A Macro-Prudential Framework for Liquidity Regulation

Arvind Krishnamurthy, Kellogg School, Northwestern University and NBER

10/15/2011

This note discusses liquidity mismatch regulation. For purposes of this note, think of a financial institution that owns an illiquid asset that is funded by a liquid liability (i.e. overnight debt) and earns a liquidity premium on running this mismatch. Liquidity regulation is about targeting the quantity of this mismatch that is taken at the institution level and at the financial sector level (i.e., aggregated across institutions).

In the wake of the financial crisis, there has been considerable interest among policymakers and academics in regulations that enable macroeconomic and financial stability. Thus far, most of the discussion has centered on capital regulation, with liquidity regulation receiving less attention. This is because frameworks for macro-prudential capital regulation are well developed, with many of the main issues well understood. But, it is important to not diminish the importance of liquidity regulation.[1] Indeed, there are reasons why it may be more beneficial to implement liquidity regulations than capital regulations.

All regulation creates a migration problem to unregulated sectors. The tradeoff in financial stability due to increasing regulation has to balance against this migration cost. The migration problem is less severe for liquidity regulations than for capital regulations. This is because unregulated entities inherently have unstable funding so that unregulated institutions are less able to intermediate illiquid assets and earn the liquidity premium. For example, if one looks across the hedge fund sector, the main strategy that owns illiquid assets funded by short term debt is fixed income, which is on the order of a few $100s of billions. This is relative to a regulated sector that runs a mismatch in the $10s of trillions. Another way of saying this is that institutions that are regulated get a “carrot” – access to central bank and other government funding facilities that enables them the latitude to run a liquidity mismatch – which can balance the “stick” of higher liquidity regulations. Capital regulation involves the stick, but activities that require capital can be as well done by regulated institutions as unregulated institutions.

[1] Arguably, the most dramatic episode of the recent crisis in the fall of 2008 was due to liquidity problems and not capital problems.

Liquidity Measurement

Liquidity regulation requires measurements of "liquidity." I lay out what these measures should look like, but considerable work remains to be done in optimizing the measurements.

a. It is necessary to measure, at an institution level, the dollar amount of liquidity mismatch. Brunnermeier, Gorton and Krishnamurthy (2011) propose such a measure. Briefly, the mismatch is measured by summing across assets multiplied by asset-specific "liquidity" weights and netting the sum against a similarly constructed measure from liabilities, including contingent liabilities such as derivatives, which are likewise "liquidity"-weighted. It is important to measure the mismatch in dollar terms, because only such a measure can be aggregated across institutions. The Brunnermeier, Gorton and Krishnamurthy scheme is similar in many ways to the liquidity ratio measures proposed in Basel 3, but with the key difference being that ratios do not aggregate. Let us denote the aggregate mismatch across the financial sector by $\mathcal{L}^{\text{private}}$.
b. It is also necessary to construct a comprehensive measure of a liquidity premium. Drawing from the finance literature, such a measure can be constructed from bond market spreads such as on-the-run/off-the-run Treasury spreads, Aaa-Treasury spreads, Agency-Treasury spreads, Swap spreads, the TED spread, Treasury yield curve slope at short-maturities, as well as from equity market measures of liquidity (bid-asks, turnover, etc). Obviously there is much work that needs to be done to settle on the right measure of a premium. Call this premium, $\mathbf{P}$.

Equilibrium and Empirical Relations

There are equilibrium relations, analogous to money demand relations, which describe appropriately constructed $\mathcal{L}^{\text{private}}$ and $\mathbf{P}$. Following Krishnamurthy and Vissing-Jorgensen (2010), households and corporations have a well-defined demand for financial liquidity that is satisfied by $\mathcal{L}^{\text{private}}$ and $\mathcal{L}^{\text{public}}$. The largest item in the latter category is U.S. Treasury securities. We can describe the liquidity market equilibrium as,

$$\mathcal{L}^{\text{demand}}(\mathbf{P}) = \mathcal{L}^{\text{private}}(\mathbf{P}) + \mathcal{L}^{\text{public}},$$

where, $\mathcal{L}^{\text{demand}}$ is decreasing in $\mathbf{P}$ while $\mathcal{L}^{\text{private}}$ is increasing in $\mathbf{P}$. That is, the financial sector has an incentive to increase its liquidity mismatch when the liquidity premium, $\mathbf{P}$, is high.

