

THE LIQUIDITY PREMIUM OF NEAR-MONEY ASSETS ^{*}

STEFAN NAGEL

This paper examines the link between the opportunity cost of money and time-varying liquidity premia of near-money assets. Higher interest rates imply higher opportunity costs of holding money and hence a higher premium for the liquidity service benefits of assets that are close substitutes for money. 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 over periods going back to the 1920s. Once the opportunity cost of money is taken into account, Treasury security supply variables lose their explanatory power for the liquidity premium, except for transitory short-run effects. These findings indicate a high elasticity of substitution between money and near-money assets. As a consequence, a central bank that follows an interest-rate operating target not only elastically accommodates and neutralizes shocks to money demand, but, effectively, also shocks to near-money asset supply and demand.

JEL codes: E43, G12, E41

^{*}I am grateful for comments from John Campbell, John Driscoll, Sam Hanson, Arvind Krishnamurthy, Francis Longstaff, Paolo Pasquariello, George Pennacchi, José-Luis Peydró, Jeremy Stein, Annette Vissing-Jorgensen, Randy Wright, Yao Zeng, seminar participants at the Bank of Canada, Banque de France, Cambridge, Cornell, Cheung Kong Graduate School of Business, Duke, the Federal Reserve Board, Federal Reserve Bank of New York, Georgetown University, Harvard, Hong Kong University of Science and Technology, University of Hong Kong, HEC Paris, University of Illinois Urbana/Champaign, Indiana, UMass Amherst, Michigan, Michigan State, Northwestern University, Purdue, Shanghai Advanced Institute of Finance, Stanford, UT Austin, UT Dallas, Toulouse, Warwick, Wisconsin, conference participants at the American Economic Association Meetings, the Bank of International Settlements, the NBER Monetary Economics Meetings, the LBS Safe Assets Conference, and the Western Finance Association Meetings. Mike Schwert and Yesol Huh provided excellent research assistance. I thank Henning Bohn, Daniel Hanson, Allan Mendelowitz, and Juan Pablo Nicolini for providing data.

I. INTRODUCTION

Investors pay a premium for the liquidity service flow provided by near-money assets. Highly liquid assets such as Treasury bills or recently issued (“on-the-run”) US Treasury Bonds trade at a liquidity premium: their yield is lower than the yield of comparably safe but illiquid assets of similar maturity. While the cross-sectional relationship between liquidity premia and various characteristics of these assets is well documented in the literature, our understanding of time-variation in near-money asset liquidity premia is still very limited. Recent policy initiatives underscore the need to fill this knowledge gap. Unconventional monetary policies and regulatory changes in the financial sector affect near-money asset supply and demand. Whether such policies affect liquidity premia or, in the extreme case, generate shortages of near money assets is unclear without an understanding of the determinants of time-variation in liquidity premia.

My approach in this paper builds on the premise that the liquidity service flow offered by near-money assets is closely related to the benefits from holding money. To the extent that Treasury bills, for example, serve as a substitute for money, the price premium that investors are willing to pay to enjoy their liquidity service flow should be linked to the price of the substitute, i.e., the opportunity cost of holding money. With non-interest bearing money, the opportunity cost of money is the forgone interest income. Thus, liquidity premia should be linked to the level of short-term interest rates.¹ The higher the elasticity of substitution, the tighter should be this link. This substitution relationship gains added importance from the fact that the opportunity cost of money is under the control of the central bank. For example, with an interest-rate operating target, the central bank endogenously supplies money to fix the opportunity cost of money at a level determined by, say, the inflation rate and the output

1. This opportunity cost of money theory of the liquidity premium goes back to Kessel (1971) and Cagan (1969). This earlier literature used the opportunity cost of money theory to analyze the dynamics of the yield curve and it measured liquidity premia by comparing assets of different maturity. The difficulty with measuring liquidity premia in this way is that yield differences across the maturity are confounded by expectations about future yield curve movements and risk premia.

gap. Depending on the elasticity of substitution between money and near-money, the central bank's interest-rate policy could then also affect the liquidity premium of near-money assets.

While there has been a surge of recent interest in time-variation of near-money liquidity premia, the literature has not explored the implications of substitution between money and near-money for the behavior of liquidity premia. One branch of the literature focuses on the rise of liquidity premia during flight-to-liquidity episodes, but without analyzing the connection between money and near-money assets (Longstaff [2004]; Vayanos [2004]; Brunnermeier [2009]; Krishnamurthy [2010]; Musto, Nini, and Schwarz [2014]). A second branch analyzes the role of near-money asset supply, but without considering money as a substitute (Bansal, Coleman, and Lundblad [2010], Krishnamurthy and Vissing-Jorgensen [2012], and Greenwood, Hanson, and Stein [2015]). As a consequence, time-variation in liquidity premia in these models is driven by changes in near-money asset supply, but the opportunity cost of money does not play any role.

My objective in this paper is to study all three channels in a common parsimonious framework: the opportunity cost of money, near-money asset supply, and potential flight-to-liquidity effects. To guide the empirical analysis, I set up a model in which households obtain liquidity service flow from deposits and T-bills. Banks create deposits, but to do so, they require reserve balances at the central bank. This allows the central bank to control nominal short-term interest rates through open market operations. Deposits pay interest equal to a fraction of the market interest rate. The behavior of the T-bill liquidity premium in this model depends on the elasticity of substitution between deposits and T-bills and a liquidity share parameter that captures how useful T-bills are for the purposes of holding liquidity relative to the same quantity of deposits.

In the case that deposits and T-bills are perfect substitutes, the liquidity premium is perfectly positively correlated with the level of the short-term interest rate. A linear aggregate of deposit and T-bill supply determines the short-term interest rate, but given the level of the aggregate, which the central bank can control through open-market operations, its

composition is irrelevant for liquidity premia. In the other extreme case, if T-bills and deposits are not substitutes at all, the liquidity premium is a function on the supply of T-bills and it is independent of the interest rate level. In the intermediate case of imperfect substitution, both the interest rate and T-bill asset supply affect the liquidity premium. In addition, in all cases, the liquidity premium varies with the liquidity share parameter. Time-variation in this parameter can be thought of as a representation of flight-to-safety effects which cause T-bills to rise in attractiveness relative to deposits during a crisis.

The goal of my empirical analysis is to infer the elasticity of substitution between money and near-money and to characterize the drivers of time-variation in the liquidity premium. I start with the perfect substitution case as the baseline and I look for co-movement between liquidity premia and interest rates. The cleanest measure of the liquidity premium of T-bills is the spread between the interest rate on three-month general collateral repurchase agreements (GC repo, effectively an interbank loan collateralized with Treasury securities) 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 and the bid-ask spreads between lending and borrowing rates are relatively wide compared with T-bills. 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. I show that this liquidity premium is closely related to the level of short-term interest rates. For every percentage point increase in the interest rate, the liquidity premium rises roughly 6 basis points.

These baseline 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 back to the 1970s, or the spread between three-month banker's acceptance rates and T-bills back to the 1920s exhibit a similarly strong positive correlation with the level of short term interest rates. 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, 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.

To allow for imperfect substitution, I then include near-money asset supply variables in the regression. According to the model, if the elasticity of substitution is not very high, near-money asset supply should have explanatory power for the liquidity premium after controlling for short-term interest rates. This is not what I find. While there is a strong negative univariate relationship between the T-bill liquidity premium and the Treasury supply variable of Krishnamurthy and Vissing-Jorgensen (2012) or the T-bill supply variable of Greenwood, Hanson, and Stein (2015), this relationship disappears once the short-term interest rate is included in the regression. In contrast, the coefficient on the short-term interest rate hardly changes when the near-money supply variables are added as explanatory variables. These empirical results indicate a very high elasticity of substitution between money and near-money assets.²

These findings leave open the possibility that the elasticity of substitution could be lower in the short-run. For example, if a T-bill supply shock hits the market, it may be absorbed initially by a small set of intermediaries. The high elasticity of substitution might only apply once the supply has been dispersed throughout the financial system. In this case, the supply

2. My analysis focuses exclusively on liquidity premia at the short end of the Treasury yield curve while Krishnamurthy and Vissing-Jorgensen (2012) additionally investigate premia in long-term Treasuries, too. Long-term Treasuries are not as liquid as T-bills, but they may command a convenience yield for reasons other than moneyness (e.g., for “long-term safety”). Neither money nor any private-sector assets are likely to be good substitutes for these properties of long-term Treasuries. As a consequence, substitution with money may be less relevant than for T-bills and Treasury supply could be more important for explaining premia in long-term Treasuries.

shock could have a transitory effect on the liquidity premium. To check for transitory effects, I look at differenced regressions of the repo/T-bill spread on monthly changes in interest rates and near-money asset supply. These differenced specifications also allow me to employ instrumental variables to address the reverse causality concern that interest rate or near-money asset supply changes could be caused by changes in the liquidity premium during the same time period. I instrument supply changes with seasonal dummies, as in Greenwood, Hanson, and Stein (2015), and interest-rate changes with lagged federal funds futures prices. These differenced regressions uncover transitory supply effects: Changes in T-bill supply affect liquidity premia in the same month even controlling for changes in interest rates, but the effect reverts during the next month. Due to their transitory nature, these supply effects don't have persistent effects on the level of liquidity premia.

While my conclusions about the determinants of the liquidity premium, i.e., the *price* of liquidity are sharply different, the model and the empirical results here are not inconsistent with the *quantity* evidence in Krishnamurthy and Vissing-Jorgensen (2015) and Greenwood, Hanson, and Stein (2015) that government supply of T-bills crowds out the private-sector supply of short-term debt, or the evidence in Sunderam (2015) that changes in the liquidity premium are associated with a subsequent short-term debt supply response by the private sector at weekly frequencies. In my model, a central bank with an interest-rate operating target would have to respond to a reduction in T-bill supply or a rise in liquidity demand with an elastic expansion in reserves supply—otherwise the interest rate would move away from the target. The expansion in reserve supply in turn facilitates an expansion in deposit supply, i.e., private-sector short-term debt.³

Knowledge of the elasticity of substitution between money and near-money is important for predicting the effects of changes in the conduct of monetary policy on liquidity premia. My finding that the elasticity of substitution is very high suggests that liquidity premia would

3. In addition to facilitating deposit creation, reserve expansion could also stimulate issuance of other types of private-sector short-term debt that come with some form of liquidity support from commercial banks. See, e.g., Acharya, Schnabl, and Suarez (2013) for a discussion of liquidity-support in the case of asset-backed commercial paper.

shrink dramatically if money became interest-bearing. The experiment of giving the public access to interest-bearing central bank money has not been implemented anywhere, but some central banks have moved to paying interest on excess reserves (IOR) held by banks in their accounts at the central bank. To what extent IOR 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. To shed light on this question, I examine the introduction of IOR in the UK and Canada in 2001 and 1999, respectively. I find that the IOR introduction 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.

