The Imperfect Intermediation of Money-Like Assets*

Jeremy C. Stein
Harvard University and
NBER

Jonathan Wallen
Harvard Business School
Harvard University

July 2024

Abstract:
We study supply-and-demand effects in the U.S. Treasury bill market by comparing the returns on T-bills to the policy rate on the Federal Reserve's reverse repurchase (RRP) facility. We develop and test a simple model, where the RRP-bill spread is policed by both: (i) heterogeneously elastic money funds; and (ii) corporate treasurers who derive collateral benefits from holding T-bills. When T-bills initially become scarce and the spread widens, money funds act as the front-line arbitrageurs, and the effect of T-bill supply shifts on the spread depends on the elasticity of the marginal money fund. Eventually, when T-bills become extremely scarce and all money funds have left the T-bill market, the RRP-bill spread is entirely pinned down by corporate treasurers' collateral demand function.

*emails: jeremy_stein@harvard.edu and jwallen@hbs.edu. We are grateful to Wenxin Du, Robin Greenwood, Sam Hanson, Sebastian Hillenbrand, Victoria Ivashina, Adi Sunderam, Luis Viceira, seminar and conference participants at HBS, UMass Amherst, and the 2024 AFA for helpful feedback, and to Viral Acharya for discussing the paper.

I. Introduction

A growing literature emphasizes how the frictions and constraints associated with financial intermediation can influence the behavior of asset prices. In this paper, we apply a frictional-intermediation lens to study pricing anomalies in the U.S. Treasury bill market. More specifically, we compare the returns on T-bills to those on another safe money-like government claim, namely an investment in the Federal Reserve's reverse repurchase (or RRP) facility. The RRP facility allows a pre-approved group of eligible counterparties, including primary dealers, government-sponsored enterprises, and a set of over 100 money funds, to lend to the Fed on an overnight basis against Treasury collateral, at an administered policy rate. Both of these markets are large, with the RRP facility reaching a peak usage of $2.55 trillion in December 2022, at a time when the private supply of T-bills (excluding T-bills held by the Fed) was $3.4 trillion.

Figure 1 displays the one-month realized and ex-ante expected returns on the RRP facility alongside the rate on one-month T-bills, over the period from June 2021 to June 2024.[1] A priori, one might have guessed that the expected RRP rate would be a lower bound on the T-bill rate. Both instruments have the full backing of the U.S. government and hence are free of credit risk. Moreover, in addition to having zero duration and hence no exposure to interest-rate risk, an investment in RRP can also be said to be more purely money-like than one in T-bills, since it liquidates and hence can be monetized at zero cost on an overnight basis; by contrast, to monetize a T-bill overnight would require selling it and thereby accepting some bid-ask cost. These differences should, in principle, translate into a weakly lower rate of return on the RRP as compared to T-bills.

Yet, as Figure 1 shows, this hypothesis is rejected by the data, sometimes dramatically so. While T-bill and expected RRP rates were nearly identical from June 2021 through early 2022, after that they began to diverge, with one-month T-bill rates falling well below expected RRP rates. This gap frequently exceeded 50 basis points in 2022, before spiking to over 160 basis points during the initial period of uncertainty over the debt ceiling in March and April of 2023. After the

[1] Because we are comparing one-month T-bills to overnight RRP, and because interest rates were rising during this period, we need to be careful to control for the difference in maturities. Looking at realized one-month returns on a strategy that invests continuously in the RRP facility is one way to do this. Alternatively, we can proxy for the ex-ante expected return on the RRP over a one-month horizon by adding the one-month OIS rate, and subtracting the current Fed Funds rate, with this adjustment intended to capture expected changes in overnight rates over the coming month. Thus we have: expected one month RRP return = RRP Rate + one month OIS – Fed Funds. As can be seen in Figure 1, our measures of realized and expected one-month RRP returns track each other closely, suggesting that our proxy for expected RRP returns is a sensible one.

debt ceiling was raised, an additional $1.8 trillion dollars of T-bills was issued over the period June 3, 2023, to June 30, 2024, and the spread collapsed back to zero.

In an effort to understand this puzzling behavior, we consider three types of frictions that could potentially impede the seemingly simple arbitrage that would appear to be available whenever T-bills offer a lower return than that on the RRP facility—namely, one would expect anyone holding the lower-yielding T-bills to sell these T-bills and to replace them with an investment in the RRP. The first friction is a form of market segmentation created by the fact that the Fed's RRP facility is not available to all investors. Rather, as noted above, access is limited to primary dealers, GSEs, and a select group of money funds, with money funds being the dominant participants in the facility by far, holding on average about 90 percent of the outstanding volume of RRP (Afonso, Cipriani, and La Spada, 2022). Money funds are also an important player in the T-bill market, where their share of outstanding T-bills has averaged 15 percent over the period of September 2013 to June 2024.

By definition, if an investor without access to the RRP facility—whom we henceforth label a "corporate treasurer" for concreteness—is unhappy with the relatively low yield on her holdings of T-bills, she cannot directly substitute into RRP. She can attempt to substitute indirectly, by selling her T-bills and putting the proceeds in a money fund that then invests primarily in the RRP. But as we will see, such indirect substitution via inflows into RRP-intensive money funds is a weak force in the data. So, it turns out to be a reasonable approximation to think of there being a wall between the corporate treasurers of the world and the money funds, or alternatively to think of money fund assets under management as exogenously determined and unaffected by movements in the RRP-bill spread.

An alternative strategy that is available to a corporate treasurer would be to sell her T-bills and put the proceeds into *private repo*, i.e., to lend against Treasury collateral, but to a private party, as opposed to the Federal Reserve. This strategy does not require access to the RRP facility, yet it turns out to be economically similar, since as we show below, the rates on the RRP facility and on private repo are generally very close to one another, implying that the two are nearly perfect substitutes. Therefore, the failure of this second type of arbitrage to close the RRP-bill spread implies that—perhaps surprisingly—limited access to the Fed's RRP facility may in fact not be the pivotal friction, and that broadening the facility's access to other money-market investors might therefore not be sufficient to eliminate the spread.[2] Simply put, corporate treasurers appear to prefer holding T-bills to either money-fund shares or to *any form* of repo.

This observation leads us to consider a second friction, which is a difference in pledgeability. T-bills can be more readily pledged as collateral for a variety of purposes, including, e.g., margin requirements in derivatives transactions.[3] In other words, the RRP-bill spread may reflect an equilibrium collateral premium that gets larger when the supply of collateral (i.e., T-bills) available to corporate treasurers declines relative to their demand for such collateral. Although we cannot establish that this mechanism is the whole story behind the limited responsiveness of the non-money-fund sector to the RRP-bill spread, we are able to provide suggestive evidence that collateral considerations play some role in shaping their relative demand for Treasury securities. In a sample of large non-financial corporates, we find that more intensive use of derivatives is associated with a larger share of Treasury holdings relative to other types of securities and cash-like instruments.

This collateral-demand story would appear to leave the money funds themselves as the front-line players who are in the best position to police the RRP-bill spread, both because they have direct access to the RRP facility, and because money funds do not typically use derivatives or engage in other types of transactions where collateral is required. Thus, one might expect any money fund to respond very elastically to an opportunity to substitute away from lower-yielding T-bills and into the RRP facility, which for all practical purposes is available in unlimited supply from the Fed at a fixed rate.[4] However, this is where we encounter our third friction: the evidence sharply contradicts this presumption. Money funds are slow and inelastic in their substitution from T-bills to RRP. Even when the spread is large, money funds continue to actively buy T-bills. This inelasticity in turn leads the money funds to cede fewer T-bills to the corporate treasurers when T-bill supply shrinks relative to demand, thereby widening the equilibrium collateral premium that

[2] This differs from other settings where partial segmentation has been found to be a pivotal friction. For example, Bech and Klee (2011) show that limited access caused a wedge between the interest on reserves (IORB) rate and the fed funds rate because only banks could earn IORB, while other lenders in the fed funds market could not.

[3] Interestingly, the Money Funds Association has recently requested that regulators permit the use of money fund shares as collateral for derivatives transactions, suggesting that the current lack of pledgeability is a consequential difference between money fund shares and Treasuries. See their letter to the Commodity Futures Trading Commission at https://www.isda.org/a/FHPTE/ISDA-Letter-to-US-Regulators-Cash-and-Money-Market-Funds-as-Initial-Margin-8.1.19.pdf. The CFTC has also drafted a rule permitting this type of pledgeability for government money funds. See https://www.cftc.gov/media/9721/gmac_TechnicalSubcommitteeRecommendationsMMF110623/download.

[4] The program has a fund-level cap, but funds are hardly ever near the cap. Also, the cap has been raised several times, from 30 billion to 80 billion in March 2021, and from 80 billion to 160 billion in September 2021.

the treasurers end up paying.

Digging deeper, we document that there is substantial heterogeneity in the cross-section of money funds with respect to their elasticity of substitution: some funds exhibit a persistently (and thus predictably) lower elasticity than others. As we show with the help of a simple model, this heterogeneity can lead to interesting time variation in how the RRP-bill spread responds to various supply and demand shocks, such as a shock to the supply of T-bills. To see the intuition, imagine that we start out with the RRP-bill spread at zero. Now there is a negative shock to the supply of T-bills that creates some T-bill scarcity and thus begins to drive the spread positive. The most elastic money funds will readily accommodate this shock, ceding their holdings of T-bills and moving into RRP as a substitute. So the impact of the shock on the spread will initially be muted. But if the shock is large enough, these highly elastic funds will eventually hit a corner, where they are out of T-bills, and hence can no longer buffer any further shocks. At this point, there are fewer remaining non-cornered money funds, and these funds are also less elastic, so a subsequent supply shock of the same magnitude will tend to have a larger impact on the RRP-bill spread.

In the limit, if the T-bill supply shrinks enough, and the RRP-bill spread gets sufficiently large, we can get to a point where the entire money fund sector is at a corner, holding a large quantity of RRP and essentially no T-bills. Indeed, this is what happened around the time of the debt-ceiling-related turmoil in March and April of 2023, when there was an especially large and inelastic demand from outside the money fund sector for those T-bills that matured before the June 1 date at which point the government was predicted to run out of money.[5] At this point, our model predicts that the RRP-bill spread can become extremely sensitive to even modest shocks in T-bill supply and demand. As we discuss further below, this mechanism appears to be part of the story behind the sharp gyrations in one-month T-bill yields seen during this period.

It might be argued that some of the inelastic behavior that we find among money funds can be attributed to a form of gradual learning. According to this story, the RRP facility was a relatively new and untested instrument for money-fund managers in our sample period, and they may have been wary of ramping up their use of it too rapidly the first time around. If so, it is possible that their behavior will look different, and more elastic, in subsequent episodes. In this scenario, our model implies that future modest changes in T-bill supply will initially have a reduced impact on the RRP-bill spread. However, the same logic implies that the money fund sector as a whole will

[5] On May 1st, 2023, the Treasury informed Congress that it would be unable to meet its obligations by June 1st.

hit the corner sooner in the face of shrinking T-bill supply, at which point the only players left in the market are the corporate treasurers, whose elasticity of substitution is shaped by their collateral-demand considerations. And it seems plausible to think that this sort of friction is more structural and durable in nature. Which would imply that even if money funds do become more elastic next time there is a large shock to T-bill supply, the ultimate implications for the RRP-bill spread may be similar to those we document below.

The primary aim of this paper is to shed light on the factors that influence rate dynamics in one of the most important money markets in the world, the market for Treasury bills. The fact that T-bill rates can become disconnected from the rate paid by the Fed on its RRP facility would at first glance seem to suggest that the RRP facility as currently designed is not fully living up to the Fed's articulated goal for it, which is to put a floor on all short-term money market rates. On the Federal Reserve Bank of New York's website, a description of the program says that: "The [RRP] provides a floor under overnight interest rates by offering a broad range of financial institutions that are ineligible to earn [interest on reserve balances] an alternative risk-free investment option."

However, one message that emerges here is that no RRP-like facility, no matter how broadly accessible it is made, can be expected to put an absolute floor on T-bill rates if T-bills provide a form of collateral services that repo does not—e.g., if it is easier to pledge T-bills as collateral to satisfy margin requirements in derivatives transactions. Indeed, once the money fund sector is at a corner, and is no longer holding any T-bills, our model implies that the RRP-bill spread is entirely shaped by collateral demand on the part of other players.

Perhaps more surprising is that the money funds themselves, who would not appear to have any obvious source of collateral demand, are also quite inelastic in terms of their response to the RRP-bill spread. Of course, the literature is full of examples wherein frictions to financial intermediation can leave a noteworthy imprint on asset prices (Duffie 2010; Adrian et al. 2014; He et al. 2017; Du et al. 2018; Du et al. 2020; Ivashina et al. 2015; Bech and Klee 2011; Haddad and Muir 2021; d'Avernas and Vandeweyer 2023).[6] The novelty here is that when it comes to the money funds, we are looking at a setting that is about as simple and transparent as one can imagine. The spread between RRP and T-bills is directly observable and money funds have uncomplicated and regularly disclosed balance sheets.[7]

---

[6] For a more comprehensive review of the literature see He and Krishnamurthy (2018).

