For release on delivery
1:15 p.m. EST
January 3, 2014

Banks as Patient Debt Investors

Remarks by

Jeremy C. Stein

Member

Board of Governors of the Federal Reserve System

at the

American Economic Association/American Finance Association Joint Luncheon

Philadelphia, Pennsylvania

January 3, 2014

I'm delighted to be speaking to this group; it's really an honor. Given the audience--and because I may never again have so many economists in front of me at one time--I thought that rather than giving a policy-oriented speech, I would test-drive a new research project. The project is ongoing joint work with Samuel Hanson and Andrei Shleifer of Harvard and with Robert Vishny of the University of Chicago. We haven't quite finished writing up the paper, but we have most of the raw material and a pretty good idea of where we are trying to go, so I figured I would give it a shot.[1]

Our aim is primarily positive, as opposed to normative: We're interested in better understanding the economic role played by commercial banks, as well as the interplay between the traditional commercial banking sector and the so-called shadow-banking sector, which includes various non-bank intermediaries such as broker-dealers, money-market funds, and hedge funds.

It would be an understatement to say that the literature on the role of banks is vast. Without attempting to do it justice, let me just note three prominent classes of theories. The first focuses on the liability side of banks' balance sheets--that is, their deposit-taking role. In this view, what is most important about banks is that they create safe claims which, precisely because of their safety and immunity from adverse-selection problems, are useful as a transactions medium.[2] In other words, banks are special because they are the institutions that create private, or "inside," money. Although it does not necessarily follow as a logical matter, this view has led some observers to advocate

[1] The views expressed here are my own and are not necessarily shared by other members of the Federal Reserve Board and the Federal Open Market Committee.

[2] See Gorton and Pennacchi (1990) for an early articulation of the idea that bank liabilities are designed to minimize adverse-selection problems and hence facilitate exchange between uninformed parties.

narrow-banking proposals, whereby bank-created money is backed by very safe liquid assets, such as Treasury bills.[3]

A second class of theories emphasizes the asset side of banks' balance sheets and their role as delegated monitors in the lending process.[4] Here banks are seen as a mechanism for dealing with the information and incentive problems that would otherwise make it difficult for credit to be extended to opaque borrowers such as small businesses. Because this work is silent on the precise structure of the liability side, it does not draw much of a distinction between banks and other, nonbank lending intermediaries, such as finance companies.

Finally, a third class of theories explicitly addresses the question of what ties together the asset and liability sides of banks' balance sheets--that is, why do the same institutions that create private money choose to back their safe claims by investing in loans and other relatively illiquid assets?[5] What is the nature of the synergy between the two activities?

Our work fits into this third category, with two twists relative to previous research. First, we stress the fact that on the asset side, banks hold not only loans, but like their shadow-banking counterparts, they also hold securities, often in very substantial amounts. Moreover, these securities holdings have a particular pattern. Banks tend to stay away from the safest and most liquid securities, such as Treasury securities, and

---

[3] See Pennacchi (2012) for a detailed discussion of narrow-banking proposals.
[4] See Diamond (1984).
[5] Diamond and Dybvig (1983) is the classic reference in this area. They argue that banks allow households that are unsure of the timing of their consumption needs to more efficiently invest in long-lived projects. Other papers include Diamond and Rajan (2001), who suggest that the fragility inherent in runnable deposit finance acts as a useful disciplinary device for bank management, and Kashyap, Rajan, and Stein (2002), who emphasize the similarities between deposits and loan commitments, and the cost savings that accrue to an institution that offers both products.

instead concentrate their holdings in securities that are less liquid and whose market prices are more volatile, including agency mortgage-backed securities (MBS), collateralized mortgage obligations (CMOs), asset-backed securities and corporate bonds. We argue that this pattern is an important clue to the business of banking more generally.

Second, while we follow previous work in assuming that the safety of bank liabilities is an important part of what makes them special, we depart from much of the rest of the literature by downplaying the vulnerability of bank deposits to runs. Indeed, we emphasize precisely the opposite aspect of deposit finance: Relative to other forms of private-money creation that occur in the shadow-banking sector--notably, short-term collateralized claims such as broker-dealer repurchase agreements (repos)--bank deposits are noteworthy because, in the modern institutional environment, they are highly sticky and not prone to run at the first sign of trouble.

In its simplest terms, our story is as follows: There are different private technologies for creating safe money-like claims. The "banking" technology involves meaningful amounts of capital as well as deposit insurance and thus leads to deposits that are both safe and relatively stable. The "shadow banking" technology uses less capital and manufactures safety by, instead, giving repo investors collateral and the right to seize the collateral on a moment's notice. So shadow banking money is much more run prone than bank money. Given its relatively stable nature, the banking model is better suited to investing in assets that are illiquid and subject to interim price volatility--that is, to fire-sale risk. These assets can be loans that involve significant amounts of monitoring, or they can be securities that require less monitoring. What is essential is the synergy between issuing stable types of money claims and investing in assets that have some

degree of exposure to fire-sale risk. That synergy, in our view, is at the heart of the business of traditional banking.

We have developed a simple theoretical model that captures the main ingredients of this story and makes some further testable predictions. It's probably ill-advised over lunch, but I will take a crack at sketching this model for you here. However, let me start with three stylized facts to motivate the theory.

