

A Comparative-Advantage Approach to Government Debt Maturity*

Robin Greenwood
Harvard University and NBER

Samuel G. Hanson
Harvard University

Jeremy C. Stein
Harvard University and NBER

First draft: July 2010
Current draft: January 2012

Abstract

We study optimal government debt maturity in a model where investors derive monetary services from holding riskless short-term securities. In a simple setting where the government is the only issuer of such riskless paper, it trades off the monetary premium associated with short-term debt against the refinancing risk implied by the need to roll over its debt more often. We then extend the model to allow private financial intermediaries to compete with the government in the provision of money-like claims. We argue that if there are negative externalities associated with private money creation, the government should tilt its issuance more towards short maturities. The idea is that the government may have a comparative advantage relative to the private sector in bearing refinancing risk, and hence should aim to partially crowd out the private sector's use of short-term debt.

* We thank Robert Barro, Markus Brunnermeier, John Campbell, Martin Eichenbaum, Ken Froot, Kenneth Garbade, Julio Rotemberg, Andrei Shleifer, Erik Stafford, Adi Sunderam, Matt Weinzierl, three referees, and seminar participants at Harvard, LSE, the NBER 2010 Corporate Finance Meetings, the University of Miami, Ohio State University, NYU, UCSD, UCLA, UC Berkeley, and Bocconi for helpful comments and suggestions.

I. Introduction

In this paper, we study the question of how the government should optimally determine the maturity structure of its debt. We focus on situations where there is no question about the government's ability to service its obligations, so the analysis should be thought of as applying to countries like the U.S. that are seen to be of high credit quality.¹ The primary novelty of our approach is that we emphasize the monetary benefits that investors derive from holding riskless securities, such as short-term Treasury bills. These benefits lead T-bills to embed a convenience premium, i.e., to have a lower yield than would be expected from a traditional asset-pricing model.

We begin with the case where the government is the *only entity* able to create riskless money-like securities. In this case, optimal debt maturity turns on a simple tradeoff. On the one hand, as the government tilts its issuance to shorter maturities, it generates more in the way of monetary services that are socially valuable; this is reflected in a lower expected financing cost. On the other hand, a strategy of short-term financing also exposes the government to rollover risk, given that future interest rates are unpredictable. As a number of previous papers have observed, such rollover risk leads to real costs insofar as it makes future taxes more volatile.²

This tradeoff yields a well-defined interior optimum for government debt maturity, unlike traditional tax-smoothing models which imply that government debt should be very long term. It also implies a number of comparative statics that are borne out in the data. Most notably, it predicts that government debt maturity will be positively correlated with the ratio of government debt to GDP, a pattern which emerges strongly in U.S. data. The intuition is that as the aggregate debt burden grows, the costs associated with rollover risk—and hence with failing to smooth taxes—loom larger.

¹ This is by contrast to a literature (e.g., Blanchard and Missale (1994)) that argues that countries with significant default or inflation risk may have a signaling motive for favoring short-term debt, or at the extreme, may have little choice but to issue short-term securities.

² See, e.g. Barro (1979), Lucas and Stokey (1983), Bohn (1990), Angeletos (2002), Aiyagari, Marcet, Sargent, and Seppälä (2002), and Nosbusch (2008).

The simple tradeoff model also captures the way in which Treasury and Federal Reserve practitioners have traditionally framed the debt-maturity problem. According to former Treasury Secretary Lawrence Summers:³

“I think the right theory is that one tries to [borrow] short to save money but not [so much as] to be imprudent with respect to rollover risk. Hence there is certain tolerance for [short term] debt but marginal debt once [total] debt goes up has to be more long term.”

Our focus on the monetary services associated with short-term T-bills is crucial for understanding Summers’ premise that the government should borrow short to “save money”.⁴ As we demonstrate formally below, if short-term T-bills have a lower expected return than longer-term Treasury bonds simply because they are less risky in a standard asset-pricing sense (i.e., because they have a lower beta with respect to a rationally priced risk factor) this does not amount to a coherent rationale for the government to tilt to the short end of the curve, any more so than it would make sense for the government to take a long position in highly-leveraged S&P 500 call options because of the positive expected returns associated with bearing this market risk.

After fleshing out the simple tradeoff model, we go on to examine the case where the government is not the only entity that can create riskless money-like claims, but instead competes with the private sector in doing so. Following Gorton and Metrick (2011), Gorton (2010), and Stein (2012), we argue that financial intermediaries engage in private money creation, thereby capturing the same monetary convenience premium, when they issue certain forms of collateralized short-term debt—e.g., overnight repo, or asset-backed commercial paper. As Stein (2012) observes, the incentives for such private money creation can be excessive from a social point of view, as individual

³ Private email correspondence April 28, 2008, also cited in Greenwood, Hanson and Stein (2010).

⁴ In a similar spirit, Bennett, Garbade and Kambhu (2000) explain the appeal of short-term financing by saying: “Minimizing the cost of funding the federal debt is a leading objective of Treasury debt management...liquidity is an important determinant of borrowing costs...Longer-maturity debt is inherently less liquid than short-term debt...”

intermediaries do not fully take into account the social costs of the fire sales that can arise from a heavy reliance on short-term financing.

In the presence of these fire-sale externalities, there is an additional motivation for the government to shift its own issuance towards short-term bills. By doing so, it reduces the equilibrium money premium, thereby partially crowding out the private sector's socially excessive issuance of short-term claims. This is desirable as long as the marginal social costs associated with government money creation—making future taxes more volatile—remain lower than the marginal social costs associated with private money creation, which stem from fire sales. In other words, the government should keep issuing short-term bills as long as it has a *comparative advantage* over the private sector in the production of riskless money-like securities.

This line of reasoning adds what is effectively a regulatory dimension to the government's debt-maturity choice. An alternative way to address the fire-sales externalities associated with private money creation would be to try to control the volume of such money creation directly, e.g., with either a regulatory limit or a Pigouvian tax on short-term debt use by financial intermediaries. However, we show that to the extent that such caps and taxes are difficult or costly to enforce—say because some of the money creation can migrate to the unregulated “shadow banking” sector—there will also be a complementary role for a policy that reduces the incentive for private intermediaries to engage in money creation in the first place, by lowering the convenience premium that money commands. The more costly are such direct regulations, the larger is the role for government debt management in reducing private money creation.

To be clear, we intend for this comparative-advantage argument to be taken in a normative, rather than positive spirit. That is, unlike with the simple government-only model, we don't mean to suggest that the comparative-advantage aspect of the theory provides further testable predictions regarding how governments have historically chosen their debt maturity structures. Rather, we offer it

as a framework for thinking about policy going forward—albeit one grounded in an empirically-relevant set of premises. In this sense, it is like other recent work on financial regulatory reform.

The ideas here build on five strands of research. First there is the prior theoretical work on government debt maturity, especially that which has put forward a tax-smoothing motive for long-term finance (Barro (1979), Lucas and Stokey (1983), Bohn (1990), Angeletos (2002), and Nosbusch (2008))⁵. Second, a series of empirical studies examine how and why government debt maturity structure varies over time and across countries (Blanchard and Missale (1991) and Missale (1999)). Third, there is a literature that documents significant deviations from the predictions of standard asset-pricing models—patterns which can be thought of as reflecting money-like convenience services—in the pricing of Treasury securities generally, and in the pricing of short-term T-bills more specifically (Krishnamurthy and Vissing-Jorgensen (2010), Greenwood and Vayanos (2010), Duffee (1996), Gurkaynak, Sack and Wright (2006), Bansal and Coleman (1996)). Fourth, there is the set of recent papers alluded to above, which emphasize how private intermediaries try to capture the money premium by relying heavily on short-term debt, even when this creates systemic instabilities (Gorton and Metrick (2011), Gorton (2010), and Stein (2012)). And finally, there is evidence that changes in government debt maturity influence private-sector debt-maturity choices, consistent with a crowding-out view: when the government issues more short-term debt, private firms issue less, and substitute towards long-term debt instead (Greenwood, Hanson and Stein (2010)).

In Section II, we further motivate our theory by laying out a set of key stylized facts, drawing on the papers cited just above, as well as on some new empirical work of our own. In Section III, we develop the simple tradeoff model of optimal debt maturity when the government is the only entity that can create riskless money-like securities. In Section IV, we add financial intermediaries and

⁵ See also Calvo and Guidotti (1990), Barro (2002), Benigno and Woodford (2003), and Lustig, Sleet, and Yeltekin (2006). More closely related to our work is Guibaud, Nosbusch, and Vayanos (2007), who propose a clientele-based theory in which the optimal debt maturity structure helps overcome imperfections in intergenerational risk sharing.

private money creation to the mix, and pose the comparative-advantage question: to what extent should the government actively try to crowd out private money creation. Drawing on our empirical work, we develop a simple calibration which suggests that the potential crowding-out benefits of short-term government debt may be of the same order of magnitude as the direct monetary services it provides. Section V discusses some further practical implications of our framework. Specifically, we explore how the answer to the comparative advantage question changes if the government has alternate regulatory tools at its disposal and discuss an extension in which we allow for three, rather than just two debt maturities. Section VI concludes.

II. Stylized Facts

A. Convenience Premia in Treasury Securities

Krishnamurthy and Vissing-Jorgensen (2010) argue that Treasury securities have some of the same features as money, namely liquidity and “absolute security of nominal return.” They find that these attributes lead Treasuries to have significantly lower yields than they would in standard asset-pricing models—their estimate of the money premium on Treasuries from 1926-2008 is 72 basis points. Their identification is based on measuring the impact of changes in Treasury supply on a variety of spreads. For example, they show that an increase in the supply of Treasuries reduces the spread between Treasuries and AAA-rated corporate bonds—arguably because the money premium falls as the quantity of money-like claims rises.

Krishnamurthy and Vissing-Jorgensen (2010) treat all Treasury securities as having similar money-like properties, and do not distinguish between Treasuries of different maturities. However, other work (e.g. Amihud and Mendelson (1991), Duffee (1996), Gurkaynak, Sack and Wright (2006)) has documented that the yields on short-term T-bills are often strikingly low relative to those on longer-term notes and bonds. Gurkaynak et al write: “...bill rates are often disconnected from the

rest of the Treasury yield curve, perhaps owing to segmented demand from money market funds and other short-term investors.”

Panel A of Figure 1 provides an illustration. We plot the average spread, over 1983-2009, between actual T-bill yields (with maturities from 1 to 26 weeks) and fitted yields, where the fitted yields are based on a flexible extrapolation of the Treasury yield curve from Gurkaynak, Sack, and Wright (2006) that is calibrated using only notes and bonds with remaining maturities greater than three months.⁶ In other words, the n -week “ z -spread” in Figure 1, $z_t^{(n)} \equiv y_t^{(n)} - \hat{y}_t^{(n)}$, represents the extent to n -week T-bills have yields that differ from what one would expect based on an extrapolation of the rest of the yield curve. The differences are large: four-week bills have yields that are roughly 40 basis points *below* their fitted values; and for one-week bills, the spread is about 60 basis points.⁷

Our preferred interpretation of these z -spreads is that they reflect the extra “moneyness” of short-term T-bills, above and beyond whatever money-like attributes longer-term Treasuries may have. For example, short-term bills offer not only absolutely certain ultimate nominal returns, as Krishnamurthy and Vissing-Jorgensen (2010) stress for Treasuries as a whole, but also are completely safe at short horizons since they have no interest-rate exposure. Furthermore, while long-term Treasuries are highly liquid, short-term T-bills are even more liquid (Amihud and Mendelson (1991)). Presumably, these attributes are what makes T-bills so attractive to money-market mutual funds and desirable as collateral for backing repurchase agreements and other financial contracts.

This interpretation is supported by the work of Greenwood and Vayanos (2010). They find that the returns on short-maturity Treasuries go up (as compared to those on longer-maturity Treasuries) when the government does a greater proportion of its issuance at the short end of the yield

⁶ Gurkaynak, Sack, and Wright (2006) estimate a parametric model of the instantaneous forward rate curve that is characterized by six parameters. Zero coupon yields are then derived by integrating along the estimated forward curve.

⁷ Because all T-bills have slightly lower yields than notes and bonds with similar remaining maturities, and because our fitted yields are based solely on notes and bonds, the z -spreads in Figure 1 do not converge to zero as maturity rises.

curve. In other words, when there are more of the most money-like short-term securities in the system, the convenience premium on these securities shrinks. Panel B of Figure 1 shows that this logic applies especially to short-term T-bills. Each quarter from 1983-2009, we plot the average 4-week “z-spread” alongside the ratio of T-bills to GDP. As can be seen, there is a positive relation between the two series ($R^2 = 0.49$).

Table 1 shows weekly univariate regressions of the n -week z -spread on the ratio of Treasury bills to GDP for $n = 2, 4$ and 10 weeks:⁸

$$z_t^{(n)} = a^{(n)} + b^{(n)} \cdot (BILLS / GDP) + \varepsilon_t^{(n)}. \quad (1)$$

To compute $BILLS/GDP$ precisely each week, we use detailed data on the size and timing of Treasury auctions. Consistent with Figure 1, Table 1 shows a strong response of z -spreads to the supply of short-term Treasuries. The coefficient of 5.60 in column (1) of Panel A (for the sample period 1983-2009) means that a one-percentage-point increase in the ratio of T-bills to GDP (roughly half of a standard deviation) leads to a 5.6 basis point increase in the 2-week z -spread. Table 1 also shows that the effect is strongest for very short-term T-bills: the coefficient for the 2-week spread is more than twice that for the 10-week spread.

