Central Banking and Financial Stability*

Jeremy C. Stein
Harvard University and NBER

October 2011

* Prepared for Paolo Baffi lecture, Rome, November 30, 2011. This lecture is based on Stein (2011), and also draws from Kashyap and Stein (2011). I am grateful to Anil Kashyap for allowing me to use material from our joint work.

I. Introduction

The modern literature on monetary policy emphasizes the central bank's role in fostering *price stability*.[1] Historically, however, a dominant concern for central bankers has been not just price stability, but also *financial stability*. Indeed, Goodhart (1988) argues that the original motivation for creating central banks in many countries was to temper the financial crises associated with unregulated "free banking" regimes:

> "In the nineteenth century, the advocates of free banking argued that the banking system could be trusted to operate effectively without external constraints or regulation….[But] experience suggested that competitive pressures in a milieu of limited information (and, thence, contagion risks) would lead to procyclical fluctuations punctuated by banking panics. It was this experience that led to the formation of noncompetitive, non-profit maximizing Central Banks." (p. 77).

A related emphasis on crisis mitigation is evident in Bagehot's (1873) famous discussion of the lender-of-last-resort function.[2] And certainly, events of the last few years have served to underscore the importance of the central bank's role in preserving financial stability.

In this lecture, I outline a conceptual framework that can be used to think about both the goals and methods of central-bank financial-stability policies. In its simplest form, the overarching message of the lecture is that central banks have an interest in using various tools to control the aggregate quantity of "money" created by private financial intermediaries. Of course, the idea that central banks should control the money supply has been standard textbook fare for generations. However what distinguishes my approach is that it is based not on the premise that too much money is inflationary, but rather on the premise that too much privately-created money (in a sense I will make precise) can lead to a socially-excessive degree of financial fragility. In other words, monetary control in this setting is a particular form of financial-stability regulation.

---

[1] See, e.g., Goodfriend (2007) for a recent articulation of this view.

[2] Tucker (2009) paraphrases Bagehot's (1873) dictum as follows: "to avert panic, central banks should lend early and freely (i.e., without limit) to solvent firms, against good collateral, and at 'high rates'".

Some motivation for this point of view comes from looking at the behavior of various measures of the stock of money-like claims in the years leading up to the current financial crisis. For example, in the 10-year period 1997-2007, the relatively narrow monetary construct M2 grew at a modest 6.4% annual rate. Over the same period, however, other measures of private money creation grew explosively. Balances in institutional money market funds—which had been a primary component in M3—rose at a 16.9% rate, reaching nearly $2 trillion by year-end 2007.[3] The volume of asset-backed commercial paper (ABCP) grew at a similarly rapid rate, peaking at $1.2 trillion in July of 2007, before that market's dramatic collapse.[4] Given the prominent role that money market funds and ABCP played in the unfolding of the crisis, it seems fair to say that, at least with the benefit of hindsight, the rapid growth of these money-like claims was deserving of more policy attention than it received at the time.

In what follows, I develop the economic logic for regulating private money creation by financial intermediaries. The argument proceeds in several steps. The first step is to define the fundamental market failure that gives rise to a need for regulation. I begin by considering an unregulated financial system in which "banks" raise financing from households to invest in projects.[5] Banks can raise this financing in the form of either short-term or long-term debt. Households are risk-neutral with respect to fluctuations in their consumption, but derive additional monetary services from holding any claim that is entirely riskless—with the notion

[3] Interestingly, the Federal Reserve discontinued its publication of the M3 aggregate in March of 2006, saying that "M3 does not appear to convey any additional information about economic activity that is not already embodied in M2 and has not played a role in the monetary policy process for many years."

[4] There is some double-counting involved if one wants to think of both money-market funds and ABCP as separate forms of money that are both ultimately held by the non-financial sector, since a moderate fraction of total ABCP outstanding is held by money market funds. At the peak in mid-2007, this fraction was approximately 21% (Krishnamurthy, Nagel and Orlov (2011)).

[5] While I use the term "banks" throughout for convenience, I mean this term to cover any type of financial intermediary that engages in a similar type of maturity-transformation activity.

being that riskless claims are easy to value and hence facilitate exchange among households. I argue that banks can manufacture some amount of riskless private "money" of this sort, thereby lowering their financing costs. Moreover, they can do so in greater quantity by issuing short-term debt, since it is harder for long-term bank debt to be made risk-free.

The role for financial-stability policy arises because the private choices of unregulated banks with respect to money creation are not in general socially optimal. When banks issue cheaper short-term debt, they capture its social benefits, namely the monetary services it generates for households. However, they do not always fully internalize its costs. In an adverse "financial crisis" state of the world, the only way for banks to honor their short-term debts is by selling assets at fire-sale prices. I argue that in equilibrium, the potential for such fire sales may give rise to a negative externality. Thus left to their own devices, unregulated banks may engage in excessive money creation, and may leave the financial system vulnerable to costly crises.

There are a variety of ways for a regulator to address this externality. One possibility is the use of conventional monetary-policy tools, i.e. open-market operations. To see how monetary policy might be of value, note that a crude approach to dealing with the externality would be for the regulator to just impose a cap on each bank's total money creation. However, when the regulator is imperfectly informed about banks' investment opportunities, he will not know where to set the cap, since it is desirable for banks with stronger investment opportunities to do more money creation. In this setting, the regulator can do better with a flexible "cap-and-trade" system in which banks are granted tradable permits, each of which allows them to do some amount of money creation.[6] The market price of the permits reveals information about banks' investment opportunities to the regulator, who can then adjust the cap accordingly—when

[6] Kashyap and Stein (2004) suggest using an analogous cap-and-trade approach to implement time-varying bank capital requirements.

the price of the permits goes up, this suggests that banks have strong investment opportunities, and so the regulator should loosen the cap by putting more permits into the system.

All of this may sound a bit like science fiction; we don't observe cap-and-trade regulation of banks in the real world. However if banks' short-term liabilities are subject to reserve requirements, it turns out that monetary policy can be used as a mechanism for implementing the cap-and-trade approach. When the central bank injects reserves into the system, it effectively increases the number of permits for private money creation. And the nominal interest rate, which captures the cost of holding reserves, functions as the permit price. Thus open-market operations that adjust aggregate reserves in response to changes in short-term nominal rates can be use to achieve the cap-and-trade solution.

An interesting benchmark case is where reserve requirements apply to the money-like liabilities of *all* lenders in the economy. This allows the central bank to precisely control private money creation *with monetary policy alone*. While this case may roughly capture the situation facing central banks at an earlier period in history, it is less realistic as a description of modern advanced economies. Nowadays there are a range of short-term financial-intermediary liabilities that are not subject to reserve requirements, and yet may both: i) provide monetary services; and ii) create fire-sale externalities. For example, Gorton and Metrick (2011), and Gorton (2010) argue that an important fraction of private money creation now takes place entirely outside of the formal banking sector, via the large volumes of short-term collateralized claims created in the "shadow banking" sector.

In this richer environment, monetary policy as it is conventionally practiced is generally not sufficient to rein in excessive money creation. Continuing with the above example, it may in addition be necessary to regulate the volume of activity in the shadow-banking sector, either by

expanding the reach of reserve requirements, or by some other means. Thus the theory helps to make clear the circumstances under which monetary policy needs to be supplemented with other measures. Moreover, it suggests that these other measures lie squarely in the central bank's traditional domain, to the extent that they are also targeted at the fundamental externality associated with excessive private money creation.

