

# **The Optimal Conduct of Monetary Policy With Interest on Reserves<sup>\*</sup>**

Anil K Kashyap  
Chicago Booth School of Business and NBER

Jeremy C. Stein  
Harvard University and NBER

December 2010

---

<sup>\*</sup> We are grateful to Daniel Xie for research assistance, and to Marvin Goodfriend, Sam Hanson and Elizabeth Klee for helpful discussions.

## I. Introduction

In October of 2008, the U.S. Federal Reserve announced that it would begin to pay interest on depository institutions' required and excess reserve balances, having just been authorized by Congress to do so. The Fed thereby joined a large number of other central banks who were already making use of interest on reserves (IOR) prior to the onset of the global financial crisis. Given the Fed's current policy of keeping the federal funds rate near zero, IOR has not been a quantitatively important tool thus far; as of this writing, the rate being paid is only 25 basis points. However, IOR may turn out to be extremely useful going forward, given the expansion of the Fed's balance sheet as a result of its quantitative easing policies.

To see why, consider the following decomposition for the nominal federal funds rate, denoted by  $i$ :

$$i = r_{IOR} + y_{SVR} \tag{1}$$

where  $r_{IOR}$  is the level of interest on reserves, and the  $y_{SVR}$  is the scarcity value of reserves—i.e., the convenience yield that accrues to the marginal holder of reserves. This decomposition implies that there are two ways for the Fed to start raising the funds rate once it decides that economic conditions warrant such an increase: i) it can increase  $r_{IOR}$ ; or ii) it can drain reserves from the system to the point that they become scarce again, so that  $y_{SVR}$  rises above zero.

The latter approach is only viable to the extent that the Fed is willing to greatly reduce the size of its balance sheet. As of November 2010, total reserves were approximately \$1 trillion; this compares to a pre-crisis value (in mid-2007) of approximately \$40 billion. Thus to make reserves sufficiently scarce that  $y_{SVR}$  once again becomes meaningfully positive would require draining almost \$1 trillion of reserves from the system. Absent other forms of non-reserve financing (such as repurchase agreements), this would require the Fed to make a

commensurate reduction in its holdings of assets such as mortgage-backed securities and long-term Treasury bonds—in other words, to unwind much of its policy of quantitative easing.

By contrast, simple arbitrage logic suggests that, no matter how large the quantity of reserves in the system, the funds rate should be bounded below by  $r_{IOR}$ —for if it were not, a bank could make a riskless profit by borrowing in the federal funds market at  $i$ , and depositing the reserves at the Fed to earn  $r_{IOR}$ .<sup>1</sup> Thus the ability to pay IOR should allow the Fed to decouple its funds-rate policy from the management of its asset holdings. For this reason, it seems likely that IOR will play an important role in monetary policy over the next several years.

In this paper, we take up a longer-run question about the use of IOR as a policy tool. Suppose that we reach a point where the Fed has reduced the asset side of its balance sheet to roughly pre-crisis levels, so that large values of reserves are no longer needed simply as a means of financing its investments. At this point, the Fed will once again have a choice as to how to manipulate the funds rate—whether, at the margin, it should do so by altering  $r_{IOR}$  or by varying the quantity of reserves in the system so as to induce movements in  $y_{SVR}$ . Given this choice, what is the optimal way to proceed? When the Fed seeks to tighten monetary policy, should it raise the rate paid on reserves, contract the quantity of reserves, or some combination of the two?

This question can be further motivated by observing the diversity of central-bank practices before the 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 from quantity-mediated

---

<sup>1</sup> In practice, there have been some small deviations from this arbitrage relationship—there have been instances where the funds rate has fallen below  $r_{IOR}$ . These deviations appear to reflect two technical complications. First, some holders of reserves—notably the government-sponsored enterprises—are not eligible to receive interest on reserves. And second, the arbitrage that involves banks borrowing large quantities of reserves and holding them on their balance sheets may bump up against a regulatory leverage constraint, which specifies an upper limit on total bank assets relative to regulatory capital. We abstract from these issues in what follows, but see Bech and Klee (2010) for a detailed treatment.

changes in  $y_{SVR}$ . At the other extreme was the 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. In such systems the quantity of reserves is typically 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.<sup>2</sup>

Note that 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. So it is natural to ask whether this approach can be thought of as an optimal one, and if so, under what circumstances.