Stylized Facts

The first fact, and the one that I suspect will surprise you the least, is the strong homogeneity of the liability side of banks' balance sheets: Banks are almost always and everywhere largely deposit financed. For example, in the cross section, and using year-end 2012 data, a bank at the 10th percentile of the distribution had a ratio of deposits to assets of 73.6 percent, while a bank at the 90th percentile had a ratio of 88.9 percent.[6] A similar homogeneity is apparent in the time series: Over the past 115 years, deposits have averaged 80 percent of bank assets, with a standard deviation of only 8 percent. These patterns are in sharp contrast to those for nonfinancial firms, for which capital structure tends to be much less determinate, both within an industry and over time. They suggest that for banks--unlike non-financials, and counter to the spirit of Modigliani and Miller (1958)--an important part of their economic value creation happens on the liability side of the balance sheet, via deposit-taking.

[6] In this calculation and those that follow, we eliminated from the sample the very smallest banks--those with less than $1 billion in assets. We did so because otherwise, these banks--which are extremely numerous but account for only a small fraction of total banking system assets--would dominate any statements we might make about percentiles of the population distribution.

The second fact, which is perhaps just a bit more surprising, is that the asset side of banks' balance sheets--and, in particular, their mix of loans versus securities--is considerably more heterogeneous. In the 2012 cross section, a bank at the 10th percentile of the distribution had a ratio of securities to assets of 6.9 percent, while, for a bank at the 90th percentile, the ratio was almost six times higher, at 40.7 percent.[7] One interpretation of this heterogeneity is as follows: While lending is obviously very important for a majority of banks, it need not be the case that a bank's scale is pinned down by the nature of its lending opportunities. Rather, at least in some cases, it seems that a bank's size is determined by its deposit franchise, and that, taking these deposits as given, its problem then becomes one of how best to invest them. Again, this liability-centric perspective is very different from how we are used to thinking about nonfinancial firms, whose scale is almost always presumed to be driven by its opportunities on the asset side of the balance sheet.

The third stylized fact, illustrated in figure 1, is the one that I found most eye-opening: While banks might be quite heterogeneous in their mix of loans and securities, within the category of securities, they appear to have well-defined preferences. As can be seen in the figure, banks hold very little in the way of Treasury securities and securities issued directly by the government-sponsored agencies: These two categories accounted for just 7.7 percent and 5.8 percent of total securities holdings, respectively, on a value-weighted basis in 2012. The large bulk of their holdings is in MBS guaranteed by the agencies and other types of mortgage-linked securities--such as CMOs and commercial

[7] These figures on securities holdings do not include banks' holdings of cash and reverse repos, which averaged 10.2 percent and 4.1 percent of assets, respectively, on a value-weighted basis in 2012.

mortgage-backed securities (CMBS)--which collectively accounted for 57.7 percent of securities holdings in 2012. Also important is the "Other" category, which includes corporate and municipal bonds, as well as asset-backed securities, and which accounted for 29.3 percent of holdings in 2012.

This composition of banks' securities portfolio is not what you would expect if they were simply holding securities as a highly liquid buffer against, say, unexpected deposit outflows or commitment drawdowns. It also appears--superficially, at least--at odds with the narrow-banking premise that one can profitably exploit a deposit franchise simply by taking deposits and parking them in Treasury bills. Rather, it looks as if banks are purposefully taking on some mix of duration, credit, and prepayment exposure in order to earn a spread relative to Treasury bills. And indeed, over the period from 1984 to 2012, the average spread on banks' securities portfolio relative to bills was 1.73 percent.

By contrast, over this same period, we estimate that if banks had raised deposits and invested them exclusively in three-month Treasury bills, they would have earned an average net return of only 0.06 percent of deposits. This figure is based on three components. First, the spread between the rate earned on bills and that paid on deposits averaged 0.87 percent over the sample period. Second, there was additional noninterest income on deposits (for example, overdraft fees) of 0.49 percent of deposits. Third, however, deposit-taking involves considerable bricks-and-mortar costs. Using Call Report data on banks' total costs, and a hedonic regression approach to infer the portion of cost that is due to deposit-taking, we estimate this piece to have averaged 1.30 percent

of deposits over the sample period. Thus, the net profitability of narrow banking is given by 0.87 percent plus 0.49 percent minus 1.30 percent, which equals 0.06 percent.

Overall, our synthesis of these stylized facts is that banks are in the business of taking deposits and investing these deposits in fixed-income assets that have certain well-defined risk and liquidity attributes but which can be either loans or securities. The information-intensive nature of traditional lending--in the delegated monitoring sense--while clearly important in many cases, may not be the defining feature of banking. Rather, the defining feature may be that, whether banks invest in information-intensive loans or relatively transparent securities, they invest in fixed-income assets that have some degree of price volatility and illiquidity and so offer a higher return than very liquid and safe Treasury securities. In this sense, small business loans and CMOs are on one side of the fence, and Treasury securities are on the other.

Before proceeding, I should note a natural first reaction to these observations. You might say, "Of course banks prefer holding riskier securities to riskless ones, even if they create no monitoring value in either case. They are just taking advantage of the put option created by deposit insurance. In other words, the evidence on the patterns of banks' securities holdings just reflects a moral hazard problem and nothing more."