Here is some empirical evidence, drawn from Krishnamurthy and Vissing-Jorgensen (2010), to support the equilibrium relations that I have asserted.

a. Evidence that is $\mathcal{L}^{\textbf{demand}}(\mathbf{P})$ quantitatively large and decreasing in $\mathbf{P}$.

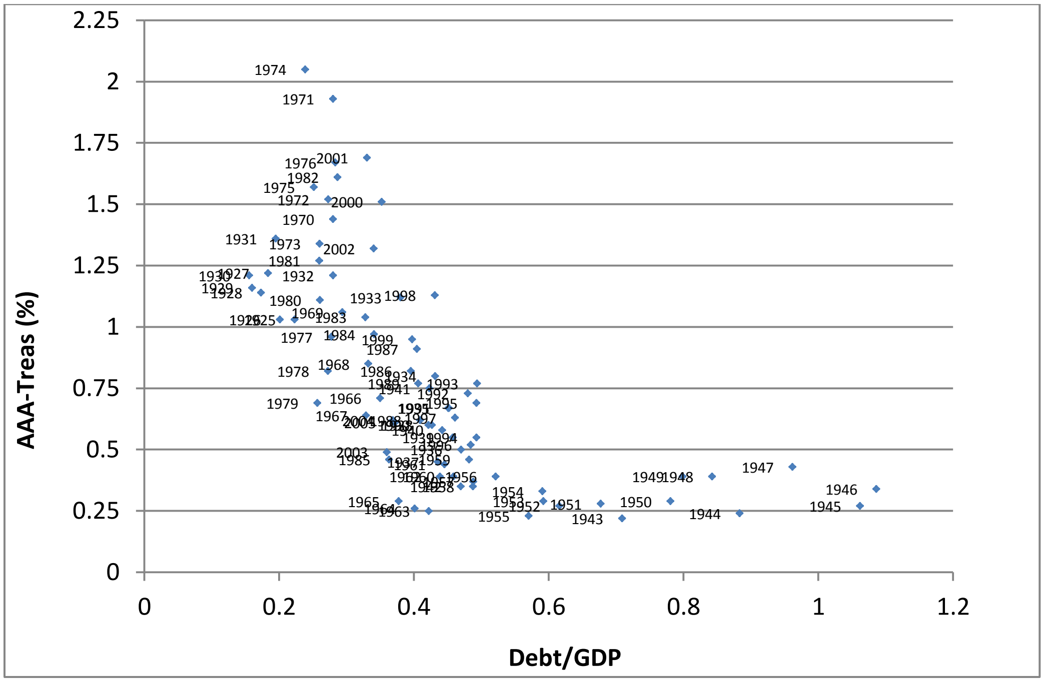

The figure is of the yield spread between AAA rated corporate bonds and US Treasury bonds against the US Debt/GDP ratio, using annual data from 1925 to 2005. When the stock of debt is low, the spread is high. The spread data in the figure is measured using long maturity bonds, but the relation holds as well based on short maturity bond spreads.

One can think of the figure as tracing out the demand function $\mathcal{L}^{\textbf{demand}}(\mathbf{P})$ by exploiting variation in the government supply $\mathcal{L}^{\textbf{public}}$ and using the Aaa-Treasury spread as a measure of the liquidity premium, $\mathbf{P}$. Thus the figure shows that the liquidity demand function is downward sloping in the liquidity premium. The figure also shows the premium is positive for values of Debt/GDP around 0.4, or near $6 trillion today. That is, investors' liquidity demand is for at least $6 trillion of liquidity (at least, because investors also hold $\mathcal{L}^{\textbf{private}}$).

b. Evidence that $\mathcal{L}^{\textbf{private}}$ is large and increasing in $\mathbf{P}$.