The results in this paper also help understand why spreads between private-sector short term debt and T-bill yields forecast future real activity. In particular, my evidence sheds light on the finding of Bernanke and Blinder (1992) that the federal funds rate subsumes much of the predictive power of the commercial paper (CP)/T-bill spread. Their interpretation is that monetary tightening, as indicated by a rise in the federal funds rate, induces a “credit crunch” that manifests itself in a rising spread between CP and T-bill yields. My analysis offers an alternative interpretation: The CP/T-bill spread rises with the federal funds rate because the spread captures a liquidity premium that rises with the opportunity cost of money.

Finally, the empirical evidence in this paper provides the underpinning for recent work that interest-rate induced variation in liquidity premia as a mechanism to explain other phenomena. In Drechsler, Savov, and Schnabl (2016), a high liquidity premium reduces the willingness of banks to invest in risky assets which results in elevated risk premia. Azar, Kagy, and Schmalz (2016) study the corporate demand for liquid assets and show that changes in the opportunity costs of holding liquid assets can explain variation over time in the level of corporate liquid asset holdings.

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 concludes.

II. A MODEL OF THE LIQUIDITY PREMIUM OF 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 exchange economy in which near-money assets can earn a liquidity premium. Prices in this economy are flexible. The only purpose of introducing nominal quantities and the price level here is to allow the central bank to influence nominal short-term interest rates. The real rate will be fixed by the endowment process. The economy consists of a representative household, a government comprising the fiscal authority and the central bank, and a banking sector, which is simply a technology to transform assets and reserve holdings at the central bank into deposits that can be held by households. In this model, the household sector should be understood broadly to represent any non-bank, non-government agents in the economy.

II.A. Households

There is a single perishable consumption good with endowment stream $\{Y_t\}$. The representative household seeks to maximize the objective

$$E_0 \sum_{t=1}^{\infty} \beta^t [u(C_t) + \alpha \log(Q_t)], \tag{1}$$

where C_t is consumption and Q_t is a real stock of liquid assets that the household draws service flow from. Liquidity services are supplied by money in the form of demandable bank deposits and near-money in the form of T-Bills issued by the government. As in Poterba and Rotemberg (1987), households' stock of liquidity is a constant-elasticity-of-substitution

(CES) aggregate

$$Q_t = \left[(1 - \lambda_t) \left(\frac{D_t}{P_t} \right)^\rho + \lambda_t \left(\frac{B_t}{P_t} \right)^\rho \right]^{\frac{1}{\rho}}, \quad (2)$$

where D_t/P_t denotes real balances of demandable deposits and B_t/P_t denotes the real value of T-bill holdings.⁴ P_t is the nominal price of the consumption good.

The parameter ρ controls the elasticity of substitution, $\sigma = 1/(1 - \rho)$, between T-bills and deposits. With $\rho = 1$, the liquidity aggregate becomes linear and T-bills and deposits are perfect substitutes. In the Cobb-Douglas limiting case $\rho \rightarrow 0$, T-bills and deposits are neither substitutes nor complements. Since money and near-money assets both provide liquidity services, they should be substitutes to some extent, and hence $0 < \rho \leq 1$. Quantifying the degree to which there is strong substitutability between money and near-money assets is a key objective of the empirical analysis below.

The parameter λ_t , which I label the liquidity share of T-bills, controls the relative contribution of T-bills and deposits to the stock of liquidity. Since deposits are money and can be used directly in transactions, their contribution to liquidity should typically be greater than that of T-bills, i.e. $\lambda_t < 1/2$. Time-variation in λ_t could arise if there are changes over time in households' perception of the usefulness of T-bills, relative to the same quantity of deposits, for the purpose of holding liquidity.

To map the model into the empirical analysis that follows, one can think of one period as lasting roughly a quarter. The T-bills have a one-period maturity. The liquidity benefits of near-money asset holdings and deposits arise from their use in (unmodeled) intra-quarter transactions, and the indirect utility from these within-period benefits enters into the felicity function in (1). Furthermore, since the model does not feature any intermediation frictions, one can abstract from institutional details of how liquid assets are held by households. For example, households' T-bill holdings do not necessarily have to represent direct holdings. One could think of the liquidity aggregate Q_t as representing households' investment in money

4. In the special case of $u(C_t) = \log(C_t)$, the felicity function in (1) becomes a (monotonic) log transformation of the felicity function in Poterba and Rotemberg (1987).

market mutual funds that invest some of their funds in T-bills and some in deposits.

Households optimize subject to the flow budget constraint

$$D_t + B_t + k_t A_t - L_t + P_t C_t = D_{t-1}(1 + i_{t-1}^d) + B_{t-1}(1 + i_{t-1}^b) + k_{t-1}(A_t + P_t Y_t) - T_t - L_{t-1}(1 + i_t) + \Pi_t, \quad (3)$$

where k_t denotes the share of total endowment stream that the household owns at the end of period t , A_t is the price of the claim to the endowment stream, T_t denotes transfers, i_t^b is the T-bill yield, i_t^d is the interest rate on demand deposits, L_t denotes bank loans, Π_t is the flow of profits from banks to households, and i_t is the nominal interest rate applicable to assets and transactions that do not produce a liquidity service flow.

The household first-order condition with respect to consumption yields the Euler equation

$$1 + i_t = \frac{1}{\beta} \left\{ E_t \left[\frac{u_c(C_{t+1})}{u_c(C_t)} \frac{P_t}{P_{t+1}} \right] \right\}^{-1}, \quad (4)$$

where u_c denotes the first derivative of $u(\cdot)$. The household's first-order conditions with respect to real liquid asset balances yield

$$\frac{\alpha(1 - \lambda_t)(D_t/P_t)^{\rho-1}}{(1 - \lambda_t)(D_t/P_t)^\rho + \lambda_t(B_t/P_t)^\rho} = u_c(C_t) \frac{i_t - i_t^d}{1 + i_t}, \quad (5)$$

$$\frac{\alpha \lambda_t (B_t/P_t)^{\rho-1}}{(1 - \lambda_t)(D_t/P_t)^\rho + \lambda_t(B_t/P_t)^\rho} = u_c(C_t) \frac{i_t - i_t^b}{1 + i_t}. \quad (6)$$

The left-hand side of these equations represents the marginal benefit of holding liquidity in the form of deposits and T-bills, respectively. The right-hand side represents the marginal cost, in terms of utility, of the (discounted) foregone interest.

II.B. Banks

For the sake of parsimony, the banking sector in this model is simply a technology, owned by households, 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.

Deposits originate in two ways. First, banks create deposits by making loans to households. Second, in open-market operations the central bank trades T-bills with households. These transactions are settled through banks whereby the bank credits (debits) household deposit accounts, while the central bank debits (credits) the banks' reserve account at the central bank. Thus, the bank balance sheet identity is

$$D_t^s = L_t^s + M_t, \tag{7}$$

with deposits D_t^s on the liability side and loans, L_t^s , and reserves, M_t , as assets. Loans earn the illiquid nominal rate i_t .

The amount of deposits that banks can create is constrained by the amount of reserves they hold according to

$$D_t^s \leq \phi M_t, \tag{8}$$

where $\phi > 1$ is a constant multiplier.⁵ The constraint (8) could be motivated by explicit reserve requirements. However, even in the absence of regulatory reserve requirements, banks may have a precautionary need for reserve holdings to hedge unexpected payment flows, as modeled, for example in Ashcraft, McAndrews, and Skeie (2011) and Bianchi and Bigio (2014). Since reserves are not interest-bearing and hence costly as long as $i_t > 0$, the constraint (8) holds with equality in equilibrium.⁶

5. It would be straightforward to let ϕ change over time to capture, for example, changes in the degree to which banks can and want to economize on holding reserves. Such time-variation in ϕ is would not be important for any of the results on the liquidity premium below as long as the central bank can track movements in ϕ in real time so that the central bank knows how a change in the supply of M translates into a change in D^s .

6. Set up in this way, the model is not suitable for thinking about a situation in which the central bank creates huge amounts of excess reserves, like the Federal Reserve did since 2008. One would need to introduce a satiation level for households' money demand. In this case, if the central bank expands reserves beyond a certain threshold, the reserves constraint on deposits ceases to be binding and the satiation level would provide an upper bound on demand deposit quantities. In this situation, the marginal benefit of a deposit is zero, interest rates hit the zero lower bound, $i_t = 0$, liquidity premia are zero, and the amount of reserves could be much larger relative to the amount of deposits than in "normal" times.

Banks pay a deposit interest rate

$$i_t^d = \delta i_t, \quad 0 \leq \delta \leq 1 - 1/\phi. \quad (9)$$

These interest payments can represent implicit interest or explicit interest. Regulation Q until July 2011 explicitly prohibited explicit payment of interest on demand deposits that are not subject to any restrictions on their use and withdrawals. Interest rates on other types of deposits (e.g., money market deposit accounts) were deregulated in the early 1980s. Before the 1980s, banks often paid implicit interest in the form of reduced service fees. For the purposes of the analysis here, implicit interest is only relevant to the extent that it contributes to the marginal rate of return on a deposit. Theoretically, the magnitude of δ is not clear. For example, with competitive banks, as analyzed in Klein (1974), $\delta = 1 - 1/\phi$, while in the monopolistic model of Startz [1983]) $\delta < 1 - 1/\phi$. As I discuss in more detail in Online Appendix B, the available evidence on implicit interest rates during this period (Barro and Santomero [1972], Startz [1979]) suggests that δ is between 1/3 and 1/2.

Following deregulation in the early 1980s, payment of explicit interest on demand deposits was still prohibited, but banks developed various alternative interest-bearing accounts (e.g., NOW accounts and money market deposit accounts) that resembled demand deposits in many ways. Online Appendix B shows, based on data from Lucas and Nicolini (2015), that interest rates on such demand-deposit substitutes, too, imply that δ is between 1/3 and 1/2.

Each period, banks' flow of profit

$$\Pi_t = L_t^s i_t - D_t^s i_t \delta, \quad (10)$$

is paid out to households who own the banks.

II.C. Government: Fiscal Authority and Central Bank

The government issues liabilities in the form of one-period T-bills and reserves, M_t^s , through the central bank (CB). To change M_t^s , the CB conducts open-market operations, exchanging T-bills against reserves. The net supply of T-bills that remains available to the public after subtracting the T-bill holdings of the central bank is B_t^s . For simplicity, I assume that $\{B_t^s\}$ is an exogenous process. In other words, if the CB purchases T-bills in open-market operations, the Treasury issues additional T-bills to leave B_t^s unaffected. This assumption is not of substantive importance for anything that follows, but it leads to much more transparent expressions for the liquidity premium.

The path of M_t^s is chosen by the CB to meet an interest-rate operating target, i_t^* , according to the policy rule

$$i_t^* = \psi(P_t/P_t^*; \xi_t), \quad (11)$$

where P_t^* represents a target path for the price level and ξ_t is an exogenous random shock.

