The Liquidity Premium of Near-Money Assets *

Stefan Nagel†
University of Michigan, NBER, and CEPR

September 2014

Abstract

This paper proposes a theory that links the liquidity premium of near-money assets with the level of short-term interest rates: Higher interest rates imply higher opportunity costs of money holdings and hence a higher premium for the liquidity service benefits of money substitutes. Consistent with this theory, short-term interest rates in the US, UK, and Canada have a strong positive relationship with the liquidity premium of Treasury bills and other near-money assets. Treasury security supply variables lose their explanatory power for the liquidity premium once short-term interest rates are controlled for, which favors the opportunity-cost-of-money theory over asset supply-driven models.

*I am grateful for comments to John Campbell, John D. Driscoll, Arvind Krishnamurthy, Francis Longstaff, Paolo Pasquariello, George Pennacchi, José-Luis Peydró, Jeremy Stein, Randy Wright, Yao Zeng, seminar participants at the Federal Reserve Board, Harvard, HEC Paris, Michigan, Stanford, Toulouse, participants at the American Economic Association Meetings and the LBS Safe Assets Conference. Mike Schwert and Yesol Huh provided excellent research assistance. I thank Henning Bohn, Daniel Hanson, Allan Mendelowitz, and Paolo Pasquariello for providing data.

†Stephen M. Ross School of Business and Department of Economics, University of Michigan, 701 Tappan St., Ann Arbor, MI 48109, e-mail: stenagel@umich.edu

I Introduction

Investors pay a premium for the liquidity service flow provided by near-money assets such as Treasury bills or recently issued ("on-the-run") US Treasury Bonds. What determines the magnitude of this liquidity premium? Does it vary over time? If so, why? It is well documented that liquidity premia rise during times of crises (Longstaff 2004; Vayanos 2004; Brunnermeier 2009; Krishnamurthy 2010; Musto, Nini, and Schwarz (2014)), but little is known so far about the economic forces that pin down liquidity premia outside of these episodes.

In this paper, I start with the basic premise that near-money assets are valued for their use as money substitutes. Holding money in the form of non-interest bearing deposits or currency is costly due to the opportunity cost of foregone interest. Near-money assets offer a lower-cost alternative for storing liquidity. The liquidity premium that market participants are willing to pay for this liquidity service flow of money substitutes should depend on the opportunity cost of holding money, that is, the level of short-term interest rates.

This opportunity-cost-of-money perspective on the liquidity premium contrasts with the asset-supply view of the liquidity premium emphasized in recent work by Bansal, Coleman, and Lundblad (2010), Krishnamurthy and Vissing-Jorgensen (2012), Krishnamurthy and Vissing-Jorgensen (2013), and Greenwood, Hanson, and Stein (2014). According to the asset-supply view, changes in the quantity of outstanding near-money assets are important drivers of the liquidity premium. These supply-driven changes in the liquidity premium are unrelated to interest-rate movements because there is no money (that pays zero or non-market determined interest) in these models and hence no substitution with money and no link to the opportunity cost of money.

To work out the relationship between the opportunity cost of money and liquidity premia, I present a model in which households obtain liquidity service flow from non-interest bearing deposits and Treasury bills. Deposits and Treasury bills are perfect substitutes, but they have different liquidity multipliers: one unit of deposits provides more liquidity service flow

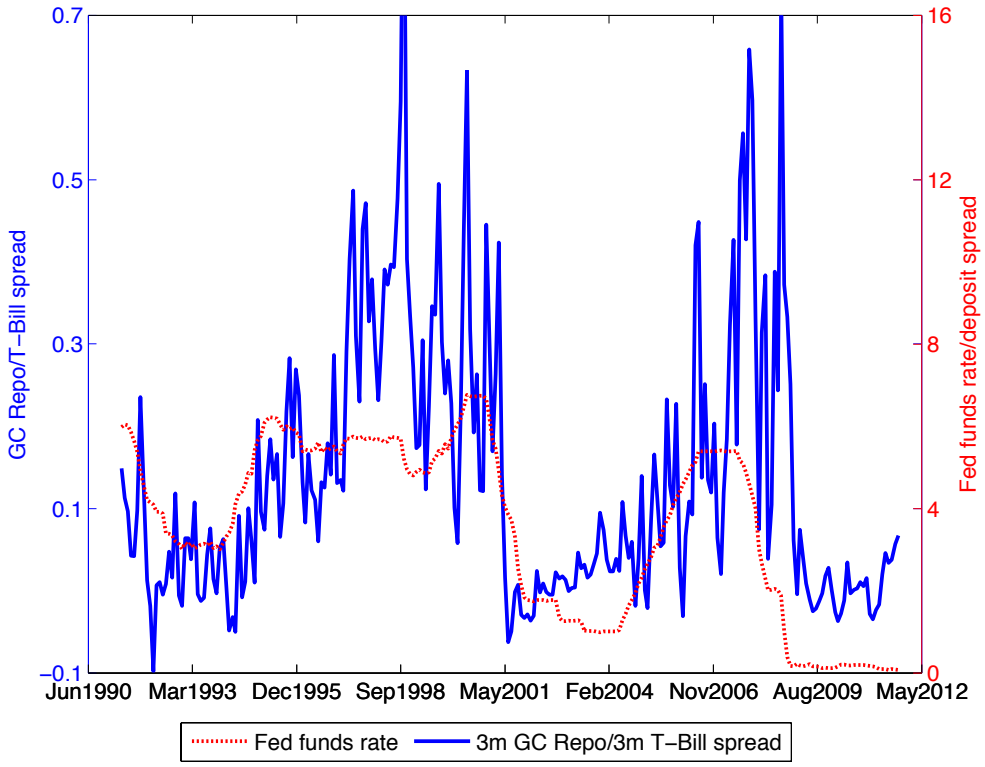


Figure 1: Spread between 3-month General Collateral Repo Rate and 3-month Treasury Bill yield

than one unit of Treasury bills. The banking sector supplies these deposits, but it requires some reserve holdings at the central bank as a liquidity buffer. In this model, the level of short-term interest rates represents households' opportunity costs of holding money (in the form of deposits). The liquidity premium of Treasury bills is proportional to the the level of short-term interest rates. This prediction is the focus of the empirical analysis in this paper.

Figure 1 illustrates the key empirical finding. The solid line shows the difference between the interest rate on three-month general collateral repurchase agreements (GC repo, a form of collateralized interbank lending) and the yield on three-month US T-bills. The GC repo term loan is illiquid, as the money lent is locked in for three months. In contrast, a T-bill investment can easily be liquidated in a deep market with minuscule bid/ask spreads. Consistent with this difference in liquidity, the GC repo rate is typically higher than T-bill

yields. This yield spread reflects the premium that market participants are willing to pay for the non-pecuniary liquidity benefits provided by T-bills. As the figure shows, there is substantial variation over time in this liquidity premium. Most importantly, the liquidity premium is closely related to the level of short-term interest rates, represented in the figure by the federal funds rate shown as dotted line. Thus, as the level of short-term interest rates changes, the opportunity cost of holding money changes, and hence the premium that market participants are willing to pay for the liquidity service flow of money substitutes changes as well.

In the opportunity-cost-of-money model, changes in near-money asset supply have effects on the liquidity premium only to the extent that the resulting change in demand for the substitute, i.e., money, leads to a change in short-term interest rates and hence the opportunity cost of money. Consistent with this prediction, I find that the Treasury supply variable of Krishnamurthy and Vissing-Jorgensen (2012) loses its explanatory power for the T-bill liquidity premium once the level short-term interest is controlled for. Greenwood, Hanson, and Stein (2014) focus more narrowly on T-bill supply, but I show that T-bill supply, too, loses its explanatory power in the presence of the short-term interest rate. Thus, what matters for the liquidity premium is the opportunity cost of money, not the supply of near-money assets. If the central bank follows an interest-rate operating target—which, according to Bernanke and Mihov (1998) seems to be a good description of the Federal Reserve's policy for much of post-WWII history—the central bank's elastic supply of reserves would neutralize any effect of near-money asset supply on the short-term interest rate and hence the opportunity cost of money and the liquidity premium.

While my conclusions about the connection between T-bill supply changes and the liquidity premium differ from the asset-supply view, the model and the empirical results here are consistent with the quantity evidence in Krishnamurthy and Vissing-Jorgensen (2013) and Greenwood, Hanson, and Stein (2014) that government supply of T-bills crowds out the private-sector supply of short-term debt, although the mechanism would be a modified one.

In the model of Greenwood, Hanson, and Stein (2014) an expansion in T-bill supply reduces household's marginal benefit of liquidity, which lowers the liquidity premium and reduces the incentive of the private sector to produce money-like short-term debt. In my model, the central bank would want to prevent this reduction in the marginal benefit of liquidity because this would push interest rates below the operating target. As a consequence, the central bank would respond by reducing the supply of reserves such interest rates and the liquidity premium would remain unchanged. The quantity of private-sector money supply would shrink, but not because of a fall in the liquidity premium, but because a lower supply of reserves limits the ability of the banking system to create deposits.

My findings are robust to a number of variations in measurement. GC repo rates are not available before the early 1990s, but I find that the spread between three-month certificate of deposit (CD) rates and T-bills exhibits a similarly strong positive correlation with the level of short term interest rates in data going back to the 1970s. A similar relationship is also evident in data from Canada and the UK. By comparing US T-bills with yields on discount notes issued by the Federal Home Loan Banks, which are illiquid, but have the same tax treatment as T-bills and are guaranteed by the US government, I am further able to rule out that the correlation between interest rates and the liquidity premium is a tax effect. The spread between illiquid Treasury notes and T-Bills (Amihud and Mendelson 1991) and between on-the-run and off-the-run Treasury notes (Krishnamurthy 2002; Warga 1992) also reflects liquidity premia, albeit of a smaller magnitude. I show that these liquidity premia, too, correlate positively with the level of short-term interest rates.

As Figure 1 shows, the liquidity premium also exhibits some higher-frequency variation around the cyclical interest-rate related component. Some of the spikes in the liquidity premium are crisis-related, such as those in September 1998 (LTCM) crisis, and in August 2007 and September 2008 (financial crisis). A plausible explanation for these deviations is that the relative liquidity service value of bank deposits is impaired during times of crises. As a consequence, market participants find the liquidity service flow of T-bills relatively more

valuable during these periods. It is also possible that part of the higher-frequency variation reflects short-lived supply effects, such as those studied by Sunderam (2013).

The opportunity-cost-of-money theory suggests that reserve remuneration policies of central banks could potentially change the relationship between the level of short-term interest rates and liquidity premia of near-money assets. Paying interest on reserves (IOR) drastically reduces the opportunity cost of one type of money (central bank reserves). However, to what extent IOR then also leads to a reduction in the opportunity cost of the monies held by households and non-financial corporations is an open question. If interest rates on the most liquid demand deposits remain close to zero, there should be little effect on liquidity premia. My empirical findings are consistent with this latter view. The introduction of IOR in the UK and Canada in 2001 and 1999, respectively, did not lead to a detectable change in the relationship between the T-bill liquidity premium and short-term interest rates. Apparently, offering IOR to a small number of financial institutions is not sufficient to induce a market-wide reduction in the opportunity cost of money and liquidity premia.

The findings in this paper contribute a new perspective on a number of research questions and policy issues related to liquidity premia. First, analyses whether there could be a "shortage" of near-money assets—e.g., due to margin requirements for derivatives transactions or other regulatory reasons (Committee on the Global Financial System 2013)—are incomplete without also taking into account the substitution relationship between near-money assets and money. As long as the central bank elastically supplies money and the banking system functions in the sense that households and corporations regard deposits as suitable for liquidity storage, even a substantial rise in demand for near-money assets does not necessarily lead to a shortage. In case that a shortage does arise (e.g., in a banking crisis), it should be apparent as a rise in the price of liquidity, i.e., the liquidity premium. My findings imply that to judge whether the liquidity premium is abnormally high, one should first remove the normal interest-rate related component.

Second, the results in this paper suggest a different interpretation of spreads between

open-market interest rates and T-bill rates as forecasters future real activity. Bernanke and Gertler (1995) note that the CD rate/T-bill spread rises during periods of monetary tightening. They interpret this finding as indicative of an imperfectly elastic demand for bank liabilities when banks respond to monetary tightening by looking for non-deposit funding. The results in this paper suggest an alternative interpretation: Monetary tightening—to the extent that it results in a rise in short-term interest rates—raises the opportunity cost of holding money and the liquidity premium of near-money assets. This alternative interpretation is also consistent with the finding of Bernanke and Blinder (1992) that much of the information in the commercial paper (CP)/T-bill spread about future real activity (as documented in Stock and Watson (1989), Bernanke (1990), and Friedman and Kuttner (1992)) is captured by the federal funds rate.

Third, the empirical evidence in this paper is relevant for applications in which liquidity premia are used as an input to explain other phenomena. In the model of Drechsler, Savov, and Schnabl (2014), the opportunity cost of (required) reserves limits the willingness of banks to invest in risky assets which results in elevated risk premia. My evidence for a tight link between the opportunity cost of money and liquidity premia suggests that a more general cost of precautionary liquidity holdings could play a similar role as the opportunity cost of reserves. Azar, Kagy, and Schmalz (2014) show that changes in the opportunity costs of holding liquid assets in the post-WW II decades can explain variation over time in the level of corporate liquid assets holdings. My findings suggest that the level of short-term interest rates can be used as a single state variable to model the evolution over time in the opportunity costs of holding liquid assets. This insight is also relevant for the construction of Divisia monetary aggregates (Barnett 1980; Barnett and Chauvet 2011). In this method, the outstanding stocks of different near-money assets are aggregated and weighted according to their degree of moneyness which is measured by each assets' liquidity premium. My finding that these liquidity premia are strongly positively correlated with the level of short-term interest rates should be useful for modeling the time series of the weights for different types

of near-money assets.