[7] In contemporaneous work, Doerr, Eren and Malamud (2023) also study the spread between RRP and T-bills, and the associated role of money funds. Their emphasis, however, is different than ours: they focus on imperfect competition among money funds, and their strategic behavior, while our model and empirical tests highlight the importance of

More broadly, we also hope to contribute to the literature on intermediary asset pricing by highlighting the role of heterogeneity among intermediaries, and the implications this has for the responsiveness of asset prices to various shocks. The theoretical literature tends to model a situation where the relevant state variable is the wealth or risk-bearing capacity of the aggregate intermediary sector (He and Krishnamurthy 2013; Brunnermeier and Sannikov 2014; Gromb and Vayanos 2018).[8] By contrast, we show that the composition of different types of intermediaries active in the market—e.g., the relative importance of money funds and corporate treasurers—is a key state variable for understanding the price impact of supply and demand shocks.

There are other circumstances where one might expect a similar mechanism to be relevant. For example, Breckenfelder and De Falco (2023) argue that successive rounds of quantitative easing by the European Central Bank (ECB) should be expected to have increasing bang-for-the-buck over time, as more elastic investors exhaust their holdings of eligible securities, leaving behind only less elastic investors.[9] What makes our setting a particularly nice laboratory for examining this hypothesis is that we know precisely who one class of the relevant intermediaries are—i.e., the money funds—and it is straightforward to see when each individual fund is at a corner, in terms of having sold off all of its T-bills. Thus, we can cleanly isolate the state variable that the model tells us is of interest, the fraction of money funds that are at a corner at any point in time. And we can then conduct a simple test to see if this state variable does in fact shape the responsiveness of T-bill rates to supply shocks, as predicted by the model.[10]

The remainder of the paper is organized as follows. In Section II, we develop a simple model of the relative pricing of T-bills and RRP that features both: (i) a heterogeneous group of money funds, who differ in their willingness to substitute between the two assets; and (ii) a set of corporate treasurers, who derive collateral benefits from holding T-bills relative to other money-market assets. In Section III, we examine a range of data that bears on both the model's key assumptions about the frictions in this sector, as well as on its implications for how T-bill rates respond to supply shocks. Section IV concludes.

---

heterogeneous preferences among a set of perfectly competitive funds, as well as the role of non-money-fund investors, i.e., the agents we are calling corporate treasurers.

[8] A notable exception is Siriwardane et al. (2022) who show that law of one price violations are weakly integrated across markets, implying that the intermediary sector is segmented across markets.

[9] For the equity market, Haddad et al. (2022) shows that demand shocks have a larger impact on stocks that have more passive investors, thereby highlighting the importance of differences in the elasticity of demand of different investors.

[10] In a somewhat similar spirit is Duffie et al. (2023) who document that liquidity in the Treasury bond market deteriorates significantly when dealers are close to hitting a capacity constraint in terms of the quantity of Treasury securities they have taken onto their balance sheets.

II. The Model

We consider a model with the following features. There is a total supply of T-bills given by S. These T-bills can be held by one of two types of investors: money funds, and corporate treasurers, (or just "treasurers" for short). Money funds also have access to the Fed's RRP facility, while treasurers do not and can therefore only invest in T-bills or in private repo. There is a mass of competitive money funds that together have aggregate assets under management (AUM) that is exogenously fixed at A. By fixing the aggregate AUM of the money-fund sector in this way, and in particular by making it insensitive to the spread between RRP and T-bills, we build in the first potential friction in the model: the hard segmentation between money funds and treasurers. Below, we will show that something close to this degree of segmentation effectively holds in the data.

We assume that treasurers' collateral-based demand for T-bills, the second key friction in the model, is given by $X(r_b - r_p)$, where r_b is the endogenously-determined T-bill rate, r_p is the RRP rate, and $X(\cdot)$ is an increasing function. In other words, as T-bill rates fall relative to the repo rate, treasurers economize on their use of T-bills as collateral, and substitute into private repo, which does not offer the same collateral services. In writing the demand of treasurers as a function of $(r_b - r_p)$, we are therefore assuming that the private repo rate is the same as the rate on the RRP facility; we show below that this assumption holds closely in the data. And because the RRP rate r_p is exogenously set by the Federal Reserve, we can simplify the notation by writing treasurers' demand for T-bills as just $X(r_b)$.

Each money fund i divides its portfolio between holding T-bills with weight T_i, and RRP with weight $(1 - T_i)$. We assume that T_i is bounded between 0 and 1, meaning that money funds cannot short T-bills or make leveraged investments. Money funds are heterogeneous in their relative preference for RRP and T-bills and have utility:

$$U_i = r_b T_i + r_p(1 - T_i) + V_i(T_i), \quad (1)$$

where $V_i(T_i)$ is an increasing concave function that allows money funds to derive some non-pecuniary benefit from holding T-bills. This non-pecuniary benefit is the third key friction in the model. Unlike with the corporate treasurers, it is less plausible to argue that it reflects collateral demand, since money funds typically do not engage in the sorts of transactions for which collateral is required. So we are somewhat more agnostic as to its origin, and hence as to the likelihood that it reflects a durable structural feature of the landscape. Nevertheless, it allows us to match the empirical fact that money funds have, at least in our sample, tended to hold non-zero allocations to both T-bills and RRP, even when the T-bill rate is lower than the RRP rate. Were it not for something that plays the role of $V_i(T_i)$ in the model, money funds would always be at a corner, with all of their assets invested in whichever of the two instruments offered the higher return.

For the purposes of generating simple closed-form solutions, we assume that $V_i(T_i)$ takes the following form:

$$V_i(T_i) = \frac{T_i - \frac{1}{2}T_i^2}{b_i}, \tag{2}$$

and where we capture heterogeneity across funds by assuming that the elasticity parameter b_i is uniformly distributed on the interval $[b_L, b_H]$, with $b_L > 0$. In this formulation, less elastic funds have smaller values of b_i and more elastic funds have larger values of b_i.

Given equations (1) and (2), the first order condition for fund *i* in an interior optimum is as follows:

$$(r_p - r_b) = V_i'(T_i) = \frac{1 - T_i}{b_i}, \tag{3}$$

which yields this expression for fund i's optimal portfolio share in T-bills T_i^* in an interior solution:

$$T_i^* = 1 - b_i(r_p - r_b). \tag{4}$$

In the cases where equation (4) applies, i.e., where fund i is at an interior solution, the fund has a positive T-bill share, $T_i^* > 0$, or equivalently, $b_i < \frac{1}{r_p - r_b}$. The linearity of the demand curves in equation (4) helps to keep things simple, and also serves to isolate our mechanism of interest: since any given fund always has a constant elasticity of demand, the aggregate elasticity of the fund sector is entirely determined by the extensive margin of fund participation, i.e. by which funds remain in the market at any point in time. However, our results are easily generalized to the case where individual money fund demand curves are non-linear.

The solution to the model can now be characterized by its behavior in three regions:

Region 1: Ample T-bill supply. In this region, $r_b \geq r_p$, i.e., the T-bill rate weakly exceeds the RRP rate. As a result, all money funds hold only T-bills, and do not hold any RRP. The T-bill rate is determined by equating the demand from treasurers to the net supply (net of money fund holdings) that is available to them:

$$X(r_b) = S - A. \tag{5}$$

This region applies when T-bill supply S is sufficiently large that it exceeds the sum of the AUM of the money funds, plus the demand of the treasurers at a rate of r_p, that is, when:

$$S \geq X(r_p) + A. \tag{6}$$

Region 2: Moderately scarce T-bill supply. In this region, some money funds hold positive amounts of T-bills, while others may be at a corner solution, holding only RRP. This region applies when $r_p - \frac{1}{b_L} < r_b \leq r_p$, that is, when the T-bill rate is sufficiently high that the least-elastic fund (that with $b_i = b_L$) is still holding a non-zero quantity of T-bills. Denote by $\theta(r_b)$ the fraction of funds that have positive holdings of T-bills at a T-bill rate of r_b. We can distinguish two sub-cases:

Case 2a: When $r_b > r_p - \frac{1}{b_H}$, all funds, including even the most elastic one (that with $b_i = b_H$) still hold some T-bills, so that $\theta(r_b) = 1$. Denoting the average T-bill share across all funds by T^*, we have that:

$$T^* = 1 - \left(\frac{b_L + b_H}{2}\right)(r_p - r_b), \tag{7}$$

and the market clearing condition is given by:

$$X(r_b) = S - AT^*. \quad (8)$$

We can use this market clearing condition to compute the sensitivity of the T-bill rate to changes in T-bill supply:

$$\frac{dr_b}{dS} = \frac{1}{X'(r_b) + A\left(\frac{b_L + b_H}{2}\right)}. \quad (9)$$

Note that the presence of the money funds in the market attenuates this sensitivity, relative to a case where they are absent, i.e., where $A = 0$. Intuitively, this is because as T-bill supply contracts, money funds help to offset the impact of the supply shock on the treasurers, by ceding T-bills to the treasurers and substituting into RRP.

Case 2b: When $r_b < r_p - \frac{1}{b_H}$, the most elastic money funds are completely out of T-bills, so that $\theta(r_b) < 1$. Now let T^{**} denote the average T-bill share across just those funds who remain in the market:

$$T^{**} = 1 - \left(\frac{b_L + b_U}{2}\right)(r_p - r_b), \quad (10)$$

where b_U is the elasticity of the fund that is just at the margin of being in the T-bill market, i.e., $b_U = \frac{1}{r_p - r_b}$.

The market clearing condition is now given by:

$$X(r_b) = S - A\theta(r_b)T^{**}. \quad (11)$$

And the share of funds active in the T-bill market can be represented as:

$$\theta(r_b) = \frac{b_U - b_L}{b_H - b_L}. \quad (12)$$

In the appendix, we show that the impact of changes to supply on the T-bill rate is now:

$$\frac{dr_b}{dS} = \frac{1}{X'(r_b) + A\theta\left(\frac{b_L + b_U}{2}\right)}. \tag{13}$$

We can then use equation (13) to show that, in Region 2b, $\frac{d^2 r_b}{dS db_U} < 0$, or equivalently, that $\frac{d^2 r_b}{dS d\theta} < 0$. That is, the sensitivity of the T-bill rate to changes in T-bill supply is greater when the marginal investor is more inelastic (i.e., when b_U is smaller), or alternatively, when fewer money funds are at an interior and are holding positive positions in T-bills (i.e., when $\theta(r_b)$ is smaller).

This comparative static forms the basis for one of our empirical tests below. The intuition can be seen from equation (13). There are two complementary mechanisms at work. First, as $\theta(r_b)$ falls, this works like a reduction in A, the AUM of the money fund sector: with more funds at a corner, there is less total capacity available to offset shocks to the supply of T-bills by substituting into RRP. Second, given that the most elastic funds leave the market first, those that remain are by construction less elastic. So we are left not only with fewer funds, but with those funds that are least willing to accommodate shocks to T-bill supply.

Region 3: Extremely scarce T-bill supply. In this region, all money funds have fully exited the T-bill market, and hold only RRP. This region applies when $r_b \leq r_p - \left(\frac{1}{b_L}\right)$, i.e., when the T-bill rate is sufficiently far below the RRP rate to have driven even the least elastic money fund out of the T-bill market entirely. In this region, the T-bill rate is given by:

$$X(r_b) = S. \tag{14}$$

And the sensitivity of the T-bill rate to changes in T-bill supply is now simply:

$$\frac{dr_b}{dS} = \frac{1}{X'(r_b)}. \tag{15}$$

Thus $\frac{dr_b}{dS}$ increases monotonically and continuously as we move from Region 2a to Region 2b, and finally to Region 3. Importantly, in Region 3, $\frac{dr_b}{dS}$ is shaped entirely by the slope of the treasurers' collateral-demand function $X'(r_b)$. Again, the intuition is that as the entire money-fund sector reaches a corner and exits the T-bill market altogether, money funds are no longer available to provide any buffering capacity against subsequent shocks to T-bill supply. All that remains to pin down spreads is the elasticity, or lack thereof, of treasurers' collateral demand.

It should be emphasized that by modeling $X(r_b)$ as reflecting collateral demand on the part of treasurers, we are implicitly taking a stand that it is likely to be a relatively structural and hence durable friction. Thus even if we are uncertain as to the primitive determinants of the $V_i(T_i)$ function that shapes money-fund preferences, and hence cannot be confident that money funds will behave in the future as they have in our sample period—e.g., perhaps a form of learning will lead them to respond more elastically to the RRP-bill spread next time around—the model's predictions about how the dynamics of the RRP-bill spread change as we move from Region 2 to Region 3 are more likely to generalize.

III. Empirical Evidence

Data Sources

To study the responsiveness of money funds to the RRP-bill spread, we collect data for a sample that spans from September 30, 2013, to June 30, 2024. Our start date is the first month-end following the inception of the RRP facility on September 23rd, 2013. From Crane Data LLC, we obtain monthly data on the holdings and characteristics of money funds. Our initial sample of money funds includes all Government, Prime, and Treasury funds that have ever made use of the RRP program at any point in time. This primarily excludes specialized tax-advantaged money funds that only invest in municipal securities. There are 140 money funds that appear in our sample at some point in time, and these funds have in aggregate \$3.2 trillion of AUM on average over our sample period. From the New York Federal Reserve, we retrieve data on the details of the RRP program, in particular the identities of those money funds that have access to the facility.

We use data from the Treasury and Federal Reserve to measure changes in the outstanding private supply of Treasuries. From Bloomberg, we obtain secondary market yields for T-bills and overnight index swap rates (OIS). From the Federal Reserve, we get the administered interest rate for the RRP program and the effective Fed Funds rate.

