size of the interbank market, and given the reticence of banks to use central bank liquidity facilities observed above in Section 2.1, any financial crisis that creates short term liquidity deficiencies may have the consequence of inducing considerable financial stress, as the Davis et al. experiment suggests.[30]

4 | POLICIES AFFECTING INDIVIDUAL FINANCIAL INSTITUTIONS

A third set of pertinent experiments considers policies that restrict the activities of individual financial institutions whose actions were widely considered to increase the likelihood and magnitude of the financial crisis. A first subsection below reviews an experiment conducted to evaluate a proposed ban on uncovered credit default swaps, the misuse of which critics contend led to an unjustified confidence in the value of mortgage-backed securities. A second subsection reviews a pair of experiments conducted to behaviorally assess policies intended to reduce ratings inflation by credit ratings agencies. A third and final subsection considers triggering mechanisms of contingent capital bonds, a new tool for helping banks satisfy increased capital requirements imposed in Basel III.

4.1 | Credit default swaps and the proposed ban on uncovered positions

Credit default swaps are the most common and arguably the most important type of credit derivative. In its simplest form, a credit default swap consists of a seller offering a commitment to compensate a buyer in the event of a default. The buyer in turn, pays the seller a periodic stream of payments until bond maturity. Credit default swaps are essentially a type of insurance contract that transfers the risk of default from bond holders to institutions better prepared to deal with adverse events (e.g., big banks or insurance companies). Unlike a standard insurance policy, however, swap purchases are not tied to the insured item. One commentary likens this feature of credit default swaps to allowing a buyer to purchase insurance on your neighbor's car, and then compensating the buyer in the event your neighbor has an accident (Noeth & Sengupta, 2012). The severed link between the insurance contract and the insured item allows credit default swaps to also be used for speculative purposes. A trader with a sufficiently large volume of credit default swaps on a security, for example, is invested in the security's failure, and may short the institution or project that the security supports in order to achieve a default. Credit default swap purchases without associated bonds, are termed uncovered or "naked" positions.

Credit default swaps became enormously popular following their creation in 1990s. By 2007 the total estimated value of credit default swap issues exceeded $61 trillion, a total larger than 2007 world GDP.[31] Some of this contract volume was unquestionably held for speculative purposes as roughly 80% of credit default swaps were held in uncovered positions.[32] The use of these derivatives for speculative purposes is controversial. As a general rule, economists regard financial derivatives as instruments that usefully improve market efficiency by increasing the flow of information. Many argue that credit default swaps are no different from other derivatives (e.g., Sultz, 2010). Others assert that swaps can create systemic risk in the banking system, and that they were a destabilizing factor in the financial crisis (e.g., McIlroy, 2010).

Concerns about the destabilizing effects of credit default swaps prompted policy reforms articulated in the Basel III Accord and in the U.S. codified in the Dodd-Frank legislation (Carlson & Jacobson, 2014). These reforms were mainly aimed at improving transparency and reducing an excessive accumulation of risk by a single issuer.[33] Subsequent regulation barring uncovered swap positions was imposed by the European Union in 2012 for sovereign debt. Broader proscriptions of uncovered purchases are being considered.[34] Concerns regarding such a ban include the possibilities that the prohibition may interfere with the market liquidity provided by speculative swap purchases, and may undermine capacity of credit default swap prices to reflect the risk of the underlying assets (Jaing, 2020).

The use of credit default swaps for speculation and the effects of regulations cannot be evaluated absent empirical information, and as Sultz (2010) observes, “there is a dearth of serious empirical studies on the social benefits and costs of credit default swaps and other derivatives – not just in the last 2 years, but in the last several decades.” In an effort to bridge this informational gap Weber et al. (2020) report a novel experiment that extends a bond market design that these same authors previously developed (Weber et al., 2018), by appending to the bond market an associated credit default swap market.

The experiment allows an initial behavioral assessment of the effects of introducing credit default swaps on bond markets, as well as the effects of restricting swap purchases to only covered positions. A baseline control treatment assesses bond market performance absent the presence of credit default swaps. The sessions consisted of six-player markets where traders bought and sold bonds over a series of 11-period rounds. Prior to the start of each round participants are exogenously given an endowment of lab currency. Then in period 0 an initial public offering for a fixed quantity of “bonds” occurs as a one-sided uniform price auction. These bonds have an announced face value which the bondholder, barring a default, will receive in period 10, and also generate a fixed coupon payment each periods 1–10. Each period, bonds are subject to a probabilistically occurring default. The default probability is determined endogenously by the initial public offering price: low prices raise project financing costs, which in turn increase the default probability.[35] In the event of a default, coupon payments cease, and the bond’s face value falls to zero. Following the initial public offering, participants trade bonds in periods 1–9 via a two-sided uniform price call market. In the final period 10, provided the bond hasn’t defaulted, a final coupon payment is made and the bond face value is repaid.

Two additional treatments modify control treatment procedures by adding a credit default swap market. For sessions in these treatments, each participant is endowed with two credit default swaps at the beginning of each round, along with a compensating cash endowment reduction made to maintain the expected value of the initial endowment across treatments. In the event of a default, the swaps pay the bond’s face value. Then in each trading period after the initial public offering, participants are given the opportunity to trade their credit default swaps in a two-sided uniform price call market that follows the bond market exchange. The treatments are distinguished by whether or not the swap purchases must be covered. In an unregulated treatment, participants could buy and sell credit default swaps without restriction. In the alternative regulated treatment traders were allowed to acquire a swap only if they had at least one “unprotected” bond. Traders were not forced to subsequently maintain a bond for every credit default swap held. A trader, for example, might sell previously held bonds after acquiring a swap. Nevertheless, in the event of a default, credit default swap holders were paid the bond face value only if they held an associated bond.

With this design, Weber et al. explore two primary research questions. First, how do credit default swaps (regulated or otherwise) affect the bond market? Second, what effect does restricting the purchase of uncovered swaps exert on either the bond market or the credit default swap market? In particular, does the restriction impede the capacity of credit default swap prices to reflect the bond’s risk? To the first question, experimental results are largely negative. In both the regulated and unregulated treatments, while bond prices rose modestly relative to the control sessions in the initial public offering market, the differences are not statistically significant in either case. Similarly, in both the regulated and unregulated treatments, bond mispricing within periods was modest and not statistically different from the degree of within period mispricing in the control sessions.