Of course, this evidence is subject to endogeneity concerns. Specifically, the government might respond to money demand shocks, increasing the proportion of short-term debt when money demand rises. Indeed Panel B of Figure 1 shows that $BILLS/GDP$ jumps in the fall of 2008 just as z -spreads plummet—the telltale signature of an endogenous supply response. The existence of large money demand shocks such as that in fall 2008 would tend to bias our OLS estimates downward. To address this concern, we first omit 2008 and focus on the 1983-2007 period in Panel B. Not surprisingly, the coefficients are twice as large and far more precisely estimated when we omit 2008.

⁸ We obtain similar results if we regress z -spreads on D_S/GDP , where D_S includes T-bills plus notes and bonds maturing within one year.

Admittedly, simply dropping the outlying 2008 observations is somewhat ad hoc. To better address endogeneity concerns, we adopt an instrumental variables strategy based on the observation that there is substantial high-frequency variation in short-term government financing patterns associated with seasonal fluctuations in tax receipts. Specifically, the Treasury expands the supply of bills ahead of statutory tax deadlines (e.g., April 15th) to meet its ongoing cash needs and these borrowings are then repaid rapidly following the deadlines.⁹ Moreover, this seasonal variation in bill supply is plausibly unrelated to business-cycle conditions or to shocks to money demand. These seasonal fluctuations in T-bill supply can be seen clearly in Panel A of Figure 2 which plots *BILLS/GDP* each week from 1983-2009.

We begin by estimating equation (1) in changes, to focus on the high-frequency variation in the data. Specifically, we regress 4-week changes in the *z*-spread on 4-week changes in *BILLS/GDP*:

$$\Delta_4 z_t^{(n)} = a^{(n)} + b^{(n)} \cdot \Delta_4 (BILLS / GDP)_t + \Delta_4 \varepsilon_t^{(n)}. \quad (2)$$

Columns (4)-(6) of Table 1 show that when estimated in changes, the slope coefficients $b^{(n)}$ are generally larger than the estimates from the levels-regressions in columns (1)-(3). However, the estimates are not significant for the full 1983-2009 period.

Next, columns (7)-(9) report instrumental variables (IV) estimates of the same changes specification. Consider, for example, the IV estimates for the full 1983-2009 sample in Panel A. In the first stage, we regress 4-week changes in *BILLS/GDP* on a series of 52 week-of-year dummy variables. Panel B of Figure 2 plots the coefficients on these 52 week-of-year dummies, which collectively explain 41% of the changes in *BILLS/GDP* from 1983-2009 (and 62% from 1983-2007).

The figure shows that T-bill supply expands significantly the weeks leading up to tax deadlines and

⁹ There are significant spikes in Federal tax receipts on the individual and corporate tax deadlines: January 15th, March 15th, April 15th, June 15th, September 15th, and December 15th. In the weeks leading up the tax deadline, the Treasury increases the size of its regularly scheduled bill auctions and issues large “Cash Management Bills” which mature just after the deadline. These borrowings are repaid following the deadline using tax receipts. As can be seen in Panel A of Figure 2, these seasonal patterns in bill supply become far more pronounced in 1992, and indeed our IV results are significantly stronger in the 1992+ period (not reported).

that these borrowings are then repaid using tax receipts. In the second stage, we regress the 4-week change in the z -spread on the fitted change in T-bill supply from the first stage.

Table 1 shows that these IV estimates of equation (2) for the full 1983-2009 period are larger and more precisely estimated than the corresponding OLS estimates. Table 1 also shows that OLS estimation of (2) yields large positive estimates of $b^{(n)}$ from 1983-2007, but that the coefficients decline substantially once we add 2008. By contrast, the IV estimates are very similar for both the 1983-2007 and the 1983-2009 periods. This makes sense if one views the large spike in $BILLS/GDP$ in the fall of 2008 as being driven by the government's endogenous response to a major money demand shock. Specifically, following the collapse of Lehman Brothers in September 2008, z -spreads fell significantly just as $BILLS/GDP$ was soaring. However, aside from this crisis period, the remaining high-frequency variation in $BILLS/GDP$ appears to be largely driven by seasonal supply shocks. As a result, the OLS and IV estimates of equation (2) are similar for the 1983-2007 period.

B. The Correlation Between Debt Maturity and the Debt-to-GDP Ratio

Figure 3 plots the weighted average maturity of outstanding U.S. government debt against the debt-to-GDP ratio from 1952-2009. As can be seen, the two series are strongly positively correlated—the correlation coefficient is 0.71 over the full sample period. This relationship between debt maturity and debt-to-GDP, also noted in Greenwood and Vayanos (2010), Greenwood, Hanson and Stein (2010), and Krishnamurthy and Vissing-Jorgensen (2010), is one of the most direct implications of the tradeoff model of government debt maturity that we develop in the next section.

To be clear, such a positive correlation could potentially arise even if the government adopted a mechanical strategy of always issuing new debt with the same average maturity: when debt-to-GDP rises, the ratio of newly-issued debt to outstanding debt rises, so even holding fixed the average maturity of new issues, the average maturity of outstanding debt can increase. However, as noted by Garbade and Rutherford (2007), the Treasury actively manages its issuance and repurchases to

achieve a target maturity of outstanding debt. To test whether issuance behavior actively adjusts in response to a target debt maturity that is itself a function of the debt-to-GDP ratio, we estimate a regression of the maturity of new debt *issues* M_t^{Iss} on the lagged maturity of debt issues M_{t-1}^{Iss} , the lagged maturity of outstanding debt M_{t-1}^{Out} , and the debt-to-GDP ratio. This yields

$$M_t^{Iss} = 0.98 + 2.86 \cdot (D / GDP)_t + 0.76 \cdot M_{t-1}^{Iss} - 0.34 \cdot M_{t-1}^{Out}, \quad R^2 = 0.55. \quad (3)$$

[$t=2.52$]
[$t=3.40$]
[$t=7.43$]
[$t=-3.60$]

One can interpret this regression in terms of a partial adjustment model of the maturity of government debt issues: when the debt-to-GDP ratio is high, the government gradually adjusts the maturity of its new issuance towards longer-term debt.

C. Private-Sector Responses to Government Debt-Maturity Choices

The comparative-advantage version of our model rests on the premise that privately-issued and government-issued short-term debt claims are partial substitutes in the sense that both provide a form of monetary services.¹⁰ As a result, the government can, by issuing more T-bills, crowd out the issuance of short-term money-like claims by financial intermediaries. Greenwood, Hanson and Stein (2010) investigate a similar crowding-out phenomenon, looking at how the maturity choices of private debt issuers respond to changes in government debt maturity over the period 1963-2005. Figure 4 reproduces the main finding of that paper, extending the time series to 2009: it shows that as government debt maturity contracts, the debt maturity of non-financial firms rises significantly.

While this result provides general support for a debt-maturity crowding-out hypothesis, here we introduce another piece of evidence that is more precisely targeted to understanding the money-creation behavior of financial intermediaries. To do so, we build on Krishnamurthy and Vissing-Jorgensen (2010), who estimate regressions of the form:

¹⁰ To be clear, when we refer to “government money” we mean the supply of short-term T-bills, and not the monetary base—i.e., currency and central bank reserves. While it seems natural that the former would be a substitute for privately-created money, the latter might actually be a complement. In particular, when reserves go up, banks subject to reserve requirements can create more demand deposits; this is the textbook “money multiplier”.

$$(PrivateMoney / GDP)_t = a + b \cdot (D / GDP)_t + u_t, \quad (4a)$$

and find $b < 0$. Recall that they are interested in the claim that *all government debt is to some degree money-like*. So they interpret their results as saying that when total government debt is higher, financial intermediaries have less incentive to create private money.¹¹

By contrast, we want to emphasize the idea that *short-term government debt is more money-like than long-term government debt*, and hence should be expected to have a more powerful crowding-out effect on private money creation. In Table 2, we compare estimates from specification (4a) with estimates from regressions of the form:

$$(PrivateMoney / GDP)_t = a + b \cdot (D_S / GDP)_t + u_t. \quad (4b)$$

Total debt, D , in (4a) is marketable government debt held by the public, whereas short-term government debt, D_S , in (4b) is marketable public debt with remaining maturity of less than one year.¹² We use two measures of private money creation. The first follows Krishnamurthy and Vissing-Jorgensen (2010) and is the difference between M2 and M1. The second is the difference between M3 and M1. We estimate (4a) and (4b) using annual data from 1952 through 2009.

We start in Table 2 by verifying the key result from Table V of Krishnamurthy and Vissing-Jorgensen—the correlation between $(M2-M1)/GDP$ and D/GDP in the first column is negative. However, the second column shows that the magnitude of the coefficient and the R^2 rise substantially when we instead regress $(M2-M1)/GDP$ on D_S/GDP . Panel B shows that similar results obtain when we use the broader measure of private money creation M3-M1. In untabulated regressions, we verify that these results are robust to business cycle controls. Furthermore, we obtain nearly identical results

¹¹ Bansal, Coleman, and Lundblad (2010) also find that an increase in the supply of government debt leads to a decline in short-term private sector debt.

¹² Thus, D_S includes both T-bills and notes and bonds with a remaining maturity of less than one year. Unlike Krishnamurthy and Vissing-Jorgensen (2010) who measure D using all government debt held by the public, we restrict attention to marketable debt held by the public (i.e., we exclude savings bonds and other miscellaneous forms of non-marketable public debt) on the theory that nonmarketable securities provide far less in the way of monetary services.

if we estimate (4b) via two-stage least squares using D/GDP as an instrument for D_S/GDP . In summary, the results in Table 2 are consistent with the idea that short-term Treasuries are especially money-like and, thus, have a particularly strong crowding-out effect on private money creation.¹³

III. A Tradeoff Model of Government Debt Maturity

The full model features three sets of actors: households, the government, and financial intermediaries (a.k.a., “banks”). In this section we begin with a stripped-down version that leaves out the intermediaries, thus focusing on the optimal maturity structure of debt when the government is the sole creator of money. This setup generates a simple tradeoff between the monetary services provided by issuing more short-term debt, and the increased rollover risk that comes as a result. In Section IV, we allow banks to compete with the government in money creation.

A. Households

There are three dates, 0, 1, and 2. Households receive a fixed exogenous endowment of one unit in each period. After paying taxes in each period, households can consume the remainder of their endowment, or invest some of it in financial assets. Households have linear preferences over consumption at these three dates.

Households can transfer wealth between periods by purchasing government bonds. At date 0, households can purchase short-term bonds $B_{0,1}$, which pay one unit at date 1, or long-term zero-coupon bonds $B_{0,2}$, which pay one unit at date 2. Households can also purchase short-term debt at date 1, $B_{1,2}$, which pays one unit at date 2. The discount factor between date 1 and date 2 is random from the point of view of date 0, and is not realized until date 1, so refinancing maturing short-term debt at date 1 introduces uncertainty over date-2 taxation and hence over consumption.

¹³ Nevertheless, one should maintain some degree of skepticism regarding time-series regressions with persistent variables. High frequency data on private money creation is available from 2001-2009 and we find that *changes* in financial commercial paper outstanding respond negatively to changes in T-bill supply in this short sample (not reported).

In addition to direct consumption, households derive utility from monetary services.¹⁴ For starkness, we assume that these services only come from short-term debt issued at date 0, although the crucial assumption is just that short-term bonds provide more in the way of monetary services than long-term bonds. The utility of a representative household is thus given by

$$U = C_0 + E[C_1 + \beta C_2] + v(M_0), \quad (5)$$

where β is the random discount factor which is realized at date 1 and where $M_0 = B_{0,1}$, the amount of short-term government bonds held by households at date 0. We assume that $E[\beta] = 1$ without loss of generality. For now, we also assume that $v'(M_0) > 0$ and $v''(M_0) \leq 0$. However, in Section IV when we analyze whether the government should try to crowd out private money creation, we must assume that there are strictly diminishing returns to holding money, i.e., that $v''(M_0) < 0$.

Equation (5) can be used to pin down real interest rates. Long-term bonds issued at date 0 have price $P_{0,2} = 1$. Short-term bonds issued at date 0 have price $P_{0,1} = 1 + v'(M_0)$, thereby embedding an additional money premium. Short-term bonds issued at date 1 have a price that is uncertain from the perspective of date 0, $P_{1,2} = \beta$.

B. Government

The government finances a one-time expenditure G at date 0, using a combination of short- and long-term borrowing from households, and taxes which it can levy in each period. The government budget constraint is given by:

$$\begin{aligned} t = 0: G &= \tau_0 + B_{0,1}P_{0,1} + B_{0,2}P_{0,2} \\ t = 1: B_{0,1} &= \tau_1 + B_{1,2}P_{1,2} \\ t = 2: B_{1,2} + B_{0,2} &= \tau_2 \end{aligned} \quad (6)$$

¹⁴ We follow a long tradition in economics, starting with Sidrauski (1967), of putting monetary services directly in the utility function. As discussed further below, our results are qualitatively unchanged if we allow long-term government bonds to also carry a convenience premium, provided that the premium is strictly less than that on short-term bonds.

where $P_{0,1}$ and $P_{0,2}$ denote the prices of short- and long-term bonds issued at date 0, and $P_{1,2}$ denotes the (uncertain) price of short-term bonds issued at date 1. At date 0, the government may levy taxes of τ_0 on households, and sell short- and long-term bonds. If the government borrows short-term, then at date 1, it must levy taxes to pay off the maturing debt, or roll over the debt by issuing new short-term bonds $B_{1,2}$. At date 2, the government pays off all maturing debt by levying taxes.