In a separate effort, we also examine the association between derivatives usage by large non-financial firms, and the composition of their cash and securities holdings. We obtain data on cash and securities holdings from Darmouni and Mota (2023), who analyze an annual sample of

193 large firms from 2001 to 2021. For these same firms, we then also collect their annual profits and losses from derivatives transactions from Compustat, following the methodology of Bonaimé et al. (2014). This derivatives data is only available starting in 2004, so we begin this analysis from that point.

Basic Facts on RRP and T-Bill Rates and Money Fund Holdings

We have already discussed Figure 1, which plots the one-month realized and ex-ante expected returns on the RRP facility alongside the rate on one-month T-bills, over the period from June 2021 to June 2024. Because we are comparing one-month T-bills to overnight RRP, and because interest rates were rising during this period, we need to control for the difference in maturities between the two assets. Looking at realized one-month returns on a strategy that invests continuously in the RRP facility is one way to do this. Alternatively, we can proxy for the ex-ante expected return on the RRP over a one-month horizon by adding the one-month OIS rate, and subtracting the current Fed Funds rate, with this adjustment intended to capture expected changes in overnight rates over the coming month. Thus, our maturity adjustment is given by: expected one month RRP return $=$ RRP rate $+$ (one month OIS $-$ fed funds).

As previously noted, T-bill and expected RRP rates are nearly identical from June 2021 through early 2022. After that, one-month T-bill rates fall significantly below expected RRP rates, with the spread exceeding 50 basis points at several points in 2022, before spiking to over 160 basis points during the initial period of uncertainty over the debt ceiling in March and April of 2023. This spike then reverses in May of 2023, with T-bill rates suddenly, albeit briefly, jumping well above the RRP rate, presumably because of temporary debt-ceiling-induced fears that these T-bills would not be repaid on time. After the debt ceiling was raised on June 3, 2023, the spread converges to nearly zero in conjunction with a large increase in the supply of T-bills.

Appendix Figure A1 further compares the RRP rate to two private-market repo rates, that on tri-party repo and that on general collateral financing (or GCF) repo. As the RRP rate is an administered rate, it is of course possible for it to fall below these private-market rates when it is set so low that there is no take-up of the RRP facility. However, when RRP take-up is significantly positive, the tri-party rate, which captures the rate at which money funds can lend to dealer banks, is typically very close to the RRP rate. For example, over the period June 2021 to June 2024 when RRP take-up was consistently well above zero, the spread between the RRP rate and the tri-party rate was on average 0 basis points, implying a high degree of substitutability between these two instruments.

The money funds in our sample hold a variety of assets, including T-bills and other government or agency debt, RRP, private repo, commercial paper and certificates of deposit. Figure 2 shows the aggregate portfolio weights over time for these assets. The combined money fund portfolio weight in all forms of repo and T-bills averages 51 percent over the full sample and 69 percent from June 2021 to June 2024. Thus, while our model directly captures an important fraction of money fund holdings, it also abstracts from a large chunk of their portfolios.

In the early portion of the sample from September 2013 to December 2017, money fund investments in RRP exhibit sharp quarter-end spikes. As noted by Anderson et al (2020), these spikes are driven by window dressing on the part of foreign banks. These banks intentionally shrink their balance sheets at quarter ends to reduce their regulatory capital requirements. This leads them to borrow less in the private repo market as well as the commercial paper (CP) and certificate of deposit (CD) markets. Faced with a reduced supply of foreign bank paper, money funds then substitute into RRP. Also noteworthy in Figure 2 is a sharp reduction in money fund holdings of CP and CD around October 2016, which coincides with the implementation of new restrictions on the riskier holdings of institutional stable-NAV money funds.

Table 1 shows some summary statistics for RRP and T-bill rates, as well as for our sample of money funds. Note that money funds on average hold assets with 33 days to maturity. This approximately matches the maturity of the one-month RRP-bill spread, which we focus our analysis on.

Figure 3 plots the RRP-bill spread over a longer horizon from the inception of the RRP facility in September 2013 to June 2024; this spread is the empirical counterpart to $(r_p - r_b)$ in the model. Additionally, Figure 3 shows the ratio of money fund holdings in T-bills to the sum of their holdings in T-bills and RRP. We plot the three-month moving average of this ratio to smooth out the quarter-end spikes in RRP that were evident in Figure 2. This ratio corresponds to the variable $T^{\star}$ in the model, i.e., the average portfolio weight in T-bills.

By putting these two lines together, Figure 3 serves to illustrate Regions 1 and 2 of the model. Recall that Region 1 is defined by ample T-bill supply, where both: (i) T-bills yield weakly more than RRP (a negative RRP-bill spread); and (ii) money funds only hold T-bills. Empirically, this corresponds to the shaded area in the figure, from March 2018 to March 2021. During this period, the average RRP-bill spread is -11 basis points and the average T-bill portfolio ratio is 96 percent. By contrast, outside of Region 1, money funds hold a significant amount of RRP. Across these non-Region-1 periods, the average RRP-bill spread is 6 basis points and the average T-bill portfolio ratio is 43 percent. Within these intervals, we have a subperiod when T-bills are especially scarce and the market approaches Region 3: this is from April 2022 to April 2023, as indicated by the hatched area in the figure. During this subperiod, the average RRP-bill spread is 37 basis points and the average T-bill portfolio ratio is down to 15 percent.

Segmentation Between Money Funds and Outside Investors

The first key friction in our model is the assumption that money funds have an exogenously fixed AUM, independent of the RRP-bill spread. Or said differently, we are assuming that flows from outside investors, i.e. the treasurers in our model, do not respond to this spread. To assess the extent to which this assumption roughly captures reality, we can do a simple decomposition. Note that when the aggregate money fund sector decreases its holdings of T-bills and increases its holdings of RRP, this may either be due to money fund managers changing their portfolio weights, or to investors reallocating capital away from funds with a high T-bill portfolio weight and toward funds with a high RRP portfolio weight. It is a straightforward exercise to estimate the relative importance of these two effects.

Denote the AUM of each fund i at time t by $A_{i,t}$ and its portfolio weight in T-bills as $w_{Bill,i,t}$. Define the change in the dollars invested in T-bills by fund i from $t-1$ to t as:

$$\Delta D_{Bill,i,t} = w_{Bill,i,t} A_{i,t} - w_{Bill,i,t-1} A_{i,t-1} \tag{16}$$

We can decompose this dollar change into a component due to investor flows and a component due to managerial rebalancing. Since money funds maintain a stable net asset value of $1 per share, investor flows are equivalent to the change in fund AUM. Therefore, investor flows into T-bills can be written as:

$$\text{IFlow}_{Bill,i,t} = w_{Bill,i,t-1}(A_{i,t} - A_{i,t-1}) \tag{17}$$

Thus the managerial rebalancing component is given by:

$$\text{MFlow}_{Bill,i,t} = \Delta D_{Bill,i,t} - \text{IFlow}_{Bill,i,t} = (w_{Bill,i,t} - w_{Bill,i,t-1})A_{i,t} \tag{18}$$

We can also do an analogous decomposition for the change in total dollars invested in RRP for any fund i.

Figure 4 constructs a counterfactual T-bill portfolio share (i.e., the ratio of T-bills to T-bills plus RRP) that would have prevailed in a hypothetical world where fund managers kept their portfolio weights constant, and changes in the share were driven only by flows from external investors. To do so, we define the counterfactual dollars invested in T-bills at time t by all money funds as:

$$\widetilde{D}_{Bill,t} = D_{Bill,Sept\ 2013} + \sum_{\tau=1}^{t} \text{IFlow}_{Bill,\tau}, \tag{19}$$

where $\tau = 0$ for September 2013, the beginning of the sample. We define the counterfactual dollars invested in RRP at time t by all money funds analogously. The counterfactual T-bill portfolio share is then given by:

$$\widetilde{\text{Ratio}}_t = \frac{\widetilde{D}_{Bill,t}}{\widetilde{D}_{Bill,t} + \widetilde{D}_{RRP,t}} \tag{20}$$

This counterfactual share is plotted in black in Figure 4, alongside the actual share in blue, which is the same series shown in Figure 3. As can be seen, the counterfactual share is an order of magnitude less variable than the actual share, and the two series are nearly uncorrelated. Thus flows from outside the money fund sector explain very little of the overall variation in the aggregate T-bill portfolio share. Notably, the significant substitution from T-bills to RRP that happens from June 2021 to February 2023 occurs nearly entirely through portfolio rebalancing by money fund managers, with external flows playing almost no role. This suggests that our modeling assumption—namely, that the AUM of individual money funds is exogenously fixed and unrelated to the RRP-bill spread—is a good approximation to reality.

The pattern seen in Figure 4 is a specific manifestation of a more general phenomenon, which is that investor flows into money funds have a relatively modest sensitivity to money fund returns. We can illustrate this pattern for our September 2013 to June 2024 sample period by running regressions of monthly or quarterly inflows into a given fund i (scaled by its assets) on the contemporaneous returns of the fund. These regressions are shown in Table 2. In columns 1 and 2, the data is monthly and equal weighted. In column 3, the regression weights observations by fund AUM. Finally, in column 4, the data is quarterly and the regression again weights observations by fund AUM. The upper-end estimate in column 4 implies that a 100 basis point increase in returns is associated with a 5.9 percentage-point increase in fund AUM over a quarter.

This degree of flow sensitivity would seem to be too small to exert much of an arbitrage discipline on the kinds of RRP-bill spreads seen in the data. For example, consider the period of greatest T-bill scarcity in our sample, April 2022 to April 2023, when the RRP-bill spread averaged 37 basis points. The point estimate of 5.9 seen in column 4 of Table 2 implies that even in response to a spread of this relatively extreme magnitude, a fund entirely invested in RRP would receive quarterly inflows of only 2.2 percentage points more than a fund entirely invested in T-bills.[11]

Corporate Treasurer Collateral Demand

The second friction in our model is that T-bills provide collateral services to corporate treasurers, while private repo and money fund shares do not. Although it is difficult to validate this assumption decisively, we can provide some suggestive evidence by showing that those firms that are more active in derivative markets—and who thus might be expected to have a greater need for pledgeable collateral—hold a larger share of their liquid cash and securities in Treasuries.

In Table 3, we estimate variations of the following basic regression:

$$w_{UST,i,t} = \alpha_t + \beta Deriv_{i,t-1} + \epsilon_{i,t} \tag{21}$$

here $w_{UST,i,t}$ is firm i's total holdings of Treasury and Agency securities, divided by the sum of their Treasury and Agency holdings and cash-like assets, which include bank deposits and money fund shares. $Deriv_{i,t-1}$ is the firm's one-year lagged *exposure* to derivatives, which we (imperfectly) proxy for by computing the absolute value of realized profit and loss (P&L) on derivatives transactions divided by the firm's total assets.[12] In our sample of 193 large corporates from 2004 to 2021, 69 percent of firm-year observations have non-zero derivatives exposure, 40 percent have an average derivatives exposure of at least 10 basis points of total assets, and 28

[11] This 2.2 percentage point effect is the point estimate of 5.9 in column 4 of table 2 multiplied by the average RRP-bill spread of 37 basis points ($2.2 = 5.9 \times 0.37$).

[12] We winsorize this variable at the 1-percent level to mitigate the effect of outlier realizations.

percent hold a non-zero amount of Treasuries. For those firms with derivatives exposure, the standard deviation of derivatives P&L is $426 million, which is about one third of their average Treasury holdings of $1.5 billion. This suggests that collateral requirements on derivatives positions could be meaningful relative to typical Treasury holdings.

The first column of Table 3 runs the regression in equation (21) with just controls for the log of firm assets and time fixed effects. The second and third columns add industry fixed effects and industry-by-year fixed effects respectively, with little impact on the coefficient of interest. The point estimate of 3.481 in the third column implies that the Treasury share variable is on average 1.8 percentage points higher when a firm's derivative exposure is one standard deviation (0.516 percent) larger; this can be compared to an average value of the Treasury share of 7.6 percent. Alternatively, in dollars, the point estimate implies that for every dollar of additional derivatives exposure, firms on average hold an additional 0.32 dollars of Treasuries.

There are several potential concerns with these regressions. The first is a form of measurement error in our Treasury-share variable. Some corporates do not report any Treasury holdings, but in these cases—and only these cases—it is possible that the data includes in the cash and cash-equivalents category short-term T-bills with a maturity of less than three months. This would lead us to understate the true Treasury share in these cases. To ensure that this measurement problem is not driving our results, in the fourth column of Table 3, we restrict the sample to firms that report positive Treasury holdings, where the measurement problem is not present. We find that the association between Treasury holdings and derivatives exposure is nearly twice as large in this sub-sample, which suggests that, if anything, our full-sample results may be too conservative.

Another limitation of the data is that our Treasury share variable is defined as *all* Treasury holdings divided by Treasury holdings plus cash-like instruments; we are unfortunately unable to break out Treasury holdings by maturity. While Treasuries of any maturity are plausibly good collateral for derivative transactions, this variable might also be picking up a preference on the part of corporate treasurers for longer-maturity securities of any kind, as opposed to just a preference for collateral. To address this possibility, we add to the regression an analogous one-year lagged corporate bond share, defined as corporate bond holdings divided by corporate bond holdings plus cash-like instruments; this corporate bond share might be expected to also capture a preference for longer-maturity securities, but without picking up collateral demand. As can be seen in the fifth column of Table 3, the inclusion of this variable, while itself significant, does little to dampen the effect of derivatives exposure on the Treasury share.