To the second question, the evidence does not suggest that credit default swap regulations impede the capacity of swap prices to reflect the risk of bonds. The percentage of covered positions in the regulated treatment did increase significantly, from 65% in the unregulated treatment to 81.4%. The restriction, however, exerted little effect on either the bond or credit default swap markets. Neither bond nor swap prices differed significantly across treatments. All that said, credit default swap prices failed to reflect bond risk in both the regulated and unregulated treatments. In stark contrast to bonds, traders paid five to six times the fundamental value of the swaps in each treatment, a curious finding with parallels to natural contexts.[36]

In summary, the Weber et al. (2020) experiment provides an important initial assessment of the effects of introducing credit default swaps on bond market performance, as well as the effects of restricting uncovered swap positions on the underlying bond market. Results suggest that while a restriction on uncovered credit default swap sales reduced the extent of uncovered positions, allowing unrestricted swap sales had no adverse effect on the underlying bond or credit default swap markets. At the same time, there was no evidence that the restriction actually improved the performance of either market. Thus, from a policy perspective, these initial results provide some support for both the policy position of the European Union to ban some credit default swaps on sovereign debt, and that of the United States to impose no restrictions.

As the authors are careful to observe, much more investigation is necessary to fully understand credit default swap markets and their regulation. Weber et al., for example, exogenously endow traders with the swap instruments. Allowing traders to take on risk by creating swaps would almost certainly generate situations where, because of a high appetite for risk and/or a failure to fully understand the liabilities associated with credit default swaps, catastrophic bankruptcies would occur with some frequency. The primary contribution of this experiment is the foundation for the further laboratory analysis of this complex market that the authors lay with their remarkably clever design.

4.2 | Credit rating agencies and policies to promote accurate ratings

Credit rating agencies provide investors with information about the credit quality of new securities. Over time credit ratings have also come to assume an increasingly important position, with high ratings used to satisfy regulatory capital requirements for banks, insurance companies, pension funds, and broker-dealers, who include major investment banks and securities firms. In the years preceding the 2007–2009 financial crisis rating agencies produced overly optimistic evaluations of the complex securities based on subprime mortgages.[37] The favorable ratings both helped inflate the housing bubble and allowed financial firms to take on more risk, which precipitated the crisis when the assets were downgraded.

Many policymakers regard the conflict of interest in the compensation scheme used by the rating agencies, where the security issuer pays for the report, as a primary driver of the inflated ratings.[38] Although credit ratings agencies are concerned about the reputational consequences of issuing a misleading report, the threat of an unsatisfied issuer to switch to another agency looms large, as a switch puts at risk not only the revenues from a current report, but future solicitations from that issuer as well. The threat of lost business takes on particular weight in the structured products market where a small number of broker-dealers accounted for almost all of the products being rated.[39]

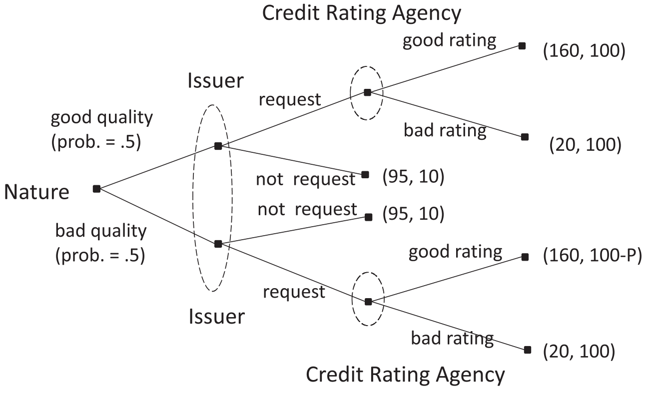

FIGURE 5 A Ratings game decision tree
(*Source*: Keser et al., 2017).

Keser et al. (2017) report an experiment that illustrates the behavioral consequences of an issuer-pays compensation scheme when the issuer has market power, and examines whether higher reputation costs improve ratings' quality. The experiment involves repeated interactions over a series of twenty periods between a bond issuer and a credit rating agency. As shown in Figure 5 each period starts with nature randomly making a quality draw, with good and bad quality being equally likely. Then the issuer, uninformed about the issue quality, chooses whether or not to request a report. If the issuer does not request a report the period ends, with the issuer earning \$95 and the credit rating agency earning \$10. If the issuer does request a report, the quality-informed credit rating agency must choose to give the asset a good or a bad rating. Issuer earnings are sensitive to the report type. In the event of a good rating, issuer earnings rise to \$160 while in the case of a bad rating issuer earnings fall to \$20. The credit rating agency on the other hand earns \$100 every time the agency provides a report, regardless of the report type or the issue's underlying quality. In the case the credit rating agency provides a good rating for a bad quality issue, however, the agency's inflated rating becomes known and generates a reputational cost, P that is deducted from its earnings. To examine the effects of a regulation that imposes higher reputation costs on the credit rating agency's propensity to make untruthful ratings Keser et al. considered two penalty levels, $P = \$10$ and $P = \$50$.

In the unique subgame perfect equilibrium for the stage-game, the issuer chooses to not request a rating. The issuer's strategy is supported by the belief that the credit rating agency will always report quality truthfully, since for any reputational cost $P > 0$, the credit rating agency will earn less making a false report than an honest one. Since finite repetition cannot create additional strategies in a game with a unique stage game equilibrium, a risk neutral (or risk averse) issuer will follow a "not request" strategy for any finitely repeated version of the game. Nevertheless, both the issuer and the credit rating agency could earn more were the agency able to convince the issuer that she would issue a good rating independent of the issue quality. In the extreme, if the credit rating agency unconditionally delivers good ratings, issuer earnings rise to \$160, while its expected earnings rise to \$100-.5P, which exceeds the credit rating agency's subgame perfect equilibrium earnings of \$10.

Contrary to equilibrium predictions, experiment results suggest that credit rating agencies frequently provided good reports for bad quality. Regardless of the penalty size, issuers solicited reports in more than 70% of periods, and in at least half of those instances the agencies subsequently inflated ratings, misreporting low quality as high quality. Importantly, reputational cost variations did not significantly affect behavior. Issuers' actions, however, varied significantly with ratings reports: the probability that an issuer sought a solicitation in the subsequent period fell following a bad report, and rose following a good report. Moreover, issuer solicitations affected ratings decisions. In the final session period, when the issuer's reactions no longer mattered the credit rating agencies overwhelmingly delivered honest reports.

