



Liquidity Crises

Understanding sources and limiting consequences: A theoretical framework

Robert E. Lucas, Jr. and Nancy L. Stokey*
University of Chicago

May 2011

ABSTRACT

Liquidity crises that induce or exacerbate deep recessions, as in 1930 or 2008, are situations in which individuals and firms want to build holdings of liquid assets. Heightened risk, or a perception of it, substantially increases demand for these assets. This reduces the supply available for normal transactions, leading to production and employment declines.

What happened in September 2008 was a kind of bank run. Creditors lost confidence in the ability of investment banks to redeem short-term loans, leading to a precipitous decline in lending in the repurchase agreements (repo) market. Massive lending by the Fed resolved the financial crisis, but not before reductions in business and household spending had led to the worst U.S. recession since the 1930s.

In this essay, we first sketch theoretical ideas that bear on the sources of liquidity crises: bank runs, sunspots and contagion effects, and the moral hazard problem created by deposit insurance. We then describe the repo market, and argue that these theoretical concepts are useful for understanding that market as well.

We conclude with several lessons for regulatory reform and for the role of Federal Reserve policy in coping with future liquidity crises:

- *Bank regulation can reduce the likelihood of liquidity crises, but cannot eliminate them entirely.*
- *During a liquidity crisis, the Fed should act as a lender of last resort.*
- *The Fed should announce its policy for liquidity crises, explaining how and under what circumstances it will come into play.*
- *Deposit insurance is part of the answer, but has a limited role.*
- *The Fed's lending in a crisis should be targeted toward preserving market liquidity, not particular institutions.*

Liquidity Crises

Understanding sources and limiting consequences: A theoretical framework

Robert E. Lucas, Jr. and Nancy L. Stokey*

May 2011

Introduction

It is hard to imagine better-motivated legislation than the Dodd-Frank Act, to date the one measure directed at preventing future financial crises. Yet it is hard to find an economist who argues that Dodd-Frank represents any appreciable progress toward this goal, nor is there anything like a consensus among its critics on what legislation should supplement or replace it. Economists cannot yet offer a complete, agreed-upon theoretical framework for thinking about liquidity crises, about the forces that precipitate them or exacerbate them or both.¹ Nevertheless, in our view, many of the main elements are in place. In this essay, we first describe these elements and then discuss how they might be combined to guide legislators and regulators.

Our interest here is in liquidity crises that induce or exacerbate deep recessions, as in 1930 or 2008. These crises are situations in which individuals and firms want to build up their holdings of liquid assets, cash and other securities that are close to cash in the sense that they can be exchanged for cash easily and at a predictable price. These assets have a special role because they are used, indeed required, for carrying out transactions. Heightened risk, or a perception of heightened risk, substantially increases the demand for these assets. This increase in demand has the effect of reducing the supply available to carry out the normal flow of transactions, leading to a reduction in production and employment.

What events are excluded by this definition? The stock market crash of 1929 and the dot-com crash of 2000 are two examples. These large, sudden changes in stock prices reflected changes in beliefs about future returns, but they did not have large, immediate effects on the inventories of cash or other liquid assets that individuals and firms wanted to hold, relative to the volume of their spending. Another example is the unexpected fall in house prices in 2007–08, which led to a reduction in construction activity. Housing construction is a large enough industry that this reduction would have shown up in a decline in overall gross domestic product (GDP), but it would have been comparable in size to other recessions of the postwar era. (Of course, mortgage-backed securities, marketed as liquid assets, did play a central role in the financial crisis, and that role will be discussed below.)

The events that followed the failure of Lehman Brothers in September of 2008 were not a modest recession. The spending declines in the fourth quarter of 2008 and the first quarter of 2009 sent U.S. GDP from 3 percent or 4 percent below trend to 8 percent or 9 percent below, where it has remained ever since. Housing was only a tangential factor in this decline.