We follow the standard assumption that taxes are distortionary (Barro (1979), Lucas and Stokey (1983), Bohn (1990)), and that the magnitude of these distortions is convex in the amount of revenue raised each period.¹⁵ For simplicity, we use the quadratic function $\tau^2 / 2$ to capture the resources that are wasted when taxes are τ . Household consumption in each period is thus given by

$$\begin{aligned} C_0 &= 1 - \tau_0 - (1/2)\tau_0^2 - B_{0,1}P_{0,1} - B_{0,2}P_{0,2} \\ C_1 &= 1 - \tau_1 - (1/2)\tau_1^2 + B_{0,1} - B_{1,2}P_{1,2} \\ C_2 &= 1 - \tau_2 - (1/2)\tau_2^2 + B_{1,2} + B_{0,2}. \end{aligned} \tag{7}$$

Substituting in the government budget constraint from (6), household consumption can be written as

$$\begin{aligned} C_0 &= 1 - (1/2)\tau_0^2 - G \\ C_1 &= 1 - (1/2)\tau_1^2 \\ C_2 &= 1 - (1/2)\tau_2^2. \end{aligned} \tag{8}$$

Since we have assumed that endowments are fixed and that the government finances a known one-time expenditure of G , there is no endowment or fiscal risk in our model. As discussed further below, this implies that tax smoothing does not give rise to the sort of hedging motive that often makes state-contingent debt optimal in models of government debt maturity. The only source of risk in our model is the random discount factor, β , which one can think of as being driven by shocks to household preferences unrelated to endowments. This setup helps to simplify the analysis and to highlight the novel forces at work in our model.

¹⁵ Bohn (1990) assumes that taxes are a linear function of endowments, and that the deadweight costs of taxation are a convex function of the *tax rate*. Given that we take endowments to be fixed, our approach amounts to the same thing.

The social planner maximizes household utility subject to the government budget constraint. Substituting household consumption (8) and money ($M_0=B_{0,1}$) into the household utility function (5) and dropping exogenous additive terms, the planner's problem can be written as

$$\max_{\{B_{0,1}, B_{0,2}, B_{1,2}\}} \left[v(B_{0,1}) - \frac{1}{2} (\tau_0^2 + E[\tau_1^2] + E[\beta\tau_2^2]) \right]. \quad (9)$$

The three right-hand terms in (9) capture the standard tax smoothing objective—the planner would like taxes to be low and constant over time. However, this objective must be balanced against the utility that households derive from holding short-term bonds.

C. Optimal Maturity Structure in the Absence of Money Demand

We first solve the planner's maximization problem in the benchmark case where households derive no utility from monetary services (i.e., $v(B_{0,1})=0$). In this case, the prices of short- and long-term bonds issued at date 0 are the same and the planner solves

$$\min_{\{B_{0,1}, B_{0,2}, B_{1,2}\}} \left[\frac{1}{2} (\tau_0^2 + E[\tau_1^2] + E[\beta\tau_2^2]) \right]. \quad (10)$$

The planner's problem can be solved by working backwards. At date 1, the discount factor between dates 1 and 2 is realized. From the government budget constraint, taxes at date 1 and date 2 are $\tau_1 = B_{0,1} - B_{1,2}\beta$ and $\tau_2 = B_{1,2} + B_{0,2}$. Substituting into the planner's date-1 problem yields

$$\min_{B_{1,2}} \left[\frac{1}{2} (\tau_1^2 + \beta\tau_2^2) \right] = \min_{B_{1,2}} \left[\frac{1}{2} (B_{0,1} - B_{1,2}\beta)^2 + \frac{1}{2} \beta (B_{1,2} + B_{0,2})^2 \right]. \quad (11)$$

The first-order condition for $B_{1,2}$ implies that

$$B_{1,2} = \frac{B_{0,1} - B_{0,2}}{1 + \beta}, \quad (12)$$

which in turn implies that $\tau_1 = \tau_2 = (B_{0,1} + B_{0,2}\beta) / (1 + \beta)$. Intuitively, the planner chooses $B_{1,2}$ to perfectly smooth taxes between dates 1 and 2 and the tax rate is such that the present value of taxes equals the present value of required debt payments. To get the quantity of short- and long-term debt

issued at date 0, we substitute (12) into (10) and solve the first-order conditions for $B_{0,1}$ and $B_{0,2}$. The solution is given by Proposition 1.

Proposition 1: *In the absence of money demand, the government perfectly smooths taxes by setting $\tau_0 = \tau_1 = \tau_2 = G/3$, $B_{0,1} = B_{0,2} = G/3$, and $B_{1,2} = 0$. This result holds even if the government can issue risky securities whose payoffs depend on the realization of the discount factor β .*

Proof: *All proofs are in the appendix.*

Proposition 1 captures the intuition that, absent money demand, the government can insulate the budget and taxes from uncertain future refinancing by never rolling over debt at date 1. With convex costs of taxation in each period, the planner sets the marginal social cost of taxation equal across dates. The government can accomplish this by issuing a long-term “consol” bond with face value of $2G/3$ that makes the same payment at dates 1 and 2.

One might wonder whether total welfare could be increased if the government were able to issue risky state-contingent securities whose payoffs depend on the realization of the discount factor β . For example, suppose the government can issue risky debt with a payoff of $X_r(\beta)$ at $t = 2$ when the realization of the discount factor is β . However, as long as these securities are fairly priced, (i.e., as long as $P_r = E[\beta X_r(\beta)]$), the government cannot improve upon the simple tax-smoothing solution. This is an important result, because it implies that absent money demand, it does not make sense for the government to try to lower its expected financing costs by selectively selling securities that have low betas with respect to priced risks.¹⁶

To see the intuition for this result, note that from (10) the planner cares about minimizing $E[\beta\tau_2^2] = Cov[\beta, \tau_2^2] + (E[\tau_2])^2 + Var[\tau_2]$. Suppose that the government reduces its issuance of 2-period riskless bonds and instead issues state-contingent securities that deliver a high payout at date 2

¹⁶ This result is reminiscent of Bohn (1999) and Missale (1997), who argue that expected return differences between bonds should only be taken into account in debt management if they are driven by market imperfections or liquidity.

when β is high. On the one hand, this would reduce expected financing costs and hence expected taxes, leading $(E[\tau_2])^2$ to fall. This is because the risky securities command a higher price than riskless ones with the same expected payout, given that they have a high payoff in states where consumption is valued most. On the other hand, the issuance of these risky securities increases $Cov[\beta, \tau_2^2]$. That is, servicing the risky debt requires the government to impose higher taxes on households in states where consumption is highly valued and hence where taxes are most painful. These two effects tend to offset one another, and we are left with the fact that issuing risky securities always increases $Var[\tau_2]$. Consequently, the effect on $E[\beta\tau_2^2]$ of shifting from riskless to risky securities is always positive, the opposite of what the planner would like to accomplish.

In summary, absent a specific hedging motive, the government should not issue a security that has a low required return simply because it is less risky in the standard asset-pricing sense. This conclusion is similar to that of Froot and Stein (1998), who argue that a financial institution cannot create value for its shareholders simply by taking on priced risks that are traded in the marketplace.¹⁷

D. Optimal Maturity Structure with Money Demand

We now turn to the case in which households derive utility from their holdings of short-term bonds. Before doing so, we introduce a notational simplification. We denote the total scale of government borrowing at date 0 as $D = B_{0,1} + B_{0,2}$, and the short-term debt share as $S = B_{0,1} / D$. Applying this notation, the benchmark optimal debt structure in the absence of money demand is given by $S = 1/2$ and $D = (2/3)G$.

We solve the planner's problem in (9) subject to the budget constraint in (6). As before, the long-term bond has price $P_{0,2}=1$, and the short-term bond issued at date 1 has uncertain price β .

¹⁷ The result that the government chooses not to issue state-contingent debt depends on our simplifying assumptions that endowments are deterministic and that there are no fiscal shocks. Otherwise, the government might have a motive to issue state-contingent debt that hedges these risks as in Bohn (1988) and Barro (1997). Since these hedging motives are well understood, we dispense with them in order to highlight the new forces at work in our model.

However, the short-term bond issued at date 0 now embeds a money premium, $P_{0,1} = 1 + v'(M_0)$. As shown in the Appendix, we can rewrite the planner's problem as

$$\min_{S,D} \left[\frac{1}{2} (G - D - DSv'(DS))^2 + \frac{D^2}{2} \left(b \left(S - \frac{1}{2} \right)^2 + \frac{1}{2} \right) - v(DS) \right], \quad (13)$$

where $b \equiv E[(\beta - 1)^2 / (1 + \beta)] \approx \text{Var}[\beta] / 2$ is a measure of the magnitude of date-1 refinancing risk.

The first-order condition for the short-term debt share S can be written as

$$\overbrace{Db(S - 1/2)}^{\text{Marginal tax-smoothing cost}} = \overbrace{v'(SD)}^{\text{Marginal benefit of money services}} + \overbrace{\tau_0 \cdot [v'(SD) + SDv''(SD)]}^{\text{Marginal tax lowering benefit}}. \quad (14)$$

Each of the three terms in (14) has a natural interpretation. The left-hand side represents the marginal tax-smoothing cost of shifting government financing towards short-term debt. Note that this cost depends on the difference between S and $1/2$, i.e., on the extent of the departure from perfect tax smoothing. It also depends on the magnitude of date-1 refinancing risk b , as well as on the raw scale of government debt D . The first term on the right-hand side of (14) reflects the direct money benefit of short-term bills—the marginal convenience services enjoyed by households. The second term on the right-hand side of (14) captures the net benefit from the *lower level of taxes* that arises when the government finances itself at a lower average interest rate. The government can raise revenue either by taxing, or by creating more money, with the marginal revenue from creating money given by $v'(SD) + SDv''(SD)$. If the latter method of revenue-raising is non-distortionary, it pushes the social planner towards further issuance of short-term bills.

Nevertheless, for much of the remainder of the paper, we ignore this latter *tax-lowering* benefit, in which case (14) reduces to:

$$Db(S - 1/2) = v'(SD). \quad (14')$$

The argument in favor of focusing on (14') rather than (14) is as follows. Given that our formulation of the deadweight costs of taxation lacks microfoundations, we don't have any basis for asserting that

one form of taxation—seignorage from money creation—is less distortionary than some other form, such as income or capital taxation. Fortunately, as we demonstrate below, our qualitative results are not sensitive to whether we derive them from (14) or (14').

The one scenario where it makes most obvious sense to include the tax-lowering benefits of short-term debt is when this debt is held by foreign investors. In this case, issuing more short-term debt corresponds to raising more seignorage revenue from parties whose utility a parochial planner may not internalize, while allowing for the reduction of other taxes on domestic households. We return to this case at the end of this section.

The solution to (14') leads directly to Proposition 2.¹⁸

Proposition 2: *Define S^* as the optimal short-term debt share which solves $Db(S-1/2) = v'(SD)$. Then $S^* > 1/2$, and S^* is decreasing in both uncertainty about date-1 short rates, as well as in G , i.e., $\partial S^* / \partial b < 0$, and $\partial S^* / \partial G < 0$. Furthermore, suppose that $v(M) = \gamma f(M)$, where $f(\cdot)$ is an increasing and weakly concave function and γ is a positive constant. Then $\partial S^* / \partial \gamma > 0$.*

Proposition 2 establishes that money demand increases the willingness of the government to issue short-term bills and thereby take on refinancing risk. In combination, Propositions 1 and 2 clarify that the government should only tilt towards a shorter maturity structure if short-term debt is “cheaper” than long-term debt *because* it benefits from a special money premium. By contrast, it would not make sense for the government to issue more short-term bills in an attempt to economize on traditional term premia that arise because long-term bonds suffer poor returns in bad states.

Proposition 2 also shows that this bias towards short-term maturities is more pronounced when either the intensity of money demand is stronger, or the variance of short rates at date 1 is lower. At the extreme, if the variance of short rates is low enough, or if the social costs of rollover

¹⁸ In establishing Proposition 2, we restrict the government to issuing only non-state-contingent securities. Since the government now incurs refinancing risk, it might want to hedge this by issuing β -contingent securities which we rule out.

risk are sufficiently small, the government may go so far as to finance its entire debt using short-term bills. A similar logic can be used to understand the relationship between the government's total debt burden and its choice of debt maturity. The greater is the size of the debt—as captured here by the parameter G —the larger is the refinancing risk in dollar terms, and thus the less willing is the planner to deviate from $S = 1/2$. Furthermore, when $v''(\cdot) < 0$, the premium on short-term debt will fall as G rises, further reducing the incentive to tilt towards short-term debt. As discussed earlier, these predictions capture the intuition used by practitioners to describe the government's approach to debt-maturity policy. And, as can be seen in Figure 2, they are clearly borne out in the U.S. data, where the correlation between debt maturity and Debt/GDP has historically been strong.