Despite the simplicity of the decision-making environment, this experiment illustrates the effects of repeated interactions on the truthfulness in ratings when the security issuer has market power. The costs of foregone solicitations from an issuer unhappy with a credit rating agency's report powerfully affects the incidence of honest reports, while increasing reputation costs in the range explored in the experiment has little effect on reporting decisions. Results thus suggest that in the absence of any change to the issuer-pays compensation scheme used in the credit rating industry, regulatory efforts to reduce the importance of ratings reports may more powerfully affect behavior than increased reputational costs of inflated reports.[40]

One prominent proposal for reducing ratings inflation is the institution of an authority to randomly assign credit ratings agencies to rate new structured security issues (Rivlin & Sourishian, 2017).[41] A related experiment Rabanal and Rud (2018), suggests while conferring monopoly power on rating agency will certainly preclude ratings shopping, it may still fail to reduce ratings inflation. The experiment, which implements the model by Bolton et al. (2012), features a more fully articulated market structure than Keser et al. Sessions consist of "rounds" in which randomly reshuffled groups consisting of a single seller, two buyers, and either one or two rating agencies interact. At the outset of each round, the seller is endowed with an asset, which may be red or blue. Buyers place a high value on blue assets and a low value on red assets. While neither the seller nor the buyers know the asset's color, the agency is endowed with a 90% accurate estimate of the color realization. Each round proceeds in four steps. First, the agency chooses both a fee to charge the seller for a report and whether to report the asset color honestly, or to always report blue. Next the seller looks at the report(s) and decides whether or not to make a report purchase. Third, buyers, either uninformed about asset color, or informed by the report if one is purchased by the seller, submit asset bids. Finally, the asset color is revealed and earnings are determined. The agency pays a penalty if it is discovered to have misreported the asset color. Rabanal and Rud consider two treatments, one with a monopoly rating agency, and the other with two rating agencies. In the case of two agencies, the seller sees both reports and selects which one to purchase (if either).

Experiment results indicate that competition promotes truthfulness in reporting. Injecting competition among rating agencies reduces the incidence of ratings inflation from 26% with one rating agency to 17% with two rating agencies. The driving force for this result is that competition forces down the price of reports by 24 percentage points, which in turn increases the costs of being caught misreporting as a percentage of total earnings.

Taken together, results of the Keser et al. and Rabanal and Rud experiments provide useful evidence regarding the effectiveness of policies intended to mitigate ratings inflation. Given an issuer-pays compensation scheme and a limited number of issuers, ratings inflation is likely. The most immediately effective remedial policies appear to be those that reduce the importance of ratings, such as eliminating the use of credit ratings for capital requirements, as imposed by the Dodd Frank Act. At the same time, regulators should exercise some caution in trying to solve market power on the part of issuers by conferring monopoly power on the part of credit rating agencies, because the reduction in the price of ratings induced by competition can reduce the value to the rating agency of an inflated report.[42]

One dimension of the structure of the credit ratings industry that has not received adequate behavioral attention regards the consequences of returning to a structure where the security purchaser pays for reports. While the potential clearly exists for security buyers to free ride off of the information purchased by others, a user pays model eliminates any incentive for credit ratings agencies to inflate earnings, and users who pay to acquire ratings have an interest in not redistributing costly reports.

4.3 | Triggering mechanisms for contingent capital

In the aftermath of the 2008 financial crisis, one focus of regulatory reform centered on insuring that banks maintain equity cushions sufficient to remain solvent in the case of a crisis. Correspondingly, the Based III agreement in 2010 introduced a schedule for substantially increasing minimum capital adequacy ratios.[43] To reduce the cost of meeting capital requirements, some countries in 2013 began allowing banks to issue a new hybrid security, termed contingent convertibles or "CoCo's" that regulators would count towards regulatory requirements.[44] CoCo's are standard fixed-term bonds unless a triggering condition is breached, in which case the bonds convert to bank equity shares at a pre-specified rate. Essentially, CoCo bonds represent a sort of pre-packaged bailout that allows the bank to raise equity in order to satisfy capital requirements in times of financial distress, just when raising capital is most problematic.

At issue is the rule used to trigger the bond-to-stock conversion. All the existing CoCo issues use accounting measures such as capital adequacy ratios as a conversion trigger, which is problematic because these measures reflect the bank's past rather than its current performance, and for this reason may fail to trigger a timely conversion. As observed by Balla et al. (2019) in the context of the FDIC's 1991 Prompt Corrective Action resolution mechanism, the costs of inaction due to the use of backward looking measures such as capital adequacy ratios as a triggering mechanism can be extraordinarily high.[45] To address this problem, many commentators advocate replacing these accounting measures with a price-based trigger, such as the value of a bank's equity, which incorporates market participants' views about a bank's current condition.[46] An example of a price-based trigger would be a pre-determined equity price, below which a bond-to-equity conversion would automatically take place. Such a "fixed trigger" rule offers obvious advantages of administrative simplicity, and clarity as to when a conversion would occur.

Despite the appeal of fixed triggers, however, they may not work. The problem, analyzed by Sundaresan and Wang (2015), arises from the feedback between bank equity prices and the fixed trigger intervention rule. The following example illustrates. Suppose that a bond-to-equity conversion occurs if the price of a share of the bank's equity falls below $5, and that the conversion is "value increasing" in the sense that the conversion terms have the effect of raising share values by clearing debt off the bank's balance sheet by a larger margin than they lower share values from the conversion-induced equity dilution. For specificity, suppose that in the case of conversion the value of bank equity increases by $2 per share.

Consider the incentives of traders in this scenario when the fundamental value of equity is, say $4. Including the value of the conversion would raise the price of equity shares to $6. A trade at $6, however, would fail to trigger the conversion. For that reason, no equilibrium price emerges. This same problem arises for any fundamental value between the trigger value (here $5) and the trigger value less the value of the conversion (here $3).

A conversion may also be "value decreasing," as would happen if the conversion terms result in a dilution in equity value that outweighs the value-increasing effects of reducing the bank's debt load. In this case, Sundaresan and Wang (2015) show the problem of equilibrium nonexistence case goes away, but is replaced by range of fundamentals for which multiple equilibria exist. To see this, suppose again that a conversion is triggered if the price of equity falls below $5, but that in the case of a conversion, the value of equity falls by $2. If the underlying fundamental value is $6, an equilibrium would exist at $6 if traders believe that other traders will continue to trade at the fundamental value. In this case no conversion will occur and the equilibrium price will remain at $6. If, however, traders believe that other traders will incorporate the conversion into their valuations, then the equilibrium price will fall to $4. Reasoning similarly, two equilibria exist for every fundamental value realization between the trigger value ($5) and the sum of the trigger value and the value of the conversion (here $7).

One possible remedial alternative would be to substitute the mechanistic fixed trigger rule with a price-informed regulator. Allowing a price-informed regulator (e.g., a bank examiner) may have the advantage of allowing the examiner to incorporate any available information not included in the market price into the intervention decision. In the case of a value decreasing conversion, the use of a price-informed regulator eliminates the problem of multiple equilibria.[47] Unfortunately, however, as shown Birchler and Facchinetti (2007) and Bond et al. (2010) in the case of a value increasing conversion, the range of equilibrium non-existence not only persists, but doubles in range relative to the case of a fixed trigger rule. With a price-informed regulator, the range of equilibrium nonexistence extends both above and below the trigger cutoff, because the regulator could not tell whether prices slightly in excess of the conversion trigger reflected the asset's true fundamental value (in which case no conversion was needed) or incorporated the value of an anticipated conversion (in which case a conversion was called for.)