The ideas developed here connect to several strands of previous work. First, the basic model of fire sales that creates the rationale for policy intervention draws on Shleifer and Vishny (1992, 1997).[7] Second, the insight that banks create a valuable transactions medium by issuing low-risk claims is formalized in Gorton and Pennacchi (1990). Third, the notion that central bank reserves can be thought of as permits that allow banks to do more of a particular kind of cheap financing appears in Stein's (1998) elaboration of the bank lending channel of monetary policy transmission.[8]

And finally, in order to focus on the financial-stability consequences of monetary policy, it helps to set aside its effects on price stability. One simple way to do so is by appealing to the fiscal theory of the price level, according to which the price level is determined not by the monetary base, but by total outstanding nominal government liabilities—i.e., by the sum of Treasury securities and the monetary base. This enables open-market operations that change the *mix* of Treasuries and bank reserves (while keeping their sum constant) to have real effects on bank investment and financing behavior, even in a world where all prices are perfectly flexible. However, I also discuss how the model's conclusions carry over to an alternative New-

[7] On fire sales, see also Kiyotaki and Moore (1997), Allen and Gale (2005), Brunnermeier and Pedersen (2009), Fostel and Geanakoplos (2008), Geanakoplos (2009), Gromb and Vayanos (2002), Morris and Shin (2004), Caballero and Simsek (2011), and Stein (2009).

[8] For early work on the bank lending channel see also Bernanke and Blinder (1988, 1992), Kashyap, Stein and Wilcox (1993), and Kashyap and Stein (2000).

Keynesian setting with sticky prices, where price stability is governed by a version of the "Taylor rule" (Taylor, 1993, 1999).

The rest of the lecture is organized as follows. In Section II, I outline the basic case for regulating private money creation by banks: I compare banks' financing choices to the social planner's solution, and clarify the conditions under which banks engage in excessive money creation. I also argue that a cap-and-trade approach to regulation can be useful when the social planner has imperfect information. Section III outlines how the cap-and-trade approach can be implemented with conventional monetary policy tools such as open market operations, and/or changes in the interest rate paid on reserves. Section IV explores a number of other complementary policy tools; these include deposit insurance and a lender-of-last-resort function, regulation of the shadow-banking sector, and the use of government debt maturity as a quasi-regulatory tool. Section V concludes.

II. The Case for Regulating Private Money Creation

A. When is There an Externality in Private Money Creation?

Stein (2011) develops a formal model in which the process of private money creation involves a negative externality, and is therefore socially excessive. In what follows, I briefly sketch the key assumptions and intuition of the model. The model features three sets of actors: households, banks, and "patient investors".

Households start with an initial endowment of the one good in the economy. They can either consume this endowment right away, or invest some of it in financial assets and consume later. In addition to consumption, over which they have linear utility, households also derive utility from monetary services. The key assumption is that monetary services can be provided by

any privately-created claim on future consumption, so long as that claim is *completely riskless*.[9] Household preferences thus pin down two real rates. The first is the real return on risky "bonds". The second is the real return on riskless "money". Because riskless money offers households a convenience yield that risky bonds do not, in equilibrium it must have a lower rate of return.

Households cannot invest their initial endowments directly in physical projects, because they do not have the monitoring expertise to do so. This investment must be undertaken by banks, who in turn issue financial claims—in the form of either riskless money or risky bonds—to households. Each bank faces the following investment opportunities. If it invests in a set of projects, and the good state prevails, it earns a strictly positive gross return on its investment. If instead the bad state prevails, the bank ultimately recovers less than its initial investment, and there is a positive probability that the value of its investment collapses all the way to zero.

At an interim date—before the final realization of project performance—there is a public signal that reveals whether the good or bad state will eventually be realized. At this interim date, it is also possible for a bank to sell any fraction of its existing physical assets to a patient investor. A central feature of the model is that the fire-sale discount associated with this sale—i.e. the discount relative to the expected value of the assets as of the interim date—is endogenous, and depends on total asset sales by all banks in the economy. The equilibrium determination of this discount will be discussed shortly.

Other than their access to investment projects, banks have no initial endowments, and hence must raise all their financing externally. They can do so by issuing debt claims to households that are either short-term (maturing at the interim date) or long-term (maturing at the final realization of investment projects). Note that if they finance with long-term debt, no

[9] This assumption is meant to capture the spirit of Gorton and Pennacchi (1990), and Dang, Gorton and Holmstrom (2009). These papers argue that information-insensitive securities are an attractive medium of exchange, because they eliminate the potential for adverse selection between transacting parties.

amount of this debt can ever be made riskless, since there is a positive probability of the assets eventually yielding zero output in the bad state of the world. By contrast, short-term debt can be made riskless, if not too much is issued. This is because by forcing an asset sale upon seeing a bad signal at the interim date—when the assets still have positive expected value—short-term creditors can escape early with a sure value equal to the proceeds from the sale.

Thus, the central tradeoff in the model is this: on the one hand, banks have an incentive to issue some short-term debt, because with sufficient collateral backing it, short-term debt can be made riskless—and hence by virtue of its money-ness, represents a cheap form of finance.[10] On the other hand, what keeps short-term debt safe is the bank's ability to sell assets in the bad state. As will become clear below, these sales of existing assets can lead to social costs that are not always fully internalized by individual banks when they pick their capital structures. As a result, there may be excessive private money creation by banks.

Patient investors (PIs) are another type of intermediary. They have fixed resources and are active only at the interim date, when the banks may be trying to sell assets, depending on the signal that has been received about the future prospects for the economy. PIs can do one of two things with their resources at this interim date. First, they can invest in new, late-arriving real investment projects. Alternatively, PIs can absorb the assets being sold by banks. If the good signal is realized at the interim date, there are no asset sales, so the PIs invest all their resources in new projects. If the bad signal is realized, banks have to sell enough assets to repay their short-term creditors. In equilibrium, the PIs must absorb these fire-sold assets, which leaves them with less to invest in new projects. For the PIs to be willing to allocate their resources in

[10] Other theories of short-term financing include Flannery (1986), Diamond (1991), and Stein (2005), who stress its signaling properties, and Diamond and Rajan (2001) who argue that short-term debt is a valuable disciplining device, particularly for financial intermediaries.

this way, it must be that the marginal return on new projects is the same as the marginal return from buying existing assets from banks. This is what pins down the fire-sale discount.

It follows that the greater is the quantity of money created by the banks, and hence the more bank assets that the PIs have to absorb when the bad signal is realized at the interim date, the less the PIs have left over for investment in new projects. With scarce PI capital, the return on secondary-market arbitrage opportunities (buying up fire-sold assets) becomes the hurdle rate for new investment, a point also emphasized by Diamond and Rajan (2011) and Shleifer and Vishny (2010). This observation makes it clear why fire sales can have real effects—as opposed to just being a transfer between the banks and the PIs—and why a social planner might be concerned with mitigating these effects.

As it turns out, however, such underinvestment by the PIs in real projects is a necessary, but not sufficient condition for there to be an externality associated with banks' private choices as to the quantity of private money creation. In order for there to be an externality, it also needs to be the case that the *banks face a binding collateral constraint in their production of private money*. In other words, given the equilibrium rates of return on money and bonds, the banks' desire to engage in money creation must be sufficiently strong that they are manufacturing as much of it as they can, subject to the constraint that they have enough collateral (i.e. physical assets) backing the money that it is indeed rendered riskless. Naturally, this collateral constraint is more likely to bind when the spread between the rates on bonds and money is relatively high.

The following proposition summarizes the discussion to this point.