E. Allowing for Monetary Services from Long-term Bonds

To keep things simple, we have assumed that long-term bonds provide no monetary services whatsoever. However, all that we really need is for short-term bills to be more money-like than long-term bonds. Suppose instead that short-term bills offer one unit of monetary services and that long-term bonds offer $0 < q < 1$ units of monetary services. In this case, while the government tilts less toward short-term debt than if $q = 0$, our basic results remain qualitatively the same.

Digging deeper, one can also think about how the magnitude of q might be derived from first principles. Suppose that for a security to provide monetary services, it must be completely riskless between dates 0 and 1. Long-term government bonds are not riskless, since their date-1 value depends on the realization of β . However, the private market may still be able to create some amount of riskless claims by using long-term bonds as collateral for short-term borrowing—as is done in the repo market. Following Geanakoplos (2010), the quantity of riskless collateralized claims that can be manufactured is given by the *minimum* period-1 price of the long-term bond, i.e., $q = \beta^{\min} < 1$.

A version of this argument is likely to hold even in a more elaborate long-horizon model where the interval between the periods becomes arbitrarily short, so long as $v''(M_0) < 0$. To see the

intuition, think about the present value of the expected stream of future monetary services provided by a long-term bond that is originally issued at a market value of 100. Suppose that from one day to the next, the bond's price can rise or fall by at most one percent. Thus on the first day, it is possible to borrow 99 on a riskless overnight basis against the bond, i.e. to generate almost the same amount of monetary services as would come from 100 of short-term bills. However, over time, as the price of the bond fluctuates, the quantity of money that it can be used to collateralize will rise or fall. Given that $v''(M_0) < 0$, the value of such a risky stream of monetary services is less than the value of the sure stream of monetary services that would come from the 100 of short-term bills being rolled over repeatedly. The ratio of the value of the risky stream to that of the safe stream is equivalent to the concept of q in our simpler model.

F. The Tax-Lowering Benefits of Short-Term Debt

In the above analysis, we assumed that the social planner internalizes the monetary benefits enjoyed by households who invest in short-term debt, but does not put any weight on the tax savings that short-term debt generates—because these savings ultimately reflect a (potentially distortive) seignorage tax on its own citizens. Now we explore the opposite configuration, where the planner cares about the tax-lowering benefits of short-term debt, but not about the monetary services. As suggested above, this case is most naturally interpreted as corresponding to a situation where *all* the short-term debt is sold to foreign investors, and where a nationalistic planner looks out only for the interests of domestic households.

The nationalistic planner's date-1 problem is the same as before, since all monetary services are consumed at date 0. The planner's date-0 problem is similar to that in (9), except that we drop the direct utility of money services. Instead, as shown in the Appendix, there is a term which reflects the fact that the planner can lower date 0 taxes on domestic households by raising seignorage revenue from foreign investors:

$$\min_{\{B_{0,1}, B_{0,2}, B_{1,2}\}} \left[\frac{1}{2} (\tau_0^2 + E[\tau_1^2] + E[\beta \tau_2^2]) - v'(M)M \right]. \quad (15)$$

Expression in (15) can be rewritten as

$$\min_{S,D} \left[\frac{1}{2} (G - D - R(DS))^2 + \frac{D^2}{2} \left(b \left(S - \frac{1}{2} \right)^2 + \frac{1}{2} \right) - R(DS) \right], \quad (16)$$

where we make use of the notation $R(M) = v'(M)M$ to denote seignorage revenue—the interest savings from issuing more short-term bills to foreign investors. We restrict attention to money demand functions for which $R'(M) > 0$ and $R''(M) \leq 0$. As shown in the Appendix, the solution takes the form

$$S^* = \frac{1}{2} + \frac{R'(D^*S^*)}{b} \frac{(G - R(D^*S^*) + 3)}{R'(D^*S^*)(G - R(D^*S^*) + 3) + 2(G - R(D^*S^*) - R'(D^*S^*))}, \quad (17)$$

and

$$D^* = \frac{(2 + R'(D^*S^*))(G - R(D^*S^*)) + R'(D^*S^*)}{3 + R'(D^*S^*)}. \quad (18)$$

Equations (17) and (18) yield the following proposition.

Proposition 3: *Let $R(M) = v'(M)M$ and suppose that $R'(M) > 0$ and $R''(M) \leq 0$. Then, with foreign investors holding all the short-term debt, and with a nationalistic planner, we have that, as in Proposition 2: $S^* > 1/2$, $\partial S^* / \partial b < 0$, and $\partial S^* / \partial G < 0$.*

In summary, the case with foreign investors and a nationalistic planner works similarly to our baseline closed-economy case. The government still finds it optimal to issue more short-term debt than in the perfect tax-smoothing benchmark, and the comparative statics are directionally the same.

IV. Adding Private-Sector Money Creation

We now extend the tradeoff model from Section III to allow private financial intermediaries to compete with the government in the provision of money-like securities. Our treatment of private-sector money creation follows Stein (2012). As in Stein's (2012) model, banks invest in real projects,

and can choose whether to finance these projects by issuing short-term or long-term debt. However, given the structure of the risks on their projects, only short-term bank debt can ever be made riskless. Hence if they wish to capture the convenience premium associated with money-like claims, and thereby lower their financing costs, banks must issue short-term debt. While this has the same social benefits as government-created money, it can also lead to forced liquidations and fire sales. These fire sales in turn create social costs which the banks themselves do not fully internalize.

A. Bank Investment and Financing Choices

There are a continuum of banks in the economy with total measure one. Each bank invests a fixed amount I at date 0, financed entirely with borrowing from households—i.e., banks have no endowment of their own. With probability p , the good state occurs and the investment returns a certain amount $F > I$ at date 2. With probability $(1 - p)$, the bad state occurs. In the bad state, expected output at date 2 is $\lambda I < I$, and there is some downside risk, with a positive probability of zero output. Importantly, the potential for zero output at date 2 in the bad state means that no amount of long-term bank debt can ever be made riskless, no matter how much seniority it is granted.

At date 1, a public signal reveals whether the good or bad state will prevail at date 2. Given the linearity of household preferences over consumption, the realization of the banks' investment risk has no direct impact on the price of new government bonds issued at date 1, which continues to be given by $P_{1,2} = \beta$. To simplify the analysis we assume that this investment risk is independent of the realization of the discount factor β at date 1. However, the results are unchanged if we allow investment risk to be correlated with the realization of β .

As demonstrated in Stein (2012), if the bad signal about investment output is observed at date 1, banks will be unable to roll over their short-term debt, and will be forced to sell assets to pay off departing short-term creditors. A bank that sells a fraction Δ of its assets obtains date-1 proceeds of $\Delta\beta k\lambda I$, where $0 \leq k \leq 1$ denotes the endogenous discount to fundamental value (i.e., $\beta\lambda I$)

associated with the fire sale; we discuss the equilibrium determination of k momentarily. A bank that finances itself by issuing M_p dollars of short-term debt must therefore sell a fraction $\Delta = M_p / (\beta k \lambda I)$ in the bad state when the discount factor is β .

Banks make an initial capital structure decision at date 0. They can finance their investment by issuing either short-term or long-term debt. The advantage of short-term debt is that it can be sufficiently well-collateralized as to be rendered riskless, so as long as not too much is issued. This allows short-term bank debt to provide monetary services to households, and hence lowers its required rate of return. The disadvantage of short-term debt is that it forces banks to sell assets at discounted prices in the bad state.

Since $\Delta \leq 1$, the upper bound on the quantity of private money M_p that banks can create is given by $M_p \leq \beta^{\min} k \lambda I$. In other words, banks can fully collateralize an amount of short-term debt equal to what they can obtain by selling off all of their assets at date 1 when the discount factor β is at its lowest possible value. No bank will ever wish to issue an amount of short-term debt greater than this upper bound, since in this case the debt is no longer riskless, and hence does not sell at a premium, yet it still causes the bank to bear fire-sales costs.

A bank that finances itself with an amount of private money M_p realizes total savings of $M_p \cdot v'(M_0)$ relative to the case where it issues only long-term debt. Note that M_0 is now the *total* amount of private plus government money, i.e., $M_0 = M_p + M_G = M_p + SD$. Thus, we are assuming that government money and private money are perfect substitutes in household utility functions.

B. Fire Sales

To pin down the fire-sale discount k , we follow Stein (2012) and assume that when a bank is forced to sell assets in the bad state at date 1, these assets are purchased by a set of “patient investors”. Patient investors have a war chest of W , but cannot access capital markets at date 1 to raise more money in the event that the bad state occurs—in other words, their resources cannot be

conditioned on the realization of the state. In addition to buying any assets sold by the banks, the patient investors can also allocate their war chest to investing in new physical projects at date 1. Given an investment of K , these new projects generate a gross *social* return of $g(K)$ at time 2, where $g' > 0$ and $g'' < 0$. However, these social returns are not fully pledgeable; only a fraction $\phi < 1$ can be captured by the patient investors. Thus for an investment of K , the gross *private* return available to the patient investors at time 2 is $\phi g(K)$. This imperfect pledgeability assumption is crucial in what follows, because it implies that the market fire-sale discount of k does not reflect the full social costs of underinvestment by patient investors in the bad state.

In the good state, banks do not sell any of their assets, so patient investors invest all of their capital in new projects, i.e., $K = W$. In the bad state, banks are forced to sell assets to pay off the short-term debt that they have issued, and these assets sales are absorbed by the patient investors in equilibrium, so $K = W - M_p$. The fire-sale discount is determined by the condition that, in the bad state, patient investors must be indifferent at date 1 between buying assets liquidated by the banks and investing in new projects.¹⁹ This implies that the fire-sale discount to fundamental value is

$$1/k = \phi g'(K) = \phi g'(W - M_p). \quad (19)$$

C. Private Incentives for Money Creation

Individual banks maximize the expected present value of project cashflows, net of financing costs. Thus, banks' objective is

$$\begin{aligned} \Pi &= (pF + (1-p)E[\beta\lambda I(1-\Delta) + (\Delta\beta\lambda kI)]) - I + M^p v'(M_0) \\ &= [pF + (1-p)\lambda I - I] + M_p [v'(M_0) - (1-p)(1/k - 1)]. \end{aligned} \quad (20)$$

Each bank treats total money M_0 and hence $v'(M_0)$ as given when choosing their capital structure, and similarly for the fire-sale discount k . Because investment I is independent of financing, we can

¹⁹ See Shleifer and Vishny (2010) and Diamond and Rajan (2009) for similar formulations of the real costs of fire sales.

focus just on the right-hand terms in (20), ignoring the first expression in brackets. These latter terms capture the tradeoff that banks face when creating more private money: doing so lowers their financing costs by an amount $M_p \cdot v'(M_0)$, but with probability $(1 - p)$ leads them to have to sell their assets at a discount to fundamental value. As noted earlier, the analysis is unchanged if we allow investment risk to be correlated with the realization of β .²⁰

Substituting (19) into (20), equilibrium private money creation M_p^* is pinned down by

$$v'(M_p^* + M_G) = (1 - p)(\phi g'(W - M_p^*) - 1). \quad (21)$$

Note that this interior solution is only valid if M_p^* is below its technological upper bound, i.e., if $M_p \leq \beta^{\min} k \lambda I = (\beta^{\min} \lambda I) / (\phi g'(W - M_p))$. As long as we do have an interior optimum, M_p^* is greater when the pledgeability parameter ϕ is smaller. The intuition is that when ϕ is smaller, banks do not internalize as much of the fire-sale costs associated with private money creation. The social costs of fire sales are given by the underinvestment in new date-1 projects that they ultimately displace; these projects have a marginal social value of $g'(W - M_p^*)$. But the private costs to the banks of fire sales are only felt to the extent that they result in a discount on the assets they sell, and this discount is related to $\phi g'(W - M_p^*)$.²¹

D. The Social Planner's Problem

The social planner now maximizes total household utility, plus the net present value of date-1 investment by the patient investors:

²⁰ Let $X \in \{B, G\}$ denote the realization of the investment risk and $\beta_x = E[\beta | X=x]$. It is easy to see that discounted expected fire-sales costs are $E[\beta \lambda I (k - 1) \Delta | X=B] = M_p (1 - 1/k)$ irrespective of the value of β_B . Intuitively, discounted fire-sale costs do not depend on β_B because the fraction of assets sold $\Delta = M_p / (\beta k \lambda I)$ declines with β . Thus, private money creation is pinned down by (21) even if investment risk is correlated with the preference shock.

²¹ When $\phi < 1$ markets are incomplete and private money creation is associated with a pecuniary externality that leads to a violation of the first welfare theorem. The idea that, with incomplete markets, pecuniary externalities may lead to over-borrowing is a common theme in a recent literature in macroeconomics and finance, e.g., Lorenzoni (2008).

$$U_{SOCIAL} = E[g(K) - K] + v(M_0) - \frac{1}{2} [\tau_0^2 + E[\tau_1^2] + E[\beta\tau_2^2]]. \quad (22)$$

We begin by considering a first-best case in which the planner is able to directly control private money creation M_P , in addition to total government debt D , and the short-term government share S . Denoting the social planner's first-best values with two asterisks we have the following result.