The remainder of the paper is organized as follows. Section II presents a model that clarifies the relationship between interest rates and liquidity premia. Empirical evidence on the time-variation in liquidity premia follows in Section III. Section IV examines the effect of changes in reserve remuneration policies of central banks in the UK and Canada. Section V concludes.

II A Model of the Liquidity Premium for Near-Money Assets

To clarify the relationship between the liquidity premium, the opportunity cost of money, and central bank policy, I begin by setting up a model of an endowment economy in which near-money assets can earn a liquidity premium. The economy is populated by households, a government comprising the fiscal authority and the central bank, and a banking sector, which is simply a technology to transform loans and reserve holdings into deposits that can be held by households.

II.A Households

Households derive utility from holding a stock of liquid assets as in Poterba and Rotemberg (1987) (see also, Woodford (2003), Chapter 3). Liquidity services are supplied by government securities and deposits created in the banking sector.

There is a single perishable consumption good, and the representative household seeks to maximize the objective

$$E_0 \sum_{t=1}^{\infty} \beta^t u(C_t, Q_t; \xi_t), \tag{1}$$

subject to the budget constraint

$$W_t = W_{t-1} r_t^W + P_t Y_t - T_t - P_t C_t, \tag{2}$$

where Y_t is the endowment of the consumption good, C_t is consumption, r_t^W is the nominal

gross return on wealth, and T_t denotes taxes paid to the government. P_t is the price, in terms of money, of the consumption good. Prices in this economy are flexible and adjust without frictions. Introducing nominal quantities and the price level serves the limited purpose of allowing some scope for the central bank to influence nominal short-term interest rates. The real rate will be fixed by the endowment process and preference shocks ξ_t. Q_t represents an aggregate of liquid asset holdings that provide households with utility from liquidity services. For each value of ξ_t, $u(C_t, Q_t; \xi_t)$ is concave and increasing in the first two arguments. I further assume that utility is additively separable in the utility from consumption and utility from liquidity services.

To map the model into the empirical analysis that follows, one can think of one period as lasting roughly a quarter. The liquidity benefits of near-money asset holdings arise from their use in (unmodeled) intra-quarter transactions, but the indirect utility from these benefits throughout the period enters directly in the objective function in (1).

Households can borrow from and lend to each other at a one-period nominal interest rate i_t. Households can also borrow from banks and hold demand deposits, D_t, at banks. Households perceive loans from other households and loans from a bank as perfect substitutes, and hence the interest is the same i_t. Demand deposits with banks are special, however. Unlike loans to other households, deposits with banks provide liquidity services. Treasury bill holdings, B_t, also provide liquidity services to some extent. The T-bills mature in one period and trade at yield i_t^b. Household T-bill holdings do not necessarily have to represent direct holdings. They could also include money market mutual funds that invest in Treasury bills. In contrast, money market mutual funds that invest in private-sector debt claims are better thought of in this model as part of the direct loans from household to household.

The households' total stock of liquidity is a linear aggregate of the real balances of demandable deposits and T-Bills

$$Q_t = \ell_d(\xi_t)(D_t/P_t) + \ell_b(\xi_t)(B_t/P_t), \tag{3}$$

i.e., D and B are perfect substitutes, but with different liquidity multipliers ℓ_d, ℓ_b. The liquidity multipliers do not depend on the level of D_t and B_t, but they are potentially subject to shocks ξ_t. As I show below, the assumption perfect substitutability of T-bills and deposits leads to some strong empirical implications. The central objective of this paper is to find out whether this stark model provides a useful characterization of the empirical behavior of the liquidity premium.

The representative household's wealth portfolio is

$$W_t = A_t + B_t + D_t - L_t, \tag{4}$$

where L represents loans from banks. A denotes the household's position in assets other than the ones discussed above.

The household first-order conditions with respect to consumption yield the Euler equation

$$1 + i_t = \frac{1}{\beta} \left\{ E_t \left[\frac{u_c(C_{t+1}; \xi_{t+1})}{u_c(C_t; \xi_t)} \frac{P_t}{P_{t+1}} \right] \right\}^{-1}, \tag{5}$$

where u_c denotes the partial derivative with respect to consumption. The household first-order conditions with respect to real liquid asset balances yield

$$\frac{u_q(Q_t; \xi_t)}{u_c(C_t; \xi_t)} \ell_d(\xi_t) = \frac{i_t - i_t^d}{1 + i_t}, \tag{6}$$

$$\frac{u_q(Q_t; \xi_t)}{u_c(C_t; \xi_t)} \ell_b(\xi_t) = \frac{i_t - i_t^b}{1 + i_t}, \tag{7}$$

where u_q denotes the partial derivative with respect to Q, and ℓ_d and ℓ_b the partial derivatives of the liquidity aggregate with respect to real balances D_t/P_t and B_t/P_t, respectively. Combining the two equations, we obtain a relationship between the liquidity premium of T-bills, $i_t - i_t^b$, and the spread between i_t and deposit rates:

$$i_t - i_t^b = \frac{\ell_b(\xi_t)}{\ell_d(\xi_t)} (i_t - i_t^d). \tag{8}$$

The liquidity premium priced into T-Bill yields is commensurate with the liquidity benefits of T-bills relative to deposits, as captured by the ratio of the liquidity multipliers ℓ_b and ℓ_d.

II.B Banks

Banks supply deposits and hold liquidity in the form of reserves M_t. I do not model the banking sector in detail. Instead, I just view it as a technology for transforming liquidity in the form of central bank reserves (which households cannot access) into deposits (which households can access). The banks' assets beyond reserve holdings are invested in loans to the household sector at rate i_t. Any profits flow to households who own the banks.

Two assumptions characterize the banking sector. First, the interest paid on demand deposits is zero,[1]

$$i_d = 0. \tag{9}$$

To map this zero-interest assumption to the real world, it is useful to think of interest-bearing deposits and money-market accounts offered by actual financial institutions as packaged portfolios of treasury bills yielding i_t^b, loans yielding i_t, and possibly a non-interest bearing demand deposit component. In the model, all but the demand deposit component should be thought of as residing in the household sector. The real-world counterpart of the demand deposits in this model are only the most liquid forms of deposits that are not subject to any restrictions on their use and withdrawals. In the US, Regulation Q until July 2011 explicitly prohibited payment of interest on such demand deposits. Starting in the early 1980s, banks were able to offer alternative interest-bearing accounts (e.g., NOW accounts) that resembled demand deposits in many ways. Nevertheless, as I discuss in more detail in Appendix B, households and corporations still maintained significant balances in non-interest bearing demand deposits. Evidently, interest-bearing deposits are not fully equivalent to non-interest bearing ones in

[1]This assumption is actually stronger than necessary for the key empirical predictions that follow. Similar predictions, apart from a constant term in the liquidity premium, would follow if demand deposit rates were greater than zero, but insensitive to movements in i_t. Deposit rates in a calibrated model in Ireland (2012) have this interest-rate insensitivity property.

the liquidity services that they provide.[2] In Section IV, I investigate whether it might be necessary to relax the $i_d = 0$ assumption if central banks pay interest on reserves at a rate close to i_t.

At least in the earlier part of the time periods that I study below, currency also plays an important part in households' stock of liquidity. I do not model currency separately from demand deposits, but one can think of currency as part of the non-interest bearing deposits D_t. For currency, the zero-interest assumption is uncontroversial.

If demand deposits do not pay interest, competition could lead banks to offer their customers other banking services without charge or at subsidized rates. Thus, bank customers might receive implicit interest through these other services. However, the key issue for the purposes of the analysis here is the marginal benefit from an additional dollar in deposits. The additional benefits offered in lieu of interest rates are typically not increasing in the level of deposit balances (apart from some minimum balance thresholds that may be in place to limit eligibility for these benefits). Hence, even including these forms of implicit interest, the marginal interest payment received on a dollar of deposits is still zero.

The second assumption that characterizes the banking sector in this model concerns the relationship between banks' reserve holdings, M_t, and their supply of deposits, D_t^s. I assume that

$$D_t^s = \phi(i_t, i_t^m; \xi_t) M_t, \tag{10}$$

where $\phi(i_t, i_t^m; \xi_t)$ is a multiplier that is potentially subject to random shocks through ξ_t. The multiplier can depend on i_t and i_t^m, because the magnitude of the spread between i_t and i_t^m may influence the degree to which banks try to economize on holding reserves. The motivation for this specification is not necessarily that there are explicit reserve requirements, but rather the precautionary need for banks to hold a certain level of reserves as a liquidity

[2] Gatev, Schuermann, and Strahan (2007) provide a piece of evidence that is consistent with the notion that non-interest bearing deposits are, to some extent, still a preferred location for storing liquidity. Studying the deposit inflows that banks received during the LTCM crisis in 1998, they find that inflows were directed towards non-interest bearing demand deposit accounts rather than interest-bearing transaction deposit accounts.

buffer with the central bank. The multiplier could be microfounded by modeling banks' reserve demand as in Ashcraft, McAndrews, and Skeie (2011) and Bianchi and Bigio (2013), where banks hold excess reserves to hedge unexpected payment flows. In countries such as Canada or New Zealand, where legal reserve requirements no longer exist, banks still exhibit a small, but greater than zero demand for reserve holdings (Bowman, Gagnon, and Leahy 2010), even though reserves are remunerated at a rate below the interbank lending rate.

These precautionary liquidity-buffer considerations motivate an inequality, where D_t^s is constrained above by the right-hand side of (10). As the interest rate on deposits is below the market interest rate i_t, banks should find it profitable, however, to create deposits (through lending) up to the point that the constraint binds. For this reason, I write (10) directly as an equality.

II.C Government: Fiscal Authority and Central Bank

The government issues liabilities in the form of one-period Treasury bills through the fiscal authority, and reserves, M_t^s, through the central bank. The central bank pays interest on reserves at a rate $i_t^m < i_t$. The path of $\{M_t^s, i_t^m\}$ is chosen by the central bank to meet an interest-rate operating target i_t^*. The central bank's interest-rate operating target could be set, for example, on the basis of a Taylor rule. I assume that monetary policy dominates, in the sense that if the Treasury implements a change to the path of T-bill supply, this does not change the central bank's interest-rate operating target.

To change M_t^s, the central bank conducts open-market operations, exchanging T-bills against reserves. Let B_t^s denote the net supply of T-bills that remains available to the public after the central bank has conducted its open-market operations.

The government collects taxes and makes transfers of net amount T_t to satisfy the joint flow budget constraint for government and the central bank consistent with the chosen paths for B_t^s and M_t^s,

$$B_t^s + M_t^s = B_{t-1}^s(1 + i_{t-1}^b) + M_{t-1}^s(1 + i_{t-1}^m) - T_t. \tag{11}$$

II.D Equilibrium

Equilibrium in this model is given by a set of processes $\{P_t, i_t, i_t^b\}$ that are consistent with household optimization (7), (6), (5), the supply of deposits (10), and market clearing,

$$C_t = Y_t, \quad D_t = D_t^s, \quad M_t = M_t^s, \quad B_t = B_t^s, \tag{12}$$

given the exogenous evolution of $\{Y_t, \xi_t\}$, and processes $\{i_t^m, M_t^s, B_t^s\}$ consistent with the monetary policy rule.

Given the evolution of $\{Y_t, \xi_t\}$, the policy rule for i_t followed by the central bank, together with the Euler equation (5), determines the current price level and expectations of future price levels (see Woodford (2003)). For any targeted i_t^*, given i_t^m, one can solve (6), with market clearing conditions substituted in,

$$\frac{u_q(Q_t^s; \xi_t)}{u_c(C_t; \xi_t)} \ell_d(\xi_t) = \frac{i_t}{1 + i_t}, \tag{13}$$

where

$$Q_t^s = \ell(\phi(i_t, i_t^m; \xi_t) M_t^s / P_t, B_t^s / P_t), \tag{14}$$

for the M_t^s that the CB must supply to achieve $i_t = i_t^*$. Concerning the conditions under which a solution exists, and the question of determinacy of the price level, see the discussion in Woodford (2003). For the purposes of this analysis, we can assume that the CB's action have resulted in a specific path for i_t and P_t, and ask what these paths imply for liquidity premia.

With a target-consistent money supply, and (9) substituted into (8), the liquidity premium of Treasury bills can be written as

$$i_t - i_t^b = \frac{\ell_b(\xi_t)}{\ell_d(\xi_t)} i_t, \quad \text{where} \quad i_t = i_t^*. \tag{15}$$

Thus, time-variation in the liquidity premium is driven by changes in the opportunity cost

of holding liquidity in the form of deposits, i_t, and by changes in the relative magnitude of the liquidity multipliers ℓ_b and ℓ_d, i.e., the relative usefulness of T-bills as a store of liquidity compared with deposits. If T-bills become relatively more useful, the ratio ℓ_b/ℓ_d rises, resulting in a lower yield on T-bills and hence a bigger liquidity premium.