Alternatively, we can put the corporate bond share on the left-hand side of a version of equation (21). We do this in Appendix Table A1, which has a structure exactly analogous to Table 3. When controlling for the Treasury share, we find that the corporate bond share is wholly unrelated to derivatives exposure; the coefficient on the latter is very small and statistically insignificant. Thus there appears to be something fundamentally different about the relationship between Treasury holdings and derivatives usage that is not shared by corporate-bond holdings. This suggests that our findings for Treasury holdings are not simply driven by something having to do with a preference for longer-maturity instruments.

A final potential concern is that there could be a mechanical relationship between our proxy for derivatives usage and the Treasury share. This could arise, for example, if the profits and losses from derivatives transactions flowed directly in and out of say bank deposits, which would in turn affect our measure of the Treasury share. Note however, that this mechanical effect would presumably be stronger if we used the realized *linear* P&L on derivatives, rather than the *absolute value* of the P&L. And comfortingly, we show in Appendix Table A2 that the former linear variable has no discernible correlation with the Treasury share.

Money Fund Elasticity of Substitution

The third friction in our model is that money funds in the aggregate substitute imperfectly elastically between T-bills and RRP when the return on the former falls below that on the latter. Moreover, this aggregate elasticity varies over time in a predictable way because the model assumes that individual funds are heterogeneous in their elasticity of substitution. We now discuss the evidence for each of these premises in turn.

Aggregate elasticity of substitution

In Table 4, we estimate aggregate regressions of the form:

$$\Delta w_{Bills,t} = \alpha - \beta \Delta Spread_t + \epsilon_t \tag{22}$$

where $\Delta w_{Bills,t}$ is the log difference in the aggregate portfolio weight of money funds in T-bills over month t, $\Delta Spread_t$ is the change in the spread over the month in basis points, and β is the estimated elasticity.[13] We estimate this over September 2013 through March 2021: this period spans Region 1 as well as a portion of Region 2 where T-bills are not scarce.

Note that if one estimates equation (22) by OLS, there is an obvious endogeneity problem: The RRP-bill spread may widen either due to supply or demand shocks for T-bills. In particular, our estimate β is likely to be biased if there are demand shocks for T-bills coming from money funds, which would both increase the spread and the funds' portfolio weight in T-bills. In an effort to address this endogeneity concern, we instrument for the RRP-bill spread using changes in the private supply of Treasury securities.[14]

The Treasury states that it auctions "bills on a regular and predictable basis" and does not attempt to tactically time the market (Garbade, 2007). To the extent that this is the case, it mitigates the concern that the supply of T-bills increases precisely when money fund demand for T-bills is high. However, we further address the possibility of strategic maturity issuance by using changes in the total private supply of all Treasury securities as our instrument, rather than the supply of just T-bills. Therefore, the identifying assumption is that monthly changes in Treasury supply are not influenced by the RRP-bill spread, which seems highly plausible.[15]

In the first column of Table 4, we estimate the first-stage regression for this instrumenting procedure, namely the effect of percent changes in Treasury supply on the RRP-bill spread (in units of basis points):

$$\Delta Spread_t = \alpha + \beta \Delta \mathrm{Treasury_t} + \epsilon_t \tag{23}$$

where $\Delta \mathrm{Treasury_t}$ is the log difference in the private supply of Treasuries. The point estimate in this column implies that for a 1 percent increase in the supply of Treasuries, the RRP-bill spread decreases by 2.2 basis points over our subsample from September 2013 through March 2021 when T-bills are not scarce. The RRP-bill spread widens because the market T-bill rate falls in response to an increase in the supply of Treasuries, not because the administered RRP rate changes. We

[13] We observe portfolio weights at month-end. For the RRP-bill spread, we do not compute the change literally from one month-end trading day to the next month-end trading day, because the Fed Funds rate has a month-end negative spike of a few basis points. Instead, in an effort to reduce noise, we measure the level of the spread as the average of the five trading days prior to each month end, and use this level to compute monthly changes.

[14] The private supply of Treasuries is the total notional of outstanding Treasuries less what is held by the Federal Reserve.

[15] Appendix Figure A2 shows the monthly log differences in the private supply of Treasuries over our sample, and illustrates how the variation in this series is primarily driven by high frequency monthly changes and not low-frequency trends in Treasury supply. Indeed, these monthly changes in Treasury supply have nearly zero autocorrelation.

show in later analyses that the effect of changes in T-bill supply on the spread varies over time, depending on the overall elasticity of demand of money funds for T-bills.

The second column of Table 4 estimates equation (23) by OLS, where again, we expect the coefficient to be biased toward zero due to the endogeneity of the RRP-bill spread. Here the coefficient estimate implies that a 1 basis point increase in the RRP-bill spread is associated with money funds decreasing their T-bill portfolio weight by 0.94 percent. To be clear, this is not 0.94 percentage points, but rather 0.94 percent of the aggregate portfolio weight of money funds in T-bills. Money funds on average held 12.2 percent of their portfolio in T-bills from September 2013 to March 2021. Thus, in response to a 1 basis point increase in the RRP-bill spread, the regression implies that money funds would decrease their portfolio weight in T-bills by 11 bps ($-11 = 12.2 \times -0.94$).

The third column of Table 4 again estimates equation (23), but this time by instrumental variables, using log differences in Treasury supply as our instrument for changes in the spread. As expected, this increases the estimated elasticity of substitution substantially—the coefficient goes up by almost an order of magnitude. The coefficient now implies that in the face of a 1 basis point increase in the RRP-bill spread due to a decrease in Treasury supply, money funds decrease their T-bill holdings sharply, by 6.2 percent, which amounts to 79 basis points of portfolio weight. The fourth column of the table repeats this IV exercise at the quarterly frequency, with a broadly similar result. Thus, it appears that while the aggregate money fund sector does not respond perfectly elastically to changes in the RRP-bill spread, it is nevertheless quite responsive, and as such is clearly the first line of defense in policing movements in the spread.

One might wonder whether the imperfect elasticity of substitution that we document reflects nothing more than a reluctance on the part of money funds to sell off their existing holdings of T-bills, perhaps in an effort to avoid having to recognize capital losses. It turns out that this is not the case: during the period when the RRP-bill spread was widening significantly, many funds were not just letting their existing holdings of T-bills passively roll off, they were actually buying large quantities of newly issued T-bills.

Figure 5 shows monthly new purchases of T-bills by money funds from June 2021 through June 2024, which spans when the RRP-bill spread widened. New T-bill purchases are defined for each money fund at the cusip level, as the sum of increases in existing cusip-level positions and the initiation of new cusip-level positions. On average from April 2022 to April 2023, money funds made 102 billion dollars of new purchases of T-bills each month. Even when the RRP-spread was 45 basis points in June of 2022, money funds made 121 billion dollars of new purchases of T-bills.

New purchases of T-bills gradually decreased to about 60 billion dollars by April 2023.

Another way to make a similar point is to consider a counterfactual scenario where money funds stopped purchasing any new T-bills after December 2021 and let their existing portfolio mature. In the counterfactual scenario, after three months, 37 percent of their T-bill holdings would have rolled off, leaving the money funds with 248 billion dollars of T-bills, rather than their actual 681 billion dollars of T-bills. After six months, the counterfactual has 94 percent of T-bills rolling off, which compares to just a 33 percent decrease in actual T-bill holdings. Thus, the slow substitution by money funds out of T-bills was not simply due to an unwillingness to sell their existing holdings of T-bills. Strikingly, despite large RRP-bill spreads, money funds continued to actively buy new T-bills throughout this period.

Heterogeneity across money funds in elasticity of substitution

We now turn our attention to heterogeneities across money funds in their responsiveness to movements in the RRP-bill spread. Using exactly the same methodology (column 3 of Table 4) as we did for the aggregate fund sector, we can estimate elasticities of substitution for each individual money fund, again over the period September 2013 through March 2021. We winsorize these elasticity estimates at the 1 percent level and require that we have at least 1 year of data to estimate each fund-level elasticity. Figure 6 plots the kernel density of the estimated elasticities for our sample of funds. The median fund has an elasticity of 7.4 and the interquartile range of elasticities goes from 4.4 to 10.1.

In un-tabulated cuts of the data, we find that there are significant differences in estimated elasticities across different types of funds. Funds categorized as Treasury funds are the most inelastic with an elasticity of 5.4 and are significantly less elastic than Government funds and Prime funds with average elasticities of 7.2 and 8.6, respectively. This hints at the idea that funds with narrower mandates behave more inelastically. Treasury funds have the most restrictive mandate, and they can invest only in T-bills or in reverse repo with Treasury collateral. Government funds, by contrast, can also invest in agency debt. Prime funds have the broadest investment universe, which includes commercial paper and certificates of deposit.[16]

Of course, an obvious concern is that the variation in our fund-level elasticity estimates is driven largely by estimation error and contains little true economic signal. So, our next step is to

[16] In a separate cut, we find that bank-affiliated money funds are significantly more elastic than non-affiliated funds, with elasticities of 8.7 and 6.6, respectively.

ask whether these fund-level estimates have any out-of-sample explanatory power that aligns with our model's predictions. Bear in mind that these estimates come from the period September 2013 to March 2021. We now check to see whether they can predict subsequent behavior.

As emphasized in the model, there are two distinct effects that lead the market to become increasingly un-resilient to T-bill supply shocks as the RRP-bill spread widens. The first is that the most elastic funds become constrained and no longer have any T-bills to sell. The second is that those funds that remain away from the corner are now less elastic, and so respond less aggressively to any further widenings of the spread. Figure 7 depicts both of these mechanisms, plotting: (i) the fraction of constrained money funds; and (ii) the elasticity of the remaining unconstrained money funds, both over the period of January 2022 to June 2024.

Figure 7 shows that by the end of April of 2023, when T-bill scarcity was at its peak level, and the RRP-bill spread averaged a stunning 130 basis points over the last five days of the month, more than 90 percent of funds by AUM weight were constrained, according to our definition. Thus, the entire sector was by this point close to a corner. Or said in the language of the model, we were very close to Region 3, where money funds lose all ability to offset T-bill supply shocks.

Moreover, as the more elastic funds became constrained and left the market, the average elasticity of the remaining marginal players in T-bills fell significantly, from an elasticity of 6.3 in January 2022 to an elasticity of 1.7 by April 2023.

This situation then rapidly reversed itself in the wake of Congress raising the debt ceiling, which led to a large increase in T-bill issuance, shown in Figure 8: from June 3, 2023 to June 30, 2024, T-bill supply increased by $1.8 trillion dollars, the RRP facility experienced outflows of $1.5 trillion dollars, and the RRP-bill spread became mildly negative, on average -6 basis points. Consistent with these developments, the aggregate fraction of constrained money funds fell back down to less than 10 percent by October 2023. Furthermore, with a substantial fraction of money funds re-establishing long positions in T-bills, the average elasticity of unconstrained funds increased to 6.7 by October 2023.

The Effect of T-bill Supply Shocks on T-Bill Rates

To summarize what we have seen thus far: in response to increasing T-bill scarcity and a widening RRP-bill spread over the period January 2022 to April 2023, a growing fraction of money funds became constrained, with few or no T-bills left to sell, and the remaining unconstrained funds were those with a lower elasticity of substitution. In the context of the model, these two effects combine to drive our prediction that as the number of unconstrained funds declines, shocks to the supply of T-bills should have a more pronounced impact on T-bill rates.

Table 5 displays our tests of this hypothesis. Using a sample period that runs from September 2013 to June 2024, we estimate the following regression in monthly data.[17]

$$\Delta Spread_t = \alpha + \beta_0 \Delta \text{TBill Supply}_t + \beta_1 \text{Scarce TBills}_t + \gamma \Delta \text{TBill Supply}_t \times \text{Scarce TBills}_t + \epsilon_t \quad (24)$$

The variable Scarce TBills_t is a dummy that takes on a value of one from April 2022 to April 2023, when money fund holdings are low and the RRP-Bill spread is high, and zero otherwise.

Due to the endogeneity concerns with respect to the supply of T-bills that we discussed in the estimation of equation (23), we instrument for log differences in T-bill supply using log differences in Treasury supply. The first column of Table 5 shows that a 1 percent increase in Treasury supply is associated with a 5.2 percent increase in T-bill supply. Using this instrumented variation in T-bill supply, we find that, unconditionally over the full sample, a one percent decrease in the supply of T-bills leads a 0.66 basis points increase in the RRP-bill spread (column 2 of Table 5). In the third column of Table 5, we show the interactive specification in equation (24). The results are stark: when T-bills are scarce and a majority of money funds are already at a corner, a further 1 percent decrease in instrumented T-bill supply causes the RRP-bill spread to increase by 4.8 basis points, which is 7 times as large as the unconditional impact.

Figure 9 illustrates this result graphically, with a scatter plot of the changes in the RRP-bill spread and in instrumented changes in T-bill supply, where the black dots represent times when T-bills are not scarce, while the blue dots represent the April 2022 to April 2023 period when they are scarce. The steeper slope of the blue regression line in the figure is the visual analog to the interaction coefficient in column 3 of Table 5.

Finally, in the fourth column of Table 5, we consider the following alternative specification:

$$\Delta Spread_t = \alpha + \beta_0 \Delta \text{TBill Supply}_t + \beta_1 \text{Constrained Share}_t + \gamma \Delta \text{TBill Supply}_t \times \text{Constrained Share}_t + \epsilon_t, \quad (25)$$

[17] We exclude the months of May 2023 and June 2023 due to effects of the debt ceiling, which we discuss toward the end of this subsection.

where $\text{Constrained Share}_t$ is the AUM-weighted share of money funds in period t that have less than 5 percent of their portfolio in T-bills and a portfolio share of at least 10 percent in either RRP or T-bills. This specification, which is a literal take on a comparative static from the model, allows for the impact of supply shocks to vary continuously depending on the fraction of constrained money funds. For example, although both May of 2022 and April of 2023 lie in what we have deemed to be the scarce T-bill region, the AUM share of constrained funds is 24 percent in May of 2022, but much higher, at 91 percent in April of 2023 (as shown in Figure 7). Our model therefore suggests that the impact of a supply shock should be even greater in April of 2023 than in May of 2022, and this is what this specification allows for.