We will argue here that what happened in September 2008 was a kind of bank run. Creditors of Lehman Brothers and other investment banks lost confidence in the ability of these banks to redeem short-term loans. One aspect of this loss of confidence was a precipitous decline in lending in the market for repurchase agreements, the repo market. Massive lending by the Fed resolved the financial crisis by the end of the year, but not before reductions in business and household spending had led to the worst U.S. recession since the 1930s.

In this essay, we first sketch several theoretical ideas that bear on the sources of liquidity crises: bank runs, sunspots and contagion effects, and the moral hazard problem created by deposit insurance. We then describe the repo market and argue that these concepts are useful for

understanding that market as well. We then draw some conclusions for regulatory reform and for the role of Federal Reserve policy in coping with future liquidity crises.

Bank runs

A simple and widely used theoretical model of bank runs was developed by Douglas Diamond and Philip Dybvig (1983). It describes an economy in terms of the production and consumption of a real good, but to apply their model to actual banking practice, it is helpful to give it a monetary interpretation. In this section, we will sketch their framework, so modified.

Think of an economy as a collection of agents—individuals and firms—who are paid in cash and in turn must pay cash for the goods they buy. That is, all creditors demand payment in cash, so all transactions require the use of cash. Assume that bills for purchases arrive with unpredictable lags and that bills must be paid exactly when due, with draconian penalties for late payment. In a currency-only version of this economy, every agent would need to hold enough currency to cover the worst payment possibility. Even if this worst-case scenario rarely occurred, each individual agent would need to be prepared, and a sizable fraction of each agent's total wealth would be tied up in non-interest-bearing currency.

In this economy, there is an obvious role for deposit-taking banks, institutions that accept deposits and in return promise to pay interest on those deposits and to redeem deposits for cash at any time. If a bank has a large number of depositors and if the (random) demands for cash of those depositors are less than perfectly correlated, then the worst payment case for the bank as a whole is less than the sum of the individual worst cases for its depositors.

In fact, if the number of depositors is large and if the random cash demands of each depositor have a substantial idiosyncratic (uncorrelated) component, then the bank needs to hold in its cash reserve only something like the average payment for each depositor. If one depositor's

demand for payments is unusually large, the bank can apply the unused portion of someone else's deposit to honor the check, in effect borrowing from one depositor to finance a loan to another.

By so economizing on its cash reserves, the bank can safely invest a fraction of deposits in interest-bearing securities. Suppose for now that these investments are Treasury bills or high-grade commercial paper, so there is little risk and the bank can quickly convert the assets into cash if its own stock runs low. The return on these investments finances the interest paid to depositors, and in a competitive banking system, all of the return (net of the bank's operating cost) is used in this way. We can think of such a bank as an institution that pools payment risks, making all of its clients better off than they would be acting on their own.

The bank has promised all of its depositors cash on demand, and it can fulfill this promise if depositors make withdrawals only when they have actual payment obligations. Moreover, if depositors are confident that the bank can fulfill its promise, they have no incentive to make larger withdrawals than necessary: Funds in the bank earn interest while cash does not. Thus, in ordinary times, the bank can and does make good on its promise of cash on demand.

This arrangement, which of course is just fractional reserve banking, has a problem, however. The ability to make good on its promise is fragile. If all its depositors—or a sufficiently large number—simultaneously try to exercise the option to withdraw, the bank in fact does not have nearly enough cash in reserve to cover those withdrawals. In this case, the bank runs out of cash and cannot honor its promise to those who arrive late. If enough other depositors choose to withdraw or if such an event is merely anticipated, then any individual depositor has an incentive to withdraw his own funds as well: Each wants to be at the beginning of the line, not at the end. Diamond and Dybvig's model captures the essential nature of this second outcome, a bank run.

Notice that the run outcome is just as possible as the outcome in which the bank continues to operate normally. Whether or not a run occurs has nothing to do with mismanagement of the bank or excessive risk in the bank's portfolio. There are simply two possible outcomes, one good and one bad, and which of them is realized depends on what everyone thinks that others will do.