Without intending to dismiss the importance of moral hazard generally, I don't think it can fully account for the data in figure 1. One way to make this point is to redraw the figure, restricting the sample only to those banks with the highest levels of capital at any point in time--those above the median of the distribution by the ratio of equity to assets--as is done in figure 2. As can be seen, the basic patterns are very similar to those in figure 1. Given that these highly capitalized banks are less likely to impose losses on

the Deposit Insurance Fund, I think there is something deeper here than can be explained by a simple appeal to deposit-insurance-induced moral hazard.

So what, then, is the story? As I alluded to earlier, it begins with the observation that a deposit franchise has two important dimensions. First, it offers a bank a low-cost source of financing, given the premium that investors are willing to pay (that is, the lower yield they are willing to accept) for safe money-like claims. And, second, it offers funding stability, since with an adequate capital buffer and deposit insurance, it is rational for depositors to be "sleepy" and disregard moderate changes in the mark-to-market value of long-term fixed-income assets. Thus, a stable deposit franchise gives a bank the ability to ride out transitory valuation changes of the sort that might come from noise-trader shocks or fire sales, without being forced to liquidate assets at temporarily depressed prices.[8] As a result, traditional banks with stable funding have an advantage relative to their shadow banking counterparts in holding those assets where transitory repricing risk is high for a given level of underlying fundamental cash flow risk.

Here's one way to put the broad theme of our work: We have learned from this year's Nobel laureates, Eugene Fama and Robert Shiller, that discount rate variation--that is, transitory movements in asset prices not driven by changes in future expected cash flows--is central to understanding asset pricing. Our basic message is that this same discount rate variation may also be central to thinking about financial intermediation and, in particular, the connection between the asset and liability sides of intermediary balance sheets. In a world with transitory pricing shocks, a stable funding structure can be an important source of comparative advantage for holding certain types of assets.

[8] On noise-trader shocks, see DeLong and others (1990). For models of fire sales, see Shleifer and Vishny (1992, 1997).

The Model

We have developed a simple model to try to make these ideas precise and to flesh out further empirical implications. I won't try to walk you through it in all of its detail, but I will try to give you a flavor for the basic structure and the main results.[9]

The model has three dates--0, 1, and 2--with a timeline illustrated in figure 3. There is a collection of different long-lived fixed-income assets indexed by *i*. The payoffs on these assets are all perfectly correlated. In particular, each asset *i* pays off *R* at time 2 if the aggregate state of the world is such that there has been no default, and it pays off a lower amount, z_i, if there has been a default. At time 1, there is an interim news event. With probability *p*, the news is good, which means that all assets will definitely not default at time 2. And with probability (1 - *p*), the news is bad, which means that there is a subsequent probability of default on all assets of (1 - *q*) at time 2. Thus, in the bad-news state at time 1, the *fundamental* value of asset *i* is $F_i = qR + (1 - q)z_i$. However, if there is bad news at time 1, the *market* value of asset *i* will be depressed by a temporary fire-sale effect and will be given by $k_iF_i \leq F_i$. The value of k_i is endogenous and asset specific in our model and depends on the equilibrium quantity of asset *i* that is liquidated at time 1; I will return to this feature momentarily.

Household preferences are as follows: Households are risk neutral over fluctuations in their consumption. However, they derive additional utility--which can be thought of as transactions services--from holding completely safe claims, since these claims can be used in a money-like fashion. The upshot of these assumptions is that households require a fixed rate of return on any risky claims, no matter how risky, but a

[9] The basic structure of the model is very similar to that in Stein (2012).

lower rate of return on claims that can be made completely safe. These latter safe claims can be interpreted as privately created money.

Given these household preferences, the name of the game for financial intermediaries is to create as much in the way of safe money-like claims from each asset category *i* as possible. And given the structure of the model, there are two distinct technologies for creating safe claims.

The first technology we label "traditional banking." A traditional bank is relatively conservative in issuing safe claims. In particular, for each dollar it holds of asset *i*, it issues z_i of deposits and finances the rest with equity. Since asset *i* always pays off at least z_i, even in the worst-case scenario, a depositor in a traditional bank can sleep through whatever news comes in at time 1 and still be assured of having a safe claim. In other words, a bank backs its deposits with enough capital so as to make these deposits safe over a two-period horizon and hence endogenously sticky; there is no reason for bank depositors ever to run at time 1. Alternatively, one can think of the bank as having acquired deposit insurance--which allows depositors to sleep through time 1--and the deposit insurer as having imposed a capital requirement of $(1- z_i)$ on the bank so as to reduce its expected losses to zero.

The second technology for creating safe claims we label "shadow banking." Shadow banks are more aggressive: They rely less on capital, and more on the exit option at time 1, to create safety for their investors. Specifically, a shadow bank issues k_iF_i of short-term collateralized claims for each dollar of asset *i* it holds, where $k_iF_i > z_i$. The way these claims are kept safe is that if there is bad news at time 1, the investors seize the collateral and dump it at the fire-sale price, realizing k_iF_i. In other words, unlike

bank investors, shadow bank investors cannot afford to sleep through time 1; their ability to pull the plug at this interim date is essential to keeping their claim safe.