Proposition 4: *In the first-best outcome, the marginal costs of both public and private money are set equal to the marginal social benefit of additional money services:*

$$\overbrace{(1-p)(g'(W - M_P^{**}) - 1)}^{\text{Fire sale cost of private money}} = \overbrace{v'(M_G^{**} + M_P^{**})}^{\text{Social benefit of money}} = \overbrace{b(M_G^{**} - D^{**} / 2)}^{\text{Tax-smoothing cost of government money}}. \quad (23)$$

The latter equality in (23), that $v'(M_G^{**} + M_P^{**}) = b(M_G^{**} - D^{**} / 2)$, is just a restatement of (14') from the government-only case, generalized to allow for the existence of private money. The former equality, that $(1-p)(g'(W - M_P^{**}) - 1) = v'(M_G^{**} + M_P^{**})$, differs from the private solution in (19) to the extent that the pledgeability parameter ϕ is less than one. Simply put, for any value of government money creation M_G , the social planner always prefers a smaller value of private money creation than do the banks acting on their own, i.e., $M_P^{**} < M_P^*$. Again, this is because with $\phi < 1$, the banks do not fully internalize the underinvestment costs that accompany their money-creation activities.

While the first-best outcome in Proposition 4 is a useful benchmark, it may be difficult to implement, because it requires the government to directly control all forms of private money creation. For example, the government can try to impose a cap or a tax on short-term debt issuance by regulated banks. However, as pointed out by Gorton (2010) and others, a significant fraction of private money creation in the modern economy takes place in the unregulated "shadow banking" sector, and so may be hard to police effectively at low cost. Indeed, more stringent regulation of traditional commercial banks may simply drive a greater share of private money creation into the unregulated sector. Thus, for the time being, we assume that regulations capable of implementing (23) are either unavailable or prohibitively costly.

With this limitation in mind, an alternative way to frame the planner's problem is as a second-best one in which it still seeks to maximize (22), but where it cannot directly constrain private money creation, and hence where its only choice variables are those pertaining to the government's own debt structure, namely D and S . It is in this second-best setting that our crowding-out intuition emerges. Consider what happens if the government issues more short-term debt at the margin. The convenience premium $v'(\cdot)$ falls, making it less attractive for the private sector to cater to money demand. In particular, for any given level of government money M_G , the corresponding level of private money creation is pinned down by (21). This implicitly defines a private-sector reaction function to government short-term debt issuance, $M_p^*(M_G)$. It is straightforward to show that:

$$-1 < \frac{\partial M_p^*}{\partial M_G} = -\frac{v''(M_p^* + M_G)}{v''(M_p^* + M_G) + (1-p)\phi g''(W - M_p^*)} < 0. \quad (24)$$

Using the private-sector reaction function, we can back out the amount of public money necessary to crowd out all socially excessive private money creation. However, because it is also socially costly for the government to issue more short-term debt, it will not be optimal for the government to issue so much short-term debt as to push private money creation all the way down to M_p^{**} . Using this logic, we can derive the government's first-order condition for the optimal short-term share in the second-best case, which we denote by S^{***} .

$$\underbrace{D^{***} b(S^{***} - 1/2)}_{\text{Tax-smoothing cost of government money}} = \underbrace{v'(M_p^* + S^{***} D^{***})}_{\text{Marginal benefit of money services}} + \underbrace{(1-p)(\phi-1)g'(W - M_p^*) \frac{\partial M_p^*}{\partial M_G}}_{\text{Crowding-out benefit}}. \quad (25)$$

Relative to our previous condition in (14'), equation (25) shows that there is an additional crowding-out benefit of short-term government debt. (The last term in (25) is positive when $\phi < 1$ since $\partial M_p^* / \partial M_G < 0$.) Thus we have the following result.

Proposition 5: *When there are externalities associated with private money creation (i.e., $\phi < 1$), a government that recognizes the crowding out benefits of short-term debt issues more short-term debt than a government that ignores its impact on private money creation. Moreover, if $v'''(\cdot)$ and $g'''(\cdot)$ are not too large, the crowding-out motive grows monotonically stronger as banks' failure to internalize fire-sales costs becomes more extreme, that is, $\partial S^{***} / \partial \phi < 0$.*²²

E. Are the Magnitudes Economically Interesting?

It is not immediately apparent whether the crowding-out benefit identified in equation (25) is of an economically interesting magnitude. In an effort to speak to this question, we perform a back-of-the-envelope calibration, drawing on some of our empirical estimates from earlier in the paper. We use this calibration to make two points. First, the crowding-out-motive for issuing short-term bills—i.e., the second term on the right-hand-side of (25)—is arguably of the same order of magnitude as the direct monetary-services motive captured in the first term on the right-hand side of (25). And second, when weighed against the tax-smoothing costs of tilting towards more short-term debt, this benefit is sufficient to raise the short-term share significantly, perhaps by something on the order of five percentage points.

The crowding-out effect is given by $(1-p)(\phi-1)g'(W-M_p^*) \cdot \partial M_p^* / \partial M_G$, which means we have to estimate four parameters. The $(1-p)$ term corresponds to the annual probability of a crisis in which the financial sector's inability to roll over short-term claims leads to a significant contraction in credit. Data on the frequency of financial crises are available in Barro and Ursua (2008), and Laeven and Valencia (2010), and suggest a range of between 5% and 10%. To be conservative, we assume a value of $(1-p) = 5\%$, corresponding to a crisis occurring once every two decades.

²² The analysis is largely unchanged if, following Krishnamurthy and Vissing-Jorgensen (2010), total monetary services are $\kappa M_p + M_G$ with $\kappa < 1$ —i.e. a dollar of private short-term debt provides less monetary services than a dollar of T-bills. In this case, (24) is $\partial M_p^* / \partial M_G = -\kappa^{-1}[(\kappa^2 v''(\cdot)) / (\kappa^2 v''(\cdot) + (1-p)\phi v g''(\cdot))]$ and (25) is unchanged. Note that we can have $\partial M_p^* / \partial M_G < -1$ if κ is small enough. Thus, the results Table 2, which suggest $\partial M_p^* / \partial M_G < -1$, hint that $\kappa < 1$.

The $(1-\phi)$ term represents the fraction of an investment's return that is not appropriable by an intermediary—i.e., it is a measure of non-pledgeability of investment returns, or alternatively of the extent of financial contracting frictions. This parameter is hard to estimate based on the literature. We somewhat arbitrarily set $(1-\phi) = 0.10$.

The $g'()$ term is the marginal gross return to physical capital in a crisis state. As a practical matter, it can be equivalently thought of as the expected gross return to a deep-pocket investor who purchases assets in a fire sale. Pulvino (1998) documents that during airline-industry downturns, airplanes change hands at discounts on the order of 30% to more normal market prices. Campbell, Giglio, and Pathak (forthcoming) report that foreclosure discounts in forced home sales average 27%. With these numbers as loose motivation, we use 1.30 as our estimate of $g'()$.

Finally, $\partial M_p^* / \partial M_G$ is the equilibrium response of private money creation to shift in the supply of short-term T-bills. While our estimates in Table 2 are noisy and subject to a variety of caveats, they suggest that a value of -1.0 is a reasonable guess. Combining these numbers, we have that $(1-p)(\phi-1)g'(W-M_p^*) \cdot \partial M_p^* / \partial M_G = 0.05 \times -0.10 \times 1.3 \times -1.0 = 0.0065$, or 65 basis points.

This 65 basis-point figure can be compared to the marginal monetary benefits term in (25), given by $v'()$. As previously noted, an on-average estimate for the convenience premium in short-term T-bills is on the order of 40 basis points. Thus the crowding-out benefit and the marginal monetary benefits associated with short-term T-bills would appear to be in the same ballpark. Obviously there are a number of ways this calculation could be off, but it does suggest that the crowding-out effect is likely to be a first-order consideration.

Suppose the government has been ignoring the crowding-out benefit when choosing its debt-maturity structure—i.e., decisions have been based on (14') as opposed to (25). By how much would it raise the short-term share of debt if this additional 65 basis-point benefit were taken into account?

This is equivalent to asking how much the government would adjust the short-term share in response to a 65 basis-point shock to money demand. Unfortunately, this is a difficult question to answer since a major take-away from Section II is that much of the observed variation in both government debt maturity and the money premium is due to shocks to supply, rather than demand shocks. In a crude attempt to overcome this stumbling block we examine the government's response to the one major money demand shock observed in our sample: that in the fall of 2008. From August 14, 2008 to November 6, 2008 the premium on 4-week bills rose from 18 bps to 148 bps—a change of 130 bps. At the same time, T-bills as a share of total debt rose from 27% to 37%. If one interprets the latter 10-percentage-point increase as the government's response to a 130 basis-point money-demand shock, one might loosely argue that taking into account a 65 basis-point crowding-out motive should raise the short-term Treasury share by half this amount, or 5 percentage points. This would represent an economically significant change in the composition of the debt, though not a radical one.

V. Further Extensions

A. Using Debt Maturity and Regulation Together to Control Private Money Creation

To illustrate the crowding-out role for government debt maturity, we have made the extreme assumption that direct regulation of private money creation is impossible. This is clearly unrealistic—as underscored by the fact that much of the current financial reform agenda has focused on regulating the capital structure of intermediaries. A more appealing way to cast the argument is to say that both government debt maturity and direct regulation can be useful in limiting private money creation, but both tools have their limitations, so it is optimal to use them in combination.

To capture this intuition in the model, we assume that the regulatory tool takes the form of a Pigouvian tax at a rate θ_p per unit of private money. We assume that these taxes create their own deadweights costs of $(Y/2) \cdot \theta_p^2$. These costs might reflect efforts by banks to shift activities into

areas where the use of short-term debt is unregulated. For simplicity, we also assume that any revenues from these taxes are rebated lump-sum to the banking industry. Under these assumptions, bank profits are given by:

$$\Pi = [pF + (1-p)\lambda I] - I + M_p[v'(M_0) - \theta_p - (1-p)(1/k - 1)] + L_0, \quad (26)$$

where L_0 is the lump-sum tax rebate. Because individual banks take L_0 as fixed, private money creation is determined by:

$$v'(M_p^* + M_G) = \theta_p + (1-p)(\phi g'(W - M_p^*) - 1). \quad (27)$$

which is the same as equation (21) but for the new θ_p term. Clearly, banks issue more private money M_p^* when the tax rate θ_p on doing so is low (i.e. $\partial M_p^* / \partial \theta_p < 0$).

Armed with two tools to constrain private money creation, what should the government do?

The planner's objective is to maximize:

$$U_{SOCIAL} = E[g(K) - K] + v(M_0) - \frac{1}{2}[\tau_0^2 + E[\tau_1^2] + E[\beta\tau_2^2]] - \frac{\Upsilon}{2}\theta_p^2 \quad (28)$$

which is the same as equation (22) but for $-(\Upsilon/2) \cdot \theta_p^2$, which reflects the social losses associated with the money-creation tax. If $\Upsilon = 0$, this Pigouvian tax can be optimally set so as to make the banks fully internalize the fire-sale externality. As long as Υ is positive, however, then the optimal tax does not completely eliminate the externality.

If we maximize equation (29) with respect to government debt maturity S , we obtain the following generalization of (25):

$$\underbrace{D^{***}b(S^{***} - 1/2)}_{\text{Tax-smoothing cost of govt. money}} = \underbrace{v'(M_p^* + S^{***}D^{***})}_{\text{Marginal benefit of money services}} + \underbrace{\frac{\Upsilon}{\left|\frac{\partial M_p^*}{\partial \theta_p}\right| + \Upsilon} (1-p)(\phi-1)g'(W - M_p^*) \frac{\partial M_p^*}{\partial M_G}}_{\text{Crowding-out benefit}}. \quad (29)$$

As before, the government tilts toward short-term debt both because of the direct benefits derived from money services, as well as the crowding-out benefit. Relative to equation (25) however, the

crowding-out motive is attenuated by a multiplicative term reflecting the deadweight costs of regulation Υ and the responsiveness of private money to the regulation. In the limit when regulation is very costly to implement—which occurs as Υ goes to infinity—we recover our previous solution, with government debt maturity being the only effective policy tool. By contrast, if $\Upsilon = 0$, regulation is costless and the crowding-out benefit vanishes.

Proposition 6: *Suppose $\phi < 1$ and that the tax-smoothing costs of short-term debt and the deadweight costs of regulation are both positive (i.e. $b > 0$ and $\Upsilon > 0$). Then the government uses two tools to control private money creation, setting both a Pigouvian tax on private money ($\theta_p^{***} > 0$) and issuing additional short-term debt to crowd out private money. Moreover, if $v'''(\cdot)$ and $g'''(\cdot)$ are not too large, the government (i) does more crowding out and less regulation when the deadweight costs of regulation are higher and (ii) less crowding out and more regulation when tax-smoothing costs of short-term debt are higher. If the use of either tool in isolation is sufficient costly (i.e. both b and Υ are large enough), then the government increases its reliance on both tools as banks' failure to internalize fire-sales costs becomes more extreme (i.e., $\partial \mathcal{S}^{***} / \partial \phi < 0$ and $\partial \theta_p^{***} / \partial \phi < 0$).*

B. Multiple Maturities of Government Debt

In our simple model, there are only two maturities of government debt: short- and long-term. In reality, the government can issue at any maturity from a few days to 30 years. The presence of a wider range of maturities raises the question of whether the Treasury could increase its issuance of short-term bills—thereby providing more in the way of monetary services—without suffering much of a cost in terms of increased tax volatility. In particular, any reduction in duration brought about by an increase in short-term bills could be offset by shifting the right amount of medium-term bonds into long-term bonds. If duration is in fact an appropriate summary statistic for exposure to tax-volatility risk, it might make sense for the government to pursue a “barbell” strategy, with a lot of short and long-term debt, and relatively little medium-term debt.