A handful of other recent papers have addressed the same question that we do here. The general consensus in this literature 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 (2010)—is an application of the so-called “Friedman rule” (Friedman 1959, 1969): to the extent that central-bank reserves are a valuable transactions medium, they should be made available in

---

<sup>2</sup> See Keister, Martin and McAndrews (2008) 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.<sup>3</sup>

Our perspective here is quite different. We argue that in general, it will be optimal for the central bank to take advantage of both tools at its disposal by varying both  $r_{IOR}$  and  $y_{SVR}$ , with the mix depending on conditions in the real economy and in financial markets. The two-tools argument begins with the premise that monetary policy may have an important financial-stability role in addition to its familiar role in managing the inflation-vs.-output tradeoff. In New Keynesian models where the central bank's only job is to minimize deviations of inflation and output from desired targets, it has just one interest-rate dial to turn—a dial which it is usually assumed to adjust by following a version of the “Taylor rule” (Taylor 1993, 1999).<sup>4</sup> However, some recent research has argued that not only should central banks also be mindful of financial stability, but that *monetary policy per se* should be a key weapon in the financial-stability arsenal, above and beyond any forms of prudential regulation, such as bank capital requirements. For example, Adrian and Shin (2008) suggest that tight monetary policy should be used to head off credit booms that involve dangerous growth of intermediary balance sheets.<sup>5</sup>

If one accepts Adrian and Shin's (2008) prescription, the appeal of a second monetary-policy instrument becomes apparent: it would be difficult to manage both the inflation-output

---

<sup>3</sup> Curdia and Woodford (2010) 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...”

<sup>4</sup> 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.

<sup>5</sup> Adrian and Shin (2008) state: “In conducting monetary policy, the potential for financial sector distress should be explicitly taken into account in a forward-looking manner.”

tradeoff, and financial stability, with just a single instrument. While this logic opens the door to considering a two-instrument regime in which both  $r_{IOR}$  and  $y_{SVR}$  are actively used, it does not by itself make the case. Assuming that the nominal funds rate  $i$  is still used to address the inflation-output tradeoff (as in a Taylor rule), one needs to go further and explain why variation in  $y_{SVR}$ —holding fixed  $i$ —is an effective method for achieving specific financial-stability goals.

Stein (2010) develops the logic for using  $y_{SVR}$  to pursue financial-stability objectives. His model has three key features. First, the central bank is purely a financial regulator; it does not need to worry about the price level, which is pinned down by fiscal considerations.<sup>6</sup> Second, and relatedly, the central bank does not pay IOR, so it only has one tool, the nominal funds rate, which is thus equivalent to  $y_{SVR}$ . And finally, there is an externality in the capital-structure choices made by financial intermediaries—they have a tendency to take on too much short-term debt, because they do not fully internalize the social costs that short-term financing can create in an adverse state of the world.

In this setting, if reserve requirements are applied to intermediaries' short-term liabilities, changes in the quantity of reserves (and hence in  $y_{SVR}$ ) can be used to implement a socially efficient, time-varying regime of Pigouvian taxes.<sup>7</sup> 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

---

<sup>6</sup> On the fiscal theory of the price level, see Leeper (1991), Sims (1994), Woodford (1995), and Cochrane (1998). An alternative interpretation of Stein's (2010) model is that the price level is fixed by a commodity standard.

<sup>7</sup> Other authors have also pointed out that Pigouvian taxes on the debt issued by financial intermediaries might be helpful in curbing systemic externalities (Jeanne and Korinek (2010), Kocherlakota (2010), Perotti and Suarez (2010)). What is distinctive in Stein (2010) is an emphasis on the fact that optimal Pigouvian taxes are likely to be time-varying, and to depend on information that is not available to the regulator. In order to elicit this information, it helps to have a market price, as in a cap-and-trade system. Monetary policy is one way of implementing such a cap-and-trade system. Much of the remainder of the paper is devoted to illustrating this point.

otherwise harmful externality, and changes in the reserves tax over time represent optimal fine-tuning of this regulatory mechanism.