A crucial feature of this example is that withdrawing depositors must get in line and be served in turn. The bank gives each depositor as much of his deposit as he asks for until the money runs out, and those still in line are simply turned away. This assumption of sequential service has been criticized, and indeed the bank has other possibilities. A bank could prorate withdrawals to distribute default over more depositors, it could temporarily cease honoring demands for withdrawals to allow noncash assets to be liquidated and so on. In some financial markets, such practices are standard. For example, hedge funds typically require investors to commit their funds for a long period and require substantial advance notice for withdrawals.

In a discussion of deposit-taking banks, these criticisms simply ignore the fact that the one function unique to these banks is providing liquidity: facilitating cash payment at low cost. A bank that cannot carry out cash transfers for its depositors has reneged on an important promise, and those depositors will take their business elsewhere. The sequential service assumption in Diamond and Dybvig's theory highlights the fact that this essential function of a bank is jeopardized during a run. It seems to us an essential feature of any theory of banking.

The main conclusion of Diamond and Dybvig's theory is that in an economy where cash is required for transactions and banks help agents avoid holding excess cash by pooling their risk, there are always two possible outcomes. The run outcome can be avoided if banks are required to hold 100 percent reserves, but this requirement also makes banks superfluous.

Sunspots, contagion

It is worth stressing that the kind of multiple equilibria that Diamond and Dybvig used to account for observed bank runs can appear in a wide range of economic situations that have no particular connection to banking or the monetary system. This conclusion was established in another remarkable 1983 paper, by David Cass and Karl Shell. They showed that accepting the principle that people act rationally—in their own interest—is not, with any generality, sufficient to determine a unique economic outcome. Fractional reserve banking is but one of many examples where if people somehow come to expect a particular outcome, then that outcome will occur, but if they agree on another, the other will occur. Cass and Shell used the term *sunspot equilibrium* to emphasize that coordination of beliefs need not make any objective sense: If enough people think the occurrence of sunspots signals a run on a particular bank, it will do so. And if so, who are we to say the sunspots are unrelated to the safety of banks?

At the same time, it is hard not to see patterns in the occurrence of bank runs and currency crises, just as patterns appear in the fashionability of nightspots and in other examples where what you want to do depends on what you think others will do. Our common-sense view that the probability of banking crises can be affected by reserve or capital requirements, by regulation of bank assets or by the general state of the economy is based on real historical evidence. The Diamond-Dybvig model does not help us use this history to design better banking policies. Cass and Shell's work makes it clear that questions about the origin and influence of beliefs have to be faced.

In a series of papers dating from 2000, Stephen Morris and Hyun Song Shin (2001, 2003) offered a possible resolution of this problem. Like Cass and Shell's, their approach is abstract, focused on generic situations where people's beliefs about the beliefs of others have a central

role. Nor do Morris and Shin offer any specifics about the objective importance of the sources of their beliefs, so sunspot-based beliefs are not ruled out. The new element in their model is diversity of beliefs.

To stay with the bank run example, suppose that depositors agree on some fundamental measure of the financial health of their bank—the quality of its assets, say—but none of them has exact information about it. All a single depositor sees is a signal, information that is imperfectly correlated with asset quality. Moreover, all depositors get different information, different signals. Those who receive favorable signals—signals that the bank is solid—are content to keep their deposits in the bank. But those who receive signals that the bank is shaky will want to withdraw their funds. There is a cutoff signal value, dividing those who withdraw from those who don't. And since signals are correlated with asset quality, more people are getting unfavorable signals when asset quality is in fact low. Hence, there is a well-defined tipping point built into the bank's situation.

Asset quality in our example sounds like something solid and objective, something fundamental. The theoretical argument does not require that, however: It allows anything depositors *think* is important. But either way, Morris and Shin's model provides a framework for interpreting historical evidence on the situations that have been correlated with bank failures in the past.

One clear feature of these histories is that bank runs—and financial crises more generally—come in bunches, as though they were contagious. One source of contagion, often emphasized in the crisis of 2008, is that banks lend to each other: If one fails, its creditors can be directly injured. But bank runs can, and often have, spread where such direct connections are

minor factors. It seems to be enough that depositors in all banks think there is a useful signal, in the sense of Morris and Shin, in the distress of other banks.