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

$$B_t^s + M_t^s = B_{t-1}^s(1 + i_{t-1}^b) + M_{t-1}^s - T_t. \quad (12)$$

II.D. Equilibrium

Due to the additive separability of preferences over consumption and liquidity in this model, the marginal utility of consumption is independent of real balances of liquidity. As a consequence, the question of price level determination can be separated from the other aspects of equilibrium. Concerning price level determination and interest-rate policy, the model here is identical to Woodford (2003), Ch. 3.2. Equilibrium is given by a set of processes $\{P_t, i_t\}$ that satisfy the policy rule (11) and the Euler equation (4) with market clearing conditions

substituted in. As discussed in Woodford (2003), with a suitable specification of the policy rule functional form and parameters, one obtains a locally unique rational expectations equilibrium. For the purposes of the analysis of liquidity premia here, these questions regarding price-level determination are tangential.

To focus on the aspects of equilibrium that are of main interest here, I proceed by assuming that policy has resulted in a specific path for $\{P_t, i_t\}$ and I solve (5) and (6), given the exogenous $\{Y_t, B_t^s\}$, for the quantity M_t^s that is required to support the chosen path of i_t , and for the liquidity premium of T-bills, $i_t - i_t^b$. In the special cases where T-bills and deposits are perfect substitutes ($\rho = 1$) or not substitutes ($\rho \rightarrow 0$, i.e., Cobb-Douglas), the solution can be stated exactly in closed form.

Special case: Perfect substitutes. In the case $\rho = 1$, the solution for the liquidity premium is

$$i_t - i_t^b = \frac{\lambda_t}{1 - \lambda_t}(1 - \delta)i_t. \quad (13)$$

Since deposits and T-bills are perfect substitutes, the liquidity benefit households get from $\lambda_t/(1 - \lambda_t)$ units of T-bill balances is always exactly equal to the liquidity benefit obtained from one unit of deposits, hence, conditional on λ_t , the relative price of obtaining liquidity through these two assets is in a fixed relationship. Time-variation in the liquidity premium is driven by changes in the opportunity cost of money $(1 - \delta)i_t$, and by changes λ_t , 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 $\lambda_t/(1 - \lambda_t)$ rises, resulting in a lower yield on T-bills and hence a bigger liquidity premium.

In the case $\rho = 1$, a liquidity demand shock does not affect the liquidity premium. In this model, a positive liquidity demand shock could result, for example, from a fall in $u_c(Y_t)$, which causes a fall in the utility cost of foregone interest in (5) and (6). (Alternatively, one could also consider an increase in α , which would have equivalent effects.) However, the liquidity premium does not change, because the CB would have to offset the positive liquidity demand shock by elastically raising M_t^s (which would change the supply of deposits) to stay at the

interest-rate target prescribed by (11). Specifically, the solution for the CB's target-consistent supply of reserves in the $\rho = 1$ case is

$$M_t^s = \frac{\alpha(1+i_t)P_t}{(1-\delta)i_t\phi u_c(Y_t)} - \frac{\lambda_t}{\phi(1-\lambda_t)}B_t^s, \quad (14)$$

which shows the negative relation between $u_c(Y_t)$ and M_t^s .

Similarly, an elastic reserve supply response of the CB would also offset the effects of a change in the supply of T-bills, B_t^s , with no remaining 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.

Special case: Not substitutes. In the case $\rho \rightarrow 0$, the solution for the liquidity premium is

$$i_t - i_t^b = \frac{\lambda_t\alpha(1+i_t)}{u_c(Y_t)} \left(\frac{P_t}{B_t^s} \right). \quad (15)$$

The same result would hold if there was no money at all in the model and hence the T-bill first-order condition (6) alone determined $i_t - i_t^b$. Thus, this special case corresponds to models like Bansal, Coleman, and Lundblad (2010), Krishnamurthy and Vissing-Jorgensen (2012), and Greenwood, Hanson, and Stein (2015) in which there is no money and time-variation in the liquidity premium is driven by changes in near-money asset supply B_t^s , liquidity demand $u_c(Y_t)$, and λ_t .

General case: Imperfect substitutes. For $0 < \rho < 1$, the model cannot be solved exactly in closed-form. Defining

$$m_t \equiv M_t^s/P_t, \quad b_t \equiv B_t^s/P_t,$$

I use a log-linear approximation of the first-order conditions (5) and (6) around steady state

values of m , b , λ , and Y , denoted \bar{m} , \bar{b} , $\bar{\lambda}$, \bar{Y} , with deviations

$$\begin{aligned}\hat{Y}_t &\equiv \log\left(\frac{Y_t}{\bar{Y}}\right), & \hat{m}_t &\equiv \log\left(\frac{m_t}{\bar{m}}\right), & \hat{b}_t &\equiv \log\left(\frac{b_t}{\bar{b}}\right), \\ \hat{\lambda}_t &\equiv \lambda_t - \bar{\lambda}, & \hat{i}_t &\equiv \log\left(\frac{1+i_t}{1+\bar{i}}\right), & \hat{i}_t^b &\equiv \log\left(\frac{1+i_t^b}{1+\bar{i}^b}\right).\end{aligned}\tag{16}$$

The resulting solution for the target-consistent supply of reserves is

$$\hat{m}_t = -\gamma_i \hat{i}_t - \gamma_\lambda \hat{\lambda}_t - \gamma_b \rho \hat{b}_t + \gamma_y \hat{Y}_t,\tag{17}$$

where γ_i , γ_λ , γ_b , and γ_y are positive constants defined in Online Appendix A, and the liquidity premium of T-bills is

$$\hat{i}_t - \hat{i}_t^b = \beta_i \rho \hat{i}_t + \beta_\lambda \hat{\lambda}_t - \beta_b (1 - \rho) \hat{b}_t + \beta_y (1 - \rho) \hat{Y}_t,\tag{18}$$

where β_i , β_λ , β_b , and β_y are positive constants defined in Online Appendix A.

II.E. Empirical specification

Equation (18) is the foundation for the empirical analysis. It nests the extreme cases of perfect and no substitutability between T-bills and deposits, as well as the intermediate cases: if $\rho = 1$, the liquidity premium varies with \hat{i}_t and with $\hat{\lambda}_t$, but not with T-bill supply \hat{b}_t and liquidity demand \hat{Y}_t (which works through $u_c(Y_t)$ as discussed above); if $\rho < 1$, the latter two channels matter for the liquidity premium, too, and they are the only ones that matter if $\rho \rightarrow 0$.

I start the empirical analysis with a baseline model motivated by the perfect substitutes case $\rho = 1$, where only \hat{i}_t and $\hat{\lambda}_t$ are used to explain the liquidity premium. I then extend the model to allow for imperfect substitution, with \hat{b}_t measured as in Krishnamurthy and Vissing-Jorgensen (2012) and Greenwood, Hanson, and Stein (2015).

In the empirical implementation, I interpret the liquidity share λ_t as being driven by the

level of uncertainty in financial markets. A high level of risk can erode agents’ “trust” that bank deposits are a good store of liquidity and thereby enhance the relative preference for T-bills.⁷ To measure the risk level, I use 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, including the recent financial crisis, 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 VIX index is only available from 1990 onwards, however. In samples that include earlier time periods, I use the projection of VIX on realized S&P500 index return volatilities (with projection coefficients estimated in the 1990-2011 sample), as explained in Online Appendix D.

Liquidity demand effects, represented by \hat{Y}_t in (18), are left unobserved in the residual of the regressions. Since leaving these unobserved effects in the residual could potentially lead to a correlated omitted variable problem, I also explore an instrumental variables strategy.

II.F. Extensions and Modifications

I close the description of the model by discussing how a number of extensions and modifications would affect the predictions concerning liquidity premia.

Monetary policy without interest-rate operating target. The relationship between liquidity premia, interest rates, near-money asset supply, and the other explanatory variables in (18) does not depend on the assumption that the CB follows an interest-rate operating target according to the policy rule (11). If the CB followed a different policy, e.g., a money-growth target, the solution for the liquidity premium would still be exactly the same as in (18), only the interpretation of interest-rate movements would be different. In this case, \hat{Y}_t , $\hat{\lambda}_t$, and \hat{b}_t would all influence \hat{i} , because the CB does not neutralize these effects in the way it would

7. In a severe financial crisis $\lambda_t > 0.5$ may be possible. In this case, T-bills have a higher liquidity share than deposits. This means that T-bill yields can fall into negative territory if deposits and T-bills are substitutes, as in (13). As an example for a model in which a crisis shock can impair the liquidity value of bank deposits see Robatto (2016).

with an interest-rate operating target.

Additional types of near-money assets. It would be straightforward to introduce additional varieties of near-money assets into the model. In this case, the liquidity aggregator in (2) would sum over more than two near-money assets and the assets' liquidity shares would again sum to one. The results for the liquidity premium can be anticipated directly from the results above. For example, in the perfect substitutes case, the liquidity premium of each near-money asset would depend on the ratio of the assets' liquidity share to the liquidity share of deposits multiplied by the interest rate, as in equation (13).

Currency. Currency could also contribute to households' stock of liquidity. To the extent that currency is a perfect substitute for deposits, one could introduce currency in the model by thinking of D_t as a weighted sum of currency and demand deposit balances and i^d as a weighted average of the return on deposits and the zero return on currency. However, perfect substitutability may not be realistic. Currency may be more convenient for some types of transactions and deposits for other ones. In particular, it seems plausible that currency is more useful for smaller transactions than for larger ones. As a consequence, while it would be very inconvenient for households to live with very small currency balances, the marginal benefits of additional currency balances may decline more quickly than for deposits. Online Appendix C extends the baseline model along these lines.

Interest on excess reserves. Since banks in this model are just a simple liquidity transformation technology, the model is not suited to analyze how the CB's reserve remuneration policy affects bank behavior and liquidity premia. In a model with a micro-founded banking sector, the payment of interest on excess reserves (IOR), i_t^m , could raise i_t^d to a higher level. If banks aggressively compete for deposits, trying to earn the spread between i_t^m and i_t^d , then, possibly, $i_t^d \approx i_t^m$. Since CBs that pay IOR typically keep i_t^m at a small fixed spread to the target policy rate, households' opportunity cost of holding money and hence liquidity premia could shrink substantially and stay constant.⁸

8. 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, Good-

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.⁹ 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.

Throughout its history until 2008, the Federal Reserve did not pay IOR, i.e., $i_t^m = 0$. In October 2008 the Federal Reserve started paying IOR, albeit at a very low level with little variation since then. To explore whether the introduction of IOR affects the behavior of liquidity premia, Section III.E looks at data from Canada and the UK, where IOR was introduced much earlier.