In what follows, we use a simplified adaptation of Stein's (2010) model to illustrate how monetary policy might work in an alternative, more realistic environment where, in addition to this regulatory objective, the central bank also has responsibility for price stability, i.e., for managing the inflation-output tradeoff. We do not model this latter objective explicitly, but rather just assume that the central bank pursues it by targeting the funds rate  $i$ , as in a Taylor rule. Then, given the assumption of non-zero IOR, we explore how the central bank can independently vary  $y_{SVR}$  in order to achieve its regulatory goals.

We begin in Section II by deriving the optimal regime of Pigouvian taxes levied on short-term debt claims issued by financial intermediaries. In Section III, we show how this optimal tax regime can be implemented with a system of reserve requirements, coupled with time-variation in the quantity of reserves, and hence in  $y_{SVR}$ . In Section IV, we discuss a number of practical issues that arise when monetary policy takes on a financial-regulatory orientation. For example, a key implication of our analysis is that for the regulatory dimension of monetary policy to be fully effective, reserve requirements would have to be substantially broadened in their application—i.e., they should cover essentially all forms of short-term debt issued by financial firms. Section V concludes.

Before proceeding, we should just underscore that our analysis of monetary policy in what follows is entirely normative, as opposed to positive. That is, we offer a framework for thinking about how monetary policy might optimally be conducted when the central bank can pay interest on reserves. We do not intend to suggest that our model describes how most central banks currently behave.

## II. A Simple Model of Pigouvian Taxes on Short-Term Bank Debt

The model that we develop here is extremely reduced-form in nature; see Stein (2010) for a version with more explicit microfoundations. It is also close in spirit to Perotti and Suarez (2010). There is a continuum of intermediaries with total mass normalized to one, who engage in maturity-transformation activities, i.e., who make long-term loans financed with short-term debt. We refer to these intermediaries as “banks” in what follows, but it is important to recognize that they need not be traditional commercial banks—the activity that we model could equally well take place in the shadow-banking sector.

There are three dates, 0, 1 and 2. At time 0, each bank is endowed with a fixed amount of initial assets, which it simply owns, without having to finance—these assets can be thought of as comprising its capital base. In addition to a liquidating time-2 payout, these assets also yield an interim dividend at time 1. The time-1 dividend is uniformly distributed on the interval  $[0, K]$ .

Given its existing capital base, a bank’s only choice at time 0 is how aggressively to engage in new maturity-transformation activities. In particular, it chooses a variable amount of investment  $I$  in illiquid assets. If held until time 2, these assets yield an expected payoff of  $\theta I$ . Banks finance a fixed fraction  $m$  of any new investment with short-maturity debt that is due at time 1, and the remaining  $(1 - m)$  with long-term debt. Short-term debt is cheaper than long-term debt: the gross two-period real interest rate on long-term debt is  $R$ , while the expected cost associated with rolling over short-term financing is  $(R - \Delta)$ .<sup>8</sup> Thus the payoff to doing maturity transformation at scale  $I$  is given by  $I(\theta - R + m\Delta)$ . Clearly, maturity transformation is more

---

<sup>8</sup> The spread  $\Delta$  reflects a particular violation of the Modigliani-Miller (1958) theorem. In Stein (2010), this violation arises from the “money-like” properties of short-term bank debt—the fact that some amount of short-term debt can be made riskless, and that investors place an extra value on riskless claims because they make a particularly good transactions medium (see also Gorton (2010) and Gorton and Metrick (2010)). And the parameter  $m$  emerges endogenously, as the upper bound on the amount of safe short-term debt that a bank can issue for a given amount of “collateral”  $I$ .



attractive when banks have better investment opportunities ( $\theta$  is higher) or when the spread between short and long-term financing costs is bigger ( $\Delta$  is higher).

However, there is also a downside to maturity transformation. With probability  $p$ , the economy falls into a bad state at time 1, such that the long-term solvency of the banks is called into question. If the bad state hits, a bank becomes distressed at time 1 if it cannot service its short-term debt from the dividends on its pre-existing assets. The idea is that when solvency is threatened, debt overhang (as in Myers (1977)) makes it impossible for a bank to raise new funds at time 1, so it faces a liquidity crisis if it has a cash shortage relative to its current obligations. Thus conditional on the bad state, and given the uniform distribution of time-1 dividends on the pre-existing assets, the probability that a bank experiences distress is given by  $mI/K$ , which is a proxy for its ratio of short-term debt to capital. Simply put, doing a lot of maturity transformation increases the risk that a bank falls into distress if the economy hits the bad state.