Proposition: When the spread between the rates of return on bonds and money is sufficiently high, banks will be at a corner solution in terms of money creation, and will face a binding collateral constraint. In this case, the privately-optimal quantity of money creation

exceeds the socially-optimal quantity. In contrast, when the spread between bonds and money is lower, banks will be at an interior solution, producing less than the maximal amount of money given their availability of physical collateral. In this case, private and socially-optimal outcomes coincide.

Thus banks may create a socially excessive amount of money, but this happens only if the spread between bonds and money is high enough. If the spread is so low that any individual bank chooses an interior value of money creation, there is no divergence between private and social incentives.

B. Understanding the Nature of the Externality

At first glance, it may not be clear why fire sales create a divergence between private and socially optimal outcomes. After all, the price impact of liquidations is a pecuniary externality, and pecuniary externalities by themselves need not lead to violations of the standard welfare theorems. The result in the proposition is a specific case of the generic inefficiency result in economies with incomplete markets (Geanakoplos and Polemarchakis (1986), Greenwald and Stiglitz (1986)). Perhaps the closest analogs are Caballero and Krishnamurthy (2003) and Lorenzoni (2008), who also show how there can be socially excessive borrowing in economies with various financial frictions. In the current setting, the key friction is the presence of a binding collateral constraint. When this constraint is operative, any one agent's impact on market prices affects other agents not only by altering their budget constraints, but also by loosening or tightening their collateral constraints. The first welfare theorem effectively says that pecuniary externalities that operate solely through prices in budget constraints do not lead to inefficiencies, but when prices show up elsewhere, this conclusion no longer holds.

The intuition behind this result can be understood as follows. When the collateral constraint does not bind, in deciding how much money to create, each bank trades off the lower financing cost associated with money creation against the potential for greater fire-sales discounts. But this is exactly the same tradeoff that the planner faces in attempting to balance the marginal value of monetary services to households against the marginal cost of underinvestment by the PIs. Hence in this case, everything is well-internalized.

By contrast, when the constraint binds, an incremental increase in money creation by any one bank has an added effect: by reducing the equilibrium value of assets in a fire sale, *it effectively lowers the collateral value of all other bank's assets, thereby tightening their collateral constraints and impinging on their ability to create money*. Thus when any one bank creates an additional unit of money, and captures the private benefit for doing so, the social benefit is less than that one unit of money, since other banks can no longer produce as much money for a given level of collateral.[11]

C. A "Cap-and-Trade" Approach to Regulating Money Creation

Given the divergence between private and social interests, it would seem that in order to achieve a socially optimal outcome, a regulator would want to impose either: (i) a cap on short-term debt issuance by banks; or (ii) alternatively, a system of Pigouvian taxes on such short-term debt. However, the complication is that, as demonstrated in Stein (2011) and Kashyap and Stein (2011), the optimal level of the tax depends on the nature of banks' investment and financing opportunities. For example, in the setting considered by Kashyap and Stein (2011), when the

[11] Think of two banks A and B as factories that each have a technology for producing money out of physical assets. When the collateral constraint binds, an incremental increase in money production by A is equivalent to a form of pollution that gums up B's production technology, since it reduces the amount of money that B can manufacture out of a given stock of physical assets.

quality of banks' investment projects increases, so does their desire to engage in maturity transformation—i.e. to issue short-term debt to fund long-term investment projects. As a result, the optimal tax goes up.

This reasoning suggests that not only will the optimal tax tend to vary over time with economic conditions, but crucially, that it may depend on information that is not directly observable to regulators—since it seems plausible that banks may have private information about the nature of their investment opportunities. Thus, for example, a regulator who does not observe the quality of banks' lending opportunities does not know enough to set the correct tax rate. Nor can such a regulator implement the optimum by setting a cap on the quantity of maturity transformation activity done by a bank, since the optimal quantity depends on the same privately-observed parameter.[12]

These informational difficulties can be addressed with a system of cap-and-trade. To see the logic, suppose that a regulator knows all the parameters of the model except the productivity of the banks' investment opportunities; this one parameter is privately observed by the managers of the banks. Suppose further that the regulator endows banks with permits that allow them to issue short-term debt—for example, each permit might allow a bank to issue one unit of short-term debt. These permits can be freely traded among banks, and are issued in total quantity Q. Then for any trial value of Q such that banks in aggregate are held below their privately optimal level of activity, the permits will trade for a non-zero price $P(Q)$ that reflects the shadow value of the constraint.

[12] The "net stable funding ratio" concept recently advocated by the Basel Committee is effectively a cap on bank's use of short-term debt relative to longer-term debt. As the above discussion suggests, a potentially big stumbling block for this approach is that it is very difficult for regulators to pick the right level of the cap.

Thus given the regulator's knowledge of Q, and of the other parameters of the model, he can simply read off from the market price of permits the key item of interest, namely the quality of banks' lending opportunities. This is intuitive: the more attractive are a bank's maturity-transformation activities, the more it will pay for permits that allow it to expand these activities. And once the regulator knows this parameter, he can adjust the quantity of permits in the system to their *optimal* level. In other words, the cap-and-trade system uses market prices to generate the information that enables a regulator to set the first-best level of taxes (or equivalently, to pick the first-best level of the cap on short-term debt issuance). For example, in a dynamic setting, if the price of permits suddenly spikes up, the regulator can infer that banks' lending opportunities have improved, and that he should therefore relax the cap by injecting more permits into the system.

Interestingly, in the simple formulation in Kashyap and Stein (2011), the optimal regulatory policy is *partially accommodative* with respect to shocks to lending opportunities: when they improve, the regulator puts more permits into the system, thereby allowing banks to expand their maturity-transformation activities. However, this increase in permits is sufficiently small that their price—or equivalently, the Pigouvian tax on short-term debt—actually rises.

III. Implementing the Optimal Regulatory Scheme with Monetary Policy

A. When the Price Level is Pinned Down Elsewhere

The cap-and-trade approach to bank regulation outlined above may seem alien—it does not have any direct counterpart in the real world. However, I now argue that the cap-and-trade approach can be implemented with something that looks very much like conventional monetary policy—with open-market operations in which the central bank adjusts the quantity of nominal

reserves in the banking system. In this setting, reserves play the role of permits for money creation, given the existence of a binding reserve requirement. And the nominal interest rate corresponds to the price of the permits.

In drawing this analogy, one wrinkle is that I have so far been describing an entirely real economy. To introduce a central bank and a role for monetary policy, I need to bring in a set of nominally-denominated government liabilities, and then pin down the price level. One simple way to do so is by relying on the fiscal theory of the price level (Leeper (1991), Sims (1994), Woodford (1995), Cochrane (1998)).[13] In particular, the government is assumed to issue two types of nominal liabilities: Treasury bills, and bank reserves. According to the fiscal theory, the *sum* of these two nominal liabilities is what is relevant for determining the price level.[14] And given the sum, the *composition* of these liabilities is a real variable, since only reserves can be used to satisfy reserve requirements. Thus holding fixed total government liabilities, when there are more reserves, banks are able to create more money, i.e. to finance a greater fraction of their operations with short-term debt. Hence reserves correspond exactly to the concept of regulatory permits in the real model. And an open-market operation that increases the supply of reserves relative to T-bills is isomorphic to an increase in the regulatory limit in the all-real cap-and-trade version of the model.

Moreover, as noted above, the analog to the price of permits is the current setting is the nominal interest rate. This is because when banks want to create money, they are forced to hold non-interest bearing reserves, and the nominal interest rate represents the opportunity cost of

[13] An alternative approach that yields similar results is to assume that the price level is anchored by a commodity standard.