The estimates in column 4 imply that when 24 percent of money funds are constrained, as was the case in May of 2022, a 1 percent decrease in instrumented T-bill supply causes the RRP-bill spread to increase by 1.5 basis points. By contrast, when 91 percent of funds are constrained, as was the case in April of 2023, a 1 percent decrease in T-bill supply causes the RRP-bill spread to increase by 6.0 basis points. Thus, even within this already-stressed period, a further increase in the fraction of funds moving to the corner greatly magnifies the rate impact of supply shocks.

When the money fund sector becomes entirely constrained, as in Region 3 of the model, there is no intermediary left that can actively substitute between T-bills and RRP. Therefore, the model suggests that further shocks to the T-bill market can have a very large effect on spreads. This became apparent during the debt ceiling crisis of 2023. In January 2023, the Treasury reached its statutory debt limit of $31.381 trillion and warned that it would exhaust its funding by early June. After Congress failed to raise the debt limit, on May 1, 2023, the Treasury stated that it expected that it would be unable to meet its obligations beginning on June 1, 2023.[18] This caused investors seeking safety to shift outward their demand for T-bills maturing prior to June 1st. At the same time, we have seen that more than 90 percent of money funds by AUM were constrained, and therefore unable to help accommodate these demands.

Figure 10 shows the average yield on T-bills maturing from May to August of 2023. The average is taken over the interval from May 2, after the June deadline was announced, to May 30th, just before Congress passed a bill raising the debt ceiling. The discontinuity around the June 1st maturity date is sharp: the T-bill maturing on May 30 yields an average of 153 basis points

[18] This projection was later revised to June 6th on May 26th, 2023.

less over the course of the month than the T-bill maturing on June 1. The average RRP-bill spread for T-bills maturing prior to June 1 is 81 basis points. At the same time, T-bills maturing in June on average have a *positive* RRP-bill spread, reflecting the risk of being caught in a technical default.

It is unfortunately impossible to do the counterfactual exercise, and ask how the debt-ceiling drama would have played out in the T-bill market had money funds not already been so constrained going into it. Nevertheless, the unprecedented pricing anomalies observed during this episode are certainly consistent with the spirit of Region 3 of our model, according to which supply and demand shocks can have very large impacts on T-bill yields when the money fund sector as a whole is at a corner, and can no longer accommodate these shocks.

IV. Conclusions

We have seen that even in an extremely simple and transparent setting, frictions in financial intermediation can give rise to economically significant and time-varying spreads between similar money-like assets. Moreover, our results suggest that heterogeneities in the behavior of the relevant intermediaries—in our case, money funds and the non-money-fund sector that we have labeled corporate treasurers—are crucial to understanding market dynamics. In particular, when T-bills became relatively scarce in 2022, the more elastic money funds were driven to a corner, selling all of their T-bills and substituting into RRP. This left only the less elastic money funds in the market, and dramatically reduced the ability of the money fund sector as a whole to respond to further shocks to T-bill supply. In the limit, as virtually all money funds exited the T-bill market, the behavior of T-bill spreads was shaped almost entirely by the collateral demand of corporate treasurers. As a result, during this period of heightened T-bill scarcity, shifts in supply had an order of magnitude larger effect on T-bill yields than in prior periods.

Our setting is an especially convenient one for studying how intermediary heterogeneity can create time-varying market dynamics. This is because we know exactly the identity of one important category of intermediary—namely, the money funds—and because we can cleanly observe when each one individually, as well as the sector as a whole, is at a corner.

While other settings may prove more empirically challenging in this regard, we believe that the insights from our study can nevertheless provide helpful qualitative guidance. As one example, consider the large body of work that seeks to measure the impact of quantitative easing (QE) programs on long-term bond-market yields.[19] Much of this work implicitly assumes that the market impact of a given-sized intervention will be constant over the period of study in question. However, our results suggest that this market impact is likely to be state-dependent in important ways. As a result, it may be possible to make more accurate judgments about the effect of future interventions by thinking carefully about who the marginal players are in the market at any point in time, and what their behavior has looked like in the past. Indeed, Breckenfelder and De Falco (2023) highlight this point when analyzing the effects of the ECB's large-scale asset purchase program. It seems likely that there are a variety of other settings where a similar approach may prove fruitful.

[19] Bernanke (2020) provides a survey of much of this literature.

References

Adrian, T., E. Etula, and T. Muir (2014). Financial intermediaries and the cross-section of asset returns. *Journal of Finance* 69 (6), 2557–2596.

Afonso, G., M. Cipriani, and G. La Spada (2022). Banks' balance-sheet costs, monetary policy, and the ON RRP. *FRB of New York Staff Report* 1041.

Bech, M. L. and E. Klee (2011). The mechanics of a graceful exit: Interest on reserves and segmentation in the federal funds market. *Journal of Monetary Economics 58* (5), 415–431.

Bernanke, B. (2020). The new tools of monetary policy. *American Economic Review* 110 (4), 943-983.

Bonaimé, A., Ö. Öztekin, and R. S. Warr (2014). Capital structure, equity mispricing, and stock repurchases. *Journal of Corporate Finance* 26: 182-200.

Breckenfelder, J. and V. De Falco (2023). Large scale asset purchases and investors' heterogeneity: direct and spillover effects. *Working Paper.*

Brunnermeier, M. K. and Y. Sannikov (2014). A macroeconomic model with a financial sector. *American Economic Review 104* (2), 379–421.

d'Avernas, A., and Vandeweyer, Q. (2023). Treasury bill shortages and the pricing of short-term assets. *Journal of Finance*, forthcoming.

Darmouni, O., and L. Mota (2023). The savings of corporate giants. *Working Paper*.

Doerr, S., E. Eren, and S. Malamud (2023). Money market funds and the pricing of near-money assets. *Swiss Finance Institute Research Paper* 23-04.

Du, W., Hébert, B., and Huber, A. W. (2023). Are intermediary constraints priced? *Review of Financial Studies*, 36(4), 1464-1507.

Du, W., A. Tepper, and A. Verdelhan (2018). Deviations from covered interest rate parity. *Journal of Finance 73* (3), 915–957.

Duffie, D. (2010). Presidential address: Asset price dynamics with slow-moving capital. *Journal of Finance 65* (4), 1237–1267.

Duffie, D., M. J. Fleming, F. M. Keane, C. Nelson, O. Shachar, and P. Van Tassel (2023). Dealer capacity and U.S. Treasury market functionality. *FRB of New York Staff Report* 1070.

Garbade, K. (2007). The emergence of "regular and predictable" as a Treasury debt management strategy. *Economic Policy Review 13* (1).

Gromb, D. and D. Vayanos (2018). The dynamics of financially constrained arbitrage. *Journal of Finance 73* (4), 1713–1750.

Haddad, V., P. Huebner, and E. Loualiche (2022). How competitive is the stock market? Theory, evidence from portfolios, and implications for the rise of passive investing. *Working Paper*.

Haddad, V. and T. Muir (2021). Do intermediaries matter for aggregate asset prices? *Journal of Finance 76* (6), 2719–2761.

He, Z., B. Kelly, and A. Manela (2017). Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics 126* (1), 1–35.

He, Z. and A. Krishnamurthy (2013). Intermediary asset pricing. *American Economic Review 103* (2), 732–770.

He, Z. and A. Krishnamurthy (2018). Intermediary asset pricing and the financial crisis. *Annual Review of Financial Economics 10*, 173–197.

Ivashina, V., D. S. Scharfstein, and J. C. Stein (2015). Dollar funding and the lending behavior of global banks. *Quarterly Journal of Economics 130* (3), 1241–1281.

Siriwardane, E., A. Sunderam, and J. Wallen (2022). Segmented arbitrage. *Working Paper*.

Figure 1
Realized and Expected RRP Rates vs. T-Bill Rates, June 2021 to June 2024

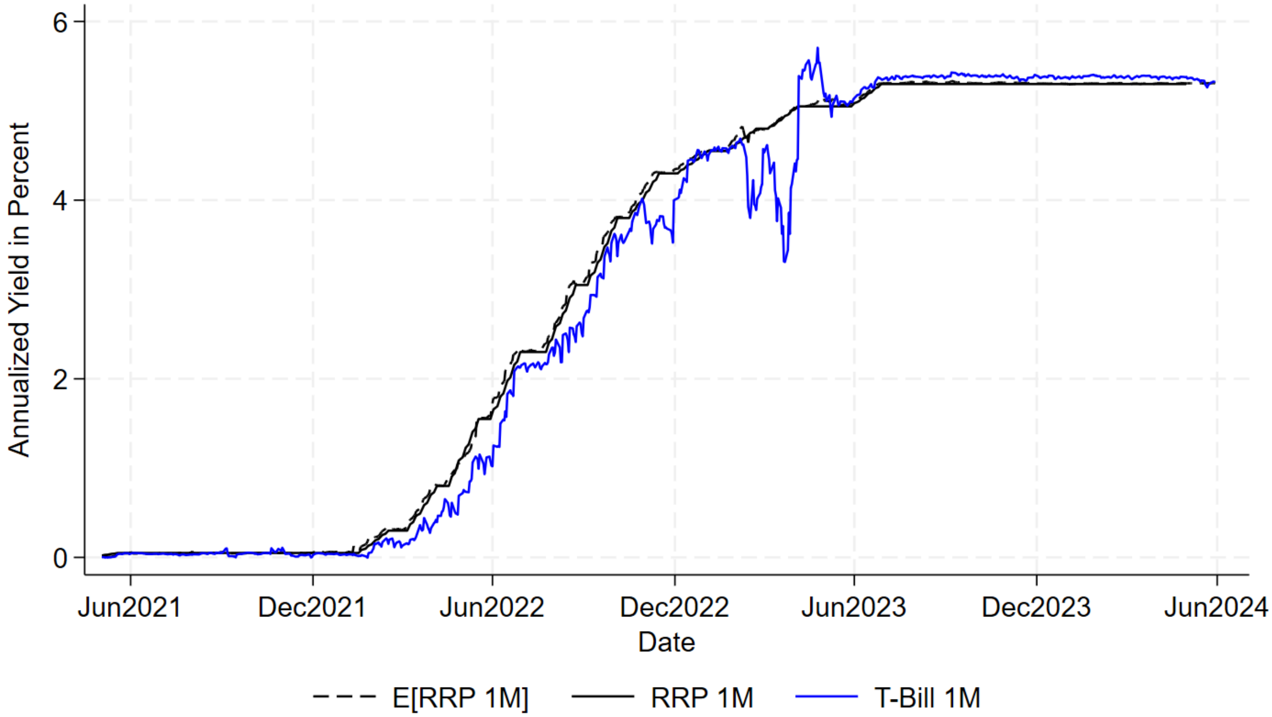

Figure 1 shows the one-month expected and realized returns on the Fed's RRP facility, as well as the one-month T-bill rate. The expected one-month RRP rate (dashed black line) is measured as the overnight RRP rate, plus the one-month OIS rate, minus the overnight Fed Funds rate. The realized 1-month RRP rate (solid black line) is the annualized rate earned from rolling over an investment in the RRP for a month. The plot spans from June 1st, 2021, to June 30, 2024.

Figure 2
Money Fund Portfolio Composition
September 2013 to June 2024

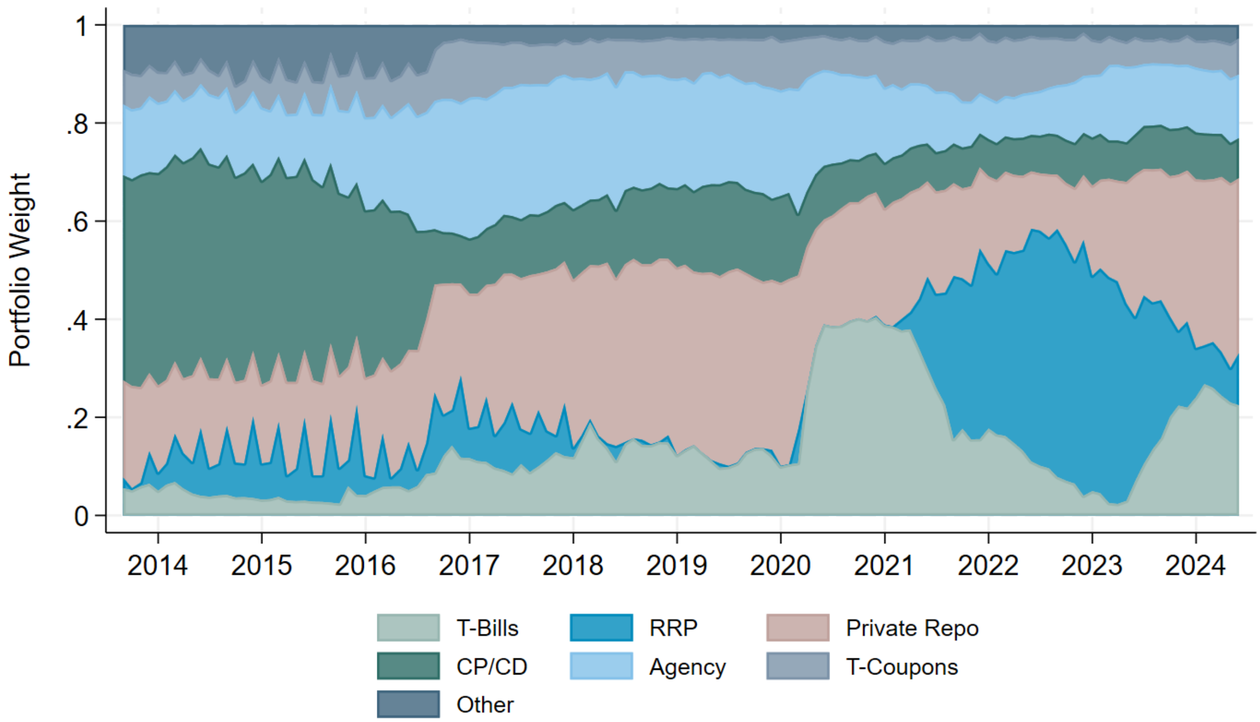

Figure 2 shows the portfolio weights by asset class for the sum of all money funds in our sample. The asset classes include Treasury Bills, RRP, reverse repo with the private sector, commercial paper and certificates of deposits, agency bonds, Treasury coupons, and other. The sample spans from the beginning of the RRP program, Sept 2013, to June 2024.