The next figure graphs M2 minus M1 scaled by GDP against the US Debt/GDP ratio, based on annual data from 1920 to 2005. The money aggregate thus includes savings deposits, small time deposits and money market deposit accounts. The figure shows that when there is less Treasury debt outstanding, and hence $\mathbf{P}$ rises, the banking sector responds to the increased $\mathbf{P}$ by increasing deposit liabilities and earning the liquidity premium.

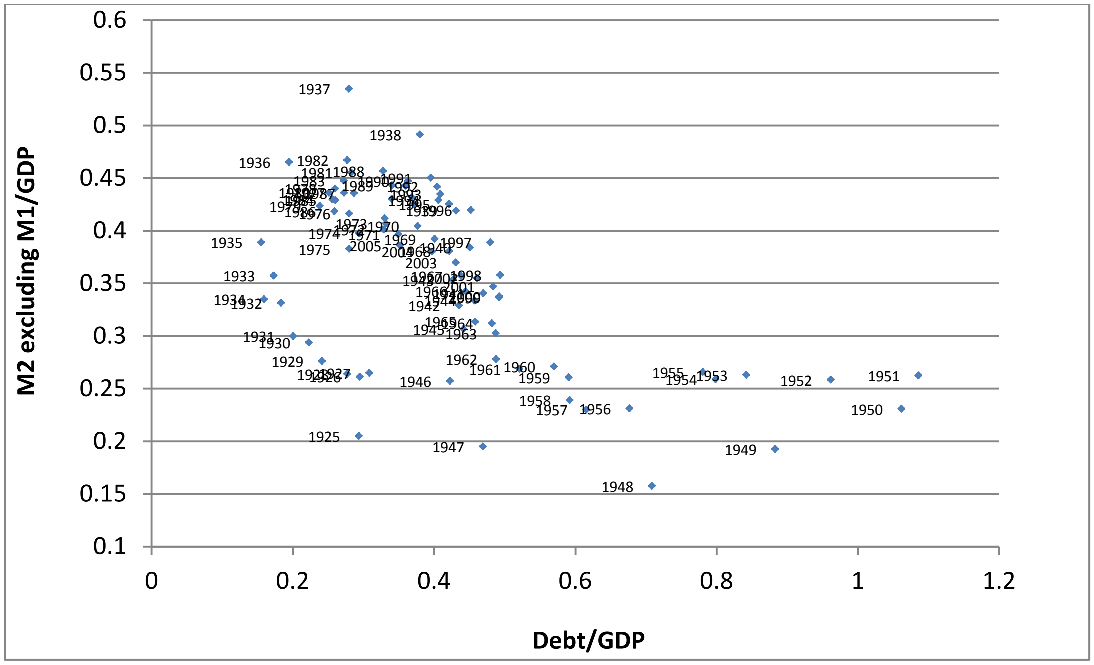

The liquidity mismatch measure in this figure is bank deposits because there is long time-series data on bank deposits. As we have discussed, an appropriate measure of liquidity mismatch is more complex than simply summing bank deposits. For example, $\mathcal{L}^{\text{private}}$ includes a weighted function of bank deposits as well as other liquidity claims issued by the financial sector, and netted against liquid assets held by the financial sector. The fact that we can see the basic supply relation even with a crude construction of liquidity mismatch suggests that a research program of accurately measuring $\mathcal{L}^{\text{private}}$ and $\mathbf{P}$ is likely to be fruitful.

<u>Liquidity Regulation</u>

The principle behind liquidity regulation is that the aggregate mismatch ($\mathcal{L}^{\text{private}}$) of the financial sector is a stability cost to the economy, and a direct potential liability to the Federal Reserve during a financial crisis in which the Federal Reserve is called to act as lender of last resort. As a result, liquidity regulation should be geared towards keeping this mismatch from growing too large (say as a ratio to GDP, or ratio to total intermediary capital). I next describe how one could regulate liquidity.