[14] To operationalize the fiscal theory, one can assume that the government anticipates real tax revenues that are exogenously fixed, and total nominal liabilities outstanding composed of Treasury bills and bank reserves. The price level is then determined by the requirement that the real value of the government's total obligations must equal the present value of its future tax revenues.

doing so. For example, suppose the parameters of the model are such that, in equilibrium, the optimal tax on short-term debt is 20 basis points. If the reserve requirement 10 percent, this tax can be (approximately) implemented with a nominal interest rate of 2.0 percent.[15]

B. Using Interest on Reserves

I have thus far assumed that the price level is determined outside the central bank, by the fiscal-theory mechanism. While this is a convenient simplification, it is not an essential piece of the story. An alternative approach, in the New-Keynesian spirit, would be to model prices as being anchored by the central bank's adherence to a "Taylor rule" (Taylor 1993, 1999) which dictates its path for the short-term nominal rate.

However, this raises a potential problem of there being more objectives than tools. If the short-term nominal rate must satisfy a Taylor rule in order to maintain price stability, how can it also be set equal to its optimal value from a regulatory perspective? One possible way out of this box is via the payment of interest on reserves (IOR), which many central banks around the world have been doing for years, and which the U.S. Federal Reserve first took up in October of 2008. As Goodfriend (2002) points out, with IOR, there are two distinct methods for raising short-term nominal rates: either by increasing the interest paid on reserve balances, or by draining reserves from the system, thereby increasing their scarcity value. These methods are not equivalent, since only the latter scarcity-based approach increases the effective "reserves tax" paid by banks, which has been the focus of the analysis above.

Building on this observation, Kashyap and Stein (2011) argue that the use of IOR allows the central bank to simultaneously accomplish two goals: i) set the short-term nominal rate in

[15] This calculation is not exact, because while one unit of reserves enables ten units of short-term debt to be issued, it only enables nine units of net financing for new loans—since one unit of the funding has to be used to support the reserves themselves. See Stein (2011) for the precise formula.

accordance with a Taylor rule; and ii) implement an optimal regulatory scheme of the sort described in this paper.[16] They note that in a regime with IOR, one can decompose the nominal federal funds rate *f* as follows: $f = r_{IOR} + y_{SVR}$, where r_{IOR} is the level of interest paid on reserves, and the y_{SVR} is the quantity-mediated scarcity value of reserves. Only the latter term is relevant from a regulatory perspective, as only it reflects the opportunity cost to a bank of holding reserves.

For example, suppose that an analysis of the sort described above yields the conclusion that, for regulatory purposes, the optimal value of y_{SVR} is 2.0%, while an application of the Taylor rule implies that the optimal value of *f* is 5.0%. In this case, the central bank should set r_{IOR} to 3.0%, and then adjust the quantity of reserves in the system until *f* equilibrates at 5.0%. In other words, the central bank should in principle stand ready to vary both IOR and the quantity of reserves in the system over time in order to meet both of its objectives. If inflation fears become more pronounced, the central bank should raise r_{IOR} so as to increase the funds rate while holding constant the quantity of reserves—and hence the implied tax on short-term debt issuance. If by contrast the worry is that maturity transformation is becoming excessive and threatening financial stability, the correct response is to drain reserves so as to increase the implied tax, while simultaneously cutting r_{IOR} so as to neutralize the effect on the funds rate.

It is interesting to contrast these normative prescriptions with the observed diversity of central-bank practices before the onset of the recent financial crisis. At one extreme of the spectrum was the Federal Reserve, which set r_{IOR} to zero, so that any variation in the funds rate had to come exclusively from quantity-mediated changes in y_{SVR}. At the other extreme was the

[16] See Woodford (2011) for a more complete treatment of these issues in a dynamic New-Keynesian model.

Reserve Bank of New Zealand, which in July of 2006 adopted a "floor system" in which reserves were made sufficiently plentiful as to drive y_{SVR} to zero, meaning that the policy rate was equal to r_{IOR} And in between were a number of central banks (e.g., the ECB, and the central banks of England, Canada and Australia) which used variants of a "corridor" or "symmetric channel" system. One approach to operating such a system is for the quantity of reserves to be adjusted so as to keep y_{SVR} at a constant positive level—100 basis points being a common value—with r_{IOR} then being used to make up the rest of the policy rate.[17]

The corridor systems share a key feature with the floor system used by New Zealand: in either case, all *marginal variation* in the policy rate comes from variation in r_{IOR}, with no need for changes in quantity of reserves. In this sense, the pre-crisis U.S. approach was fundamentally different from that in many other advanced economies. And none of these regimes can be said to have had the feature that they made purposeful use of variation in *both* r_{IOR} y_{SVR}.

To the extent that the academic literature has taken up the question of how best to use these two monetary policy tools, the general consensus appears to be a preference for the New-Zealand-style floor approach in which the banking system is satiated with reserves at all times, and in which the policy rate is controlled entirely by the level of IOR. The basic logic—as articulated by Goodfriend (2002), Keister, Martin and McAndrews (2008), and Curdia and Woodford (2011)—is an application of the so-called "Friedman rule" (Friedman 1959, 1969): if central-bank reserves are a valuable transactions medium, they should be made available in

[17] See Keister, Martin and McAndrews (2008), and Friedman and Kuttner (2011) for a more detailed discussion of central-bank operating practices.

elastic supply and not taxed. This corresponds to the case where $i = r_{IOR}$, and where $y_{SVR} = 0$, i.e., where reserves are so plentiful that there is no opportunity cost to holding them.[18]

Clearly, the perspective taken here is quite different. The heart of the difference is that this other work is rooted in a New Keynesian modeling framework where the central bank's only job is to minimize deviations of inflation and output from desired targets. In such a setting it is sufficient to give it just one interest-rate dial to turn.[19] However, if one believes that monetary policy should also be attentive to considerations of financial stability, the appeal of a second monetary-policy instrument becomes apparent: it would be difficult to manage both the inflation-output tradeoff, and financial stability, with just a single instrument.[20] In effect, the Friedman-rule logic is turned on its head: a "reserves tax" is seen not as distortionary, but as a way of internalizing an otherwise harmful externality, and changes in the reserves tax over time represent optimal fine-tuning of this regulatory mechanism.

C. Implications for the Structure of Reserve Requirements

1. Breadth of Reserve Requirements

In the United States, reserve requirements are currently applicable to only a subset of commercial banks' short-term liabilities, namely their transactions deposits. This relatively

[18] Curdia and Woodford (2011) write: "… an increase in reserves is unambiguously desirable, in any period in which they remain below the satiation level." They then go on to say: "There are possible arguments (relating to considerations not reflected in our simple model) according to which the optimal spread might be larger than zero, but it is likely in any event to be desirable to maintain a small constant spread, rather than treating the question of the interest rate paid on reserves as a separate discretionary policy decision…"

[19] See, e.g. Gali and Gertler (2007) for a survey. Interestingly, in most New Keynesian models, the central bank is implicitly assumed to follow a New-Zealand-style floor system in setting rates—that is, it simply picks the level of the nominal interest rate, with no reference to quantities of reserves.

[20] For example, Adrian and Shin (2008) argue that: "In conducting monetary policy, the potential for financial sector distress should be explicitly taken into account in a forward-looking manner." However, they do not explain how the central bank can pursue both this objective and its traditional inflation-output goals with a single instrument.

narrow focus of reserve requirements is perfectly adequate for the purposes of conventional monetary policy. For as long as the induced demand for reserves is non-zero, the policy rate can be manipulated by adjusting the quantity of reserves in the system. Indeed, if this is the only goal, it is not really necessary to have any reserve requirements at all, given that some amount of reserves would still be demanded for, e.g., interbank payment and settlement purposes.