III. EMPIRICAL DYNAMICS OF THE LIQUIDITY PREMIUM

I now turn to an empirical evaluation of time-varying liquidity premia. After describing the data, I start with the pure opportunity-cost of money model of the liquidity premium (i.e., the perfect substitutes case $\rho = 1$) and study how much of the time-variation of the liquidity premium is explained by the level of short-term interest rates. In a second step, I broaden the perspective to allow near-money asset supply to potentially affect the liquidity premium. In the last part, I analyze whether the introduction of interest on reserves in Canada and the UK had a measurable effect on liquidity premia.

friend [2002]; Cúrdia and Woodford [2011]).

9. In the model of Ireland (2014), 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.A. Measurement of Liquidity Premia

Most of the analyses below use three-month T-bills as the near-money asset, but I also present some evidence with other highly liquid assets. Online Appendix D describes the data. All interest rates in the empirical analysis are monthly averages of daily annualized effective yields.

To measure the liquidity premium of T-bills, we need to match the T-bill with an asset that is ideally similarly safe, of similar maturity, but less liquid. I use three-month general collateral (GC) repo rates for repurchase agreements with Treasury collateral as such an illiquid rate. 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, the repo rate is virtually free of any credit risk component. However, an investment into a repo term loan is illiquid, because the investment is locked in during the term of the loan. Furthermore, offsetting an existing repo term loan with repo borrowing is costly because there is a substantial difference between borrowing and lending rates in the repo market. In contrast, a T-bill investment can be re-sold easily with a tiny bid-ask spread in a highly liquid market.¹⁰ The spread between the repo rate and T-bill yields reflects this liquidity differential (Duffee [1996]; Longstaff [2000]). As column (1) in Panel A Table I shows, over the sample period from May 1991 to December 2011, this yield spread averaged roughly 24 basis points (bps = 1/100s of a percent). Panel B reports estimates of quoted bid-ask spreads of T-bills (0.4 bps) and repo (7 bps), which highlights the difference in liquidity.

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.

10. Moreover, many market participants may find T-bills more convenient than a roll-over of overnight repos, which involves daily collateral flows and a (small) risk of payment delays.

Unfortunately, repo rates are available only back to 1991. Given that interest rates are quite persistent, it would be useful to extend the sample further. This is particularly important for analyses with Treasury security supply variables, which are even more persistent than short-term interest rates. For the period 1920 to April 1991, I use a three-month bankers' acceptance (BA) rate as an illiquid, but relatively safe three-month rate. The data series ends in the 1990s. To create a series until 2011, I use the GC repo/T-bill spread from 1991 onwards. BA were an important money market instrument in the pre-war period and the first post-war decades. BA can be traded in a secondary market, but the market is less liquid than the market for T-bills. BA are of relatively low risk because they are guaranteed by a commercial bank and hence represent an obligation of the (corporate) borrower and the commercial bank. Nevertheless, some credit risk does exist. The BA/T-bill spread is therefore a less perfect measure of the liquidity premium than the GC repo/T-bill spread. Panel A of Table I shows that BA earned an average premium of around 40 bps over the 1920 to 2011 sample, consistent with the bid-ask spreads in Panel B that are about 10 times as big as those of T-bills during the same time period.

From 1976 onwards, three-month certificate of deposit (CD) rates are available as an alternative measure of the illiquid rate. These are rates on uninsured CDs and hence the spread to T-bills also captures a credit risk component. However, outside of crisis periods (financial crisis 2007-09; savings and loans crisis in the 1980s), this credit risk component is small. In the periods since the early 1990s when both repo rate and CD rate data are available, there is typically only a small difference between CD rates and GC repo rates. Panel A of Table I shows that the CD/T-bill spread is on average a bit more than twice as big as the GC repo/T-bill spread.

While T-bills are the most liquid Treasury security in the U.S., other Treasury securities can also have 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]). I use the yield 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. The auction schedule for longer maturity notes and bonds is not as regular. The resulting irregular seasonal pattern in the liquidity premium between auction dates complicates an analysis of time-variation in the liquidity premium. 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, I make further adjustments following the method of Goldreich, Hanke, and Nath (2005) as described in Online Appendix D. Panel A of Table I shows that the average off-the-run/on-the-run premium is between one and two basis points, depending on the sample period, consistent with the relatively small difference in bid-ask spreads between off-the-run and on-the-run notes. This liquidity premium compares two securities that are both relatively liquid.

The final liquidity premium measure is the spread between three-month T-bills and less liquid off-the-run two-year Treasury notes that have approximately three months remaining maturity. Amihud and Mendelson (1991) argue that this spread reflects a liquidity premium. To calculate the spread, I compare the yield of each of these Treasury notes with the average yield of two T-bills that straddle the maturity of the Treasury note. As Table I shows, the average yield spread is slightly above 5 bps, while the bid-ask spread difference between the two assets is bigger than for the off-the-run/on-the-run spread.

One potential concern regarding some of these spreads 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. 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 to directly address this issue empirically. The Federal Home Loan Bank (FHLB) issues

short-term discount notes that receive the same tax treatment as Treasury Bills and have similar maturities. Moreover, FHLB debt is perceived as implicitly guaranteed by the Federal government (Ashcraft, Bech, and Frame [2010]). Thus, a spread between FHLB discount note yields and T-bill yields cannot be driven by taxation differences nor by credit risk. Online Appendix E shows that the FHLB discount note/T-bill spread is quantitatively similar to the repo/T-bill spread.¹¹

In addition, the off-the-run/on-the-run and T-note/T-bill spreads 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.

III.B. Liquidity Premium and the Opportunity Cost of Money

To provide a preliminary illustration of the time-variation in the liquidity premium of T-bills, Figure I plots the time series of the Repo/T-bill, CD/T-bill, and BA/T-bill spread over the sample periods starting in 1991, 1976, and 1920, respectively. Each panel also shows the time series of the empirical counterpart of i_t : the federal funds rate from July 1954 onwards and the Federal Reserve Bank of New York's discount rate before that date.¹² 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.

The baseline model of (18) with $\rho \approx 1$ suggests that liquidity premia should be positively related to interest rates. Figure I shows that this prediction is borne out in the data in all three sample periods. High interest rates coincide with a high liquidity premium. Comparing the spread plotted in Panel (a) with the spread in Panels (b) and (c) one can see that the

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

12. 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 (18). However, 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 where the data span longer time periods.

estimated liquidity premium behaves quite similarly in all three cases, both in terms of level and the co-movement with the federal funds rate, with the exception of periods of financial market turmoil. For example, during the financial crisis period starting in 2007-09, the CD/T-bill spread rose much more strongly than the repo/T-bill spread. This is likely a consequence of the fact that the rates on CDs and BAs have a credit risk component that dominates in times of crises.

But even for the repo/T-bill spread in Panel (a), it is apparent that the roughly proportional relationship with the interest-rate level breaks down when financial markets are in turmoil, as, for example, during the LTCM crisis in September 1998, or around the height of the financial crisis in 2008. Following equation (18), this is to be expected if λ_t rises during times of crises—for example, because bank customers have doubts about the safety and instant availability of deposits.

The regressions in Table II therefore also include the VIX index as a proxy for λ_t . Column (1) presents a regression with the repo/T-bill spread as the dependent variable. The results confirm the visual impression from Figure I: 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 6.50 bps (s.e. 0.62). The VIX also has a highly significant relationship with the liquidity premium: a rise in the VIX by 10 percentage points is associated with a rise of 9.6 bps in the repo/T-bill spread. As one can already anticipate from the plot in Panel (a) of Figure I, the regression has a very high R^2 of 52%.

Column (2) presents results for regressions with the CD/T-bill spread as the dependent variable, including estimates from the longer sample starting in 1976 and shown in Panel B. The results are similar. The coefficient on the fed funds rate is somewhat higher than in column (1) and again highly statistically significant. The coefficient on the VIX is substantially bigger in this case. This is to be expected: The VIX now picks up not only time-variation in the liquidity premium due to changes in the liquidity share of T-bills relative to deposits, but also time-variation in the credit risk component of CD rates.

As column (3) in Table II 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 in Panel A implies that a one percentage point change in the federal funds rate translates into a change of 0.38 bps (s.e. 0.08) 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. The regression coefficient of the federal funds rate is positive, but it is not statistically significant in the short sample in Panel A. In the longer sample in Panel B, the point estimate is quite similar, but now the higher statistical power due to the longer sample period also renders it highly statistically significant.

The relative magnitudes of the point estimates of the fed funds rate coefficients in Table II across columns roughly line up with the relative magnitudes of the average liquidity premia from Table I, where the off-the-run/on-the-run spread is roughly one twentieth and the T-note/T-bill spread about one quarter of the average repo/T-bill spread. This correspondence of interest-rate sensitivity and average liquidity premium is consistent with the $\rho = 1$ baseline version of the model in (13).

III.C. Allowing for Supply Effects

As the results so far have shown, the simple and stark model with $\rho = 1$ is remarkably consistent with the observed time-variation in liquidity premia. However, the version of (18) with imperfect substitution, i.e., $\rho < 1$, suggests that variation in the supply of near-money assets could also play a role in explaining time-variation in liquidity premia. In the extreme case of $\rho \rightarrow 0$, supply effects would dominate in the sense that interest rate variation would not matter at all once supply effects are accounted for. Hence, one might be worried that the apparent strongly positive correlation of liquidity premia and interest rates in Table II could

perhaps be driven by correlated omitted supply variables. After all, existing evidence seems to suggest strong supply effects. 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 (2015) find that the liquidity premium of T-bills is negatively related to the ratio of outstanding T-bills to GDP. Of course, it could also be the case that the correlation between liquidity premia and supply in these papers reflects the presence of a correlated interest-rate variable that was omitted in their analyses.

To sort this out, Table III adds supply variables to the regression. For these regressions I rely on the sample that extends back to 1920 with the BA/T-Bill spread as the dependent variable. The government debt/GDP ratio 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 comparison, column (1) in Table III first shows a regression that mirrors the earlier ones in Table II with the federal funds rate and VIX as the only explanatory variables. The results are similar to those from the shorter samples. The coefficient on the federal funds rate is positive and statistically highly significant. With 11.04, the magnitude of the point estimate is somewhat bigger in this longer sample, and the R^2 of 56% is again very high. Thus, the strong positive relationship between interest rates and liquidity premium extends robustly into this longer sample.

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 -43.61 (s.e. 10.57). 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 conjecture. Adding the federal funds rate as an explanatory variable basically eliminates the supply effect. The coefficient on the supply variable is now estimated at 2.22 (s.e. 4.41). In contrast, the coefficient on the

federal funds rate and its standard error are almost unchanged from column (1). Moreover, adding the interest rate variable raises the adjusted R^2 from 15% to 56%. I have further explored a specification that adds 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 (2015) that T-bill supply is strongly negatively related to the T-bill liquidity premium. However, as column (5) shows, the supply effect almost completely disappears when the federal funds rate is included in the regression: The coefficient on the T-bill supply variable drops by 95% from -94.57 (s.e. 24.12) to -4.99 (s.e. 9.99). The coefficient on the federal funds rate (11.33; s.e. 1.24) is similar in magnitude to the estimate in previous regressions and inclusion of the interest-rate variable raises the adjusted R^2 from 18% to 58%. 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 II 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 roughly 90-year sample period until 2011, 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 Panel (c) in Figure I), the regressions with the federal funds rate as explanatory variable have much

higher R^2 than those with only the debt/GDP variable.