Whatever their sources, these contagion effects are exactly what is systemic about bank failures. Any one bank, no matter how large and respected, can go out of business almost without a ripple. Anyone living in an American city can list the downtown banks he grew up with that vanished in the merger movement of the 1990s. Who misses them? Indeed, who misses Lehman Brothers, for generations one of the most respected financial institutions in the world? Its valuable assets, both physical and human capital, were quickly absorbed by surviving banks without notable loss of services. It was the signal effect of the Lehman failure, whether a signal about the situations of private banks or about the Federal Reserve's willingness to lend to troubled banks, that triggered the rush to liquidity and safety that followed.

Deposit insurance

The Diamond-Dybvig model and its successors also have implications for government policy directed at bank runs and panics. In all of these models, a system of deposit insurance completely eliminates the incentive to run. By insuring depositors that in the event of a run their deposits will be promptly restored and available to them, it eliminates the possibility of the run outcome altogether. Is this a practical possibility? That is an open question. But the institution of deposit insurance in the United States in 1933 was followed by 75 years without a serious bank run, and that fact must surely be taken as encouraging.

Deposit insurance brings its own problem, however. The models we have discussed thus far are theories about the behavior of bank depositors, about what they believe and what they do. The banks in these models are automatons. This description was adequate for making Diamond and Dybvig's central point, that any fractional reserve bank is vulnerable to runs even if it is

conservatively managed. Indeed, it strengthens their point to show that a run is possible even if the bank invests only in very safe assets. But other issues involve the fact the bankers make choices, too.

In the world we have described thus far, banks invested in Treasury bills and high-grade commercial paper, short-run assets with little or no risk of default. In this world, a bank experiences a run because it is short on vault cash: It is illiquid, not insolvent. But the portfolio of securities a bank holds is a matter of choice, with the usual trade-off between risk and return. Deposit insurance alters the bank's incentives when it makes that choice, introducing the possibility of insolvency.

John Kareken and Neil Wallace analyzed the incentive effects of deposit insurance in a 1978 paper that has not lost its relevance. Deposit insurance commits the government to pay depositors in the event of asset gambles that turn out badly. An insured bank that takes on risky investments can earn a higher return, and this additional return can be passed on to depositors and shareholders without passing on the added risk. The bank need not fear losing customers by holding a risky portfolio. Indeed, it can gain customers, by offering higher interest than its more cautious rivals. In effect, deposit insurance is a contingent cash transfer from the public to the creditors, depositors and owners of banks, encouraging banks to hold riskier asset portfolios.

Parts of the Dodd-Frank Act are motivated by the desire to protect depositors from unscrupulous or foolish bankers. This is surely a legitimate concern, but it is unrelated to the point of the Kareken and Wallace analysis. Their point, and it is fundamental, is that public funds are committed to banks and their depositors together, altering their joint willingness to take on risk. How they divide the surplus is a secondary matter for this point. Regulating the portfolios of insured banks is the only effective way to deal with this problem.

But regulations designed to prevent depositors from choosing banks with risky portfolios take away options that some of them prefer, without offering any new ones. Some depositors will seek alternative routes to restore these options, and financial institutions will have much to gain from providing them in new guises. This is not just a theoretical possibility. Beginning in the 1970s, as market interest rates rose in response to high inflation rates, depositors began to move their funds out of regulated, insured commercial banks that paid little or no interest and into money market mutual funds and other liquid assets. Even after inflation subsided, depositors were motivated by the higher returns these alternative forms of liquidity offered. In 1965, demand deposits at commercial banks were 18 percent of GDP. By 2005, this ratio had fallen to 5 percent of GDP—about the same as hand-to-hand currency.