What might a move in the direction of such a barbell strategy look like in practice? T-bills averaged 24% of total marketable debt between 1995 and 2009 and the weighted average maturity of

T-bills was 90 days. As shown in Figure 1A, there appears to be a significant money premium in the shortest-maturity T-bills—those with maturities of less than four weeks. And much short-term financing on the part of financial institutions is of extremely short maturity, often overnight. This suggests that by simply reshuffling maturities *within the category of T-bills*, the Treasury could better cater to money demand and have a strong crowding-out effect on private money creation. Suppose the Treasury did so, and cut the average maturity of bills by half to 45 days. To offset the change in duration, the Treasury could simply do a small swap of 10-year notes for 20-year bonds.²³

To see if this informal duration-based intuition holds up, we extend our tradeoff model of government debt maturity from Section II by adding an extra period (i.e., the dates of the model are $t = 0, 1, 2, 3$). The formal analysis becomes fairly involved, so we leave it to the Appendix and focus on a verbal exposition here. The key ingredients are as follows. As before, the government finances a one-time expenditure G at date 0 by issuing short-term (1-period) bonds $B_{0,1}$, medium-term (2-period) bonds $B_{0,2}$, and long-term (3-period) bonds to households and by levying distortionary taxes, τ_0 . At dates 1, 2, and 3, the government must repay maturing debt by levying taxes and issuing new debt.

There are now two uncertain interest rates. At time 1, households learn β_1 which pins down the short rate between periods 1 and 2. At time 2, households learn β_2 which determines the short rate between periods 2 and 3. However, at time 1, households also update their expectations of β_2 based on the realization of β_1 . For example, if all shifts in the yield curve are parallel, then $E[\beta_2 | \beta_1] = \beta_1$. This means that there are now effectively three interest rate shocks: the realization of β_1 at time 1, the “news” about β_2 at time 1, and the ultimate realization of β_2 at time 2.

To begin, consider the perfect smoothing case. As before, the government needs to finance an expenditure of 1. One option is to set $B_{0,1} = B_{0,2} = B_{0,3} = 1/4$, which corresponds to a weighted

²³ For instance, based on a duration of 8.34 for the 10-year and 13.21 for the 20-year on December 31, 2009, a 10-year for 20-year swap equal to $(24\% \times 45/365)/(13.21 - 8.34) = 0.6\%$ of total debt outstanding would hold duration constant.

average debt maturity of 2, which allows taxes to be perfectly smoothed ($\tau_0 = \tau_1 = \tau_2 = \tau_3 = 1/4$). If there is no motive to create monetary services, this is indeed the optimal debt structure.

Now, suppose that the government wishes to provide monetary services at time 0 by increasing the supply of short-term bonds. Assume for simplicity that households do not derive monetary services at time 1 and 2.²⁴ A simple version of the barbell strategy would be to set $B_{0,1} = 3/8$, $B_{0,2} = 0$, and $B_{0,3} = 3/8$, levying taxes $\tau_0 = 1/4$. Note that this keeps the average debt maturity unchanged at 2.

At time 1, the government must pay off debt of 3/8 through a combination of taxes, new short-term debt, and new two-period debt. First, suppose that the term structure remains flat at time 1, so that $\beta_1 = E[\beta_2 | \beta_1] = 1$. The desire to smooth taxes will lead the government to raise taxes of 1/4 at time 1, leaving it with 1/8 to finance using new issues. In order to smooth taxes going forward, the government should raise 1/4 in new short-term debt, using 1/8 to pay off the maturing short-term debt and the other 1/8 to *repurchase* long-term bonds due at time 3. This operation leaves the government with 1/4 of debt maturing at both time 2 and time 3 which it repays by levying taxes of 1/4 at each date. In other words, given an initial barbell maturity structure and the desire to smooth taxes, the government issues additional short-term debt at time 1 and some of the proceeds are used to repurchase long-term bonds maturing at date 3. The issuance and repayment schedule associated with such a barbell strategy is summarized in the table below:

	$t = 0$	$t = 1$	$t = 2$	$t = 3$
G	-1.000			
Taxes	0.250	0.250	0.250	0.250
Issuance of debt due $t = 1$	0.375	-0.375	-	-
Issuance of debt due $t = 2$	0.000	0.250	-0.250	-
Issuance of debt due $t = 3$	0.375	-0.125	0.000	0.250

²⁴ In the Appendix, we show that very similar results obtain if households enjoy monetary services based on the stock of short-term debt outstanding at the interim dates ($t = 1$ and 2) in addition to the initial date. Allowing for monetary services at these interim dates has little impact on the optimal maturity structure at time 0 because the government always has the option to re-optimize its debt maturity structure at times 1 and 2.

Of course, β_1 and $E[\beta_2 | \beta_1]$ are not fixed, but instead vary randomly, which creates uncertainty about future taxes and, in turn, generates tax-smoothing costs. How do taxes and debt issuance vary as a function of the interest rate shocks? Given the initial reliance on short-term debt, the government is naturally exposed to short rates at time 1: the government needs to borrow less and can tax less when β_1 is high (short rates are low) and it can roll-over its existing short-term debt on more favorable terms. However, holding fixed β_1 , the price of the extra long-term bonds the government needs to retire (with the proceeds of new short-term bonds) is increasing in $E[\beta_2 | \beta_1]$. Thus, under the barbell strategy both short-term issuance and taxes at time 1 are increasing in $E[\beta_2 | \beta_1]$. If the realizations of β_1 and $E[\beta_2 | \beta_1]$ are highly correlated—e.g., if there are only level shifts in the yield curve—then this barbell strategy means that the government has offsetting exposures to the β_1 and $E[\beta_2 | \beta_1]$ shocks, better allowing it to smooth taxes. Intuitively, the government creates a “short” position in long-term bonds at time 1. When time 1 movements in short rates and forward rates are highly correlated, this “short” position in long bonds hedges the government’s natural “long” position in short-term bonds.

The key insight from the multi-period extension is that if interest rate shocks are primarily driven by parallel shifts in the yield curve, then a barbell strategy enables the government to hedge out most of the interest rate exposure created by its initial reliance on short-term debt. In the limiting case in which $E[\beta_2 | \beta_1] = \beta_1$, the central tradeoff between tax smoothing and the production of monetary services disappears, because the government can perfectly immunize itself against interest rate shocks. In other words, the barbell strategy allows for a decoupling of money creation—which is accomplished by issuing more short-term debt—and tax smoothing, which, roughly speaking, is accomplished by keeping duration fixed at the right value.

At the other extreme, if $E[\beta_2 | \beta_1]$ and β_1 are independently distributed, then the barbell strategy no longer provides an effective hedge. Thus, if the government wants to create more in the way of

monetary services, it must do so by shortening the weighted average maturity of its debt and by accepting some loss of tax smoothing. More generally, the effectiveness of this kind of barbell strategy is increasing in the correlation between the absolute value of $E[\beta_2|\beta_1]$ and β_1 .

The following proposition, which is proved in the Appendix, summarizes the analysis.

Proposition 7: *Consider an extension of our model in which there are short, medium and long-term bonds, and in which monetary services are only generated by short-term bonds. If the only sources of interest rate shocks are parallel shifts in the yield curve, a barbell strategy—with high proportions of short-term debt and long-term debt, and with little medium-term debt—is optimal for the government, and allows for a large volume of monetary services with little exposure of taxes to interest-rate risk. By contrast, if the yield curve can change shape, the optimum will involve: i) less of a barbell structure—i.e., less short- and long-term debt and more medium-term debt; ii) a shorter average debt maturity; and iii) less in the way of monetary services.*

Although Proposition 7 is suggestive, we should emphasize a couple of qualifications. First, our model does not explain why so much of the monetary premium is concentrated at the very shortest end of the yield curve—i.e., why one-month bills have yields that are so much lower than three-month bills. Absent a better understanding of this phenomenon, any recommendation to shorten maturities within the category of T-bills must be tempered somewhat.

Second, we suspect that our formulation of the tax-smoothing motive may not fully capture all the concerns that government debt managers have in mind when they talk about the “rollover risk” associated with short-term financing. Specifically, one can imagine that even if the discount factor β were fixed, it might be imprudent for the Treasury to put itself in the extreme position of having to roll over all outstanding T-bills every night. Such a strategy might increase its vulnerability to bank-run-type problems, whereby a sudden fear about the government’s ability to service its debts leads to a sharp increase in borrowing costs and makes the run a self-fulfilling prophecy (Bohn (2011)).

However, before one invokes run risk as a reason not to issue more T-bills, it is important to remember the core message of the model: what matters is not the absolute social cost created by a given government debt-maturity structure, but the government's *comparative advantage* in bearing this cost. While it is one thing to argue that the government may face some amount of run risk when issuing a large quantity of short-maturity debt, it is quite another to argue that it is not better-suited than the private sector to bearing such run risk.

VI. Conclusions

A growing body of evidence suggests that low-risk short-term debt securities provide significant monetary services to investors. Moreover, while both the government and private-sector financial intermediaries have the capacity to produce such money-like claims, the private sector's incentives to engage in money creation may be socially excessive, because intermediaries do not fully internalize the fire-sale costs associated with their reliance on short-term funding.

These two observations can be used as a basis for thinking about government debt maturity policy. The most novel insight to emerge from our framework is that government debt maturity can be a useful complement to prudential financial regulation. Rather than addressing private-sector financial fragility solely by writing rules that attempt to constrain the use of short-term debt by intermediaries, the government can also reduce the incentives that lead to excessive private money creation by issuing more short-term debt of its own, thereby compressing the monetary premium and crowding out private issuance. This crowding-out approach may be a powerful tool in a world where it is easy for financial activity to migrate out of the reach of regulators, thereby frustrating the intentions of more traditional capital or liquidity regulations.

One can also use the model to think about how the government should respond to shocks to money demand or private supply of money-like claims that may occur over the course of the business

cycle. For example, consider a scenario in which the private sector's ability to manufacture short-term riskless claims becomes compromised, such as in the midst of a financial crisis. The model suggests that the government should respond to such a shock by expanding the supply of T-bills. Such reasoning seems to have been borne in the fall of 2008, when the Treasury issued \$350 billion of short-term bills within a week of Lehman Brothers' failure, as part of the so-called Supplementary Financing Program. The proceeds from this program were lent to the Federal Reserve, which in turn bought long-maturity assets.

This episode raises the question of the respective roles of the Treasury and the central bank in accommodating the demand for safe short-term debt. For most of the paper, we have implicitly assumed that money demand and private money supply are relatively stable, which suggests that the Treasury could satisfy this demand while still choosing a stable maturity structure for the federal debt. But to the extent that there are important high-frequency shocks to money demand and supply, it might make more sense institutionally for these shocks to be handled by the Federal Reserve. This could be done by changing the maturity structure of the Fed's asset holdings, as in the Fed's recent Maturity Extension Program. For example, the Fed might accommodate a sudden spike in money demand by selling some of its T-bill holdings and using the proceeds to buy longer-term Treasuries. More generally, our model treats the government as a single consolidated entity, and is therefore silent on the question of how to divide responsibilities for provision of money-like claims between the Treasury and the Fed. However, if the normative ideas developed here are to be taken seriously, it will be necessary to think further about this question and the many logistical and political-economy issues that it raises.