Our key systemic-risk assumption is that for any bank, the deadweight costs of distress are increasing in the fraction of the overall banking system that is in distress at the same time. Since the probability of distress for any one bank is a linear function of its level of  $I$ , and since there is a continuum of banks, the fraction of the banking system that is in distress in the bad state will depend on  $\bar{I}$ , the average level of maturity-transformation activity across all banks in the economy. Thus a simple formulation of the systemic-risk effect is to assume that for each bank, the deadweight costs of distress are given by  $Z = \gamma \bar{I}$ .<sup>9</sup> Pulling it all together, we have the following expression for the net expected profits  $\Pi$  of a representative bank:

---

<sup>9</sup> Again, this is a very reduced-form modeling approach. Stein (2010) provides microfoundations for something along these lines, based on a fire-sales mechanism. See also Shleifer and Vishny (2010), and Diamond and Rajan (2010).

$$\Pi = I(\theta - R + m\Delta) - \frac{pmI\gamma\bar{I}}{K} \quad (2)$$

The source of the externality in this setting—and hence the rationale for regulation—flows from the fact that in maximizing its expected profits, each individual bank takes  $\bar{I}$  as fixed, thereby disregarding its contribution to aggregate systemic risk. This implies that each bank will continue to increase investment so long as  $(\theta - R + m\Delta) > \frac{pm\gamma\bar{I}}{K}$ . Given that all banks are symmetric, it follows that, absent regulation, the equilibrium level of activity  $I^*$  for each bank is given by:

$$I^* = (\theta - R + m\Delta) \frac{K}{pm\gamma} \quad (3)$$

In contrast, the social planner's solution, which does not take  $\bar{I}$  as fixed, involves a lower level of activity  $I^{**}$ , given by:

$$I^{**} = (\theta - R + m\Delta) \frac{K}{2pm\gamma} \quad (4)$$

One way to achieve the socially optimal outcome is with a system of Pigouvian taxes on short-term debt. If the tax rate on short-term debt is  $\tau$ , the net-of-tax profits of a representative bank become:

$$\Pi = I(\theta - R + m(\Delta - \tau)) - \frac{pmI\gamma\bar{I}}{K} \quad (5)$$

And the corresponding private-market outcome is altered to:

$$I^* = (\theta - R + m(\Delta - \tau)) \frac{K}{pm\gamma} \quad (6)$$

Thus the optimal tax  $\tau^{**}$  is just the value of  $\tau$  that equates the expressions in (6) and (4):

$$\tau^{**} = \frac{(\theta - R + m\Delta)}{2m} \quad (7)$$

The key thing to note here is that the optimal tax depends on the appeal to banks of undertaking maturity-transformation activities. When there is an increase in either banks' investment opportunities, (as measured by  $\theta$ ), or their long-short financing spreads, (as measured for by  $\Delta$ ), the optimal tax goes up. This suggests that not only will the optimal tax tend to vary over time with economic and financial-market 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  $\theta$  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 from (4) it follows that the optimal quantity  $I^{**}$  depends on the same privately-observed parameter.<sup>10</sup>

### III. Monetary Policy as a Cap-and-Trade System

Stein (2010) argues that these informational difficulties can be addressed with a system of cap-and-trade.<sup>11</sup> Suppose a regulator endows banks with permits that allow it to issue short-term debt—for example, each permit might allow a bank to issue one unit of short-term debt, or equivalently, to do  $1/m$  units of new investment. Suppose further that these permits can be freely traded among banks, and are issued in total quantity  $Q$ . Then for any trial value of  $Q < I^* / m$ ,

---

<sup>10</sup> 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 analysis suggests, the big stumbling block for this approach is that it is very difficult for regulators to pick the right level of the cap.