Figure 3
RRP-Bill Spread and Portfolio Shares
September 2013 to June 2024

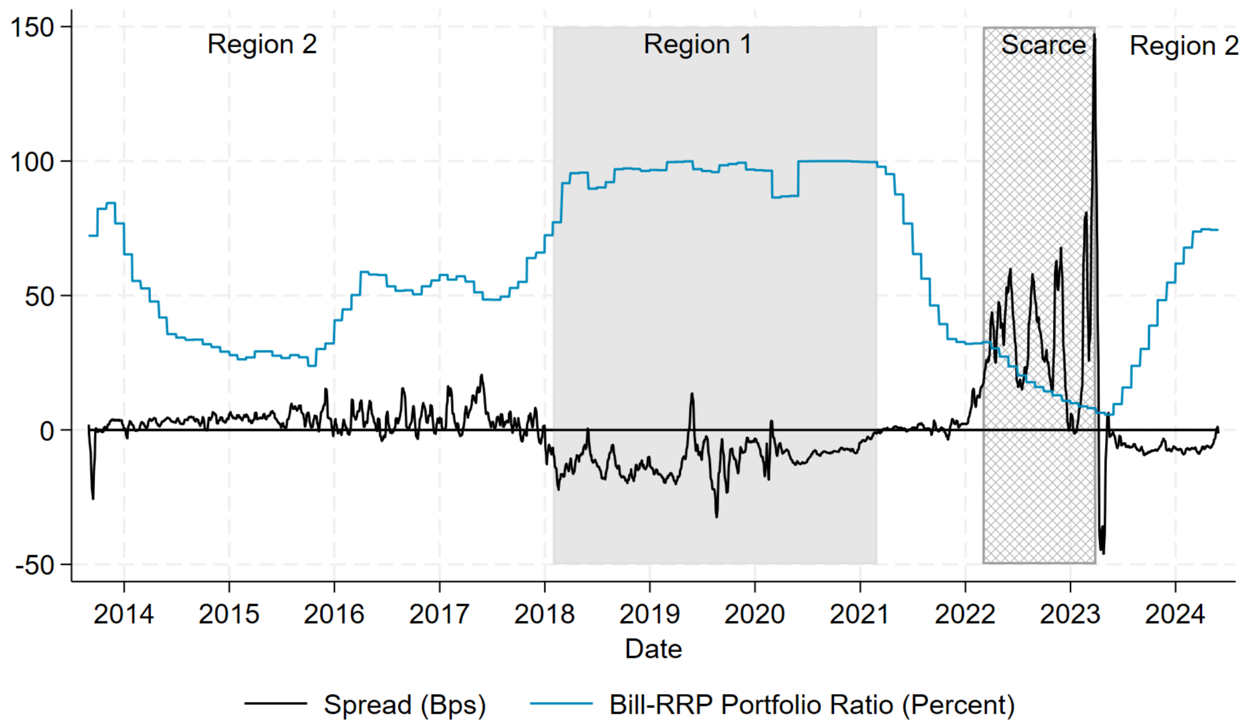

Figure 3 highlights the data corresponding to Regions 1 and 2 of the model. Spread is the one-month RRP-bill spread in units of basis points, which we show as a 1-week moving average of daily data. T-bill portfolio share is the portfolio weight in T-bills divided by the sum of the portfolio weight in T-bills and RRP for the aggregate sample of money funds. The T-bill portfolio share is plotted as a 3-month moving average of monthly data reported at month-end. This 3-month moving average smooths out the quarter-end spikes in RRP that were visible in Figure 2. Region 1, denoted by the shaded area, spans from March 2018 to March 2021. During this period, the average RRP-bill spread is -11 basis points and the average T-bill portfolio ratio is 96 percent. Region 2 spans the remainder of the sample period: from September 2013 to February 2018, from April 2021 to April 2023, and from May 2023 to June 2024. During these intervals, the average RRP-bill spread is 6 basis points and the average T-bill portfolio share is 43 percent. Within Region 2, we have a period when T-bills are especially scarce and the market approaches Region 3: April 2022 to April 2023. This is denoted by the hatched area. During this period the average RRP-bill spread is 37 basis points and the average T-bill portfolio share is 16 percent. The debt ceiling crisis causes the sharply negative spread that troughs on May 15th, 2023.

Figure 4
Variation in Bill-RRP Portfolio Ratio Due to Investor Flows
September 2013 to June 2024

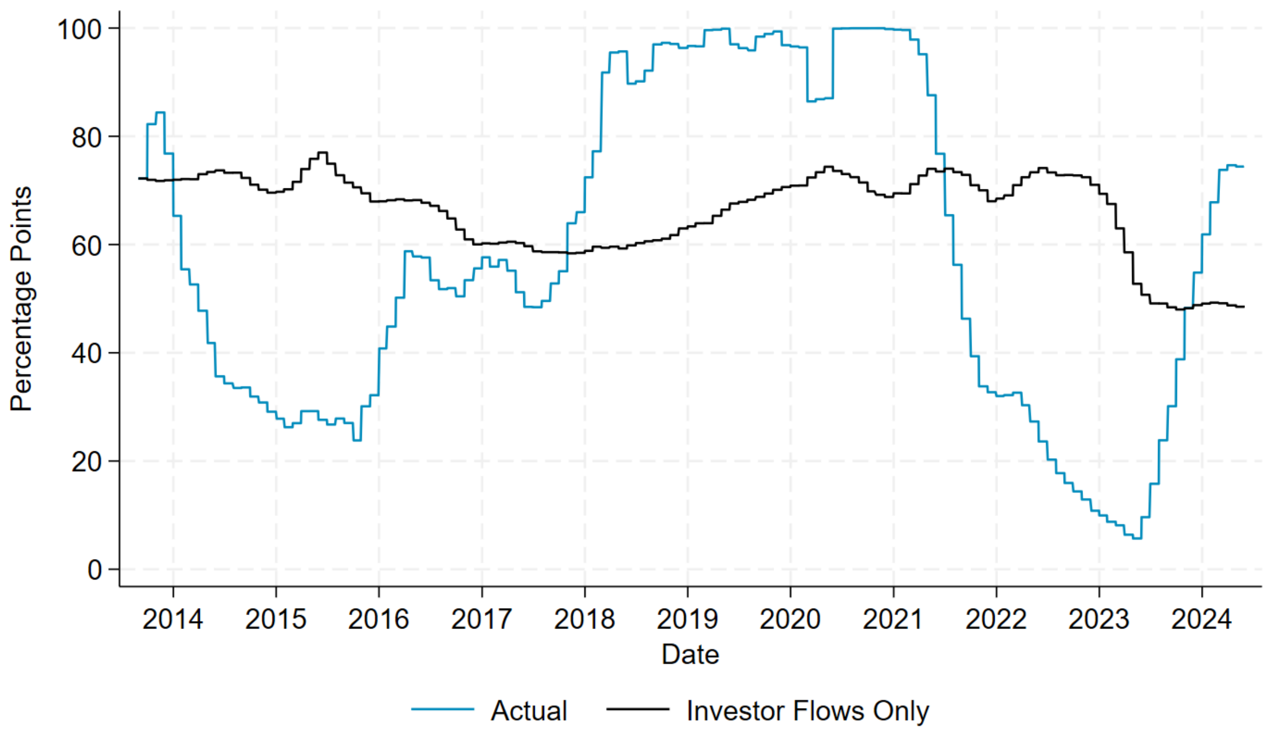

Figure 4 shows the actual Bill-RRP portfolio share (also shown in Figure 3) and a counterfactual ratio driven only by investor flows, and assuming no active rebalancing by money fund managers. Both series shown are 3-month moving averages of monthly data so as to smooth the quarter-end spikes seen in Figure 2. The details of this decomposition are described in the text.

Figure 5
Money Fund New Purchases of T-bills and the RRP-Bill Spread
June 2021 to June 2024

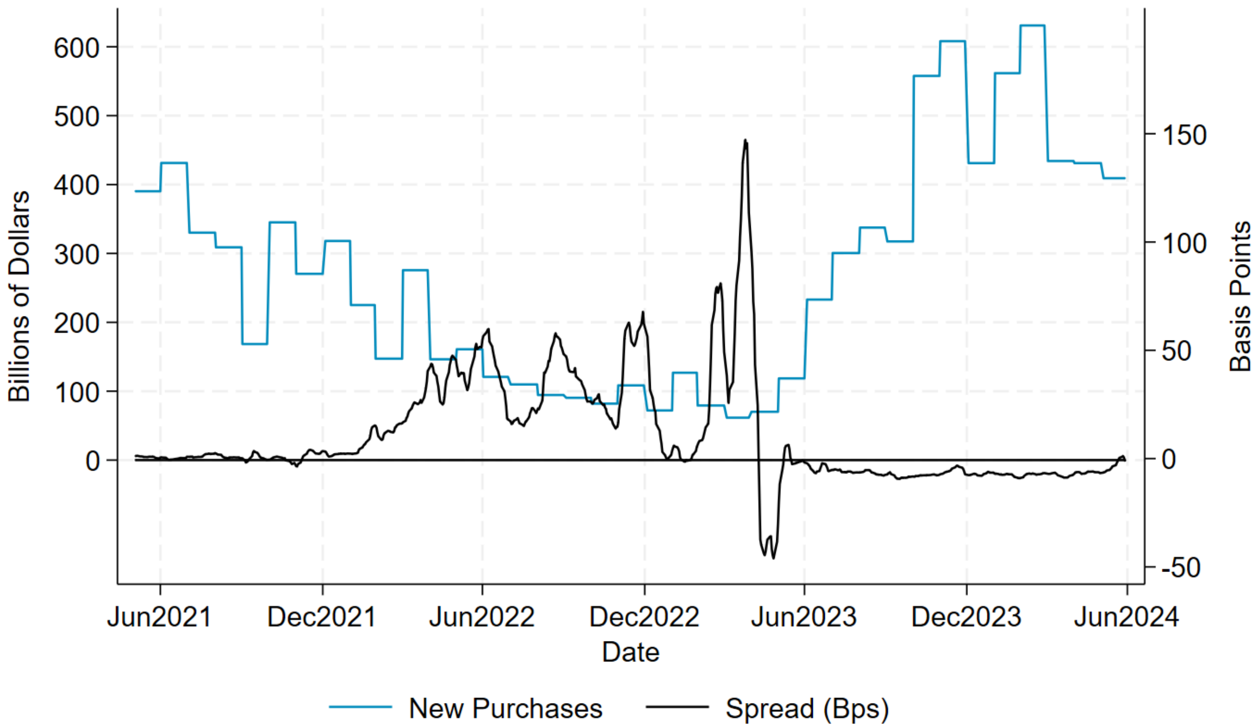

Figure 5 shows the new purchases of T-bills by money funds in our sample plotted against the RRP-bill spread from June 2021 to June 2024. New purchases are defined at the money-fund by T-bill cusip level, and are the sum of dollar increases in money fund holdings of T-bills. Despite a large RRP-bill spread from April 2022 to April 2023, money funds in sum purchased on average of 102 billion dollars of T-bills per month and new purchases never fell below 61 billion per month.