Suppose that reducing $\mathcal{L}^{\text{private}}$ increases financial stability. On the other hand, if $\mathbf{P}>0$ households value liquidity so that reducing $\mathcal{L}^{\text{private}}$ requires the household sector to manage their affairs using less liquidity. Thus there is a tradeoff that determines the optimal level of $\mathcal{L}^{\text{private}}$, which we denote $\mathcal{L}^{*}$.

How can one implement $\mathcal{L}^*$? One approach is through quantity limits as in Basel 3, adjusted to reflect that liquidity mismatch must be measured in dollar terms rather than ratios. That is, regulation can specify that institution *i* cannot exceed a liquidity mismatch of $\mathcal{L}_i^{\text{private}}$, where across all institutions $\mathcal{L}^* = \sum_i \mathcal{L}_i^{\text{private}}$.

A second approach is to use "prices" to implement the optimum. The approach, following Stein (2010), is as follows:

a. Suppose that bank-*i* is required to hold a reserve requirement at the Fed in proportion, *f*, to its liquidity mismatch, $\mathcal{L}_i^{\text{private}}$.
b. The bank earn the rate *IOR* (interest on reserves) on its reserve balance $f\,\mathcal{L}_i^{\text{private}}$.
c. The Fed sets the overnight Fed Funds rate at *FF>IOR*.
d. Then the tax on $\mathcal{L}_i^{\text{private}}$ is proportional to *f* (*FF-IOR*). Thus, a higher tax will imply that the bank will have an incentive to reduce $\mathcal{L}_i^{\text{private}}$.
e. The Fed can choose *FF* based on standard Taylor rule considerations.
f. It can independently set *(FF-IOR)* to target $\mathcal{L}^*$.
g. Alternatively, it can choose *f* to target $\mathcal{L}^*$.

Obviously this scheme is less useful if the economy is at the zero lower bound or if reserve supply is so large that *FF-IOR* is small (as is currently). On the other hand, if we envision that liquidity taxes should be procyclical, so that they are primarily used during booms when *FF* is high for standard macroeconomic reasons, then there is room to independently target *FF-IOR*. Note also that in this scheme the Fed can use the reserve ratio, *f*, as a policy tool.

A third approach to liquidity regulation is to vary $\mathcal{L}^{\text{public}}$ to target $\mathcal{L}^*$. The logic follows directly from the previous graphs: expanding Debt/GDP lowers the liquidity premium and hence lowers $\mathcal{L}^{\text{private}}$. Thus the government can alter the supply of Treasury securities to target a given $\mathcal{L}^*$.

During a non-crisis period, this approach does not seem practical because it raises the question of what the government should buy if it issues more debt. Note that it would have to purchase private sector illiquid assets, or liquidity would not expand. Thus the policies may look like "quantitative easing" with private assets, which seems less feasible outside of crisis periods.

On the other hand, during a crisis period, adjusting $\mathcal{L}^{\text{public}}$ is common central bank policy. That is, through central bank lending facilities as well as large scale quantitative easing policies, the central bank purchases illiquid assets and issues liquid liabilities.

Targeting Procedures

This section describes a rudimentary liquidity targeting regime. I am particularly interested in asking how and when the three types of policies described in the previous section (prices, quantities, and $\mathcal{L}^{\text{public}}$) can be used.

Let me start with some notation. Suppose that the household and corporate sector's valuation of liquidity can be described by a function U($\mathcal{L}$), where $\mathcal{L} = \mathcal{L}^{\text{private}} + \mathcal{L}^{\text{public}}$. That is, the private sector chooses liquidity to maximize, U($\mathcal{L}$) - $\mathbf{P}\,\mathcal{L}$, with first order condition, $\mathbf{P} = U'(\mathcal{L})$. This equation can be inverted to define the liquidity demand function $\mathcal{L}^{\text{demand}}(\mathbf{P})$ of the previous sections.