A liquidity demand shock to $u_q(C_t, Q_t; \xi_t)$ has no direct effect on the liquidity premium, because the CB would have to offset the shock by elastically changing M_t^s (which would change the supply of deposits) to stay at the interest-rate target according to (13). Similarly, an elastic reserve supply response of the CB would also offset the effects of a change in the supply of T-bills, with no effect on the liquidity premium. That asset supply shocks are neutralized by elastic money supply is strongly suggested by prior empirical findings in Bernanke and Mihov (1998). They show that Federal Reserve policy going back to the 1960s—with exception of the early 1980s—can be described well by an interest-rate operating target, even though this was not an explicitly declared policy until the 1990s.

Moreover, even if the central bank does not follow an interest-rate operating target, but instead, say, a money growth target, liquidity demand or T-bill supply shocks should have no effect on the liquidity premium once the short-term interest rate is controlled for. In this case, these shocks would not be fully neutralized by the CB's money supply response, and hence they would would now change $u_q(C_t, Q_t; \xi_t)$ and therefore also i_t via (13). However, the only modification in (15) would be that i_t is no longer equal to i_t^*. Liquidity demand and near-money asset supply shocks can now affect i_t, but liquidity premia are still remain tied to the opportunity cost of money. Thus, the empirical prediction that the level of i_t should explain time-variation in liquidity premia does not depend on the assumption about the central bank's operating target. That liquidity demand and T-bill supply shocks do not affect the liquidity premium once i_t is controlled for is a key difference of this model to others in which the substitution relationship with money is not present (e.g., Bansal and Coleman 1996; Krishnamurthy and Vissing-Jorgensen 2013).

Crisis effects. The liquidity premium is still, however, potentially subject to random

shocks ξ_t that change the relative magnitudes ℓ_d and ℓ_b. As the empirical analysis shows below, there are big spikes in liquidity premia during times of financial market turmoil. Within this model, one can think of these spikes as the consequence of a shock that destroys "trust" in bank deposits as a store of liquidity, which lowers the liquidity multiplier of deposits ℓ_d relative to ℓ_b.[3] According to (15), this raises the liquidity premium of T-bills. If the shock is sufficiently big, this could lead to T-bill yields falling into negative territory.

Alternative functional form of liquidity aggregator. For T-bills, a constant marginal liquidity contribution (in absence of shocks ξ_t) is a plausible approximation. In contrast, with a more general constant elasticity of substitution aggregator, the marginal liquidity contribution would tend to infinity as B_t/P_t approaches zero. This would imply the assumption that households and corporations are willing to pay a liquidity premium that approaches infinity if the supply of T-bills shrinks to zero, which does not seem plausible. T-bills do not perform such a special role in liquidity portfolios that they could not be substituted for by demand deposits.

On the other hand, it seems plausible that demand deposits do perform a special role that near-money assets cannot perfectly substitute for at low levels of D/P. An extremely low level of deposit balances would be a considerable inconvenience for agents and would inhibit transactions that could otherwise take place. Thus, while the linearity assumption may work well for deposits as a local approximation, it is less plausible that linearity is a good approximation for very low levels of D/P. For this reason, I also undertake some robustness checks in the empirical analysis with a quasi-linear liquidity aggregator where $\ell_d(D_t/P_t; \xi_t)$ is declining in D/P. Since demand for deposits is inversely related to the level of i_t, this implies that ℓ_d is higher in times of high i_t. Following (15), this would predict a concave relation between the T-bill liquidity premium and i_t rather than the linear one from the baseline model (15).

Interest on reserves. If the CB introduces IOR, $i_t^m > 0$, this could potentially affect

[3] As an example for a model in which a crisis shock can impair the liquidity value of bank deposits see Robatto (2013).

liquidity premia if the introduction of IOR affects the rate that banks pay on deposits. If banks aggressively compete for deposits and try to earn the spread $i_t^m - i_t^d$, then i_t^d could potentially rise close to i_t^m. If so, this would push down liquidity premia. For example, starting from a situation where $i_t^m = 0$ and $i_t^d = 0$, if the introduction of IOR leads to $i^d = i_t^m > 0$, then the liquidity premium in (15) falls from $[\ell_b(\xi_t)/\ell_d(\xi_t)]i_t$ to $[\ell_b(\xi_t)/\ell_d(\xi_t)](i_t - i_t^m)$. Because central banks that pay IOR typically keep the spread $i_t - i_t^m$ constant when they change their target for i_t, the liquidity premium in this case would also be constant over time rather than varying with the level of i_t.

The extreme case of $i_t^m = i_t^*$ corresponds to the "Friedman rule" (Friedman 1960) according to which the payment of IOR equal to market rates would eliminate the implicit taxation of reserves (see, also, Goodfriend 2002; Cúrdia and Woodford 2011). If, in addition, deposit rates rise to $i_t^d = i_t^m$, then liquidity premia shrink to zero. The payment of $i_t^m = i_t^*$ would then not only eliminate the reserves tax, but it would eliminate a liquidity tax more broadly, as T-bills and other highly liquid government liabilities would no longer trade at a liquidity premium.

On the other hand, there are a number of reasons why the introduction of IOR might not have much of an effect on deposit rates. For example, if banks do not compete on deposit rates in the market for demand deposits, then deposit rates could remain stuck at zero. Or, if the cost of holding reserves is a negligible part of the overall marginal costs of demand deposit provision, there is again not much reason to expect demand deposit rates to change.[4] Thus, whether or not i_t^d rises towards i_t^m and whether the introduction of IOR shrinks liquidity premia is, in the end, an empirical question.

[4]In the model of Ireland (2012), the introduction of IOR close to market rate has only a weak effect on deposit rates because the bulk of the marginal costs of deposit provision is accounted for by labor costs. While deposit rates are not zero in his model, they are substantially below market interest rates.

III Empirical Dynamics of the Liquidity Premium

I now turn to an empirical evaluation of the opportunity-cost-of-money theory of liquidity premia. I begin by examining the hypothesis that the liquidity premium should vary over time with the level of short-term interest rates. Most of the analyses below use T-bills as a near-money asset, but I also present some evidence with other highly liquid assets. Appendix A describes the data. All interest rates in the empirical analysis are monthly averages of daily annualized effective yields.

III.A Liquidity Premium and Short-term Interest Rates

To measure the liquidity premium in U.S. T-Bills, I compare three-month T-bill yields to a maturity-matched general collateral (GC) repo rate for repurchase agreements with Treasury collateral. This repo rate is the interest rate for a three-month term interbank loan that is collateralized with a portfolio of US Treasury securities. Due to this backing with safe collateral, there is virtually no compensation for credit risk priced into the repo rate. An investment into a repo term loan is illiquid, because the investment is locked in during the term of the loan. In contrast, a T-bill investment is more liquid, because it can be re-sold easily. The spread between the repo rate and T-bills reflects this liquidity differential (Longstaff 2000). Over the sample period from May 1991 to June 2007, this spread averaged 14 basis points (bps = 1/100s of a percent).

Furthermore, the absence of a credit risk component in the GC repo rate makes the repo/T-bill spread a more accurate measure of the liquidity premium than other frequently studied measures such as the Treasury/eurodollar (TED) spread that compare T-bills with unsecured interbank rates. The credit risk component in unsecured rates can obscure variation in the liquidity premium.

Figure 1 shows monthly averages of the repo rate minus T-bill yields since the early 1990s. As the figure shows, the repo rate typically exceeds the T-Bill yield by a substantial margin, often between 5 and 25 bps, during the sample period covered in the figure. The focus of this

paper is on understanding why this spread varies over time.

According to equation (15), the liquidity premium should, in the absence of shocks to the multipliers ℓ_b and ℓ_d, be proportional to the level of short-term interest rates. Figure 1 shows that the time-variation in the repo/T-bill spread is—outside of crisis periods such as 2007/08—remarkably consistent with (15). I measure the level of short-term interest rates with the level of the federal funds rate.[5] As the plot demonstrates, the liquidity premium co-moves strongly positively with the level of the federal funds rate. At the medium-term frequencies illustrated by the plot in Figure 1, the opportunity-cost of money mechanism that leads to (15) seems to be the dominant influence on the liquidity premium. The higher-frequency liquidity demand shocks analyzed by Sunderam (2013) may be responsible for some of the short-run variation in the repo/T-bill spread around the cyclical interest-rate related component, but the plot shows that they do not lead to persistent deviations.

When financial markets are in turmoil, however, this roughly proportional relationship with the interest-rate level is clearly broken. During the LTCM crisis in September 1998, or around the height of the financial crisis in 2008, the repo/T-bill spread is unhinged from its usual relationship with the federal funds rate. Following equation (15), this can be explained if the liquidity service value of deposits is impaired relative to T-Bills—for example, because bank customers have doubts about the safety and liquidity of deposits during a financial crisis. This leads to a fall in the liquidity multiplier of deposits, $\ell_d(\xi_t)$, relative to $\ell_b(\xi_t)$, which raises $i_t - i_t^b$.

Table I, column (1), presents the results of regressions of the repo/T-bill spread on the federal funds rate. Since the main focus of this paper is on the behavior of liquidity premia in "normal" times outside of crisis periods, these regressions use data only up to June 2007

[5] One might prefer to use the 3-month GC repo rate instead, so that the same proxy is used for i_t on the left-hand and right-hand side of (15), but for measuring the level of i_t (as opposed to yield spreads) this would make virtually no difference. There is very little difference between monthly averages of these rates, and their correlation is close to one. The advantage of using the federal funds rate is that it is available for a longer time period than the GC repo rate. This allows me to use the same type of rate to measure the level of short-term interest rates in other analyses below that use longer data samples with alternatives to the GC repo rate in the calculation of the liquidity premium in T-bills.

Table I: Liquidity Premia and Fed Funds Rate: Short Sample

The sample period is May 1991 to June 2007. The data consist of monthly averages of daily rates. The dependent variable is a yield spread expressed in basis points; the explanatory variable (federal funds rate) is expressed in percent. Newey-West standard errors (12 lags) are shown in parentheses.

	Repo/T-Bill (1)	CD/T-Bill (2)	2y Off/OnRun (3)	T-Note/T-Bill (4)
Intercept	-8.66	2.13	-0.32	0.41
(s.e.)	(3.62)	(4.14)	(0.24)	(2.38)
Fed funds rate	5.36	6.69	0.37	1.09
(s.e.)	(0.99)	(1.21)	(0.07)	(0.79)
Adj. R^2	0.34	0.34	0.08	0.05
#Obs.	194	194	194	194

so that the financial crisis period is excluded. The results confirm the visual impression from Figure 1: The spread is strongly positively related to the federal funds rate. An increase of one percentage point in the federal funds rate is associated with a rise in the repo/T-bill spread of 5.36 bps (s.e. 0.99).

While T-bills are the most liquid Treasury security in the U.S., other Treasury securities with longer maturities can also have some near-money properties. In particular, the most recently issued "on-the-run" Treasury notes and bonds are traded in a highly liquid market and there is empirical evidence that they trade at a liquidity premium compared with older "off-the-run" issues that are less liquid (see, e.g., Warga 1992; Krishnamurthy 2002).

Column (3) in Table I looks at the spread between two-year off-the-run and on-the-run notes. I focus on two-year notes because they are issued on a regular monthly auction cycle. Longer maturity notes and bonds are on a less regular auction schedule. This makes it more difficult to analyze a monthly series of the on-the-run premium, because the premium follows seasonal pattern between the auction dates. To construct the spread, I compare the yield of the most recently issued on-the-run note with the yield of the nearest off-the-run note issued one auction earlier. Because the two notes are not identical in terms of maturity and coupon

rate, some further adjustment is necessary. I follow the method of Goldreich, Hanke, and Nath (2005) and use the off-the-run zero-coupon bond yield curve of Gürkaynak, Sack, and Wright (2007) to value the cash flows of the two notes and to construct a synthetic yield difference between the two notes that reflects the shape of the off-the-run yield curve at each point in time. I adjust the off-the-run/on-the-run spread with this synthetic yield difference. This adjustment accounts for the differences in maturity and coupon rates between the two notes.

As column (3) in Table I shows, there is a statistically significant positive relationship between the level of the federal funds rate and the off-the-run/on-the-run spread. The magnitude of the liquidity premium, however, is much smaller in this case—on the order of a few basis points. Correspondingly, the magnitude of the coefficient on the federal funds rate is much smaller than in column (1). The point estimate implies that a one percentage point change in the federal funds rate translates into a change of 0.37 bps (s.e. 0.07) in the off-the-run/on-the-run spread.

Column (4) looks at the spread between T-bills and less liquid off-the-run two-year Treasury notes. Amihud and Mendelson (1991) argue that this spread reflects a liquidity premium. I construct this spread by looking for two-year Treasury notes with remaining maturity of around three months. Then I compare the yield of each of these Treasury notes with the the average yield of two T-bills that straddle the maturity of the Treasury note. As Table I shows, the regression coefficient of the federal funds rate is positive, but it is not statistically significant in this sample.