For any given asset *i*, the key question is, which type of intermediary--traditional bank or shadow bank--will end up holding the asset in equilibrium, and in what relative proportions? The answer to this question turns on the following tradeoff. On the one hand, since banks finance themselves with more equity and less cheap "money" than shadow banks, their overall cost of financing is higher, which puts them at a disadvantage. On the other hand, their more conservative capital structure means that banks' deposits never run on them. Hence, unlike shadow banks, they are never forced to liquidate assets at temporarily low prices when there is bad news at time 1; they can simply ride out this bad news and hold on to their investments until prices recover at time 2. In other words, as compared with shadow banks, traditional banks pay more to have a stabler funding structure, which is especially helpful for investing in those assets for which fire-sale discounts are high.

To close the model, we need to endogenize the fire-sale discount factor, k_i. We assume that this fire-sale discount depends in part on the amount of liquidations at time 1, which is in turn related to the extent of ownership of asset *i* by the shadow banking sector--the more the shadow banking sector owns, the more gets dumped at time 1, and hence the lower is the fire-sale price. This assumption implies that the fire-sale discount plays an important equilibrating role in the model, and that we can have interior equilibria in which the ownership of a given asset *i* is divided across the traditional banking and shadow banking sectors; as ownership migrates to shadow banks, the fire-sale discount widens, which tends to reduce the relative advantage of shadow banks.

Specifically, we adopt the following reduced form: $k_i = k(\mu_i, \varphi_i)$, where μ_i is the fraction of asset i held by the shadow banking sector and φ_i is an index of asset illiquidity, with higher values of φ_i corresponding to less-liquid assets. We assume that $dk_i/d\mu_i < 0$, meaning that a greater shadow banking share results in a lower fire-sale price, and that $d^2k_i/d\mu_i d\varphi_i < 0$, meaning that this adverse price-pressure effect is amplified in more-illiquid assets.

With all of this machinery in place, I can now state our main results, which characterize how banks' equilibrium market shares depend on the two primitive asset-level parameters, φ_i and z_i. First, we have that $d\mu_i/d\varphi_i < 0$, which means that banks have a bigger share relative to shadow banks in more-illiquid assets, all else being equal. Second, we have that $d\mu_i/dz_i < 0$, which means that banks have a bigger share in assets that have less long-run solvency risk, all else being equal.

Taken together, these two results suggest that banks have a comparative advantage in holding assets that can experience significant temporary price dislocations but, at the same time, have only modest fundamental risk. Agency MBS might be a leading example of such an asset, since they are insured against default risk but are less liquid than Treasury securities and, for a given duration, have more price volatility, as the MBS-Treasury spread varies significantly.

The model also explains why, even absent any institutional or regulatory constraints, banks would endogenously choose to avoid equities--they simply have too much fundamental risk. Because their value can fall very far over an extended period of time--that is, because their z_i is close to zero--equities cannot be efficiently used as backing to create safe two-period claims. So they are not good collateral for bank

money. By contrast, to the extent that they are highly liquid, they do make suitable collateral for very short-term repo financing. In other words, equities can be used to back some amount of shadow bank money.

It should be noted that the model admits either interior outcomes or corner solutions, depending on the asset-specific values of z_i and φ_i. So it is consistent with the possibility that some assets (say, highly illiquid loans) will be held only by banks, some (say, Treasury securities) will be owned predominantly but not exclusively by shadow banks, and some (say, agency MBS, CMBS, and CMOs) will be owned in significant quantities by each type of intermediary, with an equilibrium fire-sale discount that just balances the tradeoff between the two organizational forms. In this view of the world, loans are seen as being on a continuum with less-liquid securities, and they are held by banks because of their relative illiquidity and susceptibility to fire-sale discounts, not solely because of a need for monitoring.

Further Evidence

One way to draw out some further implications of the model is to focus on the key comparative static, $d\mu_i/d\varphi_i < 0$--namely, that banks have a bigger market share in more-illiquid assets. Now, of course in the real world, there are many different kinds of financial intermediaries, not just two as in our simple model. But a rough generalization of our theory would be to say that across intermediary types, we should expect those with relatively more-stable liabilities to hold more-illiquid assets.

To implement a test of this proposition, we use flow of funds data on eight classes of financial intermediaries: banks, broker-dealers, money market funds, real estate investment trusts, government-sponsored enterprises, property and casualty insurers, life

insurers, and finance companies.[10] The flow of funds data include breakdowns of the different assets and liabilities of each of these intermediaries. Our test then requires us to measure the "illiquidity" of each category of asset and the "stickiness" of each category of liability, which we do using homemade indexes that vary from 0 to 1. Although there is inevitably some subjectivity involved in constructing these indexes, we try to minimize this subjectivity by relying as heavily as possible on the parameters in Basel III's Net Stable Funding Ratio and Liquidity Coverage Ratio.[11] With these indexes in hand, we can then aggregate up to create values for the asset illiquidity and liability stickiness of an entire intermediary sector.[12]

Figure 4 plots the results of this exercise using data as of year-end 2012. As can be seen, there is a very tight cross-intermediary correlation between liability stickiness and asset illiquidity. At one end of the spectrum are money market funds and broker-dealers--prototypical shadow-bank-type institutions--with very low values of both, and at the other end are banks and life insurance companies. While there are only eight data points, the close fit of the regression line is consistent with the broad thrust of our story.