However, if monetary policy is to play the sort of regulatory role described above, it becomes important to expand the coverage of reserve requirements. First, within the traditional banking sector, reserve requirements should in principle apply to any form of short-term debt that is capable of creating run-like dynamics and hence systemic fragility; this would include commercial paper, repo finance, brokered certificates of deposit, and so forth. Conceptually, the aim here is very similar to that envisioned in the "net stable funding ratio" concept recently put forward by the Basel Committee on Banking Supervision (2010)—to control the total amount of short-term bank debt *of any sort*—and so the coverage should be designed accordingly.

Going further, given that essentially the same maturity-transformation activities take place in the shadow banking sector (Gorton (2010), Gorton and Metrick (2011)), it would also be desirable to regulate the shadow-banking sector in a symmetric fashion. This suggests imposing reserve requirements on the short-term debt issued by non-bank broker-dealer firms, as well as on other entities (special investment vehicles, conduits, and the like) that hold credit assets financed with short-term instruments such as asset-backed commercial paper and repo. Alternatively, to the extent that many of these short-term claims are ultimately held by stable-value money market funds that effectively take checkable deposits, a reserve requirement could be applied to these funds.[21]

[21] There are clearly some difficult issues of measurement and implementation to be addressed here, given that maturity transformation in the shadow-banking system typically takes place not under the roof of a single entity, but

2. Level of Reserve Requirements

The theory sketched above yields a unique optimum for the Pigouvian tax on short-term debt financing at any point in time. However, it is possible to achieve a given value of the tax with monetary-policy tools in one of two ways: either by adjusting y_{SVR} or by adjusting the reserve requirement. Thus, as in the example above, one can set the tax to (approximately) 20 basis points either with a reserve requirement of 10 percent, and y_{SVR} = 2.0 percent, or with a reserve requirement of 5 percent, and y_{SVR} = 4.0 percent.

Is there any reason to prefer one combination over the other? One appeal of having higher reserve requirements—and hence lower values of y_{SVR}—is that this reduces the likelihood of the regulatory and price-stability roles of monetary policy colliding with one another. In particular, suppose that, according to a Taylor rule, the optimal funds rate is determined to be 3.0 percent. With a reserve requirement of 5 percent, it is impossible to accomplish this objective and to simultaneously impose a 20 basis point tax, since the latter requires the funds rate to be at least 4.0 percent. By contrast, with a reserve requirement of 10 percent, it is possible to meet both goals.

As this discussion suggests, one can also implement the optimal time-varying Pigouvian tax on short-term debt by keeping y_{SVR} pegged at a constant value—as in a "corridor system"—and actively adjusting the reserve requirement. Interestingly, a number of central banks around the world use changes in reserve requirements as a key policy tool. For example, the Chinese

---

rather through a chain of transactions (e.g., a conduit acquires asset-backed securities and issues collateralized short-term commercial paper against these securities, a money-market fund then buys the commercial paper and issues fixed-value claims to its depositors). The challenge is to levy the proper tax on the entire chain of activity, while avoiding both gaps in coverage and double-counting.

central bank changed the level of reserve requirements six times in 2010, while moving their policy interest rate just once.[22]

India offers another intriguing case study. Since November of 2004, the Reserve Bank of India has operated a corridor system of monetary policy. In the aftermath of Lehman Brothers' bankruptcy filing, the Reserve Bank cut reserve requirements from 9.0% to 5.0% in a series of four steps between October of 2008 and January of 2009. Moreover, during this same period, it also narrowed the width of its corridor from 300 basis points to 150 basis points. This narrowing of the corridor amounts to a reduction in the difference between the rate paid on reserves and the policy rate, i.e. to a cut in y_{SVR}. Thus both policy changes had the effect of reducing the reserves tax at a time of severe stress in the financial system.[23]

Finally, Montoro and Moreno (2011) study the use of reserve requirements in three Latin American countries—Brazil, Colombia and Peru. They note that central banks in these countries raised reserve requirements in the expansion phase of the most recent credit cycle, and then, like the Reserve Bank of India, cut them sharply after the bankruptcy of Lehman Brothers. They also argue that the motivation for this approach was explicitly rooted in a financial-stability objective: "Reserve requirements could have two implications for financial stability. First, raising reserve requirements could prevent financial imbalances by restraining credit growth (and by extension, asset price increases) in the upswing of the business cycle. Second, lowering reserve requirements during a downturn can deploy the cushion of reserves built up during the expansion." (page 59).[24]

[22] See Du (2010) for a detailed discussion of the Chinese central bank's policies and procedures.

[23] See Reserve Bank of India (2011) for a recent evaluation of monetary policy practices in India.

[24] Montoro and Moreno (2011) also suggest that reserve requirements were, for these Latin American countries, a *better* tool for achieving financial-stability objectives than movements in the policy rate. Raising the latter during

It is revealing that all of the above examples of activist use of reserve requirements come from emerging economies, where non-bank financial markets are far less developed than they are in, e.g., the United States. These emerging economies mirror the assumptions of the model more closely, in that reserve requirements on bank liabilities allow the central bank to effectively regulate much of the maturity transformation in the economy as a whole. By contrast, as stressed above, making such regulation work in a more developed financial system—and in particular, one with a large shadow-banking sector—would require a substantial broadening of reserve requirements.

The suggestion that reserve requirements should be broader-based, as well as potentially higher, may at first glance strike some as tantamount to a large and distortionary tax increase on the financial sector. While this would be the case if these changes to reserve requirements were made without paying interest on reserves, this no longer need be so in the presence of IOR. Quite the opposite—the use of IOR allows the reserves tax to be targeted at precisely the level that minimizes distortions from a social planner's perspective. And as noted above, the absolute level of reserve requirements has no effect on the equilibrium reserves tax, since increase in the reserve requirement are exactly offset by reductions in y_{SVR}, or alternatively, by increases in the interest paid on reserves r_{IOR}.

3. Reducing the Cost of Holding Excess Reserves

As discussed earlier, previous authors have invoked Friedman (1959, 1969) to argue that it is undesirable to tax banks' holdings of excess reserves, since these excess reserves can be valuable for interbank payment and settlement purposes. The approach described here would

---

the expansionary phase of the credit cycle might have drawn in further capital flows from abroad and put upward pressure on the exchange rate at a time when this was seen as undesirable.

seem to run counter to this philosophy, since as I have outlined it, the tax on reserve holdings would apply to both required and excess reserves. However, this feature of the design can be altered so as to maintain the core regulatory objective, while at the same time allowing excess reserves to remain untaxed, in the spirit of Friedman.

The key modification to our scheme is to let there be a different interest rate for required reserves (IORR) and for excess reserves (IOER). The former is then set below the funds rate exactly along the lines set out above, while the latter is set equal to the funds rate. This means that there is still a tax on any form of bank funding that is regulated by being subject to reserve requirements, but there is no tax when banks choose to hold excess reserves for, say, interbank payments purposes.