Figure 2 presents a similar analysis with longer data series starting in 1976. I do not have repo rate data for time periods before the 1990s, and hence the liquidity premium in this figure is calculated by comparing T-bill yields with uninsured certificate of deposit (CD) rates rather than with GC repo rates. The spread to CD rates is an imperfect measure of the liquidity premium, because a CD rates contain a credit risk component. However, outside of crisis periods, this credit risk component is small. In the periods since the early 1990s

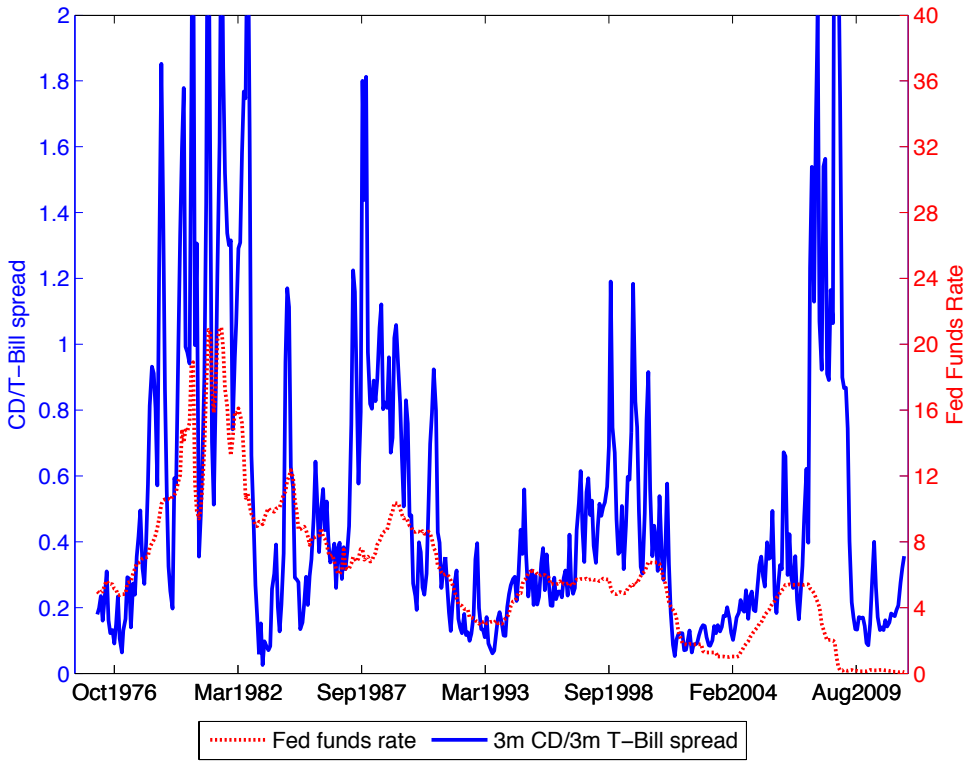


Figure 2: Spread between 3-month Treasury Bill yield and 3-month CD Rate

when both repo rate and CD rate data is available, there is typically only a small difference between CD rates and GC repo rates. Comparing the spread plotted in Figure 2 with the spread in Figure 1 one can see that the magnitude of the spread is quite similar in both cases. The big exception is the financial crisis period starting in 2007. The CD/T-bill spread rose much more strongly during the financial crisis than the repo/T-bill spread.

It is also apparent from Figure 2 that the CD/T-bill spread has a similarly strong positive relationship with the level of the federal funds rate as the repo/T-bill spread in Figure 1. This strong positive correlation is therefore not specific to the post-1990 period. The similarity of the CD/T-bill and the repo/T-bill spread can also be seen in column (2) of Table I. Regressing the CD/T-bill spread on the federal funds rate yields a slope coefficient estimate that is of similar magnitude as with the repo/T-bill spread in column (1) .

Column (1) in Table II exploits the advantage of a longer sample for the CD/T-bill spread

Table II: Liquidity Premia and Fed Funds Rate: Long Sample

The sample period is January 1976 to June 2007. The data consists of monthly averages of daily rates. The dependent variable is a yield spread expressed in basis points; the explanatory variable (federal funds rate) is expressed in percent. Newey-West standard errors (12 lags) are shown in parentheses.

	CD/T-Bill (1)	2y Off/OnRun (2)	T-Note/T-Bill (3)
Intercept	-6.31	0.52	2.89
(s.e.)	(4.53)	(0.72)	(2.17)
Fed funds rate	8.14	0.25	1.01
(s.e.)	(0.70)	(0.13)	(0.28)
Adj. R^2	0.51	0.03	0.06
#Obs.	378	378	378

and runs the regression with data from 1976 to 2007. The point estimate for the federal funds rate slope coefficient is quite similar to the estimate in column (2) of Table I. This provides further confirmation that the strong relationship between the liquidity premium of T-bills and the level of short-term interest rates is not specific to the post-1990 period.

Columns (2) and (3) further show that the off-the-run/on-the-run and the T-note/T-bill spread are positively related to the level of the federal funds rate in this longer sample, too. Moreover, the point estimate of the slope coefficient in the T-note/T-bill spread regression in column (3) is now more than three standard errors away from zero, indicating statistical significance at conventional levels.

III.B (Absence of) Supply Effects

One of the few existing alternative explanations for time-variation in the liquidity premium is the idea that changes in the supply of near-money assets cause this time-variation. Krishnamurthy and Vissing-Jorgensen (2012) show that the spread between commercial paper and T-bill yields is strongly negatively correlated with the supply of Treasuries, measured as the stock of outstanding US government debt to US GDP. Greenwood, Hanson, and Stein

(2014) find that the liquidity premium of T-bills is negatively related to the ratio of oustanding T-bills to GDP. The authors of both papers interpret these findings as evidence that the liquidity premium is declining in the supply of near-money assets.

The models that Krishnamurthy and Vissing-Jorgensen (2013) and Greenwood, Hanson, and Stein (2014) provide to explain such supply effects are similar to my model with: (i) $M^s = 0$ and deposits constrained by banks' stock of T-bill holdings rather than central bank money; (ii) the price level fixed so that the nominal rate i_t equals the real rate (which is pinned down by exogenous variation in the endowment process). In this case, (7) becomes

$$\frac{u_q(Q_t)}{u_c(C_t)}\ell_b = \frac{i_t - i_t^b}{1 + i_t}, \tag{16}$$

where Q_t is now a function of the T-bill supply only. A change in T-bill supply affects $u_q(Q_t)$, which results in a change in i_t^b because in these models i_t is, by assumption, exogenously fixed at the real rate. For example, higher supply of T-bills reduces $u_q(Q_t)$, which shrinks $i_t - i_t^b$. Thus, in this alternative theory, changes in the T-bill supply induce changes in the liquidity premium that are unrelated to movements in i_t. Furthermore, variation in i_t has virtually no effect on the liquidity premium after controlling for T-bill supply changes.[6] In contrast, the key prediction in my framework is that the liquidity premium should be highly correlated with i_t, but not related to supply variables once i_t is controlled for.

To test the alternative supply story, I construct a measure of the liquidity premium of T-bills going back to 1920. The government debt/GDP ratio—the key explanatory variable in Krishnamurthy and Vissing-Jorgensen (2012)—is highly persistent and moves up and down only a few times during the 20th century. Therefore, I follow Krishnamurthy and Vissing-Jorgensen (2012) in using a long sample to test the supply hypothesis. For the period 1920 to April 1991, I use a three-month bankers' acceptance rate as an illiquid, but relatively safe three-month rate. Bankers' acceptances (BA) are of relatively low risk because they are guaranteed by a commercial bank and hence represent an obligation of the (corporate)

[6]The effect of i_t through the denominator on the RHS is negligible.

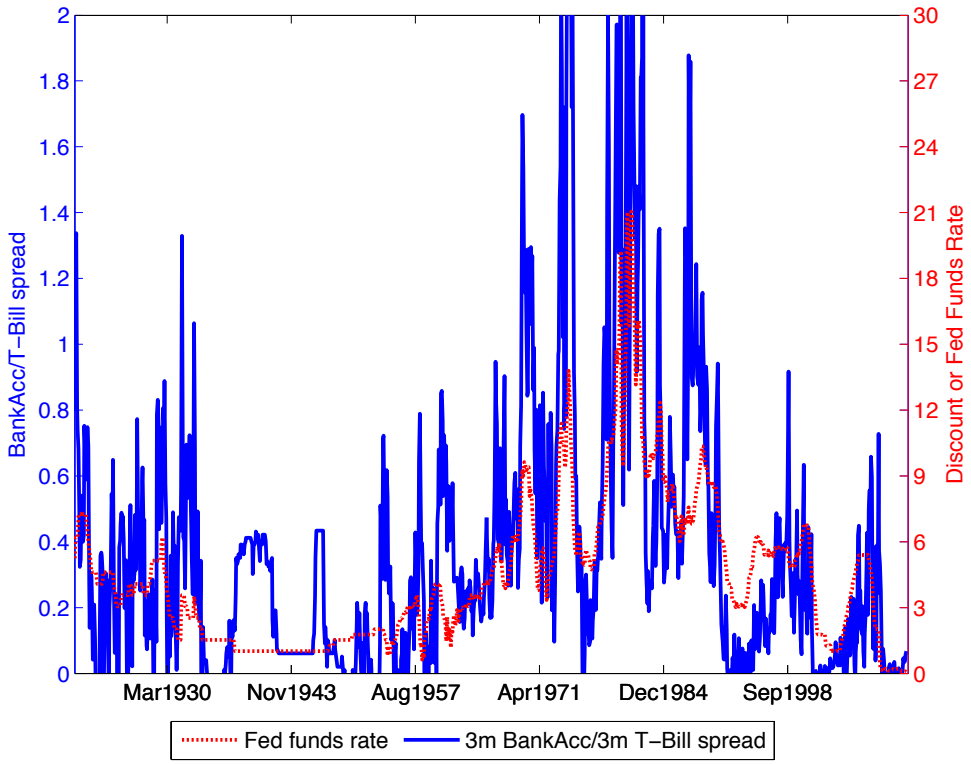


Figure 3: Spread between 3-month Treasury Bill yield and 3-month Bankers' Acceptance Rate

borrower and the commercial bank. BA can be traded in a secondary market, but the market is less liquid than the market for T-bills. From May 1991 onwards, I use three-month GC repo rates, as before. To measure i_t, I use the federal funds rate from July 1954 onwards and the Federal Reserve Bank of New York's discount rate before that date. During this early part of the sample, and unlike in later periods, banks regularly borrowed from the Federal Reserve at the discount rate, which implies that the discount rate is an appropriate proxy for the level of money market rates.

Figure 3 shows the BA/T-bill spread along with the level of the short-term interest rate. The strong correlation between the liquidity premium and the level of short-term rates is clearly apparent in this very long sample, too. Column (1) in Table III confirms this impression. A regression of the BA/T-bill spread on the level of the short-term rate produces a

Table III: Liquidity Premia and Fed Funds Rate: Controlling for Treasury Supply

The sample period is January 1920 to June 2007 in columns (1) to (3), January 1947 to June 2007 in columns (4) to (6), and January 1952 to June 2007 in columns (7) and (8). The interest-rate data consist of monthly averages of daily or weekly rates. The dependent variable is the spread between the three-month prime bankers acceptance rate and the three-month T-Bill yield until April 1991 and the spread between three-month GC repo and three-month T-bill from May 1991 to June 2007. The explanatory variables are the federal funds rate (discount rate before 1953) expressed in percent, the log government debt/GDP ratio, and the log of the ratio of outstanding amounts of T-bills (available from 1947 onwards) and GDP. Newey-West standard errors (12 lags) are shown in parentheses.

	1920 - 2007			1947 - 2007			1952 - 2007	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Intercept	-7.76	-3.52	-11.65	-192.45	-38.55	-32.89	-261.33	-48.78
(s.e.)	(4.44)	(8.05)	(6.18)	(65.08)	(30.13)	(29.22)	(76.99)	(47.02)
Fed funds rate	10.20		9.90		11.01	10.61		11.17
(s.e.)	(1.14)		(1.15)		(1.25)	(1.16)		(1.33)
log(Debt/GDP)		-44.37	-5.52			-9.25		
(s.e.)		(10.74)	(5.40)			(12.04)		
log(T-Bill/GDP)				-95.35	-9.51	-4.69	-123.90	-12.95
(s.e.)				(27.83)	(12.00)	(12.50)	(32.44)	(18.95)
Adj. R^2	0.51	0.13	0.51	0.12	0.55	0.55	0.15	0.53
#Obs.	1,050	1,050	1,050	726	726	726	666	666

slope coefficient of 10.20 (s.e. 1.14), which is quite close to the magnitudes of the coefficient in the much shorter samples in Tables I and II.

Column (2) replicates the Krishnamurthy and Vissing-Jorgensen (2012) finding for the BA/T-bill spread: Regressing the BA/T-bill spread on the log of the debt/GDP ratio yields a strongly negative and statistically significant coefficient of -44.37 (s.e. 10.74). However, it is possible that this negative coefficient arises because the debt/GDP ratio happens to be high in times when interest rates are low. Column (3) confirms this. Adding the short-term interest rate level as an explanatory variable dramatically shrinks the coefficient on log debt/GDP towards zero. At -5.52 (s.e. 5.40) it is now almost an order of magnitude smaller than in column (2) and no longer statistically significant. In contrast, the coefficient on the federal funds rate is almost unchanged from column (1) and still about nine times bigger than

its standard error. Moreover, the adjusted R^2 rises from 13% to 51%. I have further explored a specification that adds the slope the slope of the Treasury yield curve as an explanatory variable, as in Krishnamurthy and Vissing-Jorgensen (2012), but this has virtually no effect at all on the coefficients of the federal funds rate and the log Debt/GDP ratio.

Columns (4) to (6) in Table III introduce the T-bill supply as an explanatory variable. The T-bill supply variable becomes available in 1947, which is why the sample period in columns (4) to (6) starts in 1947. The regression in column (4) replicates the finding of Greenwood, Hanson, and Stein (2014) that T-bill supply is strongly negatively related to the T-bill liquidity premium. However, as column (5) shows, the supply effect largely disappears when the federal funds rate is included in the regression: The coefficient on the T-bill supply variable drops by an order of magnitude from -95.35 (s.e. 27.83) to -9.51 (s.e. 12.00). The coefficient on the federal funds rate (11.01; s.e. 1.25) is similar in magnitude to the estimate in previous regressions and the adjusted R^2 rises from 12% to 55%. Including the debt/GDP variable along with the T-bill supply in column (6) does not change this result. The federal funds rate retains its strong explanatory power.