To further highlight the importance of stable bank liabilities, figure 5 takes an approach similar to that of figure 4 but plots asset illiquidity versus the *contractual maturity* of liabilities, rather than their effective stickiness. Now the banking sector

[10] The data are from the Federal Reserve Board's Statistical Release Z.1, "Flow of Funds Accounts of the United States" (in June 2013, the release was renamed "Financial Accounts of the United States"), www.federalreserve.gov/releases/z1.

[11] See Basel Committee on Banking Supervision (2010, 2013).

[12] As an example, using this approach, we assigned values of the illiquidity index of 0 for Treasury securities, 0.15 for agency MBS, 0.5 for corporate equities, and 1.0 for unsecured commercial and industrial loans and commercial real estate loans. Similarly, we assigned values of the stickiness index of 0 for nondeposit short-term funding, 0.7 for wholesale bank deposits, 0.8 for retail time and savings deposits, and 0.9 for transactions deposits.

appears as a huge outlier, since the contractual maturity of banks' liabilities is extremely short--indeed, shorter than that of broker-dealers--even while they hold highly illiquid assets. This observation underscores a key point: If one wants to understand the divergent asset-side behavior of banks and shadow banks, one has to look not to the literal contractual maturity of their liabilities, but rather to the differing incentives that their overall financial and institutional arrangements create for short-term claims to be stable in the face of bad news.

Other Implications

Our framework may also be helpful in understanding a couple of other aspects of the business of banking. The first is banks' accounting treatment of their "available for sale" securities, according to which ongoing mark-to-market gains and losses don't appear in net income.[13] This treatment implicitly presumes that these gains and losses are temporary and that banks can ride them out by not having to sell the security in question before it matures. While I would not want to go so far as to claim that our model justifies this accounting practice, it at least provides a starting point for thinking about it.[14] By contrast, the practice would be incomprehensible in a world in which changes in asset prices reflected changes in expected future cash flows.

Second, our approach may help shed some light on the bricks-and-mortar costs associated with bank deposit-taking. As I mentioned earlier, we estimate these costs to be quite high, averaging on the order of 1.30 percent of deposits over the period from

[13] These "temporary impairments" flow through another liability account called "accumulated other comprehensive income" and affect earnings under generally accepted accounting principles only if the gains or losses are realized by selling a security.
[14] See Cochrane (2011) for a similar argument.

1984 to 2012. And of course, these costs ultimately represent an endogenous decision--banks could always choose to offer their customers fewer and less attractive branch locations, fewer opportunities for interacting with a human teller, and so forth. One view is that these amenities are simply a separable flow of services to depositors, conceptually analogous to interest income. However, an interesting alternative is that they represent a deliberate effort to build loyalty--that is, to create a form of switching costs. To the extent that other elements of their business model are also devoted to creating and exploiting deposit stickiness, it may make more sense for banks to make complementary investments of this sort in enhancing customer loyalty. By contrast, a money market fund complex--which also takes deposits, but which invests exclusively in short-term assets--has less reason to spend as heavily on a branch network.

Conclusion

Let me summarize. The creation of private money--that is, safe claims that are useful for transactions purposes--is obviously central to what banks do. But safe claims can be manufactured from risky collateral in different ways, and banks are not the only type of intermediary that engages in this activity. What makes banks unique is that they use a particular combination of financial and institutional arrangements--including capital, deposit insurance, and access to a lender of last resort--as well as substantial investments in bricks and mortar, to create liabilities that are not only safe and money-like, but also relatively stable and thus unlikely to run at the first sign of trouble. This is in contrast to shadow banks, who create money-like claims more cheaply, by relying on an early exit option, and who are therefore more vulnerable to runs and the accompanying fire-sale risk.

I have argued that there is a synergy between banks' stable funding model and their investing in assets that have modest fundamental risk but whose prices can fall significantly below fundamental values in a bad state of the world. This synergy helps explain both why deposit-taking banks might have a comparative advantage at making information-intensive loans and, at the same time, why they tend to hold the specific types of securities that they do.

References

Basel Committee on Banking Supervision (2010). *Basel III: International Framework for Liquidity Risk Measurement, Standards, and Monitoring.* Basel, Switzerland: Bank for International Settlements, December, www.bis.org/publ/bcbs188.pdf.

------ (2013). *Basel III: The Liquidity Coverage Ratio and Liquidity Risk Monitoring Tools*. Basel, Switzerland: Bank for International Settlements, January, www.bis.org/publ/bcbs238.pdf.

Cochrane, John H. (2011). "Presidential Address: Discount Rates," *Journal of Finance*, vol. 66 (August), pp. 1047-1108.

DeLong, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann (1990). "Noise Trader Risk in Financial Markets," *Journal of Political Economy*, vol. 98 (August), pp. 703-38.

Diamond, Douglas W. (1984). "Financial Intermediation and Delegated Monitoring," *Review of Economic Studies*, vol. 51 (July), pp. 393-414.

Diamond, Douglas W., and Philip H. Dybvig (1983). "Bank Runs, Deposit Insurance, and Liquidity," *Journal of Political Economy*, vol. 91 (June), pp. 401-19.