<sup>11</sup> See also Kashyap and Stein (2004) for a related discussion about the benefits of cap-and-trade in implementing capital regulation for banks.

i.e. for any trial value of  $Q$  such that banks in aggregate are held below their privately optimal level of  $I$ , the permits will trade for a non-zero price  $P(Q)$  that reflects the shadow value of the constraint:<sup>12</sup>

$$P(Q) = \frac{1}{m} \frac{d\Pi}{dI} = \frac{(\theta - R + m\Delta)}{m} - \frac{p\gamma Q}{K} \quad (8)$$

Thus given the regulator's knowledge of  $Q$ , he can simply read off from the market price of permits the key item of interest, namely  $(\theta - R + m\Delta)$ . 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  $(\theta - R + m\Delta)$ , he can adjust the aggregate quantity of permits in the system to their *optimal* level  $Q^{**} = I^{**} / m$ . Moreover, at this optimal quantity, it is easily checked that the price of the permits equals the optimal Pigouvian tax, i.e. that  $P(Q^{**}) = \tau^{**}$ . 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  $(\theta - R + m\Delta)$  has gone up, and the he should therefore relax the cap by injecting more permits into the system.

Interestingly, in this simple formulation of the model, the optimal regulatory policy is *partially accommodative* with respect to shocks in either  $\theta$  or  $\Delta$ : when either of these variables increases, 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. This is evident from inspection of equations (4) and (7).

---

<sup>12</sup> To be consistent with respect to time discounting, we need to assume that the permits are paid for at time 2, when the bank's profits are realized.

Stein (2010) notes that one particular way to implement a cap-and-trade system is by: (i) imposing reserve requirements on banks' short-term liabilities; (ii) controlling the total quantity of reserves in the financial system, and (iii) allowing these reserves to be freely traded in a federal funds market. In this case, the reserves are exactly analogous to permits, and there is a direct relationship between the scarcity value of reserves  $y_{SVR}$  and the price of permits  $P$ .

To be more precise, note that so far everything in the model has been expressed in real terms. In order to introduce nominally-denominated reserves, suppose that the price level is pinned down outside the model, say by the usual factors in a New Keynesian model. For simplicity, let the price level at time 0 be normalized to unity. This means that any bank wishing to issue one unit of short-term debt (in real terms) must hold  $\rho$  dollars of nominal reserves, where  $\rho$  is the fractional reserve requirement. Hence the *net* amount of short-term debt financing made possible by one dollar of reserves is  $(1 - \rho)/\rho$ .<sup>13</sup> It follows that if the central bank creates a amount of reserves given by  $\mathfrak{R}$ , the effective quantity of permits  $Q$  available to the banking sector is given by:

$$Q = \frac{(1 - \rho)\mathfrak{R}}{\rho} \tag{9}$$

The mapping between the permit price  $P$  and the scarcity value of reserves  $y_{SVR}$  can be understood as follows. If a bank wants to increase its net issuance of short-term debt by one unit, it must hold an incremental quantity  $\rho/(1 - \rho)$  of reserves at time 0. Since it borrows these reserves at a nominal rate of  $i$ , but only earns  $r_{IOR}$ , it has to pay a nominal financing fee at time 2

---

<sup>13</sup> As an example, suppose  $\rho = .10$ . In this case, with one dollar of reserves, a bank is allowed to raise 10 dollars of short-term debt. But given that it must hold the reserves as an asset, only 9 of these dollars represent *net* financing that is available to fund new loans.

of  $y_{SVR}$  per dollar of reserves held, or  $\frac{\rho y_{SVR}}{(1-\rho)}$  in total. Allowing for the possibility of inflation between time 0 and time 2, and denoting the time-2 price level by  $\pi$  (recall that the time-0 price level is one) the real value of the financing fee is  $\frac{\rho y_{SVR}}{\pi(1-\rho)}$ .<sup>14</sup> This is equivalent to the permit price  $P$ , as it represents the real cost that a bank must pay at time 2 in order to increase its short-term debt by one unit.

Given this equivalence, we can write the optimal value of  $y_{SVR}$  as:

$$y_{SVR}^{**} = \frac{(1-\rho)\pi}{\rho} \tau^{**} \quad (10)$$

where  $\tau^{**}$  is the optimal Pigouvian tax from equation (7) above.

**Example:** Suppose the parameters of the model are such that the optimal tax on short-term debt is 25 basis points. If the reserve requirement  $\rho$  is 10 percent, and  $\pi = 1.02$ , corresponding to inflation of 2 percent, then  $y_{SVR}^{**} = 2.30$  percent. So if the Taylor rule calls for the nominal funds rate  $i$  to be set at 5.0 percent, it would be optimal to have  $r_{IOR} = 2.70$  percent, and do the rest of the job by making reserves sufficiently scarce. If the reserve requirement  $\rho$  is raised to 20 percent, then  $y_{SVR}^{**}$  falls to 1.02 percent.