The one subtlety with this variant is that the way the central bank learns about the quality of bank investment opportunities is modified. Now, in order for the market to clear with IOER equal to the funds rate, the central bank must satiate the system with reserves. Therefore, the gap between the funds rate and IORR is no longer determined by a market-clearing condition, but rather must be set by the central bank. The problem is that, per the previous analysis, the optimal level of this gap depends on a parameter private information of the banks. However, the central bank can still infer this parameter as follows. Suppose it picks a trial value of IORR somewhere below the funds rate. At this trial value of IORR, it observes the quantity of maturity transformation that banks elect to do, and hence the quantity of required reserves they hold. This combination of price and quantity allow it to deduce how optimistic the banks are about their investment prospects, and it can then iterate to set the optimal value of IORR that takes this information into account. In other words, in the first version of the scheme, the central bank sets quantities (total reserves) and learns from market prices (the gap between the funds rate and

IOR); in this version it sets prices and learns from quantities (how much in required reserves banks choose to hold). Either way, having a reserves market where there are both observable prices and quantities allows for a better outcome than an unconditional cap on maturity transformation.

4. Monetary Policy and Credit Bubbles

In the simple model sketched above, the only divergence between individual banks and the social planner is that each bank takes the capital structure decisions of all other banks as fixed, thereby creating a systemic-risk externality. While this is perhaps the most natural starting point for thinking about monetary-policy from a regulatory perspective, there may be other effects at work that amplify this mechanism. In particular, one consideration that is often discussed informally is the idea that there may be credit bubbles—periods when private lenders make loans that would appear to have abnormally low expected returns. Greenwood and Hanson (2011) provide empirical support for the existence of something like credit bubbles, documenting that times of booming debt issuance by lower-quality (e.g., junk-rated) firms are followed by significantly reduced expected returns on corporate debt relative to Treasuries.

One simple way of incorporating a credit-bubble effect into our model is to assume that there may be a wedge between the marginal return on investment as seen by individual banks, and that perceived by the regulator. Thus a credit bubble can be thought of as a period when the former exceeds the latter, i.e., when banks are excessively bullish on their investment prospects.

Kashyap and Stein (2011) show that this case works very similarly to that studied above, with the primary difference being that the prescription is now for the central bank to be less accommodative than before in the face of incipient increases in y_{SVR}. In other words, with a

given level of reserves in the system, if bank demand for reserves appears to increase, the central bank now injects fewer additional reserves into the system, thereby allowing y_{SVR} (and hence the effective reserves tax) to go up further than it would have in the no-credit-bubble case. Intuitively, in this case, an increase in banks' demand for reserves reflects greater optimism in their *perceived* investment opportunities, but the regulator is more skeptical, and his estimate of the quality of banks' investment opportunities goes up by less than one-for-one with the banks' perceptions. Hence the regulator leans more heavily against an expansion in maturity-transformation activity brought about by such a shift in perceptions.

A number of other observers have suggested that tight monetary policy might be used to try to rein in credit bubbles.[25] However an advantage of the framework developed here is that it gives more precise guidance as to how that tightening can best be effectuated—with a contraction in reserves, and hence with an increase in y_{SVR}, rather than just by raising r_{IOR}. The intuition is straightforward: credit bubbles amplify the divergence between the private and social values of maturity-transformation activity, and hence call for a higher corrective tax to internalize the externality. In contrast, simply increasing the funds rate via the IOR channel does nothing to address the externality.

IV. Other Policy Tools

A. Deposit Insurance and Lender of Last Resort

In the baseline version of the model, the only way for banks to pay off their short-term creditors in the crisis state is by fire-selling their assets, and the only role for policy is to control the amount of short-term debt that is created ex ante. An alternative "bailout" approach would be

[25] Again, see Adrian and Shin (2008) and the references therein.

for the government to try to stem the amount of socially costly fire sales that occur for a *given* amount of short-term bank debt. This could be done either with either deposit insurance, or a lender-of-last resort policy.

Unlike in the classic framework of Diamond and Dybvig (1983), such bailout policies are not costless to the government in equilibrium, because here, in the crisis state, there is a probability that the banks' assets will turn out to be entirely worthless. So there is always a chance that taxpayers will be left on the hook. If taxpayer-financed bailouts create deadweight losses—either because of moral hazard, or because there are distortions associated with raising taxes in an adverse state of the world—the overall optimum set of policies may have the realistic feature that: i) some fraction of banks' money-like claims are insured by the government; ii) the remainder are uninsured, and hence still subject to fire-sale risk; and iii) as before, it makes sense for the regulator to control the total quantity of bank-created money.

To see this explicitly, consider a degenerate case where the deadweight costs of taxation take the following (admittedly peculiar) form: there is no cost to raising any amount less than L to pay for a bailout, but it is infinitely costly to raise anything more than L. It follows that the amount of government-insured money that can be created is bounded by L, and it will in fact always be optimal to set the quantity of insured deposits to exactly L.

Alternatively, given that the government can never put itself in a position to lose more than L, another mechanism that yields the same outcome is a lender-of-last resort facility, in which the government commits to step in and invest L alongside the PIs in the event of a fire sale. This would have exactly the same effect—it would reduce the equilibrium fire-sale discount, and thereby allow for more total money creation.

The bottom line is that one can add deposit insurance to the model in such a way as to make it more realistic, without changing any of its qualitative properties. The optimal policy mix will involve limited use of deposit insurance or equivalently, limited use of a lender-of-last-resort function. Banks will continue to issue uninsured money-like claims alongside insured deposits, and hence will continue to create some degree of fire-sale risk. Thus as before, there will continue to be a motive for regulating the creation of these uninsured short-term claims.

B. Regulating the Shadow Banking Sector and Limiting Regulatory Arbitrage

The model also assumes that all private money is manufactured by commercial banks that are subject to reserve requirements. Hence private money creation can be completely controlled by conventional open-market operations. While this may be an adequate representation of an earlier period in history, it omits an important form of money creation in the modern economy. In today's financial markets, private money—in precisely the sense meant here—is also created by the unregulated shadow banking system, via the large volume of short-term claims that are collateralized by securitized loan pools of one form or another. This observation suggests that commercial banks and shadow banks should be regulated in a symmetric fashion. As discussed above, one way to do this would be to broaden the reach of reserve requirements, so that the cap-and-trade regime covers all short-term liabilities of both commercial banks and shadow banks.

However, if either political or institutional constraints make it infeasible for the liabilities of shadow banks to be subjected to reserve requirements, there are other tools that may prove helpful. One such alternative approach might be to impose a regime of "haircut" regulation. For example, the central bank (or other regulator) could specify the maximum fraction of short-term financing that could be issued against a given pool of collateralizable assets. Moreover, just as the optimal quantity of bank-created money varies with economic conditions, optimal haircuts

would respond to these conditions as well. Stein (2011) provides an explicit analysis of haircut regulation. In the context of that model, it turns out that haircut regulation is indeed useful, but is strictly less efficient than direct control of the quantity of privately-created money via the sort of reserve-requirements-based mechanism outlined above. Intuitively, if the root of the problem is that the private sector is creating more money than is socially optimal, a system of reserve requirements allows for direct control of the quantity of money, and hence of the externality. By contrast, haircut regulation does not constrain the total quantity of money directly, but rather the *ratio* of money to collateralizable assets. Hence one potentially undesirable side-effect of haircut regulation may be that it leads to overproduction of certain kinds of hard assets (e.g. houses) as the system strains to create a larger aggregate quantity of private money.