To provide a better understanding why the federal funds rate has higher explanatory power than the supply variables, Figure 4 plots the time series of the explanatory variables used in Table III. For the purposes of this plot, the variables are de-meaned, standardized to unit standard deviation, and the sign switched, if necessary, so that a positive value implies a higher BA/T-bill spread. As the plot shows, the debt/GDP variable is very slowly moving and persistent. Over the course of the almost 90-year sample period until 2007, it has essentially only three high and two low observations. These slow movements are correlated with the federal funds rate, but, over and above this very low-frequency component, the federal funds rate also varies at higher business-cycle frequencies. Because these higher-frequency federal funds rate movements are matched with movements in the BA/T-bill spread (see Figure 3), the regression with the federal funds rate as explanatory variable have much higher R^2 than those with only the debt/GDP variable.

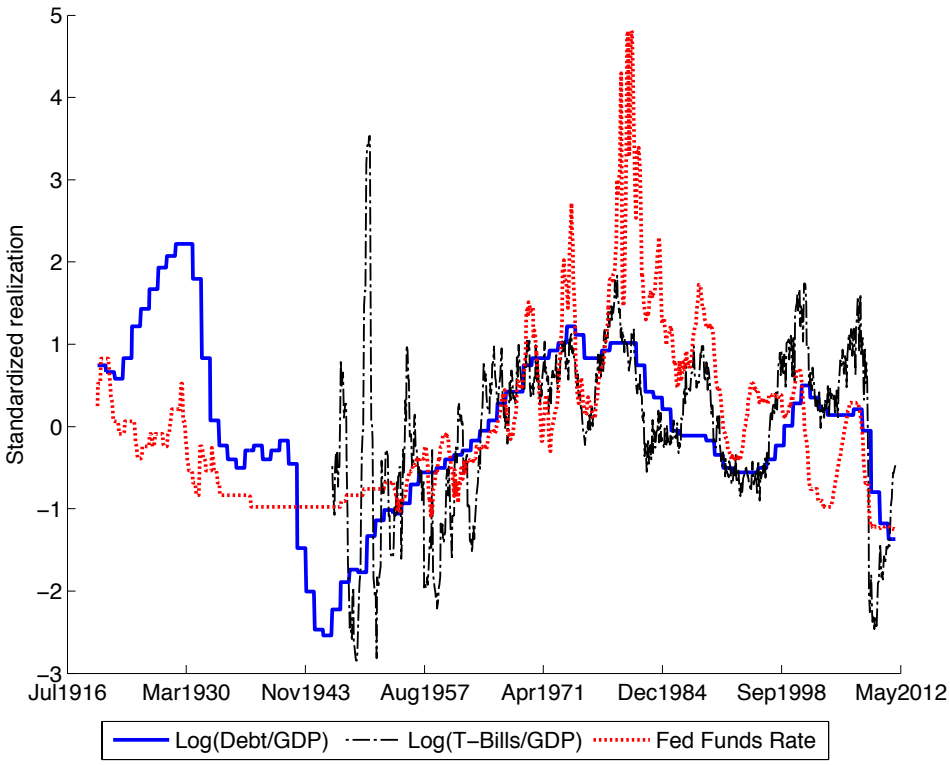


Figure 4: Standardized explanatory variables in supply regressions

The T-Bill supply variable is less persistent and more highly correlated with the level of short-term interest rates than the debt/GDP variable. But some of the movement in this variable does not line up as well with the BA/T-bill spread as the federal funds rate does. For example, the T-bill supply variable cannot explain the much higher BA/T-bill spread in the early 1980s compared with the late 1990s. It also does not explain well the big drop in the BA/T-bill spread in the early 2000s to almost zero. As a consequence, the T-bill supply variable achieves a much lower R^2 than the federal funds rate in Table III.

Figure 3 also reveals some erratic behavior of the T-bill supply variable around 1950. This may be a consequence of the Federal Reserve's policy to cap the yield of Treasury Bills, which had the consequence that much of the outstanding T-bills supply was held by the Federal Reserve. The Federal Reserve-Treasury Accord of 1951 ended this policy (Eichengreen and Garber 1991). For this reason, it is useful to check whether the results of the T-bill supply

regressions change if we limit the sample to the post-1951 period. As columns (7) and (8) in III show, this does not make much of a difference. The results are similar to those in columns (4) and (5).

Thus, once we take into account that the liquidity premium of near-money assets should be related to the opportunity cost of money, there is little evidence that Treasury supply has any effect on the level of the T-bill liquidity premium.[7] This finding supports the notion that the substitution relationship between money and near-money must be taking into account when analyzing the role of near-money asset supply. A reduction in the supply of near-money assets leads to a rise in demand for the substitute, i.e., money. This rise in money demand results either in a rise in short-term interest rates, or, if the CB follows an interest-rate operating target, in an elastic money supply response from the central bank. In either case, there is no effect of near-money asset supply on liquidity premia once the short-term interest rate is controlled for.

There is no contradiction, however, between the quantity correlations reported in Krishnamurthy and Vissing-Jorgensen (2012) and Greenwood, Hanson, and Stein (2014) and my evidence that supply-related variables are unrelated to the liquidity premium once short-term interest rates are controlled for. In fact, the elastic money supply response of the CB may be precisely what gives rise to the negative correlation between T-bill supply and private-sector near-money asset supply that these papers document. In my model, an expansion in money supply facilitates demand deposit creation. But one could extend this model to have banks create less liquid forms of near-monies, too. In practice, many types of private sector short-term debt such as commercial paper come with some form of liquidity support from commercial banks.[8] Provision of such liquidity support raises the liquidity needs of commercial banks. An expansion in the CB's supply of liquidity (in the form of reserves) to the banking system may therefore not only allow more deposit creation, but also an expansion

[7] It is possible that supply effects play a bigger role at the weekly frequency studied in Greenwood, Hanson, and Stein (2014). My regressions with monthly data do not rule this out.

[8] See, e.g., Acharya, Schnabl, and Suarez (2013) for the case of asset-backed commercial paper. Corporate commercial paper issuers typically arrange for a back-up credit line from commercial banks.

Table IV: Decomposition of Liquidity Premia: Relation to Financial Market Stress Indicators

The sample period is May 1991 to October 2011 in columns (1) to (4), where the dependent variable is CBOE VIX index expressed in percentage points and averaged within each month, and from January 1976 to December 2009 in columns (5) to (7), where the dependent variable is the financial market dislocation index of Pasquariello (2014) (multiplied by a factor of 1000). The explanatory variables are spreads of monthly averages of daily yields, decomposed into their projection on the federal funds rate and a residual. Newey-West standard errors (12 lags) are shown in parentheses.

	VIX 1991 - 2011				Dislocation Index 1976 - 2009		
	Repo/ T-Bill (1)	CD/ T-Bill (2)	2y Off/ OnRun (3)	T-Note/ T-Bill (4)	CD/ T-Bill (5)	2y Off/ OnRun (6)	T-Note/ T-Bill (7)
Intercept	21.23	21.61	23.59	23.30	-61.90	-27.25	-30.14
(s.e.)	(1.75)	(2.21)	(2.63)	(2.52)	(23.32)	(51.77)	(51.99)
proj($spread\|i_t$)	-0.16	-0.10	-3.24	-0.77	0.42	-0.99	0.03
(s.e.)	(0.10)	(0.07)	(1.87)	(0.44)	(0.38)	(18.77)	(4.31)
$spread - \text{proj}(spread\|i_t)$	0.26	0.11	0.24	0.52	1.76	2.16	2.68
(s.e.)	(0.05)	(0.02)	(0.49)	(0.12)	(0.69)	(1.72)	(0.84)
Adj. R^2	0.29	0.39	0.09	0.33	0.24	0.00	0.05
#Obs.	246	246	246	246	408	408	408

in the supply of other private-sector near-money assets.

III.C Isolating the Crisis Component in Liquidity Premia

Figure 1 suggests that the liquidity premium consists of two components with very different cyclical properties. The first is the component related to the opportunity cost of money that I focus on for the most part in this paper. The second component is one that appears predominantly in times of financial market stress like the LTCM crisis in September 1998 or the financial crisis in 2007-09. To properly interpret the magnitude of the liquidity premium at a given point in time, or to use it as predictors of other variables—e.g., as forecasters of real activity as in Stock and Watson (1989), Bernanke (1990), Friedman and Kuttner (1992) and Bernanke and Blinder (1992)—it is important to take into account that these two components carry different information about the state of the economy.

Table IV illustrates the different cyclical properties of the two components. The regressions reported columns (1) to (4) in this table show how the two components relate to the Chicago Board Options Exchange VIX index of implied volatilities of S&P500 index options. The VIX index is a widely used indicator of financial market stress. Periods of financial market turmoil and market illiquidity tend to coincide with high levels of the VIX index (Adrian and Shin 2010; Bao, Pan, and Wang 2011; Brunnermeier, Nagel, and Pedersen 2008; Longstaff, Pan, Pedersen, and Singleton 2010; Nagel 2012).

The dependent variable in columns (1) to (4) is the VIX index expressed in percentage points. The explanatory variables are the yield spreads from Table I decomposed into their projection on the federal funds rate, i.e., the fitted value from Table I, denoted proj($spread|i_t$) and the residual, denoted $spread - \text{proj}(spread|i_t)$.[9] While the regressions in Table I use only data excluding the financial crisis period after June 2007, the calculation of the fitted values for proj($spread|i_t$) applies the coefficients from Table I to the full sample period from May 1991 to October 2011.

As Table IV shows, the two components correlate very differently with the VIX. The opportunity-cost-of-money component proj($spread|i_t$) has a negative coefficient, albeit with weak statistical significance. The federal funds rate is high during booms when the VIX index tends to be low. In contrast, the residual component $spread - \text{proj}(spread|i_t)$ has a strongly positive association with the VIX index. For example, focusing on column (1), if the repo/T-bill spread widens by 10bp without a corresponding change in the federal funds rate, this is associated with a 2.6 percentage point rise in the VIX index (for comparison, the average level of the VIX is close to 20 percent). For all yield spreads except the off-the-run/on-the-run spread in column (3), the coefficient on this residual component is at least four standard errors greater than zero.

Columns (5) to (7) in Table IV use the Market Dislocation Index (MDI) of Pasquar-

[9]This residual may also be related to the common component of T-bill yields identified by Duffee (1996) that is uncorrelated with movements in other yields. Duffee conjectures that this component may reflect a time-varying convenience yield of T-bills.

iello (2014) as dependent variable instead of the VIX.[10] The MDI measures law-of-one-price violations in financial markets (covered interest parity, foreign exchange cross rates, and cross-listed stocks) and hence is a more direct indicator of abnormal levels of stress in financial markets than the VIX. This index is also available for a longer time period than the VIX, allowing me to use the long sample (as in Table II) from 1976 until the end of 2009 (when the MDI ceases to be available). As the results in Table IV show, the relationship between the MDI and the decomposed liquidity premia is broadly the same as in the case of the VIX: The coefficients of the opportunity cost of money component $\text{proj}(spread|i_t)$ are not statistically significant, whereas the coefficients of the residual $spread - \text{proj}(spread|i_t)$ are positive and, for the Repo/T-bill and T-note/T-bill spread, also more than two standard errors away from zero.

These results are consistent with the view that the residual $spread - \text{proj}(spread|i_t)$ carries information about stress levels in the financial system. Seen through the lens of the model, eq. (15), the liquidity multiplier of deposits, $\ell_d(\xi_t)$, falls in times of turmoil relative to the liquidity multiplier of T-bills, $\ell_b(\xi_t)$, which leads to an abnormally high level of the liquidity premium. While the liquidity premium can be high in good times when short-term interest rates are high, it can also be high in bad times when the liquidity service value of deposits is impaired relative to T-bills.

III.D Robustness Checks

Taxes. One potential concern with these analyses is that differences in taxation could drive a wedge between yields of T-Bills and private-sector money market rates. Earlier research, e.g. Cook and Lawler (1983), has argued that differences in state tax treatments explain the CD/T-bill rate spread. However, it is not clear whether these state-tax treatments could affect prices in a world in which some big investors are tax-exempt and taxable global financial institutions undertake elaborate efforts to minimize their tax bill. Fortunately, there is a way

[10] I thank Paolo Pasquariello for providing the data.