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.

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 incremental effect on the level of the T-bill liquidity premium, consistent with a high elasticity of substitution between money and near-money.

While the results so far—a strong relationship between liquidity premia and interest rates and a weak relationship with Treasury supply variables—are suggestive of a high elasticity of substitution between deposits and Treasury Bills, looking at the above regression results through the lens of the model helps to make this assessment more precise. In Online Appendix A, I show that the regression coefficients on the interest and supply variables in the model take, approximately, the following values,

$$\rho\beta_i \approx \rho \frac{\bar{i} - \bar{i}^b}{(1 + \bar{i}^b)\bar{i}}, \quad -(1 - \rho)\beta_b \approx -(1 - \rho) \frac{\bar{i} - \bar{i}^b}{(1 + \bar{i}^b)}, \quad (19)$$

where upper bars denote steady state values around which I log-linearized to get the baseline specification (18). Setting these steady state values equal to the sample averages, we can use the regression coefficients of the federal funds rate and T-bill supply to back out an implied ρ . Taking the point estimates of the coefficients from column (5) of Table III, we get $\rho = 1.18$ based on the federal funds rate coefficient and $\rho = 0.90$ based on the T-bill supply variable coefficient. Thus, roughly, both coefficients are consistent with ρ close to one and hence a

very high elasticity of substitution.¹³ In contrast, the coefficient on T-bill supply in column (4) would yield an implied ρ of almost exactly zero. But this estimate in column (4) is biased due to the omission of the federal funds rate. This further illustrates how the inclusion of the federal funds rate in the regression leads to a substantially different economic interpretation.

III.D. Time-differenced Specification and Instrumental Variables

One potential concern with the analysis so far is that liquidity premia, short-term interest rates and the Treasury securities supply variables could have stochastic trends that generate spurious correlation between the levels of these variables. Running the analysis in a time-differenced specification would help address this concern by removing random-walk components in these series. Taking the first difference of equation (18) we get

$$\Delta \hat{i}_t - \Delta \hat{i}_t^b = \beta_i \rho \Delta \hat{i}_t + \beta_\lambda \Delta \hat{\lambda}_t - \beta_b (1 - \rho) \Delta \hat{b}_t + \Delta \epsilon_t, \quad (20)$$

where $\epsilon_t \equiv \beta_y (1 - \rho) \hat{Y}_t$.

Table IV presents estimates corresponding to (20) with the repo/T-bill spread as the dependent variable and sample period 1991-2011. Column (1) corresponds to column (1) of Table II, Panel A, but now in first differences. With 10.74 (s.e. 4.23) the coefficient on the federal funds rate in the differenced specification is bigger than in the levels specification, but, not surprisingly, differencing also leads to a substantially greater standard error. The coefficient on VIX changes is somewhat smaller than in the levels specification. The important take-away from these results is that the strong relationship between interest rates and liquidity premia is robust to differencing and hence to removing potential spuriously correlated stochastic trends.

Column (2) corresponds to column (4) of Table III, with T-bill supply and VIX as explanatory variables, but here differenced and with a shorter sample period. The coefficient

13. Online Appendix F provides GMM estimates that identify ρ from $\rho\beta_i$ and $(1 - \rho)\beta_b$ jointly. The point estimate of ρ is 1.28 with standard error 0.11.

on the supply variable of -52.75 (s.e. 18.40) is smaller than in Table III, but still highly statistically significant. This result is consistent with Greenwood, Hanson, and Stein (2015) who examine time-differenced regressions, too, albeit without controlling for changes in the level of short-term interest rates.

In stark contrast to the levels regressions, however, adding the change in the federal funds rate in column (3) does not eliminate the supply effect. Instead, the coefficient on T-bill supply shrinks only a little to -45.34 (s.e. 19.84). One way to reconcile the finding that the supply effect is absent in levels but strong in the differenced specification would be that supply effects are transitory. If so, they would leave little trace in levels, but they could play an important role in explaining short-run changes in the liquidity premium. In the model of Section II (similar to Bansal, Coleman, and Lundblad [2010]; Krishnamurthy and Vissing-Jorgensen [2012]; Greenwood, Hanson, and Stein [2015]) the effects of a shock to supply are permanent in the sense that the effect on the liquidity premium decays only if the shock to asset supply decays itself. Transitory effects could arise in a more refined model in which the elasticity of substitution between money and near-money assets depends on the time-horizon under consideration. For example, if not all agents are instantaneously adjusting their portfolios and liquidity holdings, a T-bill supply shock may have to be absorbed initially by a small group of agents who are able and willing to adjust their portfolios immediately. In this case, perfect substitution between money and near-money may be a good approximation after all agents have adjusted their holdings, but not initially following a supply shock when a small group of agents has to make big portfolio adjustments.

Allowing for such transitory supply effects requires a modification of (18). For example, if supply effects disappear after roughly a month, then equation (18) would have $\Delta\hat{b}_t$ on the right-hand side rather than \hat{b}_t , i.e.,

$$\hat{i}_t - \hat{i}_t^b = \beta_i \hat{i}_t + \beta_\lambda \hat{\lambda}_t - \beta_b (1 - \rho^s) \Delta\hat{b}_t + \epsilon_t, \quad (21)$$

where $\sigma^s = 1/(1 - \rho^s)$ would now represent the short-horizon elasticity of substitution (while

$\rho \approx 1$ at longer horizons). Taking first-differences yields

$$\Delta \hat{i}_t - \Delta \hat{i}_t^b = \beta_i \Delta \hat{i}_t + \beta_\lambda \Delta \hat{\lambda}_t - \beta_b(1 - \rho^s) \Delta \hat{b}_t + \beta_b(1 - \rho^s) \Delta \hat{b}_{t-1} + \Delta \epsilon_t, \quad (22)$$

i.e., the lagged change in supply should show up with a coefficient of the same magnitude but opposite sign compared with the contemporaneous supply change.

Column (4) in Table IV shows that the results from estimating (22) are consistent with the transitory supply effects explanation: The coefficient on the lagged change in supply is positive and, within standard error bounds, roughly of the same magnitude, but opposite sign, as the coefficient on the contemporaneous supply change. Another notable feature is that the coefficient on the change in the federal funds rate remains big, highly statistically significant and quite close in magnitude to the point estimate from the levels specifications. If added to the regressions, further lags of the change in T-bill supply are statistically insignificant and much smaller in magnitude.

The time-differenced specification is also useful to address a potential endogeneity concern with the earlier regression results. The concern is that some unobserved shock to the liquidity premium (for example, liquidity demand shocks due to \hat{Y} , which are left in the residual in regressions based on (18) and (20)) could be correlated with interest rates or the supply variables. For example, a rise in the liquidity premium due to an unobservable liquidity demand shock may prompt the Federal Reserve to lower interest rates in the same month, leading to reverse causality from the liquidity premium to the interest rate and a downward biased estimate of the interest-rate coefficient. This concern is even more relevant for differenced regressions because the residual accounts for a bigger share of the variance than in levels regressions, but the differenced specification is also suitable for an instrumental variables approach that can address the issue.

Valid instruments should be orthogonal to $\Delta \epsilon_t$. To instrument $\Delta \hat{b}_t$, I follow Greenwood, Hanson, and Stein (2015) and use month dummies to exploit the strong seasonality in T-bill supply. This seasonality arises from seasonal fluctuations in tax receipts that are plausibly

exogenous and immune to the reverse causality problem here. In my application, I use the month dummies as instruments for the lagged supply variable $\Delta \hat{b}_{t-1}$. To construct an instrument that is highly correlated with changes in interest rates, I use data on federal funds futures. These are futures contracts that settle at the end of each month based on the average federal funds rate that prevails during that month. The futures price before expiration is a risk-adjusted forecast of the average federal funds that prevails during the expiration month (Piazzesi and Swanson [2008]). Used as an instrument, the futures price in months prior to the expiration month should therefore be highly correlated with the average federal funds rate during the expiration month.

More precisely, let f_t^{t+n} denote the average month- t price of a federal funds futures expiring in month $t+n$. I instrument $\Delta \hat{i}_t$ with $f_{t-2}^t - f_{t-2}^{t-1}$, i.e., with the (risk-adjusted) time $t-2$ expectation of $\Delta \hat{i}_t$. The main reverse causality concern is about a monetary policy (federal funds rate) or fiscal (T-bill supply) reaction to a surprise liquidity premium shock, e.g., due to the onset of a sudden crisis. Twice-lagging the instruments is necessary if these surprise liquidity premium shocks that enter the unobserved ϵ_t potentially contain a random-walk component z_t and a white noise component u_t , i.e.,

$$\epsilon_t = z_t + u_t \quad \text{with} \quad z_t = z_{t-1} + \omega_t, \quad (23)$$

and where $E_{t-1}[\omega_t] = 0$ and $E_{t-1}[u_t] = 0$. In this case,

$$\Delta \epsilon_t = \omega_t + u_t - u_{t-1}. \quad (24)$$

Thus, the presence of the stationary component u_t in ϵ_t —which gets “overdifferenced” in (20) and (22)—requires instruments dated $t-2$.

In this IV identification scheme, it would not make sense to include the change in VIX as an exogenous variable. The concern that unobserved liquidity premium shocks could reverse-cause movements applies as much to the VIX as it applies to the interest rate and supply

variables. However, unlike for Treasury securities supply and interest rates, lagged instruments with strong correlation to changes in VIX are not readily available. Many big changes in the VIX represent the surprise onset of turmoil and uncertainty, which are, by definition, unpredictable. Rather than relying on weak lagged instruments for VIX changes, a better approach may be to leave the $\Delta\hat{\lambda}_t$ component in the residual and rely on the orthogonality with the instruments (i.e., the absence of seasonality and the unpredictability of VIX changes with the twice-lagged fed funds futures spread $f_{t-2}^t - f_{t-2}^{t-1}$).

Columns (5) to (8) in Table IV present the two-stage least squares (2SLS) estimation results for the IV specification. First-stage regression estimates reported in Online Appendix G show that, as expected, $\Delta\hat{\lambda}_t$ is strongly correlated with $f_{t-2}^t - f_{t-2}^{t-1}$ and the change in T-bill supply is strongly correlated with the seasonal dummies. That the instruments are not weak is also indicated by the relatively high Cragg-Donald statistic in Table IV.