Diamond, Douglas W., and Raghuram G. Rajan (2001). "Liquidity Risk, Liquidity Creation, and Financial Fragility: A Theory of Banking," *Journal of Political Economy*, vol. 109 (April), pp. 287-327.

Gorton, Gary, and George Pennacchi (1990). "Financial Intermediaries and Liquidity Creation," *Journal of Finance*, vol. 45 (March), pp. 49-71.

Kashyap, Anil K., Raghuram Rajan, and Jeremy C. Stein (2002). "Banks as Liquidity Providers: An Explanation for the Coexistence of Lending and Deposit-Taking," *Journal of Finance*, vol. 57 (February), pp. 33-73.

Modigliani, Franco, and Merton H. Miller (1958). "The Cost of Capital, Corporation Finance, and the Theory of Investment," *American Economic Review*, vol. 48 (June), pp. 261-97.

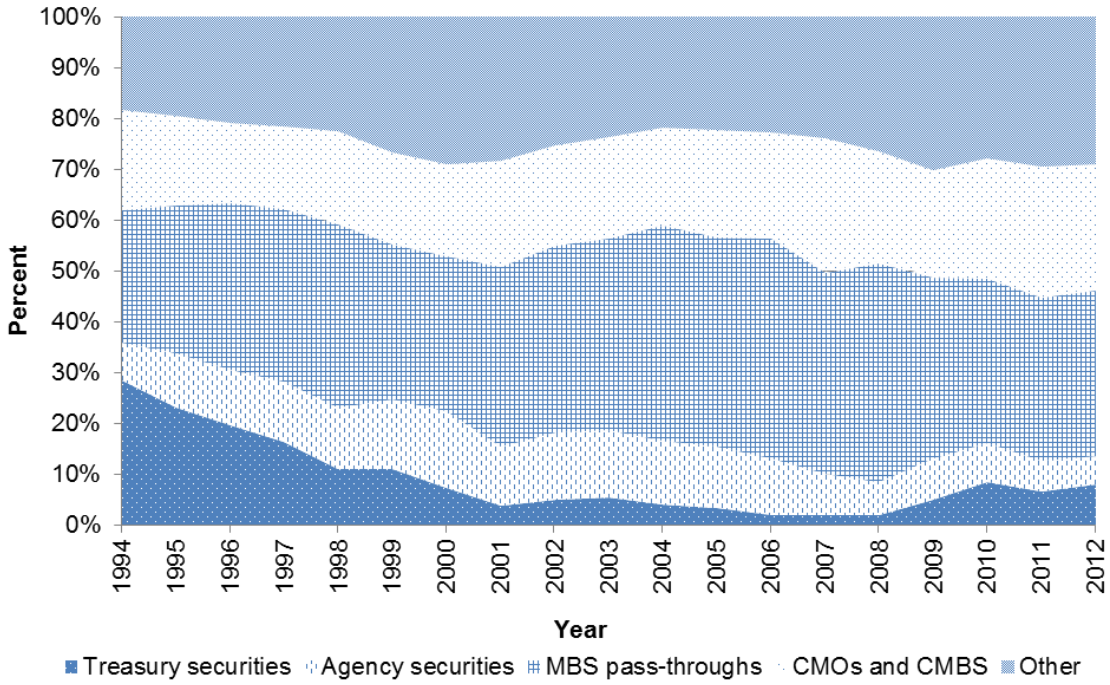
Pennacchi, George (2012). "Narrow Banking," *Annual Review of Financial Economics*, vol. 4 (October), pp. 141-59.

Shleifer, Andrei, and Robert W. Vishny (1992). "Liquidation Values and Debt Capacity: A Market Equilibrium Approach," *Journal of Finance*, vol. 47 (September), pp. 1343-66.

------ (1997). "The Limits of Arbitrage," *Journal of Finance*, vol. 52 (March), pp. 35-55.

Stein, Jeremy C. (2012). "Monetary Policy as Financial-Stability Regulation," *Quarterly Journal of Economics*, vol. 127 (February), pp. 57-95.

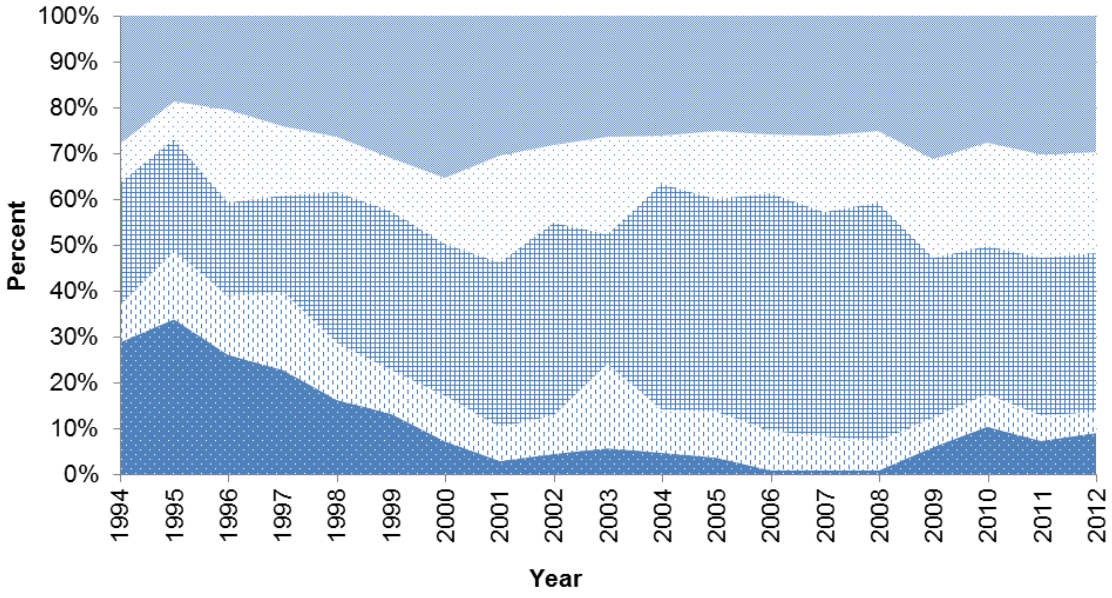
Figure 1: Composition of Bank Securities Portfolios