Davis et al. (2014) evaluate the behavioral significance of predicted equilibrium nonexistence and multiple equilibria on price-based triggering mechanisms. In the experiment groups of ten traders interact in a series of periods. In the case that a price-informed regulator makes conversion decisions, these traders also interact with three monitors. At the start of each period, each trader is endowed with two assets and a loan. For six of the traders, the asset value is determined as a common realization uniformly distributed on [$2, $8]. For the remaining four traders the asset value is set $0.60 below the common realization, to motivate trade. Participants trade assets under double auction rules, under the condition that the final value of assets would not be determined until after an intervention decision is made following the close of trade. In the fixed trigger sessions, an intervention occurs automatically if the median transaction price falls below $5. In the regulator sessions the three (accuracy incentivized) monitors are shown the median transaction price and asked to guess the asset's value and to intervene if it falls below $5. The decision of one monitor is selected at random for implementation. The experiment consisted of four treatment cells, which were combinations of value increasing and value decreasing conversion types, and fixed price or price-informed regulator triggering mechanisms.

Experiment results, summarized in Figure 6 indicate that theorists were appropriately concerned about price-based triggers. In both fixed price and price-informed regulator treatments, for both value increasing and value decreasing conversions, the ambiguous information conveyed by market prices resulted in frequent "conversion errors" or instances where an intervention occurred when intervention was not merited by underlying conditions, or where an intervention conversion did not occur when underlying conditions did merit action. Closer inspection of the intervention error rates in Figure 6, however, suggests that the relative performance of fixed price and price-informed regulator rules varies with the conversion type. Under a value decreasing conversion, the overall intervention error rate is somewhat higher in the fixed trigger treatment. Moreover, all errors under a fixed-trigger rule are "errors of commission," or socially unjustified interventions, which some commentators regard as more harmful than failures to act (e.g., McDonald, 2013). In contrast, under a "value increasing conversion" the incidence of intervention errors, particularly errors of commission are higher with a price informed regulator. A subsequent

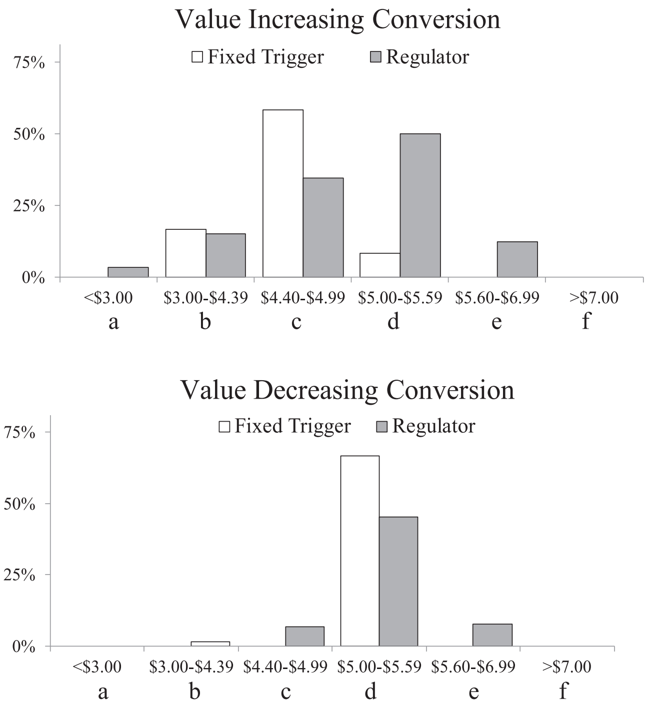

FIGURE 6 Intervention errors
Source: Davis et al. (2014).

experimental investigation by Davis and Prescott (2017), that examines a series of plausible environmental alterations that might affect a regulator's propensity to intervene, such as regulator penalties for intervention decision errors, and giving regulators the option to delay decisions (at some cost), confirm this basic finding: the choice of triggering mechanism depends critically on how the conversion affects incumbent equity holders. If CoCo bond conversions are expected to benefit incumbent equity holders (as is the case of most existing CoCo issues), a fixed-trigger rule would be more likely to result in fewer conversion errors. On the other hand, if the conversion importantly dilutes equity (as policy advocates propose), a price-informed regulator may yield better results.

As a matter of actual policy, the debate regarding the appropriate price-based triggering mechanisms has unfortunately become largely academic. Although CoCo issues exploded from their creation in 2013 to more than $350 billion in 2017, bond to equity conversion rates on the vast bulk of these issues were set at terms so unfavorable to bondholders that these CoCo bonds are little more than a complicated and poorly understood sort of bail-in debt, that banks sell to private investors at only a small premium over standard debt. As Glasserman and Perotti (2017) conclude, unless the nature of CoCo issues changes dramatically, they will fail in their purpose, to serve as a low cost way for banks to effectively satisfy capital requirements. In the event that the structure of CoCo issues are appropriately modified, however, the issue of an appropriate triggering mechanism, and the experiments reviewed here, will once again become relevant.

5 | CONCLUSION

The interrelated factors that contributed to 2007–2009 financial crisis highlighted critical weaknesses in the structure of the financial system. Addressing these weaknesses called for novel regulatory modifications of the financial system with which regulatory authorities had little experience. In just such a context the use of laboratory methods can provide important insights regarding both the efficacy and the associated risks of new, untested policies.

A critic might question the usefulness of policy experiments in a financial market context, by observing that we present no evidence that any of the experiments reviewed in this paper served directly as a basis for a new or revised policy. Such a criticism, which might equally be leveled against theoretical research on financial markets, reflects a basic misunderstanding of the role of either theory or experiments in policy analysis. Financial markets are enormously complex, and effective policies must balance an understanding of the pertinent institutional context with the implications of a new policy on the existing situation. Theory can trace through the often nonobvious consequences of a policy, given fully rational agents, and a specific set of assumptions about the pertinent environment. Experiments can assess the behavioral relevance of the resulting theoretical predictions on the domain of the theory in light of the limitations to rationality that characterize the human condition. The primary role of both policy-oriented theory and policy experiments, then, is advisory. Policymakers can (and often do) differ sharply on the conditions characterizing a pertinent environment. Even absent disagreements, the pertinent environment may not conform to any set of assumptions that lends itself to a tractable analysis. Nevertheless, theoretical and experimental research lends invaluable (if imperfectly applicable) insights about the pertinent interrelationships that in their absence would leave the policymaker very considerably less well informed.[48]

Critically, experimental investigations of financial market interactions yield useful and sometimes surprising additional insights. For example, experiments reviewed here usefully isolated the “stigma” associated with discount window access, and then demonstrated, among a menu of policy options which policy was most likely to address the issue. In other instances, such as reference price auction design and optimal information disclosure, experiments provided clear insight into the effects of policies when people are not the fully rational actors that theory presumes. Experiments also illustrate the importance of interconnectedness on financial market stability and provide evidence that supports the focus on systemically important financial institutions, as well as provide useful cautionary evidence regarding unintended effects of liquidity regulations imposed to improve the stability of financial institutions. Finally, experiments both illustrate the effects of misplaced incentives, and allow useful insight into the types of policies that might most usefully address the resulting problems. For example, experimental evidence confirms that the user-pays compensation scheme for credit reports used in the credit rating industry creates incentives for credit rating agencies to inflate earnings.