A skeptical reaction to all of this might be that a broad-based system of reserve requirements on the short-term liabilities of financial firms will naturally invite some form of regulatory arbitrage, i.e., an attempt to evade the rules by moving the borrowing to an unregulated entity. This is certainly true, in the general sense that any form of financial regulation can be expected to lead to some evasion. However, the relevant benchmark is not a world with no regulation and no evasion, but rather other, less efficient forms of regulation. Again, the net stable funding ratio concept recently put forward by the Basel Committee is an appropriate point of comparison. This rule seeks to regulate the same behavior—short-term debt issuance by financial firms—but does so by simply imposing a rigid cap on issuance, with no price-based feedback from the market. One danger with a cap of this sort is that one never gets to observe directly the shadow value of the constraint. Moreover, if the cap is set too tight, so that the shadow value of the constraint is very high, this is precisely when the incentive to evade

the rules is strongest. Thus the rigid cap approach embodied in the net stable funding ratio might be said to be particularly problematic on the evasion dimension.[26]

By contrast, it is a virtue of a cap-and-trade regime that when the price of the permits begins to move upwards, the regulator can inject more permits into the system, thereby reducing their price and the accompanying incentives to skirt the rules. In this sense, the price-based mechanism has an inherent safety valve that can help to mitigate—though never eliminate—the regulatory-arbitrage problem. This benefit of a partial-accommodation approach is likely to be especially pronounced during periods of credit expansion, whereas at such times a rigid, non-accommodating cap is most prone to drive maturity-transformation activity underground.

C. Government Debt Maturity

As was emphasized above, the magnitude of the externality associated with private money creation is related to the equilibrium level of the bond-money spread: when the spread widens, the incentive for banks to manufacture money increases, and the wedge between the social and private returns to money creation goes up. Thus an alternative way to moderate the externality would be to compress the spread.

In the plausible case where the monetary services enjoyed by households are a concave function of the supply of money—i.e., there is diminishing marginal utility of money—then it becomes possible for the government to act directly on the bond-money spread. For example, since short-term Treasury bills are riskless, they can provide the same monetary services as short-term bank debt. Hence an increase in the supply of Treasury bills will, in this modified setting, reduce the bond-money spread.

[26] Indeed, one hypothesis for why the Basel Committee has been so slow to move forward with the net stable funding ratio—it is not scheduled for implementation until 2018—is that it is difficult to calibrate the right level of the ratio absent price-based feedback, and that there are large costs to getting it wrong.

One appeal of dealing with the externality in this fashion is that unlike some other regulatory approaches, it does not invite evasion. For example, if the scope of reserve requirements were broadened, private actors might try to get around limits on their ability to use short-term debt by using various forms of hidden borrowing, e.g., by embedding the borrowing in an opaque derivative contract. In contrast, when the relative cost of short-term borrowing goes up—because the market has been saturated with riskless short-term claims—the incentive to create private money is blunted.

In Greenwood, Hanson and Stein (2010), we use this observation as the point of departure for a normative theory of government debt maturity. We argue that the government should choose a shorter debt maturity—and in particular, should issue more riskless T-bills—than it otherwise might, in an active effort to crowd out the short-term debt of financial intermediaries. The argument is based on a principle of comparative advantage. On the one hand, tilting its issuance towards short-term debt is not without cost for the government, since with stochastic interest rates this increases the variability of future interest payments and ultimately disrupts efforts to smooth tax rates over time. On the other hand, short-term government debt, unlike the short-term debt of financial intermediaries, does not create fire-sale risk. To the extent that the fire-sale externality is more costly to the economy at the margin than the disruption of tax smoothing, it can make sense for the government to take on a bigger role in providing the short-term riskless claims that the economy demands.

Of course, precisely because of tax-smoothing considerations, it will not generally be optimal for the government to tilt so strongly towards short-maturity issuance as to entirely eliminate the bond-money spread in equilibrium. Rather, optimal behavior by the government on this dimension will typically involve leaving the spread only partially compressed. So while

government debt maturity may be one helpful tool in addressing the problem of excessive private money creation, it is not a panacea, and it is unlikely to eliminate the usefulness of the other tools discussed above.

V. Conclusions

The basic message of this lecture can be summarized as follows. Banks and other financial intermediaries like to fund themselves with short-term debt. With sufficient collateral backing it, this short-term debt can be made into riskless money, which, because of the transactions services it generates, represents a cheap source of finance for banks. While society benefits from this private money creation, banks' private incentives lead them to overdo it, since they do not fully internalize the fire-sales costs that are a byproduct of their maturity-transformation activities. The externality associated with excessive private money creation provides a fundamental rationale for financial-stability regulation, and arguably, for the existence of central banks.

In a sufficiently simple institutional environment, the externality can be addressed with conventional monetary policy, complemented by either deposit insurance or a lender-of-last-resort facility. Indeed, this is one interpretation of what central banks have done for much of their history. In a more realistic modern-day setting, where a substantial shadow-banking sector exists alongside traditional commercial banks, other tools, such as expanded reserve requirements, or haircut regulation, may also be necessary. If so, central banks should not be reluctant to deploy these tools—to the extent that they do so in an effort to contain excessive private money creation, they can be said to be pursuing one of their traditional core missions in a more comprehensive and effective manner.

The above arguments put a new twist on an old notion in monetary economics—that independent of the level of nominal interest rates, it is important for the central bank to control the quantity of "money" created by the private financial sector. What is different is that here the operative definition of "money" is not just that it is a transactions medium, but that it is any form of short-term intermediary debt that has the potential to create systemic externalities. And the rationale for controlling it is not to target nominal GDP, as in a traditional quantity-theoretic model, but rather to help ensure financial stability.

References

Adrian, Tobias, and Hyun Song Shin. 2008. "Financial Intermediaries, Financial Stability and Monetary Policy." In *Maintaining Stability in a Changing Financial System*, 287-334. Federal Reserve Bank of Kansas City.

Allen, Franklin, and Douglas Gale. 2005. "From Cash-in-the-Market Pricing to Financial Fragility." *Journal of the European Economic Association* 3: 535-546.

Bagehot, Walter. 1873. *Lombard Street: A Description of the Money Market.* New York: Charles Scribner's Sons.

Basel Committee on Banking Supervision. 2010. "The Basel Committee's Response to the Financial Crisis: Report to the G20." www.bis.org/publ/bcbs179.pdf.

Bernanke, Ben S., and Alan Blinder. 1988. "Credit, Money, and Aggregate Demand." *American Economic Review* 78: 435-39.

Bernanke, Ben S., and Alan Blinder. 1992. "The Federal Funds Rate and the Channels of Monetary Transmission." *American Economic Review* 82: 901-21.

Brunnermeier, Markus K., and Lasse H. Pedersen. 2009. "Market Liquidity and Funding Liquidity." *Review of Financial Studies* 22: 2201-2238.

Caballero, Ricardo J., and Arvind Krishnamurthy. 2003. "Excessive Dollar Debt: Financial Development and Underinsurance." *Journal of Finance* 58: 867-893.

Caballero, Ricardo J., and Alp Simsek. 2011. "Fire Sales in a Model of Complexity." MIT working paper.

Cochrane, John. 1998. "A Frictionless Model of U.S. Inflation." *NBER Macro Annual* 13: 323-384.

Curdia, Vasco, and Michael Woodford. 2011. "The Central-Bank Balance Sheet as an Instrument of Monetary Policy." *Journal of Monetary Economics* 58: 54-79.

Dang, Tri Vi, Gary B. Gorton, and Bengt Holmstrom. 2009. "Opacity and the Optimality of Debt for Liquidity Provision." Yale working paper.

Diamond, Douglas W., and Philip Dybvig. 1983. "Bank Runs, Deposit Insurance and Liquidity." *Journal of Political Economy* 91: 401-419.