Note: This figure shows the value-weighted holdings of various marketable securities as a fraction of total bank securities. The sample includes all banks with assets of greater than $1 billion.

Source: Call Reports.

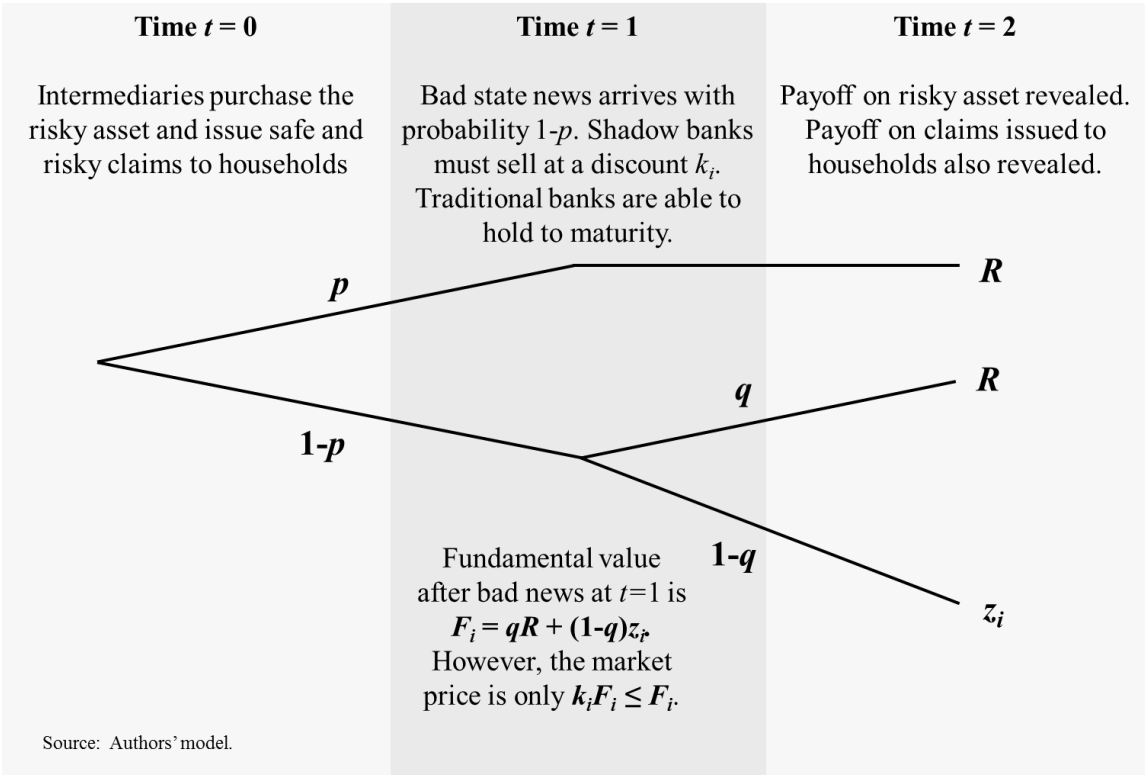
Figure 2: Composition of Bank Securities Portfolios for High-Capital Banks



Note: This figure shows the value-weighted holdings of various marketable securities as a fraction of total bank securities. The sample includes all banks with assets of greater than $1 billion and above-median values of equity to assets.