Financial markets evolve constantly, raising important new policy issues often in novel environments that have not previously existed. As the papers reviewed here suggest, in just such circumstances, laboratory methods can uniquely provide critical insights.

ACKNOWLEDGMENTS

This paper has benefited from useful comments by John Duffy, Charles Holt, Jeffrey Lacker, Andrew Kloosterman, Edward S. Prescott, Jean Paul Rabanal, and William Shobe, as well as from participants in a Virginia Commmonwealth University School of Business seminar and from an experimental economics reading group at the University of Virginia. All remaining errors remain our own. Financial Support from the National Science Foundation (SES 1949112) is gratefully acknowledged.

NOTES

[1] Good descriptions of the crisis's causes include Gorton and Metrick (2012), Allen and Carletti (2010) and Acharya and Richardson (2009).

[2] Gorton and Metrick (2012) identify 153 separate policy actions taken across 13 developed countries to mitigate the crisis. Forty-nine of these actions were in the United States alone. This count excludes actions taken subsequently to correct deficiencies in financial regulation revealed by the crisis, such as the Basel III accord issued by the Basel Committee on Banking Supervision and agreed upon in November 2010, and the Dodd Frank Act of 2010.

[3] For reviews of policy experiments, see Normann and Ricciuti (2009) and Roth (2016).

[4] We observe, however, that some of the papers we review extend the literature on financial fragility in directions related to the post 2008 reconsideration of regulatory policy. Pertinent examples include the model of discount window stigma reviewed in Section 2.1 and the analysis of interactions between the network structure of a banking system, and the system's resilance to financial shocks in Section 3.1.

[5] Some historical basis exists for this perception and it is not confined to the Fed's Discount Window. For example, a BBC leak that Northern Rock borrowed from the Bank of England was instrumental in the bank's demise. Because of concerns that stigma considerations were preventing banks from using its discount window facility, the Fed fundamentally changed its discount window policy in 2003. In particular, a primary credit program was created that allowed financially strong and well-capitalized banks to borrow from the discount window at a penalty rate above the Fed target rate, but with 'no questions asked', that is without establishing either a particular need for the funds or that the private alternative was not available. Thus after 2003 discount window access need not be motivated by pressing funding needs or signal financial weakness. Nevertheless, discount window borrowing remained sparse. See the appendix to Armentier and Holt (2020a) for further discussion.

[6] As an anonymous referee observes, simply identifying recipients of Term Auction Facility funds as auction "winners" may have also helped reduce stigma.

[7] An insolvent bank's portfolio yields a return below that needed to allow it to fulfill either its commitments to long term depositors or its immediate cash needs. On the other hand, an illiquid but solvent bank's portfolio yields a return sufficient to satisfy its commitments to long term depositors, but due to unexpectedly high demand for cash, such as a rash of withdrawals or the failure of a creditor to roll over a short-term loan, the bank finds itself with a temporary cash deficiency.

[8] The online appendix to Armantier et al. (2013) provides an historical review of the motivation for and design of the TARP auctions.

[9] This equilibrium is not unique. Since bidder 2 could bid any price above $10 without affecting her profit (of 0). Similarly, bids by bidder 1 do not affect the equilibrium as long as they are below those of bidder 2. The equilibrium presented, however, illustrates the highest possible purchase efficiency from a grand auction.

[10] In a recent related paper Armentier and Holt (2020b) evaluate a simple endogenous process for determining reference prices. Experimental results indicate that their process importantly mitigates value inaccuracies and further improves both seller profits and auction efficiency.

[11] Goldstein and Leitner (2018) motivates this bonus-induced jump in bank returns by suggesting that the bank has a project that yields a fixed return, but that requires a minimum level of cash on hand to complete. The model focuses on the capacity of a bank's short-term investments to provide liquidity sufficient for the bank to maintain its long-term investment portfolio, a focus that is eminently reasonable in the context of assessing the effects of disclosing stress test results.

[12] Actually, if banks were aware of their intrinsic value, for a bid of $9.33 they would sell only assets with a $6 fundamental realization, since only in this case does the $9.33 bid exceeds the expected value of a sale ($6+0.1*$15 = $7.50). Still, the predicted sales quantity is 0. Sales restricted to instances that turn out to involve only low intrinsic value assets would induce a lemons market effect, causing bids to fall to $6 thus eliminating sales.

[13] That is with probability 0.1 for a $6 intrinsic value, with probability .4 for a $9 intrinsic value and with probability 0.8 for a $13 basic value. Given that each intrinsic value realization is equally likely, the expected return given no sales is $(0.1) \times \frac{1}{3} + (0.4) \times \frac{1}{3} + (0.8) \times \frac{1}{3} = 0.43$.

[14] Given a certain sale for the $13 intrinsic value realization increases to 1, so the expected bonus incidence becomes $(0.1) \times \frac{1}{3} + (0.4) \times \frac{1}{3} + (1.0) \times \frac{1}{3} = .50$

[15] The bank's expected returns from holding assets with $9 and $13 intrinsic values are $9+0.4($15) = $15 and $13+0.8($15) = $25, respectively.

[16] That is, $(0.1) \times \frac{1}{3} + (1.0) \times \frac{1}{3} + (1.0) \times \frac{1}{3} = .70$.

[17] Importantly, the values used in the illustrative example (and in the experiment) are a simple case of a more general analysis by Goldstein and Leitner that allows for multiple intrinsic values and a variety of optimal poolings. The values used in the text are representative in the sense that they illustrate the basic intuition driving the analysis.

[18] The bonus incidences shown in Table 3, as well as the winning bid and bank sales decisions discussed in the text reflect a pooling of results in sessions where banks are and are not informed about their intrinsic value realizations. The presence or absence of intrinsic value information was a treatment condition, but consistent with the prediction that variations in bank information conditions does not affect predicted bonus incidences, Cox et al. found no significant treatment effects.

[19] In a policy analysis of stress test disclosure policy Goldstein and Sapra (2013) offers a similar conclusion (see in particular, pp. 43–45). These authors argue that in good economic times no disclosure is at all is the most preferred policy. Notably, however, Goldstein and Sapra observe a number of general negative consequences of full disclosure in addition to risk pooling, such as encouraging banks to hold suboptimal loan portfolios to pass the stress tests.