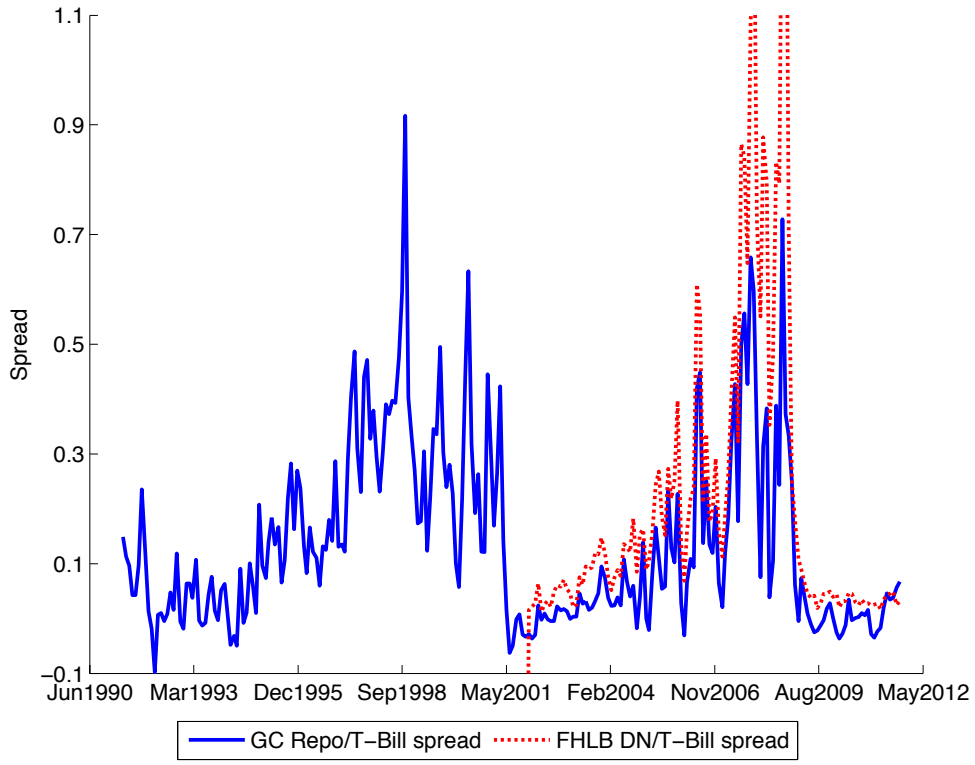


Figure 5: Spread between Repo/T-bill spread compared with T-bill/FHLB discount note spread

to directly address this issue empirically. For a number of years, the Federal Home Loan Bank (FHLB) has issued short-term discount notes with maturities in similar ranges as T-bills.[11] These discount notes receive the same tax treatment as Treasury Bills and FHLB debt is also explicitly guaranteed by the Federal government (Cowan and Petrine 2002). Thus, a spread between FHLB discount note yields and T-bill yields cannot be driven by taxation differences nor by credit risk. As Figure 5 shows, the FHLB note/T-bill spread is quantitatively similar to the repo/T-bill spread. The correlation between the two series is 0.91. This suggests that a the repo/T-bill spread cannot be explained by differential tax treatment. Instead, both the FHLB note/T-bill spread and the repo/T-bill spread reflect the superior liquidity of T-bills.

In addition, the off-the-run/on-the-run and T-note/T-bill spreads analyzed in Tables I and

[11] I am grateful to Allan Mendelowitz for providing the discount note yield data

II also compare instruments with similar tax treatment. The existence of a spread between their yields and the time-variation of this spread therefore cannot be explained by a tax story either.

Non-linear liquidity aggregator. As discussed in Section II.D, a quasi-linear liquidity aggregator may be more plausible than a linear aggregator. This would imply a concave relationship between i_t and the liquidity premium. To check this, I have experimented with a quadratic specification for the regression of the T-bill liquidity premium on the federal funds rate, but I did not find evidence for significant non-linearity.

IV The Effect of Reserve Remuneration Policies

As the evidence in the previous section shows, the assumption that $i_d = 0$, which leads to

$$i_t - i_t^b = \frac{\ell_b(\xi_t)}{\ell_d(\xi_t)} i_t \tag{17}$$

provides a good description of the behavior of liquidity premia in the US until the financial crisis. Throughout its history until 2008, the Federal Reserve did not pay IOR, i.e., $i_t^m = 0$. In October 2008 the Federal Reserve changed its reserve remuneration policy and started paying IOR. Due to the extremely low level of short-term interest rates since 2008, the IOR so far remained very close to zero and is hence unlikely to have a detectable effect on liquidity premia. A number of other countries introduced IOR at earlier points during the past two decades when short-term interest rates were higher. In those cases, effects on liquidity premia could potentially be big enough to be detectable.

Whether the assumption of $i_d = 0$ still provides a good prediction of liquidity premia after the introduction of IOR is an open question. Payment of IOR renders one type of money interest-bearing, but since reserve deposits are available only to commercial banks, it is not clear to what extent payment of IOR raises the rates of return on monies available to other market participants. It is possible that the introduction of IOR lowers banks' marginal

cost of maintaining precautionary liquidity sufficiently so that competition among banks drives demand deposit rates up to a level significantly above zero, perhaps close to i_t^m. If so, one would expect liquidity premia to shrink substantially and not vary much anymore with the level of short term interest rates following the introduction of IOR. On the other hand, it is possible that the opportunity cost of reserve holdings is not a major factor in the determination of equilibrium deposit rates. In this case, demand deposit rates may remain at zero, and liquidity premia still conform to (17).

To investigate this question empirically, I examine the UK (which introduced IOR in 2001) and Canada (which introduced IOR in 1999). For both countries, I do not have a sufficiently long series for 3-month term GC repo rates, and so I use the CD rate (UK) or the prime CP rate (Canada) as a proxy for the market rate for (illiquid) 3-month term loans.

IV.A United Kingdom

Figure 6 plots monthly averages of the CD rate/T-bill spread at 3-month maturity for the UK since 1978. The short-term interest rate shown in the figure is SONIA, an unsecured overnight interbank rate, from 1997, and the Bank of England's repo rate before 1997. As in the US, the liquidity premium of UK T-bills is positive on average, positively correlated with the level of short-term interest rates, and it exhibits positive spikes unrelated to short-term interest rate movements during the East Asian and LTCM crises in 1997 and 1998 as well as the financial crisis starting in 2007.

In June 2001, the Bank of England (BoE) introduced an overnight deposit facility. Excess reserves placed into the deposit facility earned an interest rate of 100bps below the BoE's main policy rate. Starting in March 2005, the spread to the main policy rate was changed a number of times to 25bps and 50bps (Bowman, Gagnon, and Leahy 2010). Figure 6 does not suggest that this change in reserve remuneration policy in 2001 had a substantial effect on the liquidity premium of UK T-Bills. Until the onset of the financial crisis in 2007, the CD rate/T-bill spread continued to be substantially positive, and it correlated positively with the

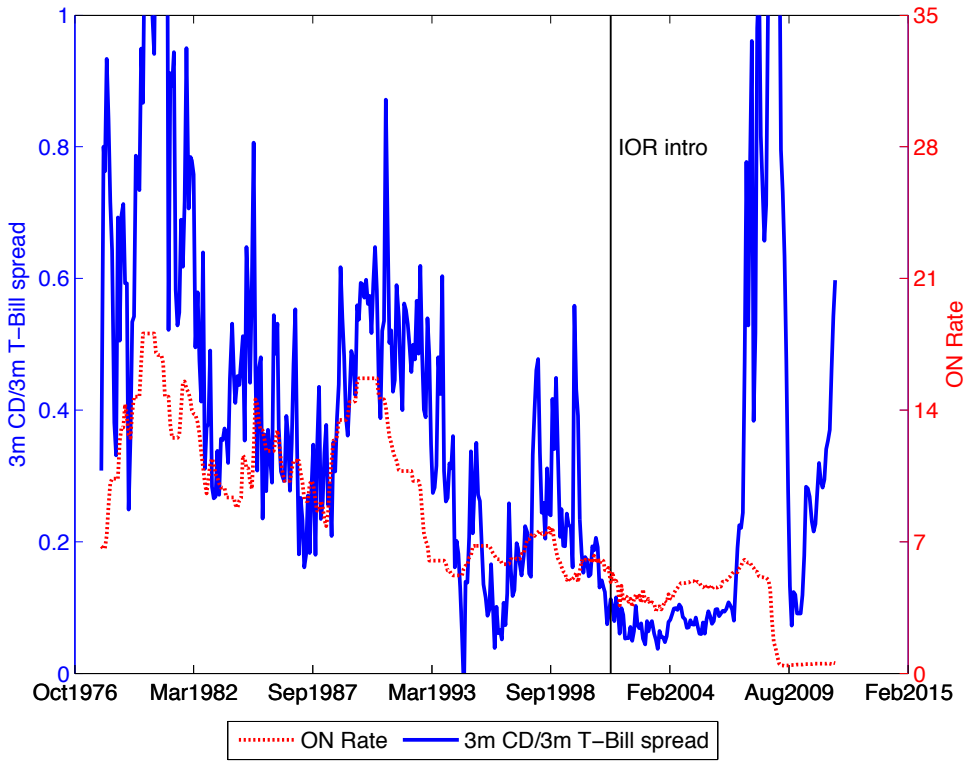


Figure 6: CD/T-Bill Spread in the UK

level of short-term interest rates.

Among the three countries examined in this study, the UK shows the most pronounced rise in the CD rate/T-bill spread towards the end of 2011 in Figure 6. A potential explanation for this rise is that it reflects a rise in perceived bank credit risk in the wake of the European debt crisis that was building up around that time and that UK banks may have exposure to.

Table V presents regressions of the UK's CD rate/T-Bill spread on the level of the short-term interest rate using data up to June 2007. Column (1) shows that there is a strong positive relationship between the level of the overnight rate and the CD rate/T-Bill spread. The magnitude of the coefficient is quite similar to the estimate in US data in Table I: a one percentage point rise in the overnight rate is associated with a 6.01 bps (s.e. 1.00) rise in the liquidity premium. Column (2) interacts the overnight rate with a dummy that equals one in the periods after the introduction of IOR. If the IOR introduction reduced liquidity

Table V: Liquidity Premium and IOR: UK

The sample period is January 1978 to June 2007. The data consist of monthly averages of daily rates. The dependent variable is a yield spread expressed in basis points; the explanatory variable (SONIA, an unsecured overnight interbank rate) is expressed in percent. Newey-West standard errors (12 lags) are shown in parentheses.

	CD/T-Bill (1)	CD/T-Bill (2)
Intercept	-16.89	-15.19
(s.e.)	(7.16)	(11.11)
ON Rate	6.01	5.87
(s.e.)	(1.00)	(1.29)
IOR dummy		6.91
(s.e.)		(12.05)
IOR dummy × ON rate		-2.11
(s.e.)		(1.71)
Adj. R^2	0.60	0.60
#Obs.	354	354

premia and dampened their co-movement with short-term interest rates, one would expect a negative coefficient on this interaction term. The estimated coefficient on the interaction term is indeed negative. However, the combined effect with the ON rate variable evaluated at the point estimates ($5.87 - 2.11 = 3.76$) is still a strong positive effect. Moreover, at conventional significance levels, one cannot reject the hypothesis that the coefficient on the interaction is zero. Thus, the T-bill liquidity premium still has a strong positive relationship with the short-term interest rate level, despite the IOR.

The behavior of interest rates on deposits and aggregate balances in different types of deposit accounts may offer additional insights into the (absent) effects of IOR. Appendix B shows that balances in non-interest bearing deposits in the UK did not change noticeably with the introduction of IOR. Average interest rates paid on transaction deposit accounts (current accounts) remained close to zero. This suggest that the IOR introduction had little effect on the way banks' conduct their deposit business, and hence little effect on the opportunity cost of money for non-banks.

IV.B Canada

Figure 7 shows the history of the liquidity premium, measured as the spread between prime commercial paper (CP) rates and Candadian T-bills at 3-month maturity. The overnight interest rate shown by the dotted line in Figure 7 is the CORRA overnight GC repo rate since December 1997, and prior to that date, the overnight Canada dollar LIBOR rate. As the figure shows, the liquidity premium of Canadian T-bills exhibits similar time-series behavior to the liquidity premia in the US and UK: It is positive on average, positively correlated with the level of short-term interest rates, and there are upward spikes during times of market turmoil.

The Bank of Canada introduced IOR in February 1999, as shown by the vertical line in Figure 7, with i_t^m set to 25bps below the target interbank lending rate (Bowman, Gagnon, and Leahy 2010). If deposit rates changed one-for-one with i_t^m, this would have a dramatic effect on the magnitude of liquidity premia according to (17): In 1999, short-term rates were at $i_t \approx 6\%$ and so with the introduction of IOR $i_t - i_t^m$ shrank from about 6% to 0.25%.

Figure 7 indicates, however, that the introduction of IOR had little effect on the magnitude of the liquidity premium of Canadian T-bills. While the introduction of IOR is preceded by some positive spikes in the liquidity premium in 1997 and 1998—presumably a consequence of the East Asian and LTCM crises occurring at the time—there is little evidence of a persistent change in the way the liquidity premium relates to the level of short-term interest rates. The time-series behavior of the liquidity premium is quite similar before and after the IOR introduction.