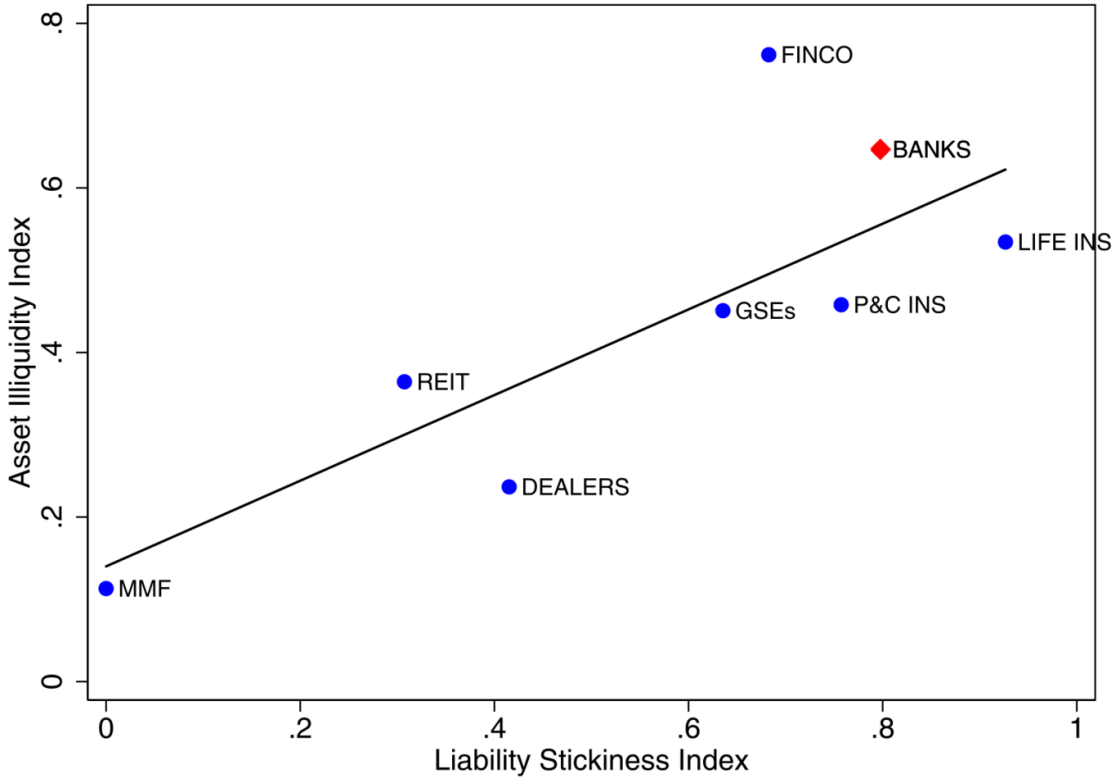
Source: Call Reports.

Figure 3: Timing of the Model



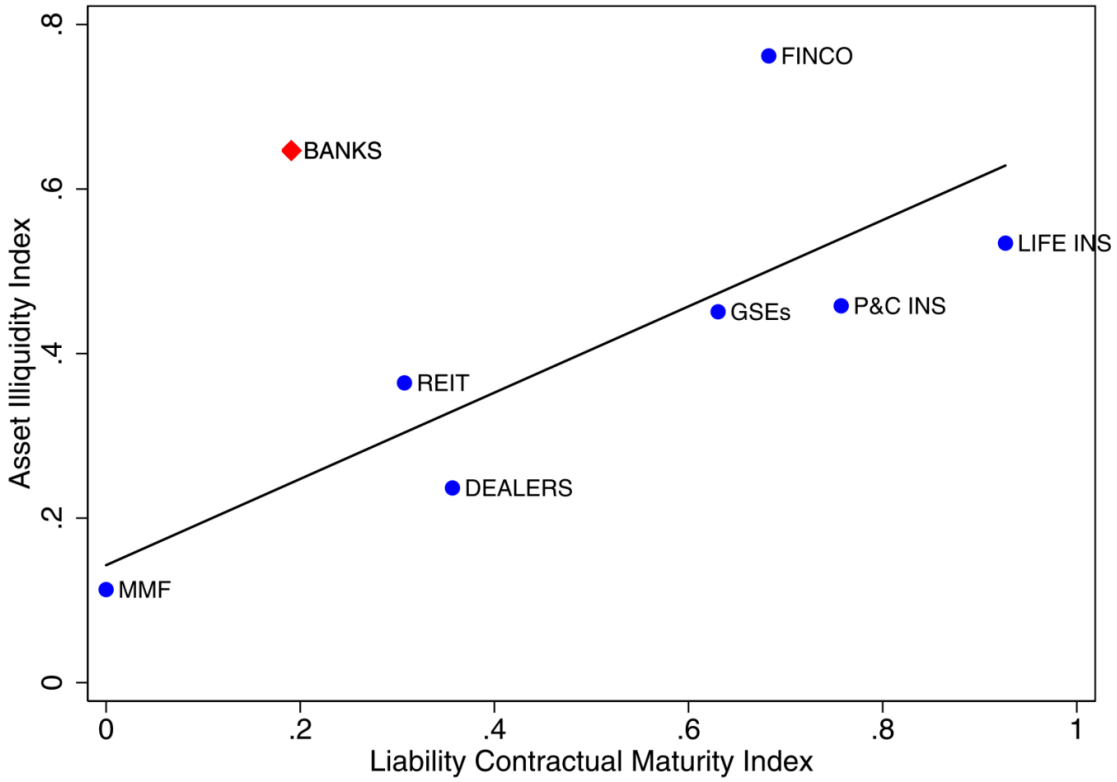
Source: Authors' model.

Figure 4: Asset Illiquidity versus Liability Stickiness across Different Intermediary Types



Source: Federal Reserve Board flow of funds data and authors' calculations.

Figure 5: Asset Illiquidity versus Liability Contractual Maturity across Different Intermediary Types



Source: Federal Reserve Board flow of funds data and authors' calculations.