[20] With a liquidation value $r = 0.2$ the contagion free outcome is an equilibrium in the complete network setting but not in the incomplete network setting. With $r = 0.4$ the contagion free outcome is an equilibrium in both the complete and incomplete network settings. The full contagion equilibrium is also an equilibrium in both networks with both liquidation rates.

[21] Entries in Table 4 are estimates from the simplest specification of a mixed effects panel logit regression analysis of withdrawal decisions reported as table 8 in Duffy et al. (2019). The authors also report results of a number of additional specifications that include factors such as own withdrawal and partner withdrawal decisions. Although these added variables do significantly affect withdrawal rates, they do not importantly affect the estimates shown above.

[22] Choi et al. consider a balanced set of small, medium and large shocks in each session. Information about the shock magnitude each period was provided as public information.

[23] Formally, the Basel committee recommended two separate liquidity requirements. The first one is the Liquidity Coverage Ratio, which requires a bank to hold enough liquid assets to meet expected net cash outflows for a 30-day period. The second one is the Net Stable Funding Ratio, which is intended to ensure that banks adequately balance the sources and uses of funds over a longer term (1 year). In the United Sates, the Liquidity Coverage Ratio was implemented in 2014 and the Net Stable Funding Ratio in 2021. Both apply only to the largest banks.

[24] Experimental results by Davis et al. (2019) suggest that liquidity restrictions may actually improve the stability of individual banks when banks can access their own reserves in times of stress. As discussed below, the issue of concern here is whether interactions in the interbank market undermine the amelioratory effects of liquidity restrictions on single banks.

[25] Allowing banks to lend required reserves to other banks because those 'stressed' banks needed liquidity would effectively be the same as having no reserve requirement.

[26] The authors also consider a 'simple shock' treatment where banks are susceptible to only the idiosyncratic first stage shock. We omit discussion of this treatment here for purposes of brevity and because it is the compound shock environment that evaluates the reason liquidity requirements were implemented – that is, to help banks address cash needs given probabilistically occurring shock to the system.

[27] The intuition for the equilibrium follows from considering the cash needed in the system to deal with both the first and second stage shocks. A response to the idiosyncratic first stage shock requires \$32 in cash reserves, and a response to the event of a second stage shock requires an additional \$16, for a total of \$48. If each of the eight banks initially holds \$6 in cash, then the four banks realizing a shock in the first stage will each need an additional \$2, which the initially unshocked banks can supply in exchange for assets, while still leaving each initially unshocked banks with \$4 in cash per bank, for a total of \$16.

[28] Maximum sustainable investment falls by \$8 because the liquidity requirement dictates that the two banks not shocked in the second stage each retain their \$4 cash reserve. This can be accomplished by each bank initially holding \$7 in cash reserves. In response to the first stage shock each initially shocked bank will need only \$1 in cash, leaving each initially unshocked bank free to buy one asset while keeping \$6 in cash. In the event of a

second stage shock, the two shocked banks will need $2 each, which the unshocked banks can supply without violating their $4 reserve requirement.

[29] Given the liquidity requirement, a second 'exposure' equilibrium also exists, in which banks to simply ignore the possibility of the second stage shock and run the risk of bankruptcy. We do not develop this equilibrium in the text because it had very little explanatory power.

[30] Yankov (2020) also observes that an important effect of the elevated liquidity requirements under Basel III is to push riskier investments with higher returns to non-bank financial institutions. In the event of a crisis, these nonbank institutions will draw on their lines of credit at banks, creating short term liquidity deficiencies.

[31] In 2007 world GDP was $58 trillion (Aldasoro & Ehlers, 2018). Since 2007 credit default swap issues fell considerably. By 2018 the market value of issues had fallen to $9 trillion.

[32] U.S. Congress, House, Committee on Financial Services. The Effective Regulation of the Over-The-Counter Derivatives market: Hearing Before the Subcommittee on Capital Markets, Insurance, and Government Sponsored Enterprises of the Committee on Financial Services U.S, House of Representatives. 111th Cong., 1s sess., 2009, 35.

[33] In 2008 credit default swap issues were very concentrated; just 15 banks accounted for 90% of total issues. Also at that time swaps were customized contracts, sold over the counter, making cross-contract comparisons difficult. The regulations on swaps imposed in the Dodd Frank Act were intended to reduce the concentration of credit default swap issues and to increase contract transparency.

[34] Germany temporarily banned all "uncovered" credit default swap purchases 2010–2011, and the European Commission proposed a more general ban in 2010. "European Credit Default Regulation" in MarketReformWiki.com

[35] Weber et al. (2018) offers a structural model that develops the specific inverse relationship between initial public offering prices and bond default probabilities used in the experiments.

[36] As Weber et al. note, evidence of severe credit default swap mispricing has also been observed in the field (Singh and Andritzky, 2005).

[37] Importantly, ratings for more standard corporate security issues did not suffer from ratings inflation.

[38] The issuer-pays compensation structure was not the only factor alleged to drive ratings inflation. In particular, the association of ratings with regulatory capital requirements mentioned in the opening paragraph motivated investors to eagerly seek high ratings for issues.

[39] In 2006 the top 10 subprime issuers were responsible for 65% of market volume and the top 25 firms were responsible for 95% of volume (Ashcraft & Schuermann, 2008) in the more traditional corporate bond market (who's valuations did not collapse in 2008) is characterized by thousands of issuers (Frenkel, 2015).

[40] The Dodd Frank Act (2010) (Title IX, subsection C) mandates reforms that both increased the reputational cost of misreports to credit rating agency (by increasing their legal liability for mis-reports), and reduced the demand for inflated ratings by eliminating the use of credit ratings for capital requirements.

[41] The Dodd Frank Act mandates that the Securities Exchange Commission and the General Accounting Office consider alternatives to the issuer-pays compensation model. As yet no alternative has been implemented, but the Act mentions that in the event that the regulatory agencies fail to recommend an alternate business model an authority should be created to randomly assign agencies to rate structured security issues. Rivlin and Sourishian (2017) also endorse the option.

[42] Further research by Rabanal and Rud confirms the importance of price competition among credit rating agencies on mitigating the incidence of ratings inflation. In an environment that excludes price competition among rating agencies, Rud et al. (2018) find that the presence of multiple agents increases the incidence of ratings inflation. On the other hand, in a replication of the Rabanal and Rud framework discussed above, Angerer et al. (2021) show that 'accountability,' or the obligation for credit rating agencies to justify their actions anonymously in a chat box, reduces ratings inflation considerably below the levels observed by Rabanal and Rud. All that said, while conferring monopoly power on credit ratings agencies may have undesirable consequences, the dramatic incidence of inflated ratings by competing credit ratings agencies in the period prior to the crisis calls into question the potential effectiveness of competition as a policy remedy.