Table VI confirms the visual impression from Figure 7. The results in this table are based on data from January 1976 to June 2007, i.e., excluding the recent financial crisis period. The regression of the prime CP/T-bill spread on the short-term interest rate in column (1) yields a positive coefficient on the overnight rate. The point estimate suggests that a one percentage point rise in the overnight rate is associated with a 2.15 bps (s.e. 0.59) rise in the liquidity premium. In column (2), the overnight rate is interacted with a dummy that equals

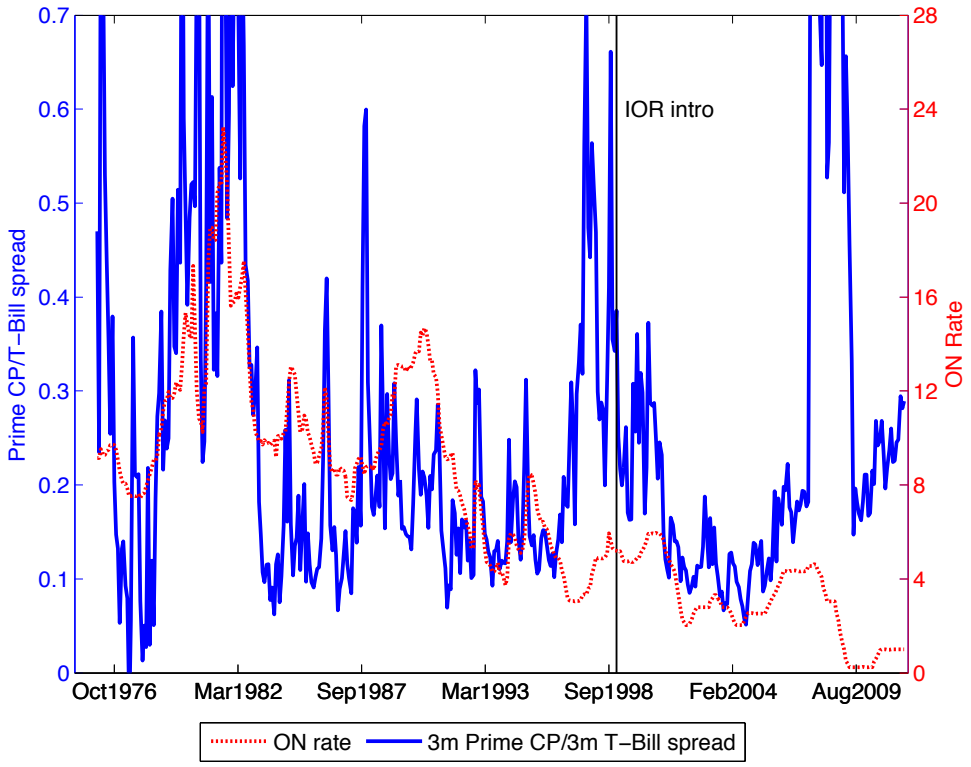


Figure 7: Commercial paper/T-Bill Spread in Canada

one in the periods after the introduction of IOR. If the IOR introduction reduced liquidity premia, one would expect a negative coefficient on this interaction term. In contrast, the estimated coefficient is positive, albeit not significantly different from zero at conventional significance levels.

One somewhat puzzling feature of the data is the elevated prime CP/T-bill spread in Figure 7 towards the very end of the sample. Possibly, this reflects a higher perceived riskiness of prime commercial paper following the financial crisis. If so, this would suggest that the prime CP/T-bill spread is not a good proxy for the liquidity premium during the last two years of the sample. Consistent with this explanation, the US CD rate/T-bill spread in Figure 2 shares this feature, but not the GC repo/T-bill spread (which is virtually free of credit risk) in Figure 1.

Table VI: Liquidity Premium and IOR: Canada

The sample period is January 1976 to June 2007. The data consist of monthly averages of daily rates. The dependent variable is a yield spread expressed in basis points; the explanatory variable (CORRA, a general collateral overnight rate since December 1997; overnight LIBOR in earlier periods) is expressed in percent. The IOR dummy is set to one in the time periods following the introduction of IOR in February 1999. Newey-West standard errors (12 lags) are shown in parentheses.

	Prime CP/T-Bill (1)	Prime CP/T-Bill (2)
Intercept	6.38	4.93
(s.e.)	(4.27)	(8.30)
ON rate	2.15	2.25
(s.e.)	(0.59)	(0.86)
IOR dummy		-5.40
(s.e.)		(8.51)
IOR dummy × ON rate		2.11
(s.e.)		(1.08)
Adj. R^2	0.26	0.26
#Obs.	378	378

IV.C Summary

The combined evidence from the UK and Canada offers little support for the conjecture that the introduction of IOR uncoupled liquidity premia from their close relationship with the short-term interest rate. Even though IOR lowers the opportunity cost of holding one form of money (central bank reserves) for banks with access to the central bank deposit facility, this does not seem to carry over into a substantial reduction in the opportunity costs of holding other types of money (deposits) faced by non-bank market participants without access to central bank deposits. As a caveat, though, the statistical power to estimate the effect of the introduction of IOR is somewhat limited.

Recent experience with IOR in the US also points towards frictions that prevent market participants from arbitraging discrepancies between IOR and open-market rates. After the Federal Reserve introduced IOR in October 2008, the federal funds rate has persistently traded below IOR. As Bech and Klee (2011) argue, this reflects the fact that some large

participants in the federal funds market are not eligible to receive IOR. Instead, they have to lend their funds in the federal funds market to banks who are eligible to receive IOR. These banks are not bidding for these funds aggressively enough to push the federal funds rate to the level of IOR. This illustrates that payment of IOR does not automatically establish the level of IOR as the floor for money market rates.

V Conclusion

The evidence in this paper suggests that liquidity premia of near-money assets reflect the opportunity cost of holding money. When interest rates are high, the opportunity costs of holding money are high, and market participants are willing to pay a big premium for highly liquid money substitutes such as T-bills. As a consequence, liquidity premia are positively correlated with the level of short-term interest rates. This interest-rate related variation is a dominant driver of liquidity premia at business-cycle frequencies, except in periods of financial market turmoil when liquidity premia are elevated relative to their normal level. The liquidity premium is unrelated to Treasury securities supply once short-term interest rates are controlled for.

Payment of IOR could potentially reduce and stabilize liquidity premia because IOR reduces the opportunity cost of holding money for at least some market participants (banks with reserve accounts at the central bank). However, the evidence from UK and Canada shows that liquidity premia remained strongly tied to the level of short-term interest rates after the introduction of IOR. Evidently, payment of IOR did not substantially affect the opportunity cost of holding money for non-bank market participants. One might conjecture that liquidity premia would indeed shrink and uncouple from the short-term interest rate if a much broader group of market participants—perhaps even including households and non-financial corporations—had direct access to interest-bearing electronic central bank money.

References

Acharya, Viral, Philipp Schnabl, and Gustavo Suarez, 2013, Securitization without risk transfer, *Journal of Financial Economics* 515–536, 515–536.

Adrian, Tobias, and Hyun Song Shin, 2010, Liquidity and leverage, *Journal of Financial Intermediation* 19, 418–437.

Amihud, Yakov, and Haim Mendelson, 1991, Liquidity, maturity, and the yields on us treasury securities, *Journal of Finance* 46, 1411–1425.

Ashcraft, Adam, James McAndrews, and David Skeie, 2011, Precautionary reserves and the interbank market, *Journal of Money, Credit and Banking* 43, 311–348.

Azar, José, Jean-François Kagy, and Martin C Schmalz, 2014, Can changes in the cost of cash explain the corporate cash puzzle?, Working paper, University of Michigan.

Bansal, Ravi, and Wilbur John Coleman, 1996, A monetary explanation of the equity premium, term premium, and risk-free rate puzzles, *Journal of Political Economy*, 1135–1171.

Bansal, Ravi, Wilbur John II Coleman, and Christian T. Lundblad, 2010, Endogenous liquidity supply, Working paper, Duke University.

Bao, Jack, Jun Pan, and Jiang Wang, 2011, The illiquidity of corporate bonds, *Journal of Finance* 66, 911–946.

Barnett, William A, 1980, Economic monetary aggregates: An application of index number and aggregation theory, *Journal of Econometrics* 14, 11–48.

Barnett, William A., and Marcelle Chauvet, 2011, How better monetary statistics could have signaled the financial crisis, *Journal of Econometrics* 161, 6–23.

Bech, Morten L., and Elizabeth Klee, 2011, The mechanics of a graceful exit: Interest on reserves and segmentation in the federal funds market, *Journal of Monetary Economics* 58, 415–431.

Bernanke, B., 1990, On the predictive power of interest rates and interest rate spreads, *New England Economic Review* Nov/Dec, 51–68.

Bernanke, Ben S., and Alan S. Blinder, 1992, The federal funds rate and the channels of monetary transmission, *American Economic Review* 82, 901–921.

Bernanke, Ben S., and Mark Gertler, 1995, Inside the black box: The credit channel of monetary policy, *Journal of Economic Perspectives* 9, 27–48.

Bernanke, Ben S., and Ilian Mihov, 1998, Measuring monetary policy, *Quarterly Journal of Economics* 113, 869–902.

Bianchi, Javier, and Saki Bigio, 2013, Liquidity management and monetary policy, Working paper, University of Winsconsin and Columbia University.

Bohn, Henning, 2008, The sustainability of fiscal policy in the United States, in R. Neck and J. Sturm, eds.: *Sustainability of Public Debt*, pp. 15–49).

Bowman, David, Etienne Gagnon, and Mike Leahy, 2010, Interest on excess reserves as a monetary policy instrument: The experience of foreign central banks, Working paper, Board of Governors of the Federal Reserve System.

Brunnermeier, Markus K., 2009, Deciphering the liquidity and credit crunch 2007–2008, *Journal of Economic Perspectives* 23, 77–100.

Brunnermeier, Markus K., Stefan Nagel, and Lasse H. Pedersen, 2008, Carry trades and currency crashes, in Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, eds.: *NBER Macroeconomics Annual 2008*, pp. 313–347, Chicago, IL (University of Chicago Press, Chicago, IL).

Committee on the Global Financial System, 2013, Asset encumbrance, financial reform and the demand for collateral assets, Working paper, CGFS Papers, no 49.

Cook, Timothy Q., and Thomas A. Lawler, 1983, The behavior of the spread between treasury bill rates and private money market rates since 1978, *Federal Reserve Bank of Richmond Economic Review* 4, 3–15.

Cowan, Adrian M., and Robert T. Petrine, 2002, Taxes and the pricing of government sponsored entity bonds, Working paper, Federal Home Loan Banks Office of Finance.

Cúrdia, Vasco, and Michael Woodford, 2011, The central-bank balance sheet as an instrument of monetary policy, *Journal of Monetary Economics* 58, 54–79.

Drechsler, Itamar, Alexi Savov, and Philipp Schnabl, 2014, A model of monetary policy and risk premia, Working paper, New York University.

Driscoll, John C., and Ruth A. Judson, 2013, Sticky deposit rates, Working paper, Federal Reserve Board.

Duffee, Gregory R., 1996, Idiosyncratic variation in treasury bill yields, *Journal of Finance* 51, 527–552.

Eichengreen, Barry, and Peter M. Garber, 1991, Before the accord: U.s. monetary-financial policy, 1945-51, in R. Glenn Hubbard, ed.: *Financial Markets and Financial Crises*).

Friedman, Benjamin M., and Kenneth N. Kuttner, 1992, Money, income, prices, and interest rates, *American Economic Review*, 472–492.

Friedman, Milton, 1960. *A Program for Monetary Stability* (Fordham University Press, New York, NY).

Gatev, Evan, Til Schuermann, and Philip Strahan, 2007, How do banks manage liquidity risk? evidence from the equity and deposit markets in the fall of 1998, in *The Risks of Financial Institutions*, pp. 105–132).

Goldreich, David, Bernd Hanke, and Purnendu Nath, 2005, The price of future liquidity: Time-varying liquidity in the u.s. treasury market, *Review of Finance* 9, 1–32.

Goodfriend, Marvin, 2002, Interest on reserves and monetary policy, *Federal Reserve Bank of New York Economic Policy Review* 8, 13–29.

Greenwood, Robin, Samuel G. Hanson, and Jeremy C. Stein, 2014, A comparative-advantage approach to government debt maturity, *Journal of Finance, Forthcoming*.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright, 2007, The U.S. treasury yield curve: 1961 to the present, *Journal of Monetary Economics* 54, 2291 – 2304.

Ireland, Peter N., 2012, The macroeconomic effects of interest on reserves, Working paper, NBER.

Krishnamurthy, Arvind, 2002, The bond/old-bond spread, *Journal of Financial Economics* 66, 463 – 506.

Krishnamurthy, Arvind, 2010, How debt markets have malfunctioned in the crisis, *Journal of Economic Perspectives* 24, 3–28.

Krishnamurthy, Arvind, and Annette Vissing-Jorgensen, 2012, The aggregate demand for treasury debt, *Journal of Political Economy* 120, 233–267.

Krishnamurthy, Arvind, and Annette Vissing-Jorgensen, 2013, Short-term debt and financial crises: What we can learn from US treasury supply, Working paper, UC Berkeley.

Longstaff, Francis A., 2000, The term structure of very short-term rates: New evidence for the expectations hypothesis, *Journal of Financial Economics* 58.

Longstaff, Francis A, 2004, The flight-to-liquidity premium in US treasury bond prices, *Journal of Business* 77.

Longstaff, Francis A., Jun Pan, Lasse H. Pedersen, and Kenneth J. Singleton, 2010, How sovereign is sovereign credit risk?, *American Economic Journal: Macroeconomics* 3, 75–103.

Musto, David, Greg Nini, and Krista Schwarz, 2014, Notes on bonds: Liquidity at all costs in the great recession, Working paper, Wharton.

Nagel, Stefan, 2012, Evaporating liquidity, *Review of Financial Studies* 25, 2005–2039.

Pasquariello, Paolo, 2014, Financial market dislocations, *Review of Financial Studies* 27, 1868–1914.

Poterba, James M., and Julio J. Rotemberg, 1987, Money in the utility function: An empirical implementation, in William A. Barnett and Kenneth J. Singleton, eds.: *New Approaches to Monetary Economics: Proceedings of the Second International Symposium in Economic Theory and Econometrics*, pp. 219–240, New York (Cambridge University Press, New York).

Robatto, Roberto, 2013, Financial crises and systemic bank runs in a dynamic model of banking, Working paper, University of Chicago.