## IV. Discussion

### A. Breadth of Reserve Requirements

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

---

<sup>14</sup> Since the nominal funds rate is  $i$ , and one plus the real short rate is  $(R - \Delta)$ , it must be that  $\pi = (1 + i) / (R - \Delta)$ .

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 (2010)), it would also be desirable to regulate the shadow-banking sector in a symmetric fashion. This would suggest 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.

### *B. Level of Reserve Requirements*

The theory sketched above yields a unique optimum for the Pigouvian tax  $\tau^{**}$  at any point in time. However, as equation (10) makes clear, it is possible to achieve a given value of  $\tau$  with

monetary-policy tools in one of two ways: either by adjusting  $y_{SVR}$  or by adjusting the reserve requirement  $\rho$ . Thus, as in the example above, one can set  $\tau$  to 25 basis points either with  $\rho = 10$  percent, and  $y_{SVR} = 2.30$  percent, or with  $\rho = 20$  percent, and  $y_{SVR} = 1.02$  percent.

Is there any reason to prefer one combination over the other? One appeal of using 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  $i$  is determined to be 2.0 percent, and  $\tau^{**} = 25$  basis points. With  $\rho = 10$  percent, it is impossible to accomplish both objectives, since we would need  $y_{SVR} = 2.30$  percent to hit the regulatory target, which obviously is incompatible with having  $i = 2.0$  percent. By contrast, with  $\rho = 20$  percent, we can set  $y_{SVR} = 1.02$  percent, and meet both goals simultaneously.

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  $\rho$ . Interestingly, the Chinese central bank (the People’s Bank of China, or PBOC) uses reserve requirements as a key policy tool. For example, the PBOC changed the level of reserve requirements six times in 2010, while moving their policy interest rate just once. This sort of behavior is hard to understand in the context of a standard New Keynesian model where the policy rate fully summarizes the stance of monetary policy, and where quantitative measures such as reserve requirements are irrelevant once one conditions on the policy rate. However, it may reflect a regulatory motive similar to the one we have outlined here, namely a desire to control certain forms of short-term-debt-financed credit creation.<sup>15</sup>

---

<sup>15</sup> See Du (2010) for a detailed discussion of the PBOC’s policies and procedures.



The proposal 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—in our setting, 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  $\rho$  are exactly offset by reductions in  $y_{SVR}$ , or alternatively, by increases in  $r_{IOR}$ .

### *C. Regulatory Arbitrage*

Another skeptical reaction 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.<sup>16</sup>

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.

#### *D. Monetary Policy and Credit Bubbles*

In the simple model developed above, the only divergence between individual banks and the social planner is that each bank takes  $\bar{I}$  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 can be credit bubbles—periods when private lenders make loans that would appear to have abnormally low expected returns. Greenwood and Hanson (2010) 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.

---

<sup>16</sup> 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 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, which we now denote by  $\theta^B$ , and that perceived by the regulator, which we denote by  $\theta^R$ . Thus a credit bubble can be thought of as a period when  $\theta^B > \theta^R$ , i.e., when banks are excessively bullish on their investment prospects.

It is easy to see that this case works very similarly to that studied above, with the optimal Pigouvian tax being modified to:

$$\tau^{**} = \frac{(2\theta^B - \theta^R - R + m\Delta)}{2m} \quad (11)$$

To put some further structure on the problem, we might posit that  $\theta^R = \alpha + \beta\theta^B$ , with  $\beta < 1$ . If we continue to assume that  $\theta^B$  is private information of the banks, this formulation captures the intuition that: i) the regulator still has something to learn about fundamentals from banks; and ii) yet when the banks are most optimistic about fundamentals, the regulator believes them to be overoptimistic. The formula for the optimal tax now becomes:

$$\tau^{**} = \frac{((2 - \beta)\theta^B - \alpha - R + m\Delta)}{2m} \quad (12)$$

As compared to equation (7), it is evident that the optimal tax now involves a higher coefficient on  $\theta^B$  than in the baseline version of the model. In the monetary-policy implementation, this implies a policy that is less accommodative than before in the face of incipient increases in  $y_{SVR}$ . As  $y_{SVR}$  begins to rise, the central bank now injects fewer additional reserves into the system, and allows  $y_{SVR}$  to go up further than it would have in the no-credit-bubble case.