Turning to the second-stage estimates in Table IV, the estimates of the coefficient on the federal funds rate change are similar to the OLS results, and the coefficient on the supply variable is somewhat bigger. Broadly, though, the results tell the same story as the OLS results. It does not seem to be the case that the OLS results are plagued by a big endogeneity bias. In particular, the 2SLS results, too, are consistent with supply effects having only a transitory effect on the liquidity premium.

The lesson from these differenced specifications is that effects of interest rates and the VIX are robust to differencing and that transitory supply effects exist that are not detected in levels regressions. To account for transitory supply effects, the regression should include supply changes as in (21). A good way to summarize the findings from the analysis in this subsection is to estimate (21) directly, without further differencing. In this case, I treat the federal funds (and the VIX) as exogenous. In the differenced specifications, instrumenting the federal funds rate had relatively little effect on the coefficient estimate. However, instrumenting did have a bigger effect on the estimate for the supply change effect. For this reason, I estimate (21) with $\Delta\hat{b}_t$ instrumented with seasonal dummies, as before (which results in an estimate of

the supply effects that is of substantially bigger magnitude than with OLS). Based on fitted values from these regressions, I decompose the time-variation in $\hat{i}_t - \hat{i}_t^b$ into the components related to \hat{i}_t , $\hat{\lambda}_t$, $\Delta\hat{b}_t$, and the residual ϵ_t .¹⁴

Figure III presents the results of this decomposition. The plots illustrate well that the interest-rate component and VIX play a much bigger role in explaining time-variation in the liquidity premium than supply changes. While supply changes play an economically significant in month-to-month changes in the liquidity premium, the lack of persistence in supply changes means that these effects revert quickly and do not leave much trace in the levels of liquidity premia.

III.E. The Effect of Reserve Remuneration Policies

As discussed in Section II.F, how payment of IOR affects liquidity premia is an open empirical 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.

Due to the extremely low level of short-term interest rates since IOR introduction in the U.S. in 2008, the IOR is unlikely to have a detectable effect on liquidity premia. However, a number of other countries introduced IOR at earlier points in time when short-term interest rates were higher. In those cases, effects on liquidity premia could be big enough to be detectable. 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 in the UK or the prime commercial paper (CP) rate in Canada as a proxy for the market rate for (illiquid) 3-month term loans.

The top panel in Figure IV 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

14. The coefficient estimates are 6.53 (s.e. 0.58), 1.06 (s.e. 0.19), -80.66 (s.e. 15.54) for \hat{i}_t , $\hat{\lambda}_t$, $\Delta\hat{b}_t$, respectively. While $\Delta\hat{b}_t$ is instrumented with seasonal dummies in estimation, the calculation of fitted values multiplies the coefficient estimate for $\Delta\hat{b}_t$ with the (de-meanned) realization of $\Delta\hat{b}_t$ (rather than its projection on the instruments).

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, 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 IV 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 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 IV. 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.

The bottom panel in Figure IV shows the spread between prime commercial paper (CP) rates and Canadian T-bills at 3-month maturity. The overnight interest rate in this figure is the CORRA overnight GC repo rate since December 1997, and the overnight Canada dollar LIBOR rate prior to that date. As the figure shows, the liquidity premium of Canadian T-bills exhibits similar time-series behavior as the liquidity premia in the US and UK: It is positive, 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 IV, 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 opportunity cost of money for the non-bank public: 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 IV 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 V presents regression evidence. Columns (1) and (3) first replicate the baseline regressions from Table II, based on equation (18) with $\rho = 1$, for the Canada and the UK. In both cases, there is a positive relationship between the level of the overnight rate and liquidity premia, which suggests that the opportunity cost of money is an important determinant of liquidity premia in Canada and the UK, too. While the coefficient estimate for Canada (1.49, s.e. 0.70) is smaller than in US data, the estimate for the UK (4.64, s.e. 1.13) is close to the US estimates in Table II.¹⁵

Columns (2) and (4) add an interaction variable between the short-term interest rate and $\mathbb{1}_{IOR}$, a dummy for the time periods following the introduction of IOR. If the introduction of IOR leads banks to raise deposit rates from δi_t to $\delta_{IOR} i_t$, with $\delta \leq \delta_{IOR} \leq 1$, this shrinks the interest-related component of the liquidity premium from $\beta_i \hat{i}_t$ prior to IOR to $\beta_i [(1 - \delta_{IOR}) / (1 - \delta)] \hat{i}_t$ after the introduction of IOR. We can express the liquidity premium for the pre- and post-IOR introduction periods jointly as

$$\hat{i}_t - \hat{i}_t^b = \beta_i \hat{i}_t - \beta_i \frac{\delta_{IOR} - \delta}{1 - \delta} (\mathbb{1}_{IOR} \times \hat{i}_t) + \beta_\lambda \hat{\lambda}_t. \quad (25)$$

15. One likely explanation why the coefficient on the short-term interest rate is smaller for Canadian T-bills than for U.S. T-bills is that Canadian T-bills are less liquid than U.S. T-bills and hence earn a smaller liquidity premium. Gravelle (1999) finds that the bid-ask spread of Canadian 90-day T-bills during the 1990s is about 1.5 basis points which is almost 4 times as big as the bid-ask spread of U.S. T-bills during the same time period (see Table I).

If it's true that paying IOR at level i_t^m close to i_t raises deposit rates close to i_t then δ_{IOR} should be close to one and the interaction variable should have a negative coefficient close in magnitude to the coefficient on \hat{i}_t . In contrast, if payment of IOR has little effect on deposit rates and $\delta_{IOR} \approx \delta$, then the interaction variable should get a zero coefficient.

The evidence in Table V is consistent with the latter case: In both countries, the point estimates of the coefficient on the interaction variable are positive, rather than negative, and not statistically distinguishable from zero. If anything, liquidity premia moved somewhat more strongly with the short-term interest rate in the post-IOR introduction period rather than less. However, the standard errors are quite large. The variation in interest rates following the introduction of IOR was quite limited and so it is difficult to pin down precisely how the behavior of liquidity premia in relation to the interest rate have changed with the introduction of IOR.

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.

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. Online Appendix H 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 the non-bank public.

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. Broad access would circumvent frictions that can prevent market participants from arbitraging discrepancies between IOR and open-market rates. Recent experience with IOR in the US illustrates the problem. After the Federal Reserve introduced IOR in October 2008, the federal funds rate 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 funds aggressively enough to push the federal funds rate to the level of IOR. This illustrates that payment of IOR to a narrow set of institutions does not automatically establish the level of IOR as the floor for money market rates. The Federal Reserve’s new Reverse Repo Facility is designed to address this problem by effectively offering IOR (minus a small spread) to a broader range of institutions that would otherwise not have access to deposits at the Federal Reserve (Frost, Logan, Martin, McCabe, Natalucci, and Remache [2015]) .

IV. 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 cost of holding money is high and market participants are willing to pay a substantial 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.

Recognizing the substitution relationship between money and near-money leads to a quite different perspective on the role of near-money asset supply compared with the recent literature on this topic. Controlling for the opportunity cost of money eliminates the otherwise seemingly strong negative relationship between Treasury securities supply and liquidity pre-

mia. Only at short horizons, supply changes have some effect in moving liquidity premia, but the effects are transitory and revert within a month. Overall, the results indicate a high elasticity of substitution between money and near-money assets.

This elasticity of substitution is important for understanding the effects of policies—fiscal, monetary, or regulatory—that change the supply of or the demand for near-money assets. With a high elasticity of substitution, a policy-induced rise in near-money asset demand (e.g., due to regulatory collateral requirements) has little effect on the relative price of holding liquidity in near-money (i.e., the T-bill liquidity premium) and money (i.e., the interest rate level). Moreover, a central bank that elastically supplies money to neutralize unwanted deviations from a short-term interest-rate target automatically neutralizes any effect of these near-money asset supply or demand shifts on the interest rate. As a result, the near-money asset supply or demand shifts may be associated with substantial changes in quantities, but little change in liquidity premia.

A high elasticity of substitution implies that making money interest-bearing could dramatically shrink liquidity premia. Payment of interest on reserves (IOR) is one step in this direction, but it directly reduces the opportunity cost of money only for a small set of market participants with access to reserve accounts at the central bank. In Canada and the UK, the introduction of IOR seems to have had little effect, so far, on liquidity premia and the opportunity cost of monies held by the non-bank public.

UNIVERSITY OF MICHIGAN, NATIONAL BUREAU OF ECONOMIC RESEARCH, AND CENTRE FOR ECONOMIC POLICY RESEARCH

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, Morten L. Bech, and W. Scott Frame, 2010, “The Federal Home Loan Bank System: The Lender of Next-to-Last Resort?,” *Journal of Money, Credit and Banking* 42, 551–583.
- 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, 2016, “Can Changes in the Cost of Cash Explain the Corporate Cash Puzzle?,” *Review of Financial Studies, forthcoming*.
- 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.
- Barro, Robert J, and Anthony M Santomero, 1972, “Household Money Holdings and the Demand Deposit Rate,” *Journal of Money, Credit and Banking* 4, 397–413.
- 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, 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 Ilian Mihov, 1998, “Measuring Monetary Policy,” *Quarterly Journal of Economics* 113, 869–902.
- Bianchi, Javier, and Saki Bigio, 2014, “Banks, Liquidity Management and Monetary Policy,” Working paper, University of Wisconsin and Columbia University.
- 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, “Carry Trades and Currency Crashes,” in *NBER Macroeconomics Annual 2008*, Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, eds., pp. 313–347 (University of Chicago Press, Chicago, IL, 2008).