The repo market

As deposits moved out of commercial banks, investment banks and money market funds increasingly provided close substitutes for the services commercial banks provide. Like the banks they replaced, they accepted cash in return for promises to repay with interest, leaving the option of when and how much to withdraw up to the lender. The exact form of the contracts involved came in enormous variety. In order to support these activities, financial institutions created new securities and new arrangements for trading them, arrangements that enabled them collectively to clear ever larger trading volumes with smaller and smaller holdings of actual cash. In August of 2008, the entire banking system held about \$50 billion in actual cash reserves while clearing trades of \$2,996 trillion per day.² Yet every one of these trades involved an uncontingent promise to pay someone hard cash whenever he asked for it. If ever a system was “runnable,” this was it. Where did the run occur?

There were no runs on commercial banks during the financial crisis of 2008. Deposit insurance through the Federal Deposit Insurance Corporation (FDIC) was effective in eliminating the incentive for depositors to withdraw funds. Indeed, as we will see below, demand deposits at commercial banks increased significantly during the crisis. There were two runs on investment banks, however. The run on Bear Stearns in March ended with its purchase by JPMorgan Chase, and the run on Lehman Brothers in September ended with its bankruptcy. In addition, there was an incipient run on money market mutual funds following the collapse of Lehman, halted only when the Treasury stepped in to provide deposit insurance for those institutions.

Of course, for the reasons discussed earlier, these events also heightened the fear of contagion for all financial institutions, altering their willingness to engage in various transactions. Gary Gorton refers to these events, aptly, as a “run on repo.” How did it work?

In economic terms a repurchase agreement (repo) is a securitized loan.³ The lender brings cash to the transaction, while the borrower supplies a T-bill or some other security to be used as collateral. The loans are short term, often one day.

Large lenders in the repo market include money market mutual funds and hedge funds. The repo market performs for these large institutions the same function that commercial banks perform for smaller depositors. In effect, it allows them to pool their cash, collectively economizing on their stocks of non-interest-bearing assets. For lenders, the repo market is attractive because the loans are very short term, so it is a way to earn a return—albeit modest—on cash reserves that would otherwise be idle. In normal times, any lender can withdraw cash by declining to roll over earlier loans. Firms that do not want liquidity do not lend in the repo market, since higher returns are available elsewhere.

What does it mean to have a run in the repo market? Consider a shock that heightens uncertainty about the soundness of financial institutions. Potential lenders will choose to hold more of their cash in reserve, anticipating possible withdrawals by their own clients. As a result, potential borrowers will find it difficult to obtain funds. Actual defaults are rare in this market, but borrowers who hoped to roll over old agreements may have to sell securities on short notice, perhaps at fire sale prices, to obtain cash elsewhere.

The role of collateral in the repo market is similar to the role of deposit insurance at commercial banks. In the cost-benefit calculus that Morris and Shin imagine depositors using when they decide whether to make a precautionary withdrawal, good collateral increases the incentive to continue rolling over short-term loans and hence reduces the likelihood of a successful run.

But while collateral reduces the likelihood of a successful run, it does not eliminate it altogether. Like other forms of fractional reserve banking, the repo market is in effect an institution for pooling cash reserves. Participants can choose to withdraw their cash from the collective pool, and in some circumstances, many will simultaneously choose to exercise this option.

To get a sense of the importance of the repo market, we can look at its size relative to other aggregates. At the end of 2007, \$774 billion was held as currency outside banks, \$511 billion in private, domestic demand deposits and \$3.033 trillion in money market mutual funds. As shown in the table below, all of these figures increased over the following year. In 2008, unlike 1930, demand deposits were a safe asset. Money market mutual funds, which are much larger than demand deposits, increased almost as much.

The Repo Market and other Monetary Aggregates January 2008 to January 2009			
	Jan. 2008 (billions)	Jan. 2009 (billions)	Change
Cash Held Outside of Banks*	\$773.9	\$832.2	7.5%
Private, Domestic Demand Deposits*	\$510.7	\$658.0	+28.8%
Money Market Mutual Funds*	\$3,033.1	\$3,757.3	+23.9%
Repos held by Primary Dealers**			
Total	\$3,699.4	\$2,585.9	-30.1%
Overnight & Continuing	\$2,543.6	\$2,005.6	-21.2%

* End of previous year. Flow of Funds Accounts, Board of Governors of the Federal Reserve System.