Stock, James H., and Mark W. Watson, 1989, New indexes of coincident and leading economic indicators, in *NBER Macroeconomics Annual*, pp. 351–409).

Sunderam, Adi, 2013, Money creation and the shadow banking system, Working paper, Harvard Business School.

Vayanos, Dimitri, 2004, Flight to quality, flight to liquidity, and the pricing of risk, Working paper, National Bureau of Economic Research.

Warga, Arthur, 1992, Bond returns, liquidity, and missing data, *Journal of Financial and Quantitative Analysis* 27, 605–617.

Woodford, Michael, 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy* (Princeton University Press, Princeton).

Appendix (for online publication)

Appendix A describes the sources of the data used in this paper. Appendix B provides information on deposit quantities and deposit interest rates in the US and UK.

A Data

All yields are expressed as effective annual yields.

A.1 United States

Treasury bill and Treasury note yields. Data for T-bills and T-notes is from the daily CRSP database. Every day I choose the T-bill closest to 91-day maturity and calculate its yield from the midpoint of the bid and ask quotes provided in the CRSP database. To match T-notes with similar maturity T-bills, I look, each day, for the two-year note with remaining maturity closest to 91 days maturity and two T-bills whose maturity straddle the T-note's maturity. To construct the T-note/T-bill spread I subtract the linear interpolation of the two T-bill yields from the T-note yield as in Amihud and Mendelson (1991). For the sample starting in 1920, the T-bill series includes, from 1920 to 1933 U.S. Treasury three- to six-month notes and certificates taken from NBER Macro History database, obtained through the FRED database at the Federal Reserve Bank of St. Louis. From 1934 to April 1991 T-bill yields are from the FRED database, and from May 1991 onwards from CRSP, as described above.

On-the-run and off-the-run Treasury notes. To construct the spread between two-year on-the-run and off-the-run notes, I compare the yield of the most recently issued on-the-run note with the yield of the nearest off-the-run note issued one auction earlier. I follow Goldreich, Hanke, and Nath (2005) and use an off-the-run zero-coupon bond yield curve to value the cash flows of the on-the-run note and the nearest off-the-run note. I adjust the off-the-run/on-the-run spread with this synthetic yield difference. This adjustment accounts for the differences in maturity and coupon rates between the two notes. The off-the-run zero-coupon yield curves used in this method are obtained from the Federal Reserve Board, and they are based on the method of Gürkaynak, Sack, and Wright (2007).

Interbank rates. Daily GC repo rates are from Bloomberg, available from May 1991. The rates represent the rates at which dealers pay interest (the rates at which dealers receive interest are a few basis points higher). CD rates are obtained from the FRED database at the Federal Reserve Bank of St. Louis. The source of these data is the H.15 Release of the Federal Reserve Board. The reported CD rates refer to average of dealer bid rates for large-denomination ($1,000,000 or greater) certificates of deposit. These large denomination CDs are not insured by the FDIC. Daily data for the effective federal funds rate based on the H.15 release is also obtained from the FRED database.

Other yields. The Federal Reserve Bank of New York discount rate and the New York three-month bankers' acceptance rate from 1920 until 1940 are from the NBER Macro History Database, obtained through the FRED database. The Banker's acceptance rate from 1941 onwards is from FRED, too.

Debt/GDP. Data on the outstanding U.S. government debt to GDP ratio is from Henning Bohn's website. The data are an updated version of the series used in Bohn (2008). Data on outstanding T-bills is from the Center for Research in Security Prices (CRSP) monthly Treasury files, starting in January 1947. The series reports the outstanding face values of Treasury Bills, Certificates of Indebtedness, Tax Anticipation Bills, and Tax Anticipation Certificates of Indebtedness. Data on quarterly nominal GDP is from the the FRED database.

A.2 United Kingdom

.

Data on UK T-bill yields is from Datastream. Before 1996, I use secondary market T-bill yields (LDNTB3M). From 1996 onwards, this series is often stale. For this reason, I use average yields from the weekly three-month T-bill tenders (UKTBTND) from 1996 until 2011. Monthly averages of CD rates and of the SONIA overnight interbank rate are from the Bank of England.

A.3 Canada

Data on three-month T-bill yields are from Global Financial Data. The data comprises daily secondary market yields from 1990 and auction yields prior to 1990. Yields on three-month prime commercial paper and Canada dollar overnight LIBOR are also from Global Financial Data. The data is weekly until 1990 and daily subsequently. When the CORRA general collateral overnight repo rate becomes available on Datastream from 12/8/1997 onwards, I use CORRA as the short-term interest rate instead of LIBOR.

B Deposit balances and deposit rates

A.1 United States

Figure A.1 sheds light on the role of non-interest bearing demand deposits in the US. The data are from the FRED database at the Federal Reserve Bank of St. Louis. These non-interest bearing demand deposits are a component of the monetary aggregate M1. This category of deposits excludes interest-bearing NOW accounts and funds swept into interest-bearing transaction deposits. Figure A.1 plots the series as a proportion of GDP. As the figure shows, the share of non-interest bearing demand deposits has been quite stable at around 5 percent of GDP. Thus, throughout the sample period that I study, households and businesses in the US have keep substantial amounts of deposits in non-interest bearing form.

Figure A.1 also shows the share of currency plus non-interest bearing demand deposits over time. Holdings of currency have dropped strongly during the post-WWII period. Together, these non-interest bearing forms of money amount to about 10 percent of GDP.

Regarding the interest rates paid on interest-bearing deposit accounts, Driscoll and Judson (2013) report that the average rate on an aggregate of liquid interest-bearing deposits that includes interest-bearing checking accounts, savings deposits, and money-market deposit

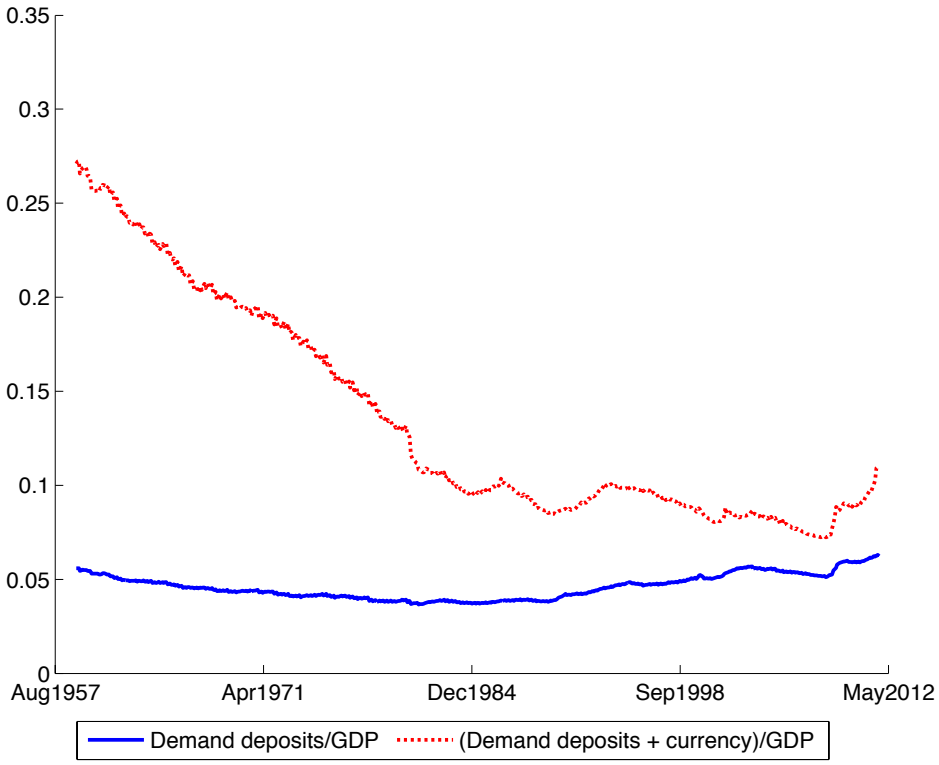


Figure A.1: US non-interest bearing demand deposits as proportion of GDP

accounts, is substantially below the federal funds rate, at roughly half of its level. Since this aggregate also includes less liquid savings and money-market deposit accounts, and it excludes non-interest-bearing accounts accounts, the average rate on the most liquid forms of deposits must be even closer to zero.

A.2 United Kingdom

.

Figure A.2 plots the quantity of non-interest bearing demand deposits in the UK, expressed as a proportion of GDP, around the time of the introduction of IOR in 2001. The data are from the Bank of England (non-interest bearing sight deposits in the Bank of England's terminology). Before the financial crisis, the share of non-interest bearing demand deposits is quite stable between 15 and 20 percent of GDP. There is no apparent break or other substantial change in the series with the introduction of IOR in 2001.

Figure A.3 shows UK deposit rates based on data from the Bank of England. The solid line shows the average rate on current accounts (i.e., checking accounts). As the figure shows, the average interest rate for this type of account is close to zero, it does not vary with the level of the short-term interest rate (which is shown in the figure as the dotted line), and it stayed close to zero following the introduction of IOR in 2001. There is a slight up-tick in

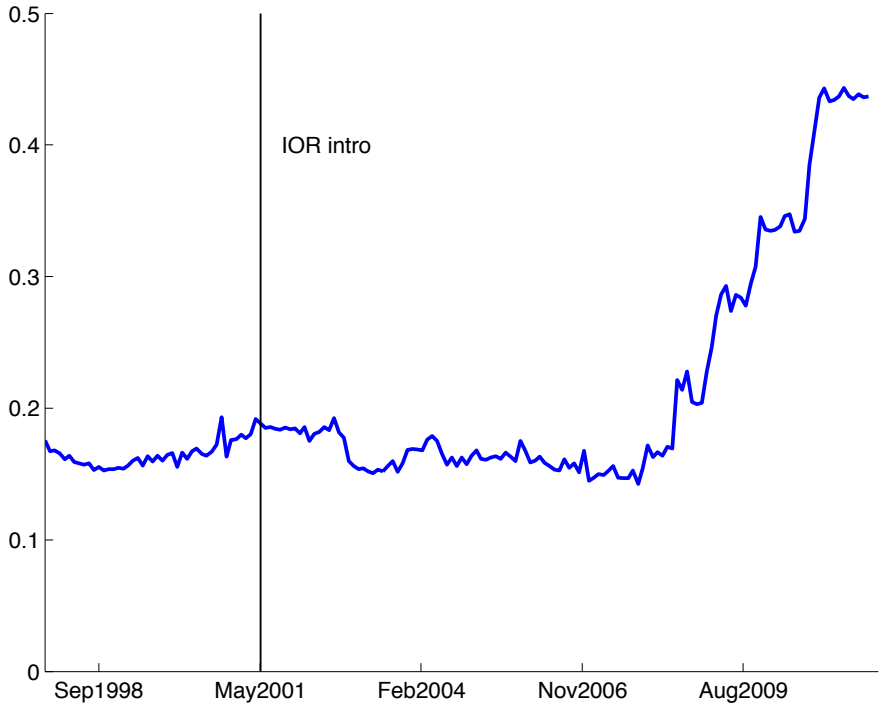


Figure A.2: UK non-interest bearing demand deposits as proportion of GDP

the rate a few months after the IOR introduction, but it is not clear whether this rise was connected to the change in the Bank of England's reserve remuneration policy. In any case, even with this slight rise, the average rate remained below one percent.

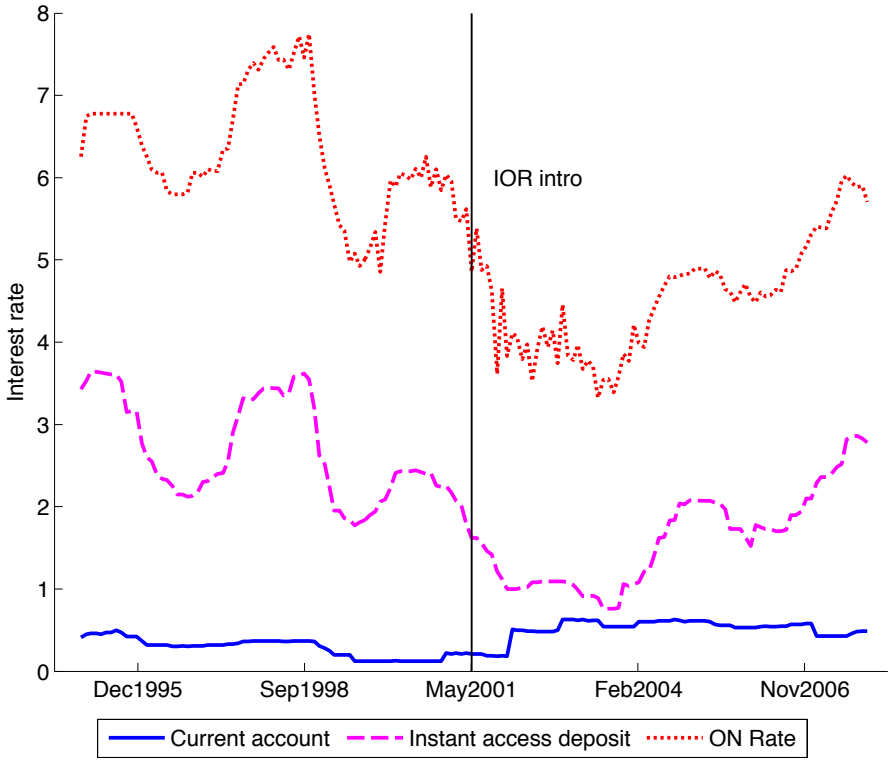


Figure A.3: UK deposit interest rates