We are certainly not the first to suggest that tight monetary policy might be used to try to rein in credit bubbles.<sup>17</sup> However, our model is more precise 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: as can be seen in equation (12), 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, increasing the funds rate via the IOR channel does nothing to address the externality.

## V. Conclusions

The introduction of interest on reserves gives the Federal Reserve a second monetary policy tool that, used properly, may prove helpful for financial-stability purposes. By adjusting both IOR and the quantity of reserves in the system, the Fed can simultaneously pursue price stability, as well as an optimal regime of regulating the externalities created by short-term bank debt. Though to be clear, the latter would also require, in addition to the use of IOR, a significant expansion in the coverage of reserve requirements, as well possibly an adjustment to their level.

Our analysis puts 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.

---

<sup>17</sup> Again, see Adrian and Shin (2008) and the references therein.

## References

- Adrian, Tobias, and Hyun Song Shin, 2008, "Financial Intermediaries, Financial Stability and Monetary Policy," in *Maintaining Stability in a Changing Financial System*, Federal Reserve Bank of Kansas City, 287-334.
- Basel Committee on Banking Supervision, 2010, "The Basel Committee's Response to the Financial Crisis: Report to the G20," < <http://www.bis.org/publ/bcbs179.pdf>>
- Bech, Morten, and Elizabeth Klee, 2010, "The Mechanics of a Graceful Exit: Interest on Reserves and Segmentation in the Federal Funds Market," Federal Reserve Board working paper.
- Cochrane, John, 1998, "A Frictionless Model of U.S. Inflation," *NBER Macro Annual*, 323-384.
- Curdia, Vasco, and Michael Woodford, 2010, "The Central-Bank Balance Sheet as an Instrument of Monetary Policy," NBER working paper 16208.
- Diamond, Douglas W., and Raghuram G. Rajan, 2010, "Fear of Fire Sales, Illiquidity Seeking and Credit Freezes," working paper.
- Du, Yang, 2010, "Monetary Policy and Bank Loan Supply in China," working paper.
- 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 (4), 25-45.
- Goodfriend, Marvin, 2002, "Interest on Reserves and Monetary Policy," Federal Reserve Bank of New York Economic Policy Review 8 (1), 13-29.
- Gorton, Gary B., 2010, *Slapped by the Invisible Hand: The Panic of 2007*, Oxford University Press.
- Gorton, Gary B., and Andrew Metrick, 2010, "Securitized Banking and the Run on Repo," *Journal of Financial Economics*, forthcoming.
- Greenwood, Robin, and Samuel Hanson, 2010, "Issuer Quality and Corporate Bond Returns," working paper.
- Jeanne, Olivier, and Anton Korinek, 2010, "Managing Credit Booms and Busts: A Pigouvian Taxation Approach," working paper.
- 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.

- 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.
- Kocherlakota, Narayana, 2010, "Taxing Risk and the Optimal Regulation of Financial Institutions," Economic Policy Paper 10-3, Federal Reserve Bank of Minneapolis.
- Leeper, Eric, 1991, "Equilibria Under Active and Passive Monetary Policies," *Journal of Monetary Economics* 27, 129-147.
- Modigliani, Franco and Merton H. Miller, 1958, "The Cost of Capital, Corporation Finance and the Theory of Investment," *American Economic Review* 48, 261-297.
- Myers, Stewart C., 1977, "Determinants of Corporate Borrowing," *Journal of Financial Economics* 5, 147-175.
- Perotti, Enrico, and Javier Suarez, 2010, "Regulation of Liquidity Risk," working paper.
- Shleifer, Andrei, and Robert W. Vishny, 2010, "Unstable Banking," *Journal of Financial Economics*, forthcoming.
- 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., 2010, "Monetary Policy as Financial-Stability Regulation," working paper.
- Taylor, John B., 1993, "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Series on Public Policy* 39, 195-214.
- Taylor, John B., 1999, "An Historical Analysis of Monetary Policy Rules," in *Monetary Policy Rules*, edited by John B. Taylor, University of Chicago Press.
- Woodford, Michael, 1995, "Price level Determinacy Without Control of a Monetary Aggregate," *Carnegie-Rochester Conference Series on Public Policy* 43, 1-46.