** First week of January. Federal Reserve Bank of New York.

The repo market behaved quite differently. At the beginning of 2008, primary dealers held total funds of \$3.70 trillion in the repo market, of which \$2.54 trillion was in overnight or continuing agreements. Those figures grew slightly during the first half of 2008. Total funds then fell to \$2.59 trillion at the beginning of 2009, a 30 percent decline, while overnight and continuing agreements fell to \$2.01 trillion, a 21 percent decline. Both figures showed further declines over the subsequent year as well.

Lessons from the panic of 2008

We began by asking what theory and evidence tell us about liquidity crises and about policies to avoid them or to mitigate their severity. The arguments above do not provide a complete answer, but they do point to some broad principles.

(a) Bank regulation can reduce the likelihood of liquidity crises, but it cannot eliminate them entirely.

Banks will fail, and these failures will make failure more likely for others. There is language in Dodd-Frank suggesting that the Fed should take responsibility for predicting and

precluding crises, but this task seems to us to be an impossible one, at least for the foreseeable future.⁴

(b) *During a liquidity crisis, the Fed should act as a lender of last resort.*

In the event of a bank run or a run on the repo market, the Fed can always add liquidity to the system, and there will be occasions—as in 1930 and in the fall of 2008—when it would be irresponsible not to do so.

(c) *The Fed should announce its policy for liquidity crises, explaining how and under what circumstances it will come into play.*

The events of 2008 illustrate the importance of an announced and well-understood policy. Over the years prior to 2008, investors came to understand that the Fed was operating under an implicit too-big-to-fail policy, in the sense that the depositors/creditors of large banks would be protected. No other policy was ever discussed, and the Fed's assistance in engineering the orderly exit of Bear Stearns in March 2008 was surely interpreted as evidence that this policy was still in place. The abrupt end of Lehman in September was then all the more shocking.

There is no gain from allowing uncertainty about how the Fed will behave. The beliefs of depositors/lenders are critical in determining the contagion effects of runs that do occur. By announcing a credible policy, the Fed can affect those beliefs, and the Fed needs to use this tool.

(d) *Deposit insurance is part of the answer.*

When introduced in the Banking Act of 1933, deposit insurance was limited to small deposits, and its role was viewed as consumer protection, not run prevention. Deposit insurance

performed this function well during the 2008 crisis: There were no runs on FDIC-insured commercial banks, although many failed or were absorbed by stronger institutions.

Deposit insurance should be retained, although for the reasons described by Kareken and Wallace, the assets held against insured deposits should be carefully regulated.

(e) Deposit insurance has a limited role.

Investment banks, money market funds and the repo market are outside the protection of the insured system, and the liquidity crisis of 2008 involved these other institutions. Could they be brought into the fold, with the relevant portion of their investment portfolio regulated in the same way that commercial banks are?

Higher returns in the uninsured sector will always be attractive for large depositors, and new institutions or arrangements would surely arise, offering liquidity provision on the old, risky terms. Clients will want it, markets will have a strong incentive to provide it and regulators will probably not be able to contain their efforts. Providers will be able to innovate around regulations or move offshore to avoid them. This dilemma leads us to our next point.

(f) The Fed's lending in a crisis should be targeted toward preserving market liquidity, not particular institutions.

There are two goals here: to have a credible policy for how liquidity will be injected in a crisis and to provide proper incentives for banks during ordinary times. Both goals are met by the Bagehot rule: In a crisis, the central bank should lend on good collateral at a penalty rate. To implement this rule, we need to know how much the Fed should lend and what assets will be regarded as good collateral.