Figure 6
Fund-Level Elasticity of Substitution
September 2013 to March 2021

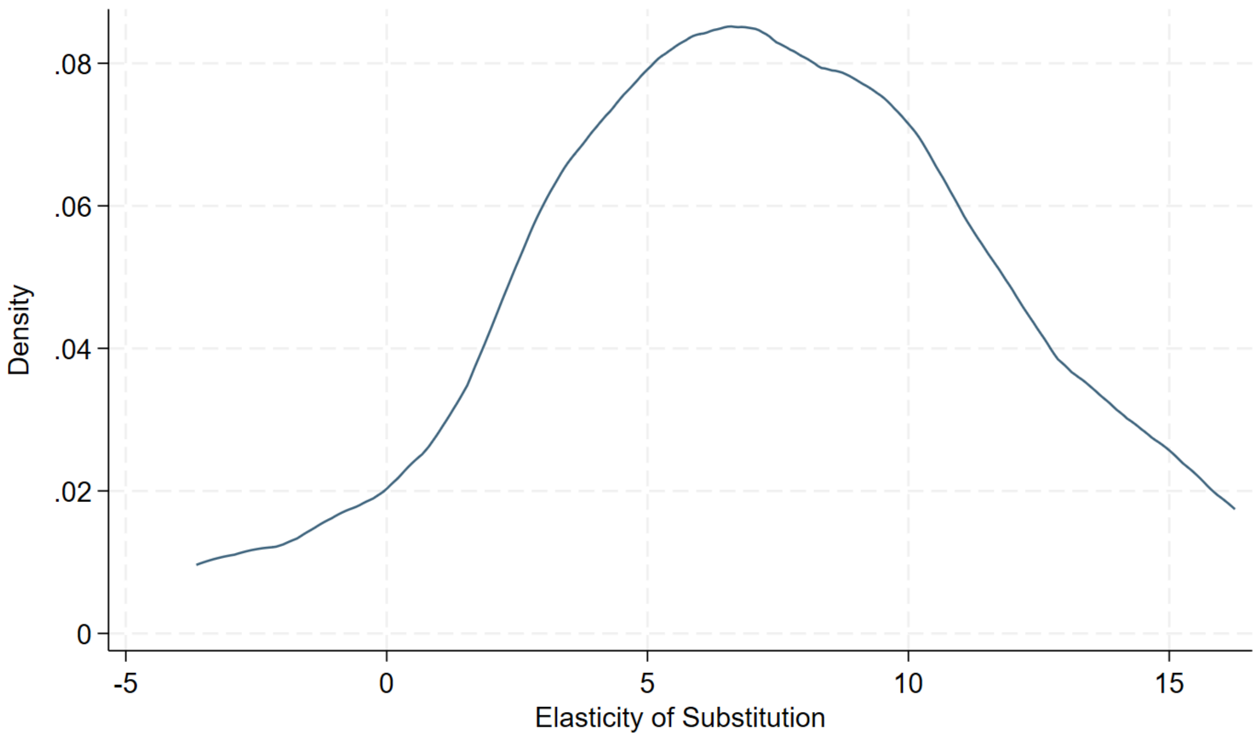

Figure 6 illustrates the distribution of fund-level elasticities. For each of the funds in our sample, we estimate the fund-level elasticity of demand for T-bills at a monthly horizon, as described in the text, over a sample running from September 2013 to March 2021. The figure shows the kernel density of these fund-level elasticities winsorized at the 1 percent level. The median fund has an elasticity of 7.4 and the interquartile range spans from 4.4 to 10.1.

Figure 7
Average Elasticity of Unconstrained Money Funds
January 2022 to June 2024

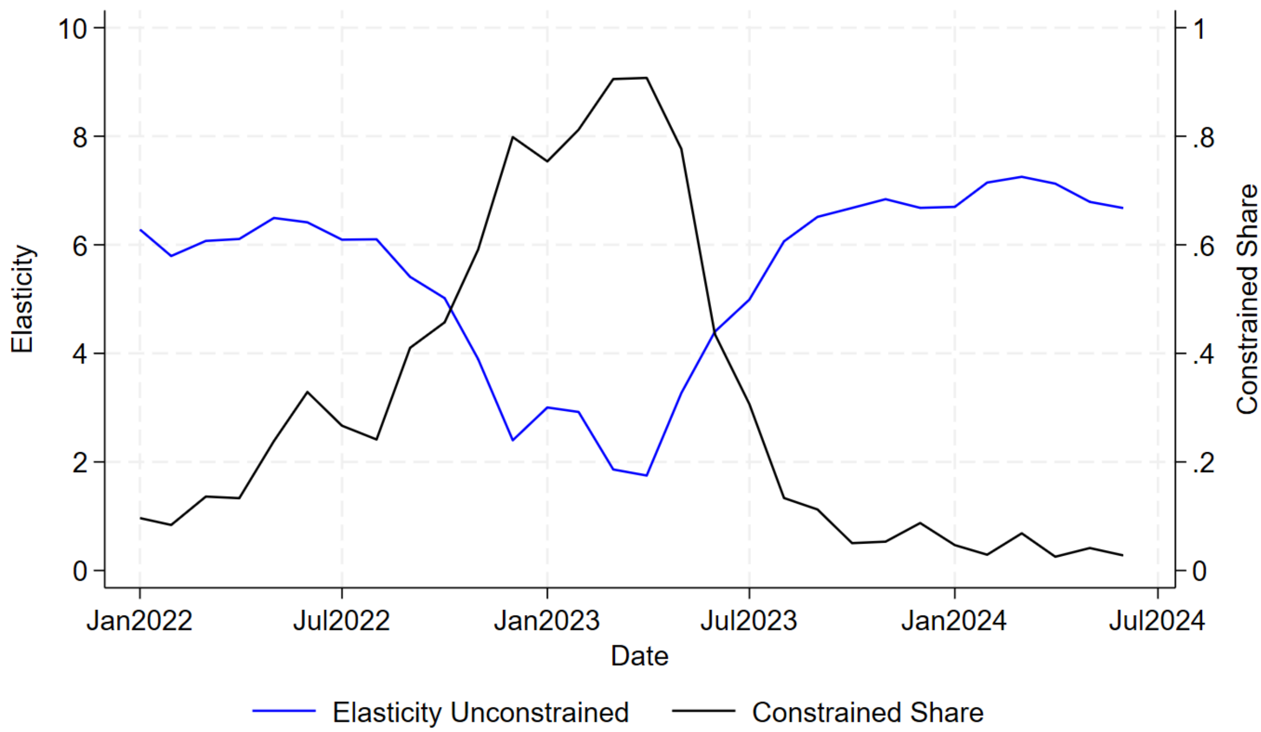

Figure 7 plots the average elasticity of unconstrained money funds weighted by their T-bill holdings (blue line and left y-axis) and the fraction of money funds that are constrained (black line and right y-axis). We estimate fund-level elasticity of demand for T-bills from September 2013 to March 2021 at a monthly horizon as described in the text and shown in Figure 6. When T-bills became scarce, the more elastic money funds sold their T-bill holdings and became constrained, leaving behind a more inelastic pool of residual unconstrained funds.

Figure 8
RRP and T-bill Supply
January 2023 to June 2024

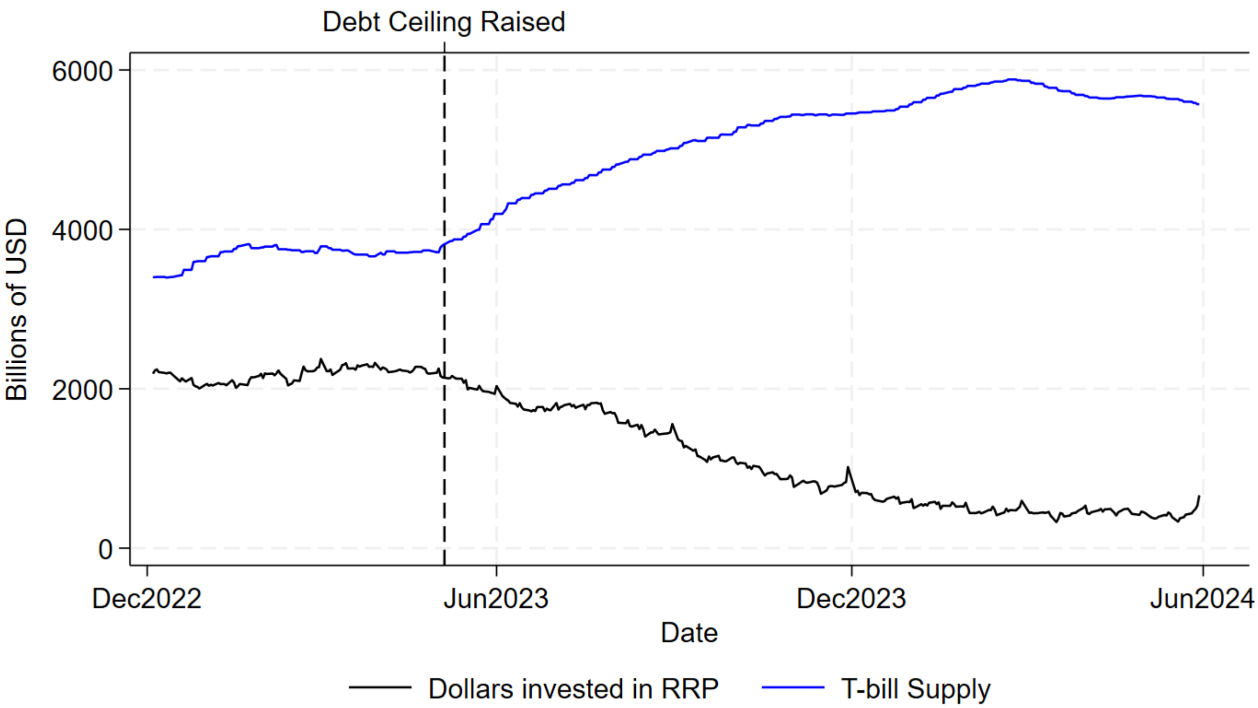

Figure 8 shows the supply of T-bills and dollars invested in the RRP facility from January 2023 to June 2024. After the debt ceiling was raised, T-bill supply increased by $1.8 trillion from June 3, 2023, to June 30, 2024, and the RRP facility experienced outflows of $1.5 trillion, which is 83 cents for every dollar of increased T-bill supply.

Figure 9
Change in RRP-Bill Spread and T-bill Supply
September 2013 to June 2024

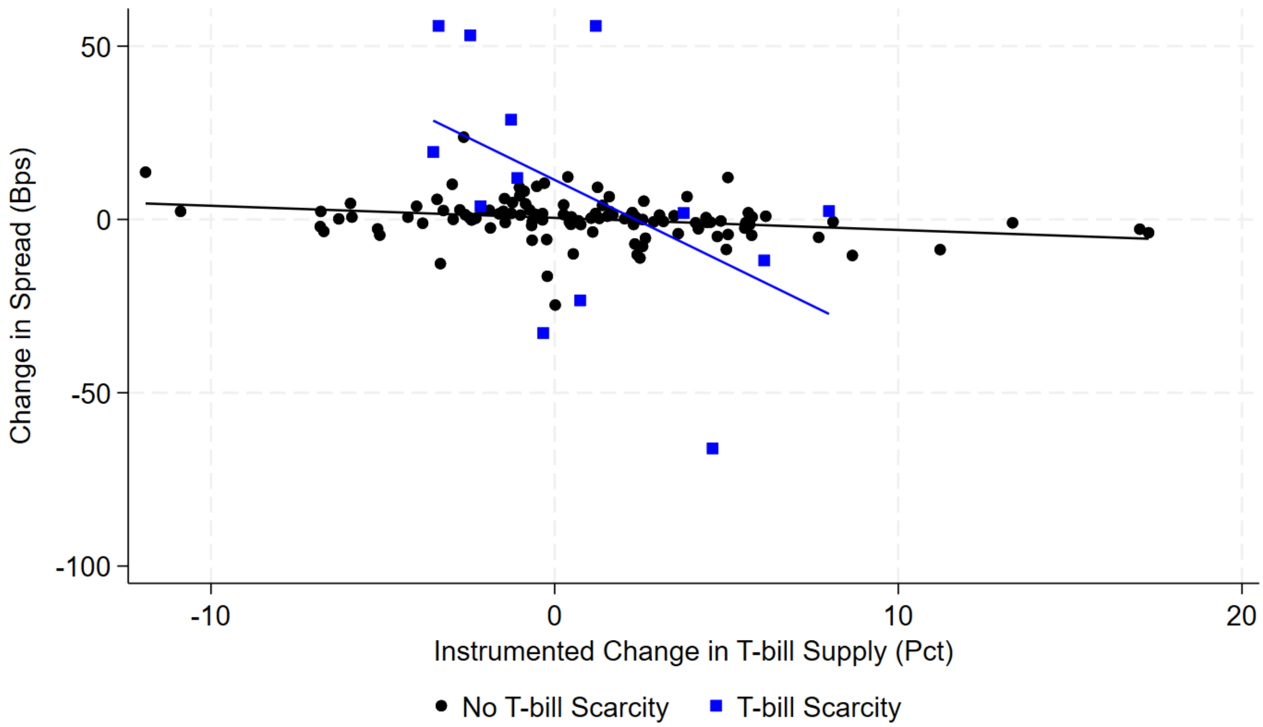

Figure 9 shows a scatter plot for the monthly change in the RRP-bill spread versus the instrumented change in the supply of T-bills (the fitted value from a regression of the change in T-bills on the change in aggregate Treasury securities.) The points are grouped by whether T-bills are scarce, as in April 2022 to April 2023 (the hatched area in Figure 3), or not scarce, as in the remainder of our sample period. We exclude the debt ceiling crisis months of May and June 2023; see Figure 10 for a plot of T-bill yields during the debt ceiling crisis. The blue line shows a linear fit of the data for when T-bills are scarce and the black line shows the analogous fit for when T-bills are not scarce.