Likewise, suppose that financial institutions incur a private cost of running a liquidity mismatch given by F($\mathcal{L}^{\text{private}}$), so that an institution maximizes $\mathbf{P}\,\mathcal{L}^{\text{private}} - F(\mathcal{L}^{\text{private}})$. The first order condition for private liquidity supply is $\mathbf{P} = F'(\mathcal{L}^{\text{private}})$.

Then, the private market no-regulation equilibrium is described by,

$$U'(\mathcal{L}^{\text{private}} + \mathcal{L}^{\text{public}}) = F'(\mathcal{L}^{\text{private}}).$$

Suppose that a planner computes that the social cost of running a liquidity mismatch is given by C($\mathcal{L}^{\text{private}}$). The cost function here reflects the instability a large liquidity mismatch may create for the financial sector as well as the cost to the central bank of having to inject liquidity in a crisis to bailout the financial sector. Thus, the planner chooses $\mathcal{L}^{\text{private}}$ to solve

$$U'(\mathcal{L}^{\text{private}} + \mathcal{L}^{\text{public}}) = C'(\mathcal{L}^{\text{private}}).$$

As noted earlier, if $\mathbf{P} > 0$ (hence $U'(\mathcal{L}^{\text{private}} + \mathcal{L}^{\text{public}}) > 0$), the solution calls for $\mathcal{L}^{\text{private}} = \mathcal{L}^* > 0$.

The solution can be implemented by direct quantity regulation or by imposing the reserves tax, measured as $\boldsymbol{\tau}$, on private sector liquidity creation. In the tax case, the private sector maximizes,

$$(\mathbf{P} - \boldsymbol{\tau})\,\mathcal{L}^{\text{private}} - F(\mathcal{L}^{\text{private}}).$$

If the tax rate is chosen so that, $F'(\mathcal{L}^{\text{private}}) + \boldsymbol{\tau} = C'(\mathcal{L}^{\text{private}})$, then the private and social optimal coincide.

Case 1: Demand Shocks

Suppose first that there are investor demand shocks. Investor utility is given by $Z^D\,U(\mathcal{L})$, where $Z^D > 1$ is a shock that increases liquidity demand. An example of such a shock is a financial crisis in which there is a flight to liquidity. I contrast how price versus quantity regulation handles the increased demand shock.

In this setup, the optimal tax rate is still $\boldsymbol{\tau}$ = $C'(\mathcal{L}^{\mathbf{private}}) - F'(\mathcal{L}^{\mathbf{private}})$. Thus, if following a tax scheme, the equilibrium will (optimally) feature a higher $\mathcal{L}^*$ and $\mathbf{P}$ in response to the positive demand shock. This can be thought of as accommodating by allowing the private sector to expand private liquidity. Another way of accommodating is for the central bank to directly purchase private sector illiquid assets and create liquid government assets in exchange, thus expanding $\mathcal{L}^{\mathbf{public}}$.

If instead, the regulator imposes a quantity limit, it will have to allow $\mathcal{L}^{\mathbf{private}}$ to expand (or expand public liquidity instead). But here it faces the difficult question of determining how much expansion to allow. Moreover, since during a crisis, the demand for liquidity is likely to be relatively inelastic, the deadweight cost of providing the wrong amount of $\mathcal{L}^{\mathbf{private}}$ will be high.

Consider another type of shock which reduces liquidity demand, but may play a role during a non-crisis period. During the 2000 to 2007 period, investors may have perceived a low risk environment in which they did not need to hold much liquidity. The optimal tax in this case is still constant, so that the tax policy implements the optimum. With the constant tax, the equilibrium will feature a lower $\mathcal{L}^*$ and $\mathbf{P}$. However, unlike the flight to liquidity case, in this case since demand for liquidity is likely to be relatively elastic the deadweight cost of providing the wrong amount of $\mathcal{L}^{\mathbf{private}}$ is low. That is, quantity regulation will work as well as price regulation.

Note that this kind of shock must not have been the dominant shock in the 2000 to 2007 period because while $\mathbf{P}$ was low in this period, $\mathcal{L}^{\mathbf{private}}$ was high. As the next case makes clear, the 2000 to 2007 period was more likely reflective of the supply side.