Time consistency requires that no upper bound be placed on crisis lending. The guidelines we have for monetary policy, whether stated in terms of monetary aggregates or

interest rates, are directed at long-term objectives and are no help in a liquidity crisis. After the Lehman failure in the fall of 2008, the Fed expanded bank reserves from \$40 billion to \$800 billion in three months, surely exceeding by far any limit that would have been imposed in August. Even with this decisive response, spending declined sharply over next two quarters.

Because crises occur too rarely for the ex ante formulation of useful quantitative rules, the Fed should have considerable discretion in times of crisis. Nevertheless, because policies should be predictable, the Fed should describe the indicators it will use to decide when lending has reached a sufficient level.

Defining good collateral is more complicated. The quality of collateral is in the eye of the lender, and it can change dramatically from week to week. In this application, though, the lender is the Fed, and it is the responsibility of the Fed to define what it will treat as good collateral. To this end, the Fed should announce an ordering of assets by their quality. The list should be long enough to cover all contingencies, and it would need to be revised from time to time.

In such a regime, banks outside the FDIC would be free to choose their portfolios, with clients, bondholders and equity holders bearing the risk that those choices entail. The lower return on lower-risk assets would be offset, at least in part, by their superior status as collateral in the event of a crisis.

Avoiding liquidity crises altogether is probably more than we can hope for. What we can do is put in place mechanisms to make such crises infrequent and to make their effects manageable.

Notes

* The authors thank Douglas Diamond, Gary Gorton, Anil Kashyap, Allan Meltzer, Edward C. Prescott, Harald Uhlig, Warren Weber and Motohiro Yogo for helpful comments. The views expressed herein are those of the authors and not necessarily of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

¹ There is a long tradition of careful historical study of financial crises. Friedman and Schwartz (1963) and Kindleberger (1978) are canonical. Recent books by Reinhart and Rogoff (2009) and Gorton (2010) enrich this literature and bring it to bear on the crisis of 2008.

² Fedwire Funds Service, Board of Governors of the Federal Reserve System.

³ See Copeland, Martin and Walker (2010), Duffie (2010, 2011) and Gorton (2010) for detailed descriptions of this market.

⁴ See Meltzer (2009) for a further discussion of this point.

References

- Cass, David, and Karl Shell. 1983. Do sunspots matter? *Journal of Political Economy* 91 (2): 193–227.
- Copeland, Adam, Antoine Martin and Michael Walker. 2010. The tri-party repo market before the 2010 reforms. Staff Report 477, Federal Reserve Bank of New York, November.
- Diamond, Douglas W., and Philip H. Dybvig. 1983. Bank runs, deposit insurance, and liquidity. *Journal of Political Economy* 91 (3): 401–19.
- Duffie, Darrell. 2010. The failure mechanics of dealer banks. *Journal of Economic Perspectives* 24: 51–72.
- Duffie, Darrell. 2011. *How Big Banks Fail and What to Do about It*. Princeton University Press.
- Friedman, Milton, and Anna J. Schwartz. 1963. *A Monetary History of the United States, 1867–1960*. Princeton University Press.
- Gorton, Gary. 2010. *Slapped by the Invisible Hand: The Panic of 2007*. Oxford University Press.
- Kareken, John H., and Neil Wallace. 1978. Deposit insurance and bank regulation: A partial-equilibrium exposition. *Journal of Business* 51 (3): 413–38.
- Kindleberger, Charles P. 1978. *Manias, Panics, and Crashes: A History of Financial Crises*. Wiley.
- Meltzer, Allan H. 2009. Regulatory reform and the Federal Reserve. Testimony, July 23, Senate Banking Committee.
- Morris, Stephen, and Hyun Song Shin. 2001. Rethinking multiple equilibria in macroeconomic modelling *NBER Macroeconomics Annual 2000*. MIT Press, 139–61.

Morris, Stephen, and Hyun Song Shin. 2003. Global games: Theory and applications, in *Advances in Economics and Econometrics*. M. Dewatripont, L. Hansen and S. Turnovsky, eds. Cambridge University Press.

Reinhart, Carmen M., and Kenneth S. Rogoff. 2009. *This Time is Different*. Princeton University Press.