[43] The implementation date for the new requirements has been repeatedly deferred, and are currently set to take effect on January 1 2023. Under the new requirements, core Tier 1 capital to risk weighted assets will rise from 2% under Basel II to 5% under Basel III, while total Tier 1 capital to risk weighed assets will increase from 4% to 6%.

[44] Since their creation, roughly 200 contingent capital issues worth a total of $212 billion have been issued (Gledhill, 2020, Bloomberg Reports, August 12, 2020). Contingent capital bonds issuances are largely confined to European banks, primarily because regulatory authorities in the U.S. have been reluctant to count contingent capital bonds issuances toward capital requirements. Additionally, contingent capital bonds have gained little traction in the U.S. because the IRS treats interest payments on contingent capital bonds as taxable dividends rather than tax deductible interest payments.

[45] In 1991 the FDIC adopted a prompt corrective action (PCA) resolution mechanism intended to force regulators to act before a bank's losses grew too large. The PCA triggering condition was based on specific regulatory capital ratios. Despite the PCA mechanism, losses to deposit insurance funds were very high following the 2007-2009 crisis. Balla et al. (2019) reports that over the years 2007–2012, FDIC losses on failed banks and thrifts equaled 24.62% of these institutions' assets.

[46] Flannery (1998) provides evidence that market prices contain information about price not contained in supervisory reports.

[47] Unlike the case of a fixed price rule, where no unilateral action by any trader could move the market away from a price that would trigger a value decreasing conversion, given the same price, a regulator incentivized to make socially optimal decisions, could profitably deviate from a conversion outcome by choosing to not convert.

[48] Hommes (2021) offers a more extensive discussion of the policy relevance of behavioral and experimental economics in complex environments such as financial systems. See also Battiston et al. (2016), which provides a concise explanation of the usefulness of experiments in such a context. Notably, the issues regarding external relevance of experiments conducted to evaluate financial systems differ somewhat from the issues arises from the use of policy experiments as testbeds for new auction mechanisms (discussed, for example, in, Cason, 2010) or factors affecting the scalability of randomized control trials, as discussed recently by Al-Ubaydli et al. (2017).

REFERENCES

Acemoglu, D., Ozdaglar, A., & Tahbaz-Salehi, A. (2015). Systemic risk and stability in financial networks. *American Economic Review*, *105*(2), 564–608.

Acharya, V., & Richardson, M. (2009). Causes of the financial crisis. *Critical Review*, *21*(2&3), 195–210.

Aldasoro, I., & Ehlers, T. (2018). The credit default swap market: What a difference a decade makes. *BIS Quarterly Review*, June, 1–14.

Allen, F., & Carletti, E. (2010). An overview of the crisis: Causes, consequences, and solutions. *International Review of Finance*, *10*(1), 1–26.

Allen, F., & Gale, D. (2000). Financial contagion. *Journal of Political Economy*, *108*, 1–33.

Al-Ubaydli, O., List, J., & Suskind, D. (2017). What can we learn from experiments? Understanding the threats to the scalability of experimental results. *American Economic Review, Papers and Proceedings*, *107*(5), 282–86.

Angerer, M., Herrmann-Meng, M., & Szymezak, W. (2021). *Losing funds or losing face? Reputation and accountability in the credit rsating industry*. Manuscript. University of Liechtenstien.

Armantier, O., Ghysels, E., Sarkar, A., & Shrader, J. (2015). Discou*nt* window stigma during the 2007–2008 financial crisis. *Journal of Financial Economics*, *118*, 317–335.

Armantier, O., & Holt, C. A. (2020a). Overcoming discount window stigma: An experimental investigation. *Review of Financial Studies*, *33*, 5630–5659.

Armentier, O., & Holt, C. (2020b). *Endogenous reference price auctions for a diverse set of commodities: An experimental analysis*. Manuscript, University of Virginia.

Armantier, O., Holt, C., & Plott, C. (2013). A procurement auction for toxic assets with asymmetric information. *American Economic Journal: Microeconomics*, *5*, 142–162.

Ashcraft, A. B., & Schuermann, T. (2008). Understanding the securitization of subprime mortgage credit. *Foundations and Trends in Finance*, *2*(3), 191–309.

Balla, E., Mazur, L. C., Prescott, E. S., & Walter, J. R. (2019). A comparison of community bank failures and FDIC losses in the 1986–1992 and 2007–2013 banking crises. *Journal of Banking and Finance*, *106*, 1–15.

Battiston, S., Farmer, J. D., Flache, A., Garlaschelli, D., Haldane, A. G., Heesterbeek, H., Hommes, C., Jaeger, C., May, R., & Scheffer, M. (2016). Complexity theory and financial regulation. *Science*, *315*(6275), 818–819.

Battiston, S., Puliga, M., Kaushik, R., Tasca, P., & Caldarelli, G. (2012). DebtRank: Too central to fail? Financial networks, the fed and systemic risk. *Scientific Reports*, *2*, 541.

Bech, M. L., & Atalay, E. (2010). The topology of the federal funds market. *Physica A: Statistical Mechanics and its Applications*, *389*, 5223–5246.

Birchler, U., & Facchinetti, M. (2007). *"Self-destroying prophecies? The endogeneity pitfall in using market signals for prompt corrective action." Manuscript*, University of Zurich.

Bond, P., Goldstein, I., & Prescott, E. S. (2010). Market-based corrective actions. *The Review of Financial Studies*, *23*, 781–820.

Bolton, P., Freixas, X., & Shapiro, J. (2012). The credit ratings game. *Journal of Finance*, *67*(1), 85–111.

Boss, M., Elsinger, H., Summer, M., & Thurner, S. (2004). Network topology of the interbank market. *Quantitative Finance*, *4*, 677–684.

Caballero, R. J., & Simsek, A. (2013). Fire sales in a model of complexity. *The Journal of Finance*, *68*, 2549–2587.

Carlson, J., & Jacobson, M. (2014). New rules for credit default swap trading: Can we now follow the risk? *Economic Commentary*, Federal Reserve Bank of Cleveland, 2014–2011.

Cason, T. N. (2010). What can laboratory experiments teach us about emissions permit market design? *Agricultural and Resource Economics Review*, *39*, 151–161.

Choi, S., Gallo, E., & Wallace, B. (2017). *Financial contagion in networks: A market experiment*. Working Paper Series no. 124, Institute of Economic Research, Seoul National University.

Cox, C., Davis, D., Korenok, O., & Lightle, J. (2021). *Stress tests and information disclosure: An experimental analysis*. Manuscript, Virginia Commonwealth University.

Davis, D. D., Korenok, O., Lightle, J. P., & Prescott, E. S. (2020). Liquidity requirements and the interbank loan market: An experimental investigation. *Journal of Monetary Economics*, *115*, 113–126.