References

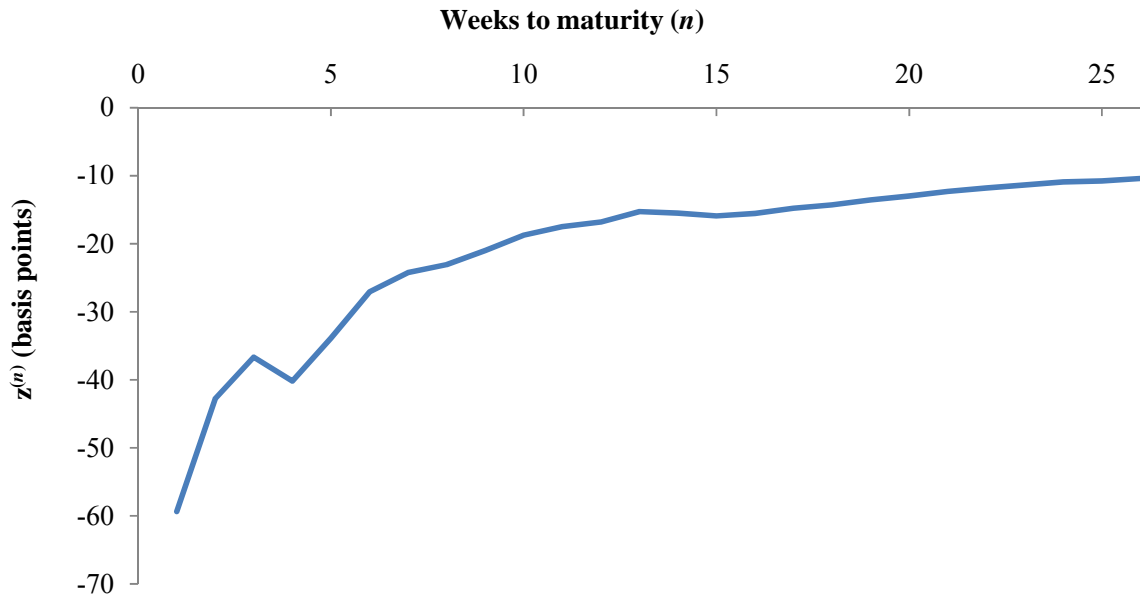
- Aiyagari, S. Rao, Albert Marcet, Thomas J. Sargent, and Juha Seppälä, 2002, Optimal Taxation without State Contingent Debt, *Journal of Political Economy*, 110, 1220-1254.
- Angeletos, Marios, 2002, Fiscal Policy with Non-Contingent Debt and the Optimal Maturity Structure, *Quarterly Journal of Economics* 117, 1105-1131.
- Amihud, Yakov, and Haim Mendelson, Liquidity, Maturity, and the Yields on U.S. Treasury Securities, *Journal of Finance*, 46, 1411-1425.
- Bansal, Ravi and Wilbur John Coleman II, 1996, A Monetary Explanation of the Equity Premium, Term Premium, and Risk-Free Rate Puzzles, *Journal of Political Economy*, 104, 1135-1171.
- Bansal, Ravi, Wilbur John Coleman II, and Christian T. Lundblad, 2010, Liquidity and Financial Intermediation, Working Paper.
- Barro, Robert J., 1979, On the Determination of the Public Debt, *Journal of Political Economy*, 87, 940-971.
- Barro, Robert J., 1997, Optimal Funding Policy, in M. King and G. Calvo, eds., *The Debt Burden and its Consequences for Monetary Policy*, London, Macmillan.
- Barro, Robert J., 2002, Optimal Management of Indexed and Nominal Debt, in F. Lefort and K Schmidt-Hebbel, eds., *Indexation, Inflation, and Monetary Policy*, Central Bank of Chile.
- Barro, Robert J., and Jose F. Ursua, 2008, Macroeconomic Crises since 1870, *Brookings Papers on Economic Activity*, Spring.
- Bennett, Paul, Kenneth Garbade, John Kambhu, 2000, Enhancing the Liquidity of U.S. Treasury Securities in an Era of Surpluses, *Federal Reserve Bank of New York Economic Policy Review*, 6, 89-125.
- Benigno, Pierpaolo, and Michael Woodford, 2003, Optimal Monetary and Fiscal Policy: A Linear Quadratic Approach, *NBER Macroeconomics Annual* 2003.
- Blanchard, Olivier, and Alessandro Missale, 1994, The Debt Burden and Debt Maturity, *American Economic Review* 84, 309-319.
- Bohn, Henning, 1988, Why Do We Have Nominal Government Debt?, *Journal of Monetary Economics*, 21, 127-140.
- Bohn, Henning., 1990, Tax Smoothing with Financial Instruments, *American Economic Review*, 80, 1217-1230.

- Bohn, Henning. 1999. Fiscal Policy and the Mehra-Prescott Puzzle: On the Welfare Implications of Budget Deficits When Real Interest Rates Are Low, *Journal of Money, Credit and Banking*, 31(1): 1-13.
- Bohn, Henning, 2011, The Economic Consequences of Rising U.S. Government Debt: Privileges at Risk, Working Paper.
- Calvo, Guillermo A. and Pablo E. Guidotti, 1990, Indexation and Maturity of Government Bonds: An Exploratory Model, in R. Dornbusch and M. Draghi eds., *Public Debt Management: Theory and History*, Cambridge University Press, Cambridge.
- Campbell, John, Stefano Giglio, and Parag Pathak, 2011, Forced Sales and House Prices, *American Economic Review* 101, 2108-31.
- Diamond, Douglas W., and Raghuram Rajan, 2009, Fear of Fire Sales and the Credit Freeze, Working Paper.
- Duffee, Gregory R., 1996, Idiosyncratic Variation of Treasury Bill Yields, *Journal of Finance* 51, 527-551.
- Froot, Kenneth A, and Jeremy C. Stein, 1998, Risk Management, Capital Budgeting and Capital Structure Policy for Financial Institutions: An Integrated Approach. *Journal of Financial Economics* 47, 55-82.
- Garbade, Kenneth and Matthew Rutherford, 2007, Buybacks in Treasury Cash and Debt Management, Staff Report 304, Federal Reserve Bank of New York.
- Geanakoplos, John, 2010, The Leverage Cycle, Cowles Foundation Discussion Paper 1715R.
- Gorton, Gary B., 2010, *Slapped by the Invisible Hand: The Panic of 2007* (Oxford University Press).
- Gorton, Gary B., and Andrew Metrick, 2011, Securitized Banking and the Run on Repo, *Journal of Financial Economics*, forthcoming.
- Guibaud, Stephane, Yves Nosbusch, and Dimitri Vayanos, 2007, Preferred Habitat and the Optimal Maturity Structure of Government Debt, Working Paper.
- Greenwood, Robin, Samuel Hanson, and Jeremy C. Stein, 2010, A Gap-Filling Theory of Corporate Debt Maturity Choice, *Journal of Finance* 65, 993-1028.
- Greenwood, Robin, and Dimitri Vayanos, 2010, Bond Supply and Excess Bond Returns, NBER Working Paper No.13806.
- Gurkaynak, Refet S., Brian Sack, and Jonathan H. Wright, 2006, The U.S. Treasury Yield Curve: 1961 to the Present, Federal Reserve Board Working Paper No. 2006-28.

- Krishnamurthy, Arvind, and Annette Vissing-Jorgensen, 2010, The Aggregate Demand for Treasury Debt, Working Paper.
- Laeven, Luc and Fabian Valencia, 2010, Resolution of Banking Crisis: The Good, the Bad, and the Ugly, IMF working paper.
- Lorenzoni, Guido, 2008, Inefficient Credit Booms, *Review of Economic Studies* 75, 809-833.
- Lucas, Robert E., and Nancy L. Stokey, 1983, Optimal Fiscal and Monetary Policy in an Economy Without Capital, *Journal of Monetary Economics* 12, 55-93.
- Lustig, Hanno, C. Sleet and S. Yeltekin, 2006, Fiscal Hedging and the Yield Curve, Working Paper.
- Missale, Alessandro, 1997, Managing the Public Debt: The Optimal Taxation Approach. *Journal of Economic Surveys*, 11(3): 235-65.
- Missale, Alessandro, 1999, Public debt management, Oxford University Press, New York.
- Nosbusch, Yves, 2008, Interest Costs and the Optimal Maturity Structure of Government Debt, *Economic Journal* 118, 477-498.
- Pulvino, Todd, 1998, Do Asset Fire Sales Exist? An Empirical Investigation of Commercial Aircraft Transactions, *Journal of Finance* 53, 939-978.
- Sachs, Lee, 1999, Statement of Lee Sachs, Assistant Secretary, Financial Markets, before the House Committee on Ways and Means, Hearing on Treasury's Debt Buyback Proposal. September 29, 1999.
- Shleifer, Andrei, and Robert W. Vishny, 2010, Unstable Banking, *Journal of Financial Economics*, forthcoming.
- Sidrauski, Miguel, 1967, Rational Choice and Patterns of Growth in a Monetary Economy, *American Economic Review* 57, 534-544.
- Stein, Jeremy C., 2012, Monetary Policy as Financial-Stability Regulation, *Quarterly Journal of Economics*, forthcoming.

Figure 1. The money premium on short-term Treasury Bills, 1983-2009. Panel A plots the average spread, over the period 1983-2009, between actual Treasury-bill yields (“on-cycle” Treasury bills with maturities from 1 to 26 weeks) and fitted yields, based on a flexible extrapolation of the Treasury yield curve from Gurkaynak, Sack and Wright (2006) as updated regularly by the Federal Reserve Board of Governors. Gurkaynak et al (2006) estimate a parametric model of the instantaneous forward rate curve that is characterized by six parameters. The set of sample securities used to estimate the curve each day includes almost all “off-the-run” Treasury notes and bonds with a remaining maturity of more than 3 months. For each quarter from 1983-2009, Panel B plots the average 4-week T-bill premium (actual yield minus fitted yield) against the ratio of Treasury bills to GDP.

Panel A: Average premium (actual Treasury-bill yield minus fitted yield) by week to maturity



Panel B: 4-week T-bill premium (actual yield minus fitted yield) and the bills-to-GDP ratio

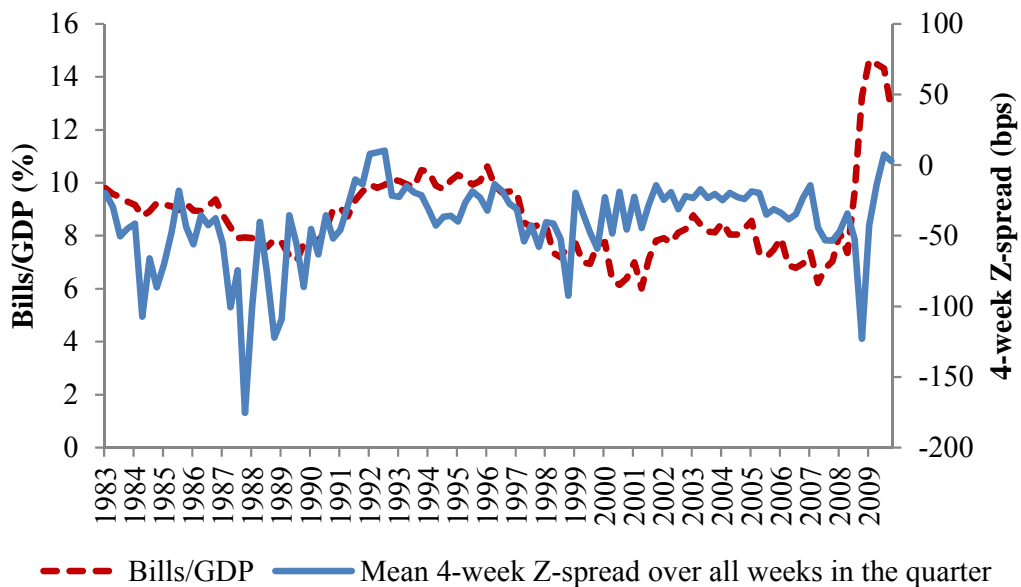
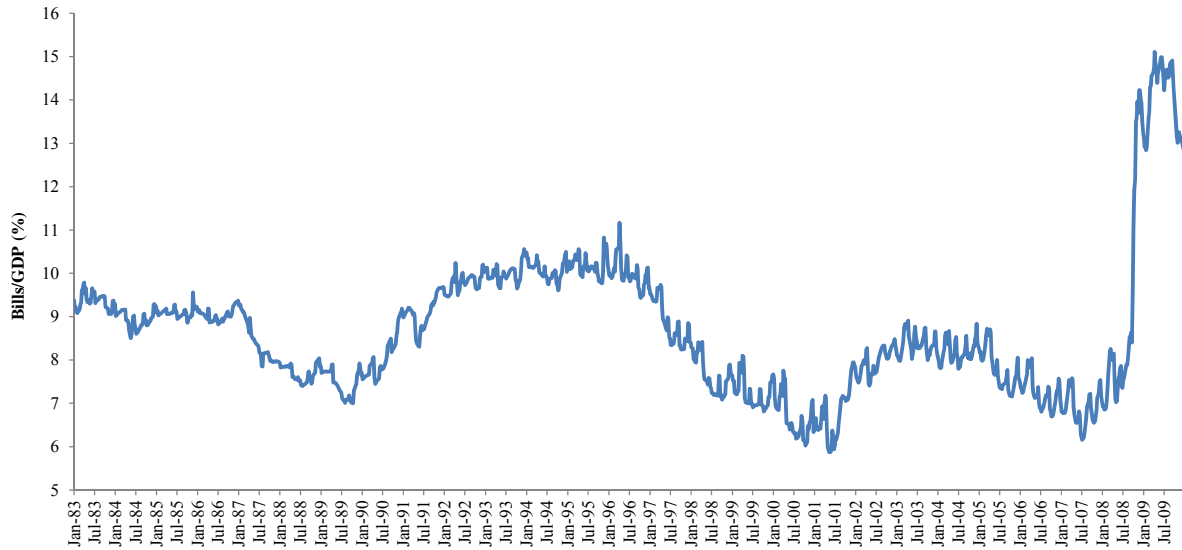


Figure 2. The supply of Treasury bills, 1983-2009. Panel A plots the ratio of Treasury bills to *GDP* each week from 1983-2009. To compute the ratio of Treasury bills to *GDP* precisely at the end of each week, we use detailed data on the size and timing of Treasury auctions. Panel B plots the coefficients from the first-stage regression from our instrumental variables estimator for 1983-2009. Specifically, we regress $\Delta_4 \text{BILLS} / \text{GDP}$ on a full set of week-of-year dummies (the omitted category is the 1st week of the year):

$$\Delta_4(\text{BILLS} / \text{GDP})_t = c + \sum_{w=2}^{53} d^{(w)} \cdot \mathbf{1}\{\text{week}(t) = w\} + \Delta_4 v_t.$$

Confidence intervals are based on Newey-West standard errors (1987) allowing for serial correlation up to 8 weeks. We mark the six major Federal tax deadlines on the figure: January 15th, March 15th, April 15th, June 15th, September 15th,

Panel A: The ratio of Treasury bills to *GDP* from 1983-2009.