<u>Case 2: Supply Shocks</u>

Suppose instead that there are shocks to the financial supply function, given by $Z^S F(\mathcal{L}^{\mathbf{private}})$, where $Z^S < 1$ is a shock reducing supply costs. An example of such a shock is securitization which led financial institutions to feel comfortable running a larger liquidity mismatch. I will assume that the shock does not shift the social cost function. In part, the assumption is driven by the common view that, ex-post, securitization did not decrease macroeconomic liquidity risk.

A fixed tax does badly with this shock. That is, the optimal tax is $\boldsymbol{\tau}$ = $C'(\mathcal{L}^{\mathbf{private}}) - Z^S F'(\mathcal{L}^{\mathbf{private}})$, which should rise if $Z^S < 1$ (Note that if the social cost was also falling this conclusion may not be valid.) A fixed tax allows $\mathcal{L}^{\text{private}}$ to expand and **P** to fall. If the social cost function of expanding liquidity too much is steep, then the deadweight cost of the fixed tax policy is large.

It turns out for this case that there is a version of the tax policy that can work. Suppose that the Fed changed the reserves tax to target a constant **P** (as an intermediate target). Since the social optimum still requires that $U'(\mathcal{L}^{\text{private}} + \mathcal{L}^{\text{public}}) = C'(\mathcal{L}^{\text{private}}) =$ **P,** which is not affected by the shock, this targeting procedure will implement the optimum.

The quantity target implements the optimum because $\mathcal{L}^*$ is constant across the shock.

What to do in Practice?

The general case of whether to use prices or quantities depends on the regulator's knowledge of the shocks and the slopes of the cost and demand function. This is the subject of a large literature clarified by Weitzman (1974).

In the context of liquidity, my conjecture is the following. During normal (i.e., non-crisis) periods, the demand curve for liquidity is relatively flat. On the other hand, the maximum amount of liquidity that the Fed can provide in a crisis is a more clearly quantifiable number (i.e. 2 or 3 trillions of dollars?). Moreover, while there may be shocks to the private liquidity-cost function, these shocks are unlikely to shift the social cost function. Thus the deadweight cost with quantity regulation is low. The central bank should choose a fixed $\mathcal{L}^*$. Basel 3-type liquidity regulations can implement this quantity limit.

There are two main issues that may develop with this type of regulation. First, while the demand curve for liquidity may be flat, it probably scales up with economic activity. If the regulator fixes a quantity target, then a slow increase in demand will raise **P.** The regulator can use the rise in **P** as information to slowly adjust the quantity target upwards.

Second, the regulator may discover that some types of private financial innovation which reduce private liquidity costs also reduce social costs, which should also push to increasing $\mathcal{L}^*$. My conjecture is that this realization can take place over time as the financial innovation is well understood, tested, and widely adopted by the marketplace. Again, the cost of reducing liquidity supply is likely to be small, so that a slow adjustment involves low deadweight costs.

In a crisis, the slope of the demand schedule is steep. Moreover, it is likely that private liquidity supply also shifts inwards (i.e. is reduced). Thus in this case the quantity target will not be desirable or even feasible. Instead, I suggest that the regulator shift to a (low) tax on liquidity, and that the central bank expands $\mathcal{L}^{\text{public}}$. Here also using the price **P** as an intermediate target may allow the central bank to determine how much to expand $\mathcal{L}^{\text{public}}$.

References:

Brunnermeier, Markus, Gary Gorton and Arvind Krishnamurthy (2010), Liquidity Mismatch, working paper, Northwestern University.

Krishnamurthy, Arvind and Annette Vissing-Jorgensen (2010), The Aggregate Demand for Treasury Debt, working paper, Northwestern University.

Stein, Jeremy (2010), Monetary Policy as Financial Stability Regulation, *Quarterly Journal of Economics*.

Weitzman, Martin (1974), Prices versus Quantities, *Review of Economic Studies.*