Davis, D. D., Korenok, O., & Lightle, J. P. (2019). Liquidity regulation and the banking environment: An experimental examination. *Journal of Economic Behavior and Organization*, (In Press). https://doi.org/10.1016/j.jebo.2019.09.005

Davis, D., Korenok, O., & Prescott, E. S. (2014). An experimental analysis of market-based contingent capital triggering mechanisms. *Journal of Money, Credit and Banking*, *46*(5), 999–1033.

Davis, D., & Prescott, E. (2017). Fixed prices and regulatory discretion as triggers for contingent capital conversion: An experimental examination. *International Journal of Central Banking*, *13*(2), 33–71.

Diamond, D. W., & Dybvig, P. H. (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy*, *91*, 401–419.

Dodd-Frank Wall Street Reform and Consumer Protection Act, 12 U.S.C. §5301 (2010).

Duffy, J., Karadimitropoulou, A., & Parravano, M. (2019). Financial contagion in the laboratory: Does network structure matter? *Journal of Money, Credit and Banking*, *51*, 1097–1136.

Dufwenberg, M. (2015). Banking on experiments? *Journal of Economic Studies*, *42*, 943–971.

Falk, A., & Heckman, J. J. (2009). Lab experiments are a major source of knowledge in the social sciences. *Science*, *326*(5952), 535–538.

Flannery, M. J. (1998). Using market information in prudential bank supervision: A review of the U.S. empirical evidence. *Journal of Money, Credit and Banking*, *30*, 273–305.

Frenkel, S. (2015). Repeated interaction and rating inflation: A model of double reputation. *American Economic Journal: Microeconomics*, *7*(1), 250–280.

Gale, D., & Yorulmazer, T. (2013). Liquidity hoarding. *Theoretical Economics*, *8*(2), 1555–1561.

Glasserman, P., & Perotti, E. (2017). The unconvertible CoCo bonds. Chapter 23 *Achieving financial stability: Challenges to prudential regulation*, World Scientific Publishing Co. Pte. Ltd. pp. 317–329.

Gledhill, A. (2020). Contingent convertibles. Bloomberg Reports, August 12. https://www.bloomberg.com/quicktake/contingent-convertible-bonds

Goldstein, I., & Leitner, Y. (2018). Stress tests and information disclosure. *Journal of Economic Theory*, *177*, 34–69.

Goldstein, I. (2013). Should banks' stress test results be disclosed? An Analysis of the costs and benefits. *Foundations and Trends in Finance*, *8*(1), 1–54.

Gorton, G. (2010). *Slapped by the invisible hand: The panic of 2007*. Oxford University Press.

Gorton, G., & Metrick, A. (2012). Getting up to speed on the financial crisis: A one-weekend reader's Guide. *Journal of Economic Literature*, *50*(1), 128–150.

Hirshleifer, J. (1971). The private and social value of information and the reward to inventive activity. *American Economic Review*, *61*, 561–574.

Hommes, C. (2021). Behavioral and experimental macroeconomics and policy analysis: A complex systems approach. *Journal of Economic Literature*, *50*(1), 149–219.

Jaing, W. In (2020). *Defense of naked credit default swaps. The FinReg Blog, Global Financial Markets Center*, Duke University School of Law.

Keser, C., Özgümüs, A., Peterlé, E., & Schmidt, M. (2017). An experimental investigation of rating-market regulation. *Journal of Economic Behavior and Organization*, *144*, 78–86.

Kiss, H., Rodriguez-Lara, I., & Rosa-Garcìa, A. (2021). *Experimental bank runs*. Manuscript, Universidad de Granada.

Mcdonald, R. L. (2013). Contingent capital with a dual price trigger. *Journal of Financial Stability*, *9*, 230–241.

Mcilroy, D. (2010). The regulatory issues raised by credit default swaps. *Journal of Banking Regulation*, *11*(4), 303–318.

Noeth, B., & Sengupta, . (2012). A look at credit default swaps and their impact on the European debt crisis. *The Regional Economist*, St. Louis Fed. April.

Normann, H. -. T., & Ricciuti, R. (2009). Laboratory experiments for economic policy making. *Journal of Economic Surveys*, *23*(3), 407–432.

Nuzzo, S., & Morone, A. (2017). Asset markets in the lab: A literature review. *Journal of Behavioral and Experimental Finance*, *13*(c), 42–50.

Palan, S. (2013). A review of bubbles and crashes in experimental asset markets. *Journal of Economic Surveys*, *27*(3), 570–588.

Rabanal, J. P., & Rud, O. A. (2018). Does competition affect truth telling? An experiment with rating agencies. *Review of Finance*, *22*(4), 1581–1604.

Rivlin, A., & Soroushian, J. (2017). Credit rating agency reform is incomplete. *Brookings Research Report*, March 6.

Roth, A. (2016). Experiments in market design, In J. H. Kagel & A. E. Roth (Eds.), *Handbook of experimental economics*, vol. 2, pp. 290–346. Princeton University Press.

Rud, O. A., Rabanal, J. P., & Horowitz, J. (2018). Does competition aggregate moral hazard? A multi-principal-agent experiment. *Journal of Financial Intermediation*, *33*, 115–121.

Singh, M., & Andritzky, J. (2005). Overpricing in emerging market credit default swap contracts: Some evidence from recent distress cases. IMF Working Paper No 05/125.

Sprague, O. (1910). *History of crises under the National Banking System*. U.S. Government Printing Office.

Smith, V. L., Suchanek, G. L., & Williams, A. W. (1988). Bubbles, crashes and endogenous expectations in experimental spot asset markets. *Econometrica*, *56*(5), 1119–1151.

Sultz, R. (2010). Credit default swaps and the credit crisis. *Journal of Economic Perspectives*, *24*(1), 73–92.

Sundaresan, S., & Wang, Z. (2015). On the design of contingent capital with a market trigger. *Journal of Finance*, *70*, 881–920.

Weber, M., Duffy, J., & Schram, A. (2020). *Credit default swap regulation in experimental bond markets*. Manuscript, UC Irvine.

Weber, M., Duffy, J., & Schram, A. (2018). An experimental study of bond market pricing. *Journal of Finance*, *73*(4), 157–1892.

Winters, B. (2012). Review of the Bank of England's framework for providing liquidity to the banking system. www.bankofengland.co.uk/publications/Documents/news/2012/cr2winters.pdf

Yankov, V. (2020). The liquidity coverage ratio and corporate liquidity management. FEDS Notes 2020-02-26. Board of Governors of the Federal Reserve System (U.S.).

How to cite this article: Davis, D. D., & Korenok, O. (2023). Policy experiments and financial regulation: Using laboratory methods to evaluate responses to the 2007–2009 financial crisis. *Journal of Economic Surveys*, *37*, 657–687. https://doi.org/10.1111/joes.12517