- Cagan, Phillip, “A Study of Liquidity Premiums on Federal and Municipal Government Securities,” in *Essays on Interest Rates*, Jack M. Guttentag and Phillip Cagan, eds., Volume 1, pp. 107–142 (NBER, New York, NY, 1969).
- 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.
- 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, 2016, “A Model of Monetary Policy and Risk Premia,” *Journal of Finance*, *forthcoming*.
- Duffee, Gregory R., 1996, “Idiosyncratic Variation in Treasury Bill Yields,” *Journal of Finance* 51, 527–552.
- Fleming, Michael, and Asani Sarkar, 1999, “Liquidity in U.S. Treasury Spot and Futures Markets,” Working paper, Federal Reserve Bank of New York.
- Friedman, Milton. *A Program for Monetary Stability* (Fordham University Press, New York, NY, 1960).
- Frost, Joshua, Lorie Logan, Antoine Martin, Patrick E McCabe, Fabio Massimo Natalucci, and Julie Remache, 2015, “Overnight RRP Operations as a Monetary Policy Tool: Some Design Considerations,” Working paper, Federal Reserve Board.
- 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.
- Gravelle, Toni, 1999, “Liquidity of the Government of Canada Securities Market: Stylized Facts and Some Market Microstructure Comparisons to the United States Treasury Market,” Working paper, Bank of Canada.
- Greenwood, Robin, Samuel G. Hanson, and Jeremy C. Stein, 2015, “A Comparative-Advantage Approach to Government Debt Maturity,” *Journal of Finance* 70, 1683–1722.
- Ireland, Peter N., 2014, “The Macroeconomic Effects of Interest on Reserves,” *Macroeconomic Dynamics* 18, 1271–1312.
- Kessel, Reuben A., “The Cyclical Behavior of the Term Structure of Interest Rates,” in *Essays on Interest Rates*, Jack M. Guttentag, ed., Volume 2, pp. 337–390 (NBER, New York, NY, 1971).
- Klein, Benjamin, 1974, “Competitive Interest Payments on Bank Deposits and the Long-Run Demand for Money,” *American Economic Review*, 931–949.
- 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, 2015, "The Impact of Treasury Supply on Financial Sector Lending and Stability," *Journal of Financial Economics*, forthcoming.
- 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, 511–526.
- 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.
- Lucas, Robert E., and Juan Pablo Nicolini, 2015, "On the Stability of Money Demand," *Journal of Monetary Economics* 73, 48–65.
- Melton, William C, and Jean M Mahr, 1981, "Bankers' Acceptances," *Federal Reserve Bank of New York Quarterly Review*, 39–55.
- Morris, Marc D., and John R. Walter, "Large Negotiable Certificates of Deposit," in *Instruments of the Money Market*, Timothy Q. Cook and Robert K. Laroche, eds., (1998).
- 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.
- Piazzesi, Monika, and Eric T Swanson, 2008, "Futures Prices as Risk-Adjusted Forecasts of Monetary Policy," *Journal of Monetary Economics* 55, 677–691.
- Poterba, James M., and Julio J. Rotemberg, "Money in the Utility Function: An Empirical Implementation," in *New Approaches to Monetary Economics: Proceedings of the Second International Symposium in Economic Theory and Econometrics*, William A. Barnett and Kenneth J. Singleton, eds., pp. 219–240 (Cambridge University Press, New York, 1987).
- Robatto, Roberto, 2016, "Flight to Liquidity and Systemic Bank Runs," Working paper, University of Wisconsin.
- Startz, Richard, 1979, "Implicit Interest on Demand Deposits," *Journal of Monetary Economics* 5, 515–534.
- Startz, Richard, 1983, "Competition and Interest Rate Ceilings in Commercial Banking," *Quarterly Journal of Economics* 98, 255–265.
- Stock, James H., and Motohiro Yogo, "Testing for Weak Instruments in Linear IV Regression," in *Identification and Inference for Econometric Models: Essays in Honor*

of *Thomas Rothenberg*, Donald W.K. Andrews and James H. Stock, eds., pp. 80–108 (Cambridge University Press, Cambridge, 2005).

Sunderam, Adi, 2015, “Money Creation and the Shadow Banking System,” *Review of Financial Studies*.

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

TABLE I
SUMMARY STATISTICS: LIQUIDITY PREMIA AND BID-ASK SPREADS

	Repo/T-Bill (1)	BAcc/T-Bill (2)	CD/T-Bill (3)	2y Off/OnRun (4)	T-Note/T-Bill (5)
Panel A: Mean and standard deviation of yield spreads (basis points)					
1991 - 2011	23.65		48.12	0.99	5.54
(s.d.)	(18.19)		(51.24)	(2.20)	(7.92)
1976 - 2011			69.43	1.68	8.26
(s.d.)			(104.49)	(4.76)	(13.83)
1920 - 2011		40.19			
(s.d.)		(48.58)			
Panel B: Typical quoted bid-ask spreads (basis points of yield)					
1990s - 2000s	7.0/0.4		5.0/0.4	1.0/0.4	> 1.0/0.4
1970s - 1980s		10.0/1.3	10.0/1.3		

Notes: Bid-ask spreads reported in Panel B are expressed in terms of basis points of annualized yield for the two assets that are used to calculate the yield spread in each column. The bid-ask spread estimates are obtained from the following sources: Three-month general collateral repo: Average difference between dealer receives and dealer pays rate on Bloomberg from May 1991 to June 2007; T-bills: Fleming and Sarkar (1999) and Melton and Mahr (1981); Bankers' acceptances: Melton and Mahr (1981); CDs: Morris and Walter (1998) and Melton and Mahr (1981); Two-year Treasury notes (on-the-run and just off-the-run, and close to maturity): Goldreich, Hanke, and Nath (2005)

TABLE II
LIQUIDITY PREMIA AND THE OPPORTUNITY COST OF MONEY: BASELINE SPECIFICATION

	Repo/T-Bill (1)	CD/T-Bill (2)	2y Off/OnRun (3)	T-Note/T-Bill (4)
Panel A: May 1991 - December 2011				
Fed funds rate	6.50 (0.62)	7.56 (1.75)	0.38 (0.08)	0.80 (0.53)
VIX	0.96 (0.18)	4.16 (1.35)	0.02 (0.04)	0.54 (0.11)
Intercept	-18.38 (4.09)	-62.80 (24.16)	-0.75 (0.77)	-8.26 (3.09)
Adj. R^2 (%)	52.33	39.92	10.82	27.46
#Obs.	248	248	248	248
Panel B: January 1976 - December 2011				
Fed funds rate		9.85 (1.65)	0.33 (0.13)	1.11 (0.27)
VIX		4.20 (1.12)	0.05 (0.05)	0.61 (0.16)
Intercept		-74.07 (20.71)	-0.90 (1.19)	-8.81 (3.42)
Adj. R^2		34.79	5.56	14.21
#Obs.		432	432	432

Notes: Newey-West standard errors (12 lags) are shown in parentheses. 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. VIX refers to the CBOE S&P500 implied volatility index. From January 1990 onwards, the data is the monthly average of the daily VIX index. Prior to 1990, the VIX variable refers to the linear projection of the monthly average VIX index on the monthly average of daily squared returns on the S&P500 during the same month with the projection coefficients estimated with data from January 1990 to December 2011.

TABLE III
LIQUIDITY PREMIA AND THE OPPORTUNITY COST OF MONEY: INCLUDING
NEAR-MONEY ASSET SUPPLY

	1920 - 2011			1947 - 2011		
	(1)	(2)	(3)	(4)	(5)	(6)
Fed funds rate	11.04 (1.15)		11.18 (1.22)		11.33 (1.24)	10.96 (1.20)
VIX	1.22 (0.22)	0.59 (0.31)	1.23 (0.23)	1.66 (0.70)	1.56 (0.44)	1.55 (0.44)
log(Debt/GDP)		-43.61 (10.57)	2.22 (4.41)			-7.87 (10.89)
log(T-Bill/GDP)				-94.57 (24.12)	-4.99 (9.99)	-1.23 (10.67)
Intercept	-30.62 (7.60)	-12.32 (9.68)	-29.50 (7.63)	-217.34 (62.18)	-50.42 (27.86)	-46.20 (27.18)
Adj. R^2 (%)	55.90	14.86	55.89	17.82	58.05	58.10
#Obs.	1,104	1,104	1,104	780	780	780

Notes: The dependent variable is the spread, in basis points, between the three-month bankers acceptance rate and the three-month T-Bill yield from January 1920 until April 1991 and the spread between three-month GC repo and three-month T-bill from May 1991 to December 2011. The interest-rate data consist of monthly averages of daily or weekly rates. VIX refers to the CBOE S&P500 implied volatility index. From January 1990 onwards, the data is the monthly average of the daily VIX index. Prior to 1990, the VIX variable refers to the linear projection of the monthly average VIX index on the monthly average of daily squared returns on the S&P500 during the same month with the projection coefficients estimated with data from January 1990 to December 2011. The other explanatory variables are the federal funds rate (Federal Reserve Bank of New York's 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.

TABLE IV
MONTHLY CHANGES IN THE T-BILL LIQUIDITY PREMIUM

	OLS				2SLS			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Fed funds rate	10.74 (4.23)		8.20 (4.08)	11.04 (4.84)	10.46 (4.75)		7.05 (4.33)	12.46 (4.47)
Δ VIX	0.52 (0.27)	0.65 (0.28)	0.67 (0.26)	0.56 (0.26)				
$\Delta \log(\text{T-Bill}/\text{GDP})_t$		-52.75 (18.40)	-45.34 (19.84)	-40.08 (19.41)		-90.92 (21.37)	-87.00 (22.01)	-68.06 (24.98)
$\Delta \log(\text{T-Bill}/\text{GDP})_{t-1}$				50.12 (19.95)				98.05 (24.90)
Intercept	0.19 (0.38)	-0.03 (0.39)	0.16 (0.38)	0.19 (0.39)	0.20 (0.42)	-0.00 (0.45)	0.16 (0.47)	0.23 (0.47)
Adj. R^2	4.79	6.36	7.72	11.67				
<i>Instruments</i>								
Fed funds futures					Y	Y	Y	Y
Seasonal dummies					Y	Y	Y	Y
<i>Weak instruments test</i>								
CD statistic					11.50	12.53	10.46	9.29
SY critical value					[6.53]	[6.53]	[6.22]	[5.90]
#Obs.	247	247	247	246	246	246	246	246

Notes: The dependent variable is the first difference of the monthly averages of the daily GC repo/T-bill spread (in basis points) from Table II, column (1) with sample period from May 1991 to December 2011. Instruments in columns (5) to (8): Seasonal dummies for calendar month and the average price difference in month $t - 2$ of federal funds futures for month t and $t - 1$. Newey-West standard errors (6 lags) are shown in parentheses. A CD (Cragg-Donald) statistic greater than the critical values in brackets from Stock and Yogo (2005) rejects the hypothesis of weak instruments (with bias greater than 20% of the OLS bias) at a significance level of 5%.

TABLE V
LIQUIDITY PREMIUM AND INTEREST ON RESERVES: CANADA AND UK

	Canada		UK	
	CP/T-Bill (1)	CP/T-Bill (2)	CD/T-Bill (3)	CD/T-Bill (4)
ON rate	1.49 (0.70)	1.77 (0.67)	4.64 (1.13)	4.69 (1.23)
VIX	1.48 (0.47)	1.45 (0.47)	1.96 (0.99)	1.97 (1.01)
ON rate \times post-IOR dummy		1.60 (1.84)		0.28 (2.18)
Intercept	-12.38 (8.04)	-15.37 (8.94)	-33.02 (18.34)	-33.84 (21.78)
Adj. R^2 (%)	22.33	23.18	37.86	37.72
#Obs.	432	432	408	408

Notes: In columns (1) and (2), the dependent variable is the spread (in basis points) between Canadian 3-month prime commercial paper yields and Canadian 3-month T-bill yields from January 1976 to December 2011. The dependent variable in columns (3) and (4) is the UK 3-month CD/T-bill spread (in basis points) January 1978 to December 2011. Both spreads are measured as monthly averages of daily rates. The ON rate for Canada is the Canadian Overnight Repo Rate Average (CORRA) and Sterling Over Night Index Average (SONIA) for the UK, both in percent. VIX refers to the CBOE S&P500 implied volatility index. From January 1990 onwards, the data is the monthly average of the daily VIX index. Prior to 1990, the VIX variable refers to the linear projection of the monthly average VIX index on the monthly average of daily squared returns on the S&P500 during the same month with the projection coefficients estimated with data from January 1990 to December 2011. Newey-West standard errors (12 lags) are shown in parentheses.