Diamond, Douglas W. 1991. "Debt Maturity Structure and Liquidity Risk." *Quarterly Journal of Economic* 106: 709-737.

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

Diamond, Douglas W., and Raghuram G. Rajan. 2011. "Fear of Fire Sales, Illiquidity Seeking and Credit Freezes." *Quarterly Journal of Economics* 126: 557-591.

Du, Yang. 2010. "Monetary Policy and Bank Loan Supply in China." Williams College undergraduate thesis.

Flannery, Mark J. 1986. "Asymmetric Information and Risky Debt Maturity Choice." *Journal of Finance* 41: 19-37.

Fostel, Ana, and John Geanakoplos. 2008. "Leverage Cycles and the Anxious Economy." *American Economic Review* 98: 1211-1244.

Friedman, Benjamin M., and Kenneth N. Kuttner. 2011. "Implementation of Monetary Policy: How Do Central Banks Set Interest Rates?" In *Handbook of Monetary Economics*, Friedman and Woodford (eds.). Amsterdam: North-Holland.

Friedman, Milton. 1959. *A Program for Monetary Stability*. New York: Fordham University Press.

Friedman, Milton. 1969. "The Optimum Quantity of Money." In *The Optimum Quantity of Money and Other Essays*. Chicago: Aldine.

Gali, Jordi, and Mark Gertler. 2007. "Macroeconomic Modeling for Monetary Policy Evaluation." *Journal of Economic Perspectives* 21: 25-45.

Geanakoplos, John. 2009. "The Leverage Cycle." *NBER Macro Annual* 24: 1-65.

Geanakoplos, John, and Heraklis Polemarchakis. 1986. "Existence, Regularity and Constrained Suboptimality of Competitive Allocations When the Asset Market is Incomplete." In *Essays in Honor of Kenneth Arrow, Volume III*, edited by W. Heller, R. Starr, and D. Starrett, 65-95. Cambridge: Cambridge University Press.

Goodfriend, Marvin. 2002. "Interest on Reserves and Monetary Policy." *Federal Reserve Bank of New York Economic Policy Review* 8: 13-29.

Goodfriend, Marvin. 2007. "How the World Achieved Consensus on Monetary Policy." *Journal of Economic Perspectives* 21: 47-68.

Goodhart, Charles. 1988. *The Evolution of Central Banks.* Cambridge, MA: MIT Press.

Gorton, Gary B., and George Pennacchi. 1990. "Financial Intermediaries and Liquidity Creation." *Journal of Finance* 45: 49-72.

Gorton, Gary B. 2010. *Slapped by the Invisible Hand: The Panic of 2007.* Oxford University Press.

Gorton, Gary B., and Andrew Metrick. 2011. "Securitized Banking and the Run on Repo." *Journal of Financial Economics*, forthcoming.

Greenwald, Bruce, and Joseph E. Stiglitz. 1986. "Externalities in Economies with Imperfect Information and Incomplete Markets." *Quarterly Journal of Economics* 101: 229-264.

Greenwood, Robin, Samuel Hanson, and Jeremy C. Stein. 2010. "A Comparative Advantage Approach to Government Debt Maturity." Harvard University working paper.

Greenwood, Robin, and Samuel Hanson. 2011. "Issuer Quality and the Credit Cycle." Harvard Business School working paper.

Gromb, Denis, and Dimitri Vayanos. 2002. "Equilibrium and Welfare in Markets with Financially Constrained Arbitrageurs." *Journal of Financial Economics* 66: 361-407.

Kashyap, Anil K, Jeremy C. Stein, and David Wilcox. 1993. "Monetary Policy and Credit Conditions: Evidence from the Composition of External Finance." *American Economic Review,* 83: 78-98.

Kashyap, Anil K, and Jeremy C. Stein. 2000. "What Do a Million Observations on Banks Say about the Transmission of Monetary Policy?" *American Economic Review* 90: 407-28.

Kashyap, Anil K, and Jeremy C. Stein. 2004. "Cyclical Implications of the Basel-II Capital Standards." *Federal Reserve Bank of Chicago Economic Perspectives,* 28, Q1: 18-31.

Kashyap, Anil K, and Jeremy C. Stein. 2011. "The Optimal Conduct of Monetary Policy with Interest on Reserves." Forthcoming, *American Economic Journal: Macroeconomics*.

Keister, Todd, Antoine Martin, and James McAndrews. 2008. "Divorcing Money from Monetary Policy." *Federal Reserve Bank of New York Economic Policy Review* September: 41-56.

Kiyotaki, Nobuhiro, and John Moore. 1997. "Credit Cycles." *Journal of Political Economy* 105: 211-248.

Krinamurthy, Arvind, Stefan Nagel, and Dmitry Orlov. 2011. "Sizing Up Repo." Stanford GSB working paper.

Leeper, Eric. 1991. "Equilibria Under Active and Passive Monetary Policies." *Journal of Monetary Economics* 27: 129-147.

Lorenzoni, Guido. 2008. "Inefficient Credit Booms." *Review of Economic Studies* 75: 809-833.

Montoro, Carlos, and Ramon Moreno. 2011. "The Use of Reserve Requirements as a Policy Instrument in Latin America." *BIS Quarterly Review* March: 53-65.

Morris, Stephen, and Hyun Song Shin. 2004. "Liquidity Black Holes." *Review of Finance* 8: 1-18.

Reserve Bank of India. 2011. "Report of the Working Group on Operating Procedure of Monetary Policy."

Shleifer, Andrei, and Robert W. Vishny. 1992. "Liquidation Values and Debt Capacity: A Market Equilibrium Approach." *Journal of Finance* 47: 1343-1366.

Shleifer, Andrei, and Robert W. Vishny. 1997. "The Limits of Arbitrage." *Journal of Finance* 52: 35-55.

Shleifer, Andrei, and Robert W. Vishny. 2010. "Unstable Banking." *Journal of Financial Economics* 97: 306-318.

Sims, Christopher A. 1994. "A Simple Model for the Determination of the Price Level and the Interaction of Monetary and Fiscal Policy." *Economic Theory* 4: 381-399.

Stein, Jeremy C. 1998. "An Adverse-Selection Model of Bank Asset and Liability Management with Implications for the Transmission of Monetary Policy." *RAND Journal of Economics* 29: 466-86.

Stein, Jeremy C. 2005. "Why are Most Funds Open-End? Competition and the Limits of Arbitrage." *Quarterly Journal of Economics* 120: 247-272.

Stein, Jeremy C. 2009. Presidential address: "Sophisticated Investors and Market Efficiency." *Journal of Finance* 64: 1517-1548.

Stein, Jeremy C. 2011. "Monetary Policy as Financial-Stability Regulation." *Quarterly Journal of Economics* forthcoming.

Taylor, John B. 1993. "Discretion versus Policy Rules in Practice." *Carnegie-Rochester Series on Public Policy* 39: 195-214.

Taylor, John B. 1999. "A Historical Analysis of Monetary Policy Rules." In *Monetary Policy Rules*, edited by John B. Taylor. University of Chicago Press.

Tucker, Paul. 2009. "The Repertoire of Official Sector Interventions in the Financial System: Last Resort Lending, Market-Making, and Capital." Remarks at the Bank of Japan 2009 International Conference, May 27-28.

Woodford, Michael. 1995. "Price Level Determinacy without Control of a Monetary Aggregate." *Carnegie-Rochester Conference Series on Public Policy* 43: 1-46.

Woodford, Michael. 2011. "Monetary Policy and Financial Stability." Presentation at NBER Summer Institute, Cambridge, MA, July.