Panel B: High-frequency variation in the supply of Treasury bills.

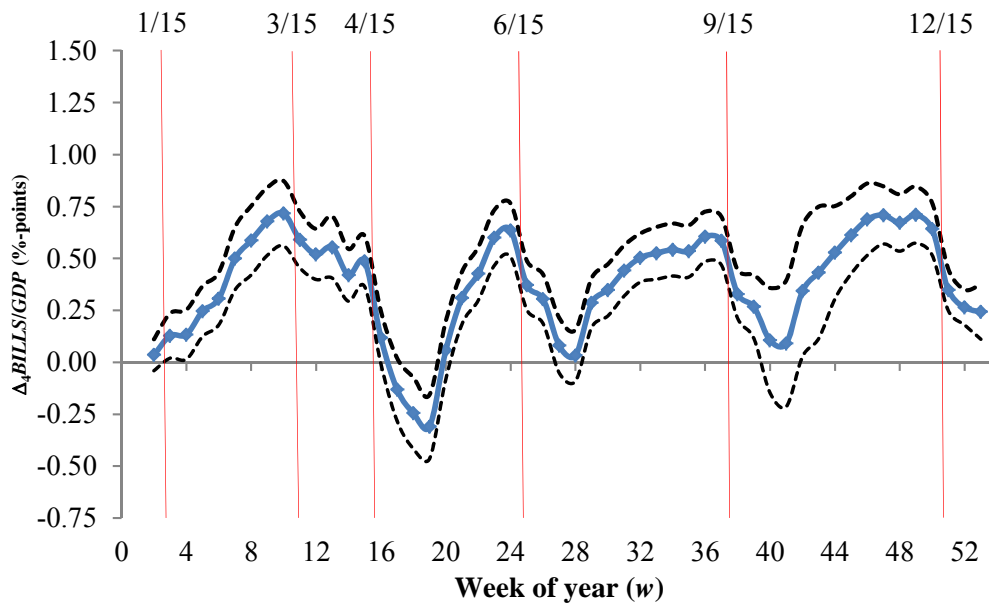


Figure 3. Debt/GDP and the maturity of U.S. Treasury debt, 1952-2009. The figure plots the weighted average maturity of all marketable Treasury debt against the debt-to-GDP ratio from 1952-2009. The weighted average maturity is calculated using issue-level data on debt outstanding from CRSP and excludes data on Treasury Inflation Protected Securities (TIPS). The debt-to-GDP ratio is obtained from Henning Bohn's website.

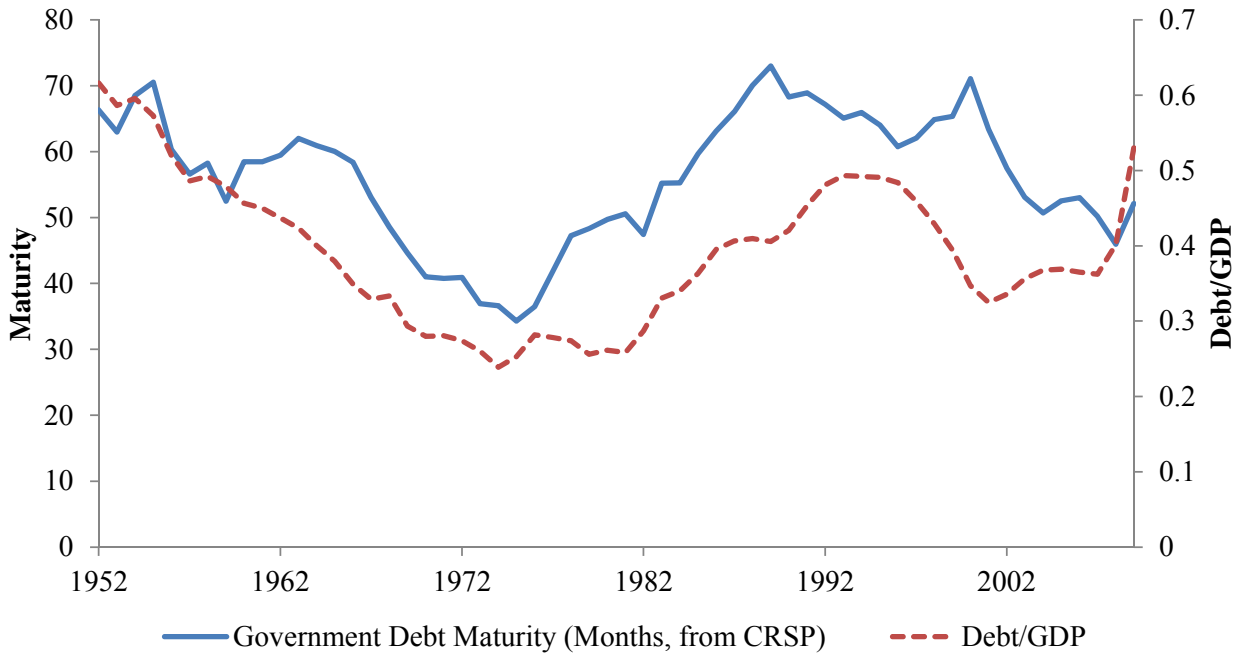


Figure 4. Corporate and government debt maturity, 1963-2009. The figure reproduces and extends Figure 1 from Greenwood, Hanson, and Stein (2010). The dashed line, plotted on the left axis, is the share of long-term corporate debt as a fraction of total debt, based on *Flow of Funds* data. The solid line, plotted on the right axis, is the share of marketable Treasury debt with maturity of one year or less based on CRSP data.

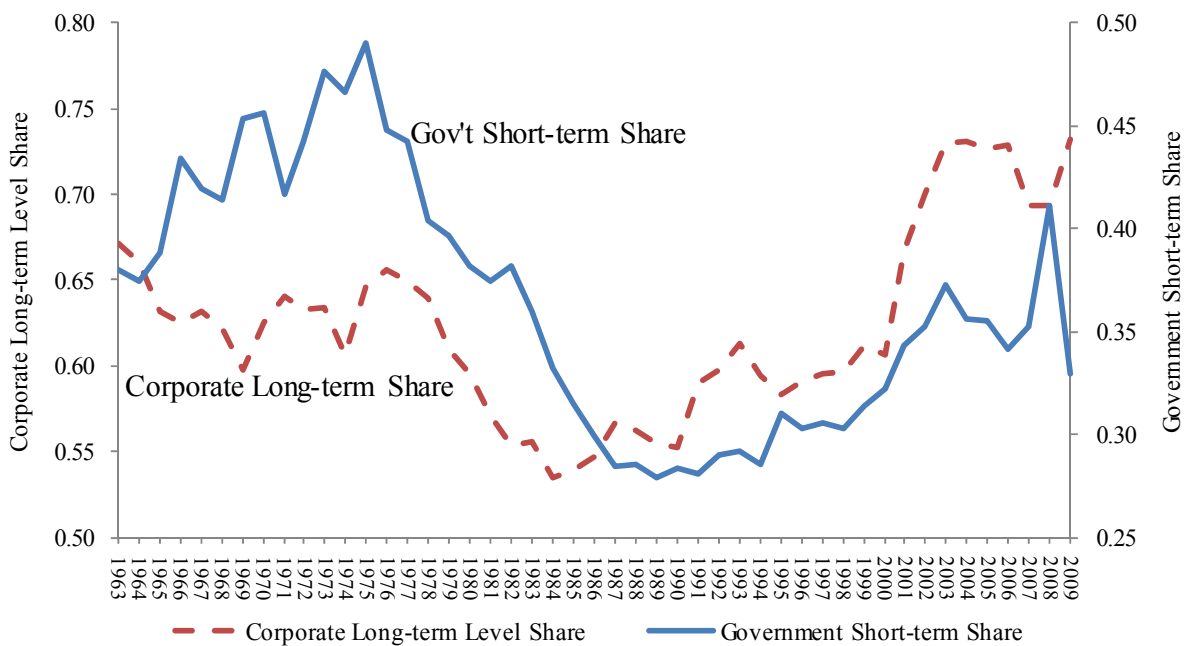


Table 1
The money premium and the supply of short-term Treasuries

Regressions of z -spreads on the supply on Treasury bills. The n -week z -spread $z_t^{(n)} = y_t^{(n)} - \hat{y}_t^{(n)}$ is the difference between the actual yield on an n -week Treasury bill and the n -week fitted yield, based on the fitted Treasury yield curve in Gurkaynak, Sack and Wright (2006) as updated regularly by the Federal Reserve Board of Governors. Each day they estimate a 6-parameter model of the instantaneous forward curve. Zero coupon yields are derived by integrating along the estimated forward curve. The parameters for each day are estimating by minimizing a weighted sum of pricing errors, where the set of sample securities includes almost all off-the-run Treasury notes and bonds with a remaining maturity of more than 3 months. Using weekly data, we regress the n -week z -spread on the supply of Treasury bills scaled by GDP. We estimate this specification in both level and in 4-week changes:

$$z_t^{(n)} = a^{(n)} + b^{(n)} \cdot (BILLS / GDP)_t + \varepsilon_t^{(n)} \quad \text{and} \quad \Delta_4 z_t^{(n)} = a^{(n)} + b^{(n)} \cdot \Delta_4 (BILLS / GDP)_t + \Delta_4 \varepsilon_t^{(n)}.$$

To compute the ratio of Treasury bills to GDP precisely at the end of each week, we use detailed data on the size and timing of Treasury auctions. The first six columns report OLS estimates in levels and changes for $n = 2, 4$, and 10-week bills. The final three columns report instrumental variables (IV) estimates which exploit seasonal variation in Treasury bill supply driven by the Federal tax calendar. In the first stage of the IV regressions, we regress $\Delta_4 BILLS / GDP$ on a full set of week-of-year dummies; in the second stage, we regress changes in z -spreads on the fitted values from the first stage. The units of the dependent variable are basis points and the units of the independent variable are percentage points. t -statistics are based on Newey-West (1987) allowing for serial correlation up to 12 weeks in the levels regressions and up to 8 weeks in the changes regressions. Panel B shows results for the 1983-2007 subsample.

	Levels			Changes			IV		
	OLS			OLS			IV		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	$z^{(l)}$	$z^{(4)}$	$z^{(10)}$	$\Delta_4 z^{(l)}$	$\Delta_4 z^{(4)}$	$\Delta_4 z^{(10)}$	$\Delta_4 z^{(l)}$	$\Delta_4 z^{(4)}$	$\Delta_4 z^{(10)}$
	2-week	4-week	10-week	2-week	4-week	10-week	2-week	4-week	10-week
Panel A: 1983-2009									
$b^{(n)}$	5.60	3.96	2.04	15.75	7.59	0.65	29.37	10.49	6.32
$[t]$	[2.67]	[2.09]	[1.29]	[1.55]	[1.20]	[0.22]	[3.39]	[1.70]	[2.69]
R^2	0.03	0.02	0.03	0.01	0.00	0.00	0.01	0.00	0.01
Panel B: 1983-2007									
$b^{(n)}$	9.58	6.50	5.90	34.89	20.29	7.41	29.99	11.51	8.09
$[t]$	[4.69]	[3.69]	[6.45]	[6.35]	[5.04]	[3.67]	[3.18]	[1.73]	[3.39]
R^2	0.04	0.03	0.15	0.03	0.02	0.02	0.01	0.00	0.02

Table 2
Private money creation and the quantity of short-term government debt, 1952-2009

Time-series regressions of private money creation on the debt-to-GDP ratio or the ratio of short-term debt to GDP:

$$(PrivateMoney / GDP)_t = a + b \cdot (D / GDP)_t + u_t \quad \text{and} \quad (PrivateMoney / GDP)_t = a + b \cdot (D_S / GDP)_t + u_t.$$

The data are annual and begin in 1952. D is marketable debt held by the public as reported in the Treasury Bulletin and D_S is marketable debt held by the public with a remaining maturity of less than one year. We measure private money creation using data from the Federal Reserve's H6 report and alternately define private money creation as non-M1 M2 divided by GDP or as non-M1 M3 divided by GDP. Non-M1 M2 consists of savings deposits and small-time deposits at banks and thrifts as well as retail money market funds. Savings deposits and small-time deposits are typically non-reservable liabilities and hence private banks can freely expand and contract the supply of these forms of money. By contrast, checking deposits and other reservable liabilities in M1 are pinned down by the quantity of central-bank controlled reserves. Non-M2 M3 consists of large time deposits at banks and thrifts, Eurodollar deposits, repurchase agreements, and institutional money market funds. M3 is published 1959-2006, but we compute the series for the missing years using data on the individual components. Nominal GDP is from the Bureau of Economic Analysis. t -statistics based on Newey-West (1987) standard errors, with 3-years of lags, are shown in brackets.

	Dep Var = (M2-M1)/GDP		Dep Var = (M3-M1)/GDP	
D/GDP	-0.515 [-2.91]		-0.659 [-1.44]	
D_S/GDP		-1.529 [-3.61]		-2.952 [-2.76]
N	58	58	58	58
R^2	0.38	0.52	0.11	0.33