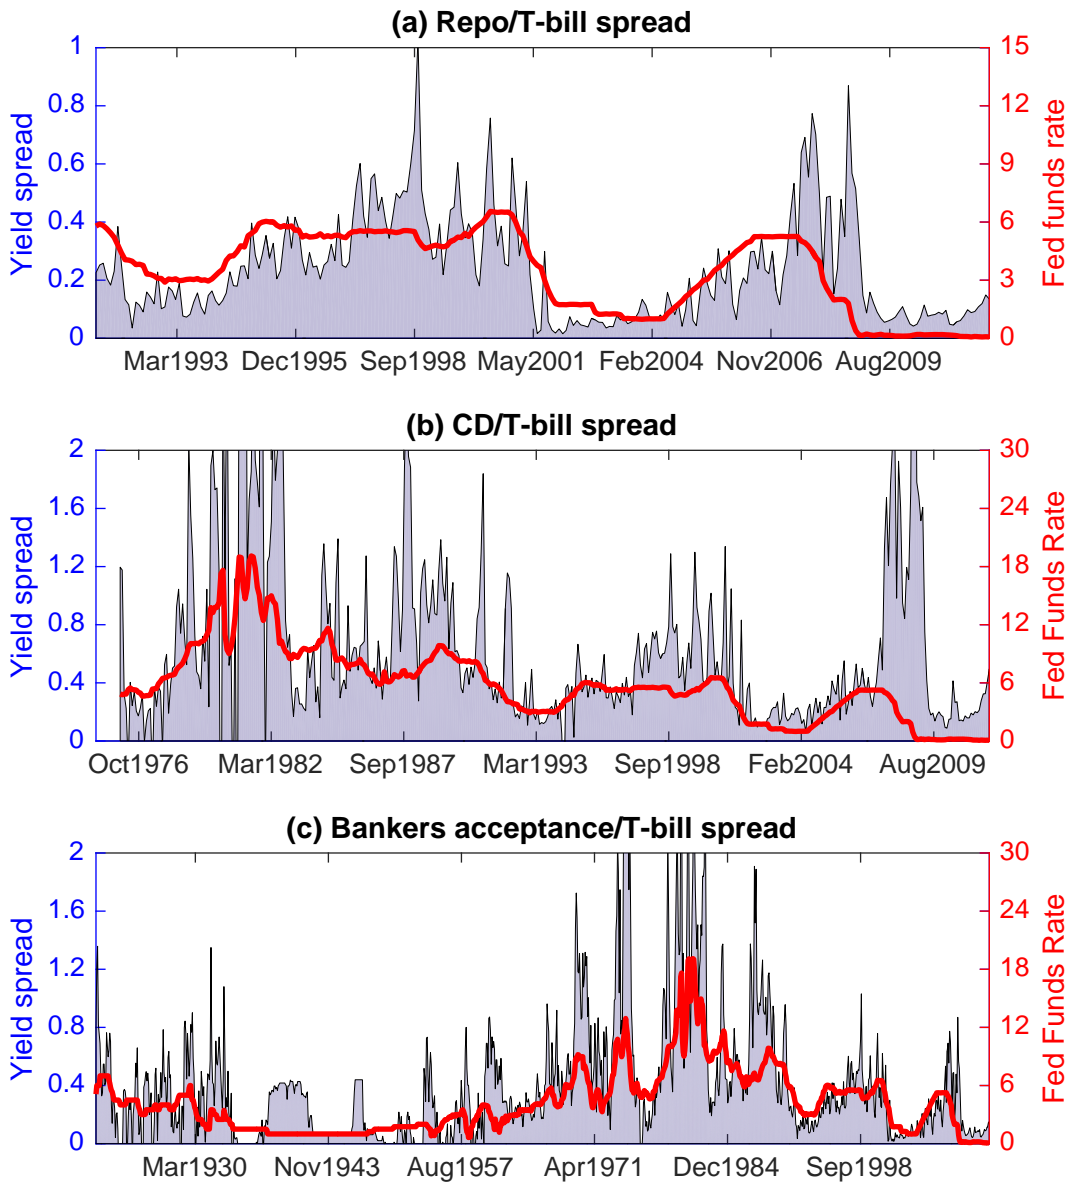


FIGURE I

Time-variation in the T-Bill Liquidity Premium and the Level of Short-Term Interest Rates
 The liquidity premium is shown as the shaded area with corresponding axis on the left; the level of the federal funds rate (Federal Reserve Bank of New York discount rate prior to July 1954) is shown as a solid line with corresponding axis on the right.

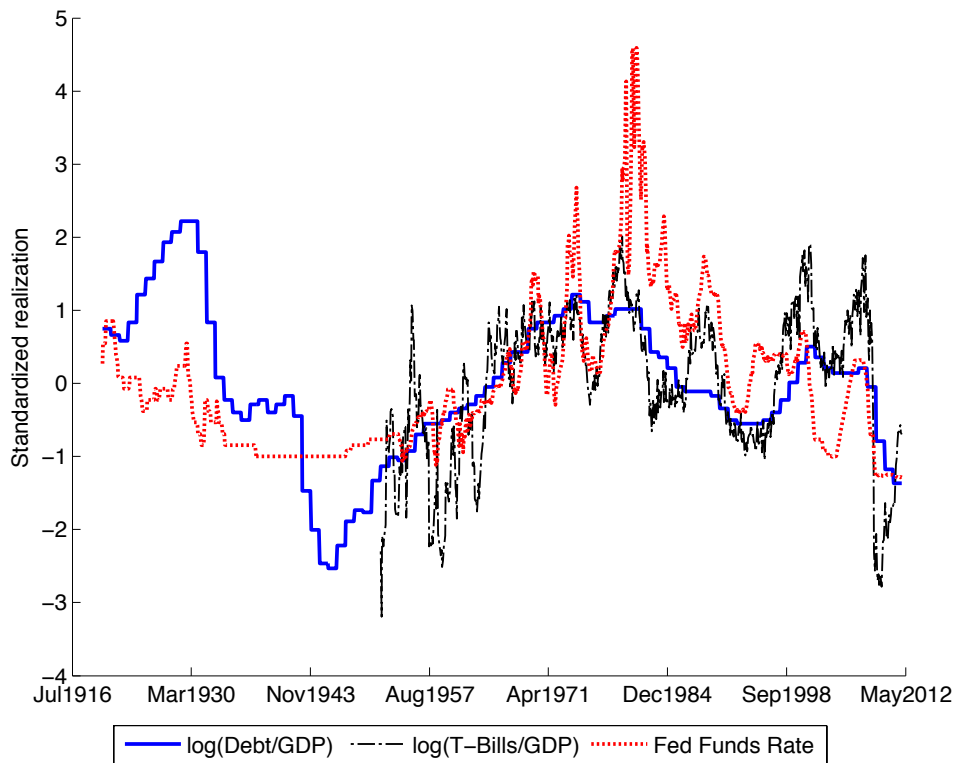


FIGURE II
Standardized explanatory variables in BA/T-bill spread regressions




FIGURE III
 Decomposition of de-meaned repo/T-bill spread

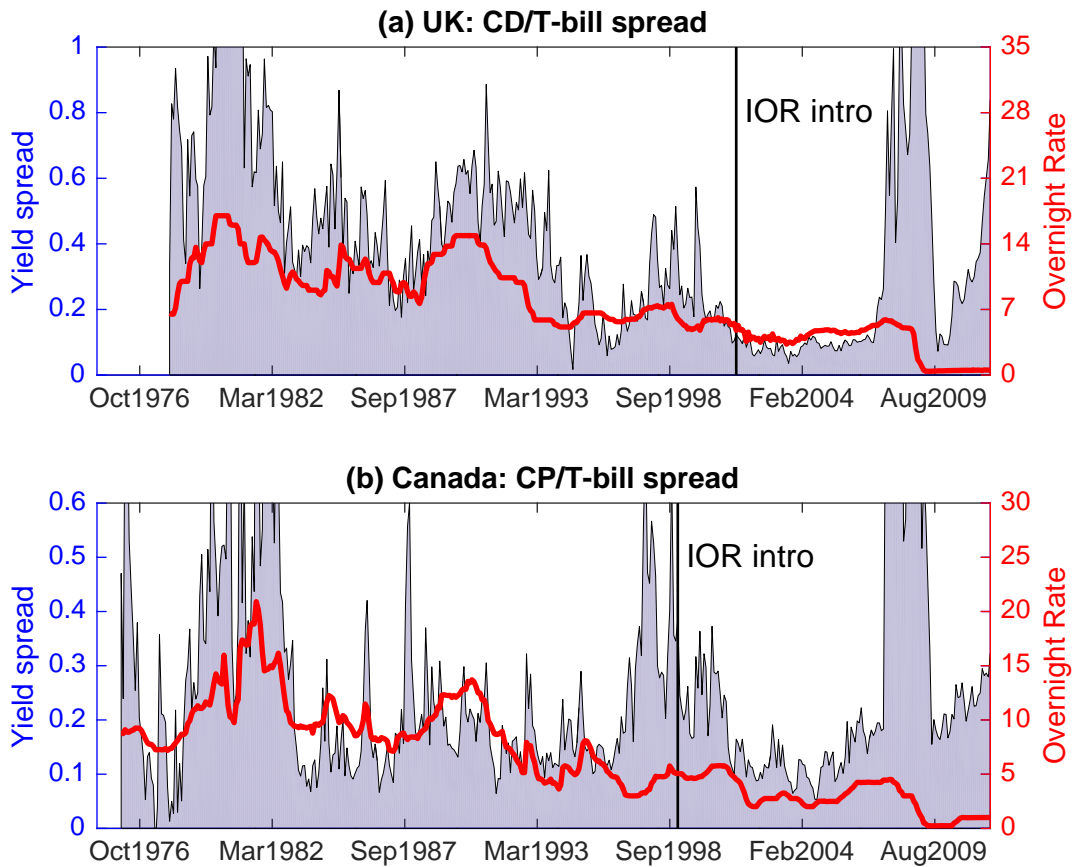


FIGURE IV

Time-variation in the T-Bill Liquidity Premium and the Level of Short-Term Interest Rates in the UK and Canada

The liquidity premium is shown as the shaded area with corresponding axis on the left; the level of the overnight interbank rate is shown as a solid line with corresponding axis on the right.