Figure 10
T-Bill and RRP Yields by Maturity
May 2023

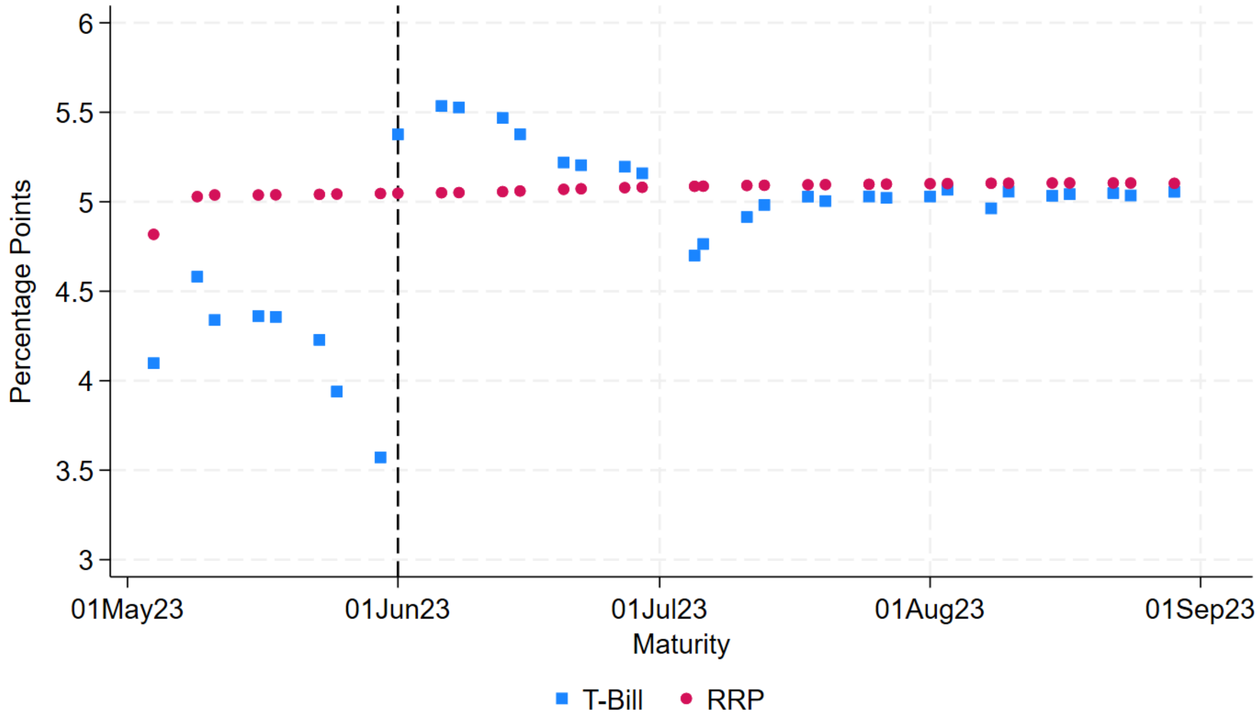

Figure 10 shows the average T-bill and RRP yield curves for May 2nd to May 30th, 2023, spanning T-bills that mature from May to August of 2023. June 1st, 2023 is demarcated as the projected date when the Treasury would be unable to meet its obligations under the debt ceiling. Each observation is a T-bill cusip and we have the date of maturity of the T-bill on the x-axis and the average yield to maturity on the y-axis. Average yield to maturity is the average of the given T-bill's yield to maturity on all the days in May of 2023 for which it traded. Corresponding to each T-bill cusip, we plot a maturity-matched expected RRP rate, given by the prevailing overnight RRP rate, plus a maturity-matched OIS rate, minus the overnight Fed Funds rate.

Table 1
Summary Statistics
September 2013 to June 2024

	Mean	p50	Std. Dev	Min	Max	*N*
E[RRP Rate] 1M in Bps	141	52	177	-13	534	2,608
T-Bill Rate 1M in Bps	139	47	177	-11	571	2,608
RRP-Bill Spread in Bps	2	0	17	-61	166	2,608
AUM in Millions	30,400	14,690	44,240	16	334,168	13,579
Days to Maturity	33	34	14	1	103	13,579
Number of Funds	104	103	5	98	114	130
Weight T-Bills in Pct	12	3	19	0	100	13,579
Weight RRP in Pct	10	0	17	0	99	13,579

Table 1 presents summary statistics for money market rates and for the money funds in our sample, which covers the period September 23, 2013, to June 30, 2024. Expected RRP Rate is the expected one-month RRP rate. T-Bill Rate is the one-month T-bill rate. RRP-Bill Spread is the difference between these two. AUM in millions are the assets under management of a given money fund, reported monthly. Days to Maturity is the dollar weighted time to maturity of a money fund's portfolio holdings in units of days, again reported monthly. Number of Funds is the number of funds present in the data for each month of our sample period, which comprises 130 months. Weight T-Bills and Weight RRP are the portfolio weights of a money fund in T-bills and RRP, respectively, once again reported monthly.

Table 2
Investor Flow Sensitivity to Money Fund Returns
September 2013 to June 2024

	Investor Flow Pct			
	(1)	(2)	(3)	(4)
Fund Return	2.908***	6.587***	4.878***	5.885***
	(0.66)	(1.09)	(1.05)	(0.98)
AUM Weighted	N	N	Y	Y
Quarterly Frequency	N	N	N	Y
Time FE	Y	Y	Y	Y
Fund FE	N	Y	Y	Y
Adjusted R^2	4%	5%	11%	14%
N	13,434	13,434	13,434	4,449

Table 2 estimates the sensitivity of investor flows to money funds returns for our sample from September 2013 to June 2024. Investor flows are measured as a percentage of lagged fund AUM. Money fund returns are monthly returns. For columns 1 and 2, the data is monthly and equal weighted. For column 3, the regression weights observations by fund AUM. For column 4, the data is quarterly (quarter-end) and the regression weights observations by fund AUM. The estimate in column 4 implies that a 100 basis point increase in returns is associated with a 5.9 percentage-point increase in AUM due to investor flows over a quarter. Standard errors are clustered by fund.

Table 3
Corporate Treasury Holdings and Derivatives Exposure
2004 to 2021

	Treasury Share				
	(1)	(2)	(3)	(4)	(5)
Derivatives Exposure	3.438**	3.249**	3.481**	6.054**	3.050**
	(1.34)	(1.48)	(1.57)	(2.92)	(1.49)
Firm Size	0.012	0.020	0.021	0.035	0.005
	(0.02)	(0.01)	(0.02)	(0.03)	(0.01)
Corporate Bond Share					0.577***
					(0.08)
Time FE	Y	Y	Y	Y	Y
Industry FE	N	Y	Y	Y	Y
Time × Industry FE	N	N	Y	Y	Y
Subsample Pos Treasury	N	N	N	Y	N
Adjusted R^2	0.01	0.18	0.15	0.11	0.41
N	2,730	2,730	2,730	757	2,730

Table 3 shows the association between corporate Treasury holdings and derivatives exposure. We measure corporate Treasury holdings as Treasury holdings divided by the sum of Treasury holdings and cash and cash equivalents. We measure derivatives exposure as the absolute value of derivatives P&L divided by the firm's total assets. Each observation in the data is a firm by year and covers a sample of 193 large firms from 2004 to 2021. The regression is an OLS regression where derivatives exposure is lagged by one year to avoid contemporaneous endogeneity between profits and cash holdings. Columns 1 to 3 include progressively more fixed effects: time fixed effects, industry fixed effects, and industry by time fixed effects. Column 4 subsamples to the firms with positive Treasury holdings so as to focus on the intensive margin. Column 5 controls for lagged corporate bond share holdings. A one standard deviation increase in derivatives exposure is 0.5 percent, which is associated with a 1.8 percentage point increase in the share of Treasury holdings (column 3). This is economically large compared to the average 7.6 percent Treasury holdings share (24 percent). Standard errors are clustered by firm.

Table 4
Aggregate Elasticity of Substitution
September 2013 to March 2021

	Δ Spread	Δ T-Bill Portfolio Weight		
	First Stage	OLS	IV	IV Quarterly
	(1)	(2)	(3)	(4)
Δ Treasury	-2.20*** (0.59)			
Δ Spread		-0.94*** (0.28)	-6.19*** (2.24)	-5.00*** (1.53)
N	90	90	90	30

Table 4 estimates the aggregate elasticity of substitution, at a monthly horizon, of money funds over the period from September 2013 to March 2021. This is a time-series regression where each observation is one month, and we have 90 months of changes in the RRP-Bill spread. Treasury supply is the privately available supply of Treasuries (outstanding notional minus Treasuries held by the Federal Reserve). ΔT-bill portfolio weight is the log difference in the money fund sector's portfolio weight in T-bills. The first column shows the first stage regression: a one- percent increase in Treasury supply decreases the RRP-Bill spread by 2.20 basis points. The second column presents an OLS regression: a one-basis-point increase in the spread is associated with a 0.94 percent decrease in money fund T-bill holdings, which is 11 basis points of portfolio weight. The third column shows the instrumental variables estimate, where we use the predicted variation in change in spread from changes in the supply of Treasuries to estimate a causal effect of changes in spread on T-bill portfolio weight. We find that a one basis-point increase in the spread due to a decrease in Treasury supply causes money funds to decrease their T-bill portfolio weight by 6.2 percent, which is 79 basis points of portfolio weight. The fourth column shows the estimate at a quarterly frequency, where we sample the data at quarter-ends. Standard errors are robust to heteroskedasticity.

Table 5
Effect of Changes to T-bill Supply on RRP-Bill Spread
September 2013 to June 2024

	Δ T-bill	Δ Spread		
	(1)	(2)	(3)	(4)
	First Stage	IV	IV	IV
Δ Treasury	5.155***			
	(0.88)			
Δ T-bill		-0.656***	-0.351***	0.035
		(0.21)	(0.12)	(0.21)
Scarce T-bills			0.077	
			(0.08)	
Scarce T-bills × Δ T-bill			-4.484**	
			(1.83)	
Constrained Share				0.139
				(0.09)
Constrained Share × Δ T-bill				-6.595***
				(2.34)
N	127	127	127	127

Table 5 shows the effect of changes in T-bill supply on the RRP-bill spread. This is a monthly time-series regression. We instrument for monthly changes in the private supply of T-bills using monthly changes in the private supply of Treasuries (first stage shown in column 1). Δ T-bill is the instrumented variation in the private supply of T-bills. Scarce T-bills is an indicator variable equal to one for April 2022 to April 2023 and zero otherwise. Constrained share is defined as the AUM weighted share of money funds that have less than 5 percent portfolio weight in T-bills and at least 10 percent portfolio weight in either RRP or T-bills. Column 2 shows that an unconditional 1 percent decrease in Δ T-bill causes the RRP-bill spread to increase by 0.66 basis points. Column 3 shows that the effect of changes in T-bill supply is much larger when T-bills are scarce: a 1 percent decrease in Δ T-bill causes a 4.5 basis points increase in the spread. Column 4 uses a continuous version of the scarcity measure: when half of the money funds are constrained, for a 1 percent decrease in instrumented Δ T-bill supply, the RRP-bill spread increases by 3.3 basis points. Standard errors are robust to heteroskedasticity.

Appendix A: Derivation of the Impact of T-bill Supply on T-bill Rates

In this appendix, we derive equation (13) in the text:

$$\frac{dr_b}{dS} = \frac{1}{X'(r_b) + A\theta\left(\frac{b_U + b_L}{2}\right)}$$

We start with the market clearing condition (equation 11):

$$S = X(r_b) + A\theta T^{**}$$

Taking the derivative of the market clearing condition with respect to supply S yields:

$$\frac{dr_b}{dS} = \frac{1}{X'(r_b) + A\frac{d\theta T^{**}}{dr_b}}$$

By the product rule,

$$\frac{d\theta T^{**}}{dr_b} = \frac{d\theta}{dr_b}T^{**} + \frac{dT^{**}}{dr_b}\theta$$

Recall that the definition of θ is:

$$\theta = \frac{b_U - b_L}{b_H - b_L}$$

and substitute $b_U = \frac{1}{r_p - r_b}$ to get:

$$\theta = \left(\frac{1}{r_p - r_b}\right)\left(\frac{1}{b_H - b_L}\right) - \frac{b_L}{b_H - b_L}$$

The derivative of θ with respect to r_b is therefore:

$$\frac{d\theta}{dr_b} = \left(\frac{1}{b_H - b_L}\right)\left(\frac{1}{r_p - r_b}\right)^2$$

Recall that the definition of T^{**} is :

$$T^{**} = 1 - \left(\frac{b_L + b_U}{2}\right)(r_p - r_b)$$

and substitute $b_U = \frac{1}{r_p - r_b}$ to get:

$$T^{**} = \frac{1}{2} - \frac{b_L}{2}(r_p - r_b)$$

The derivative of T^{**} with respect to r_b is:

$$\frac{dT^{**}}{dr_b} = \frac{b_L}{2}$$

Plugging in terms and simplifying, we have that:

$$\frac{d\theta T^{**}}{dr_b} = \frac{1}{2}\left(\frac{1}{b_H - b_L}\right)\left(\left(\frac{1}{r_p - r_b}\right)^2 - b_L^2\right)$$

Substituting $\frac{1}{r_p - r_b} = b_U$ and factoring, we have that:

$$\frac{d\theta T^{**}}{dr_b} = \frac{1}{2}\left(\frac{(b_U - b_L)(b_U + b_L)}{b_H - b_L}\right)$$

which we can rewrite as:

$$\frac{d\theta T^{**}}{dr_b} = \theta\left(\frac{b_U + b_L}{2}\right)$$

Thus, we have derived equation (13):

$$\frac{dr_b}{dS} = \frac{1}{X'(r_b) + A\frac{d\theta T^{**}}{dr_b}} = \frac{1}{X'(r_b) + \theta A\left(\frac{b_U + b_L}{2}\right)}$$

Appendix B: Additional Figures and Tables

Figure A1
RRP Rate and Repo Rate
September 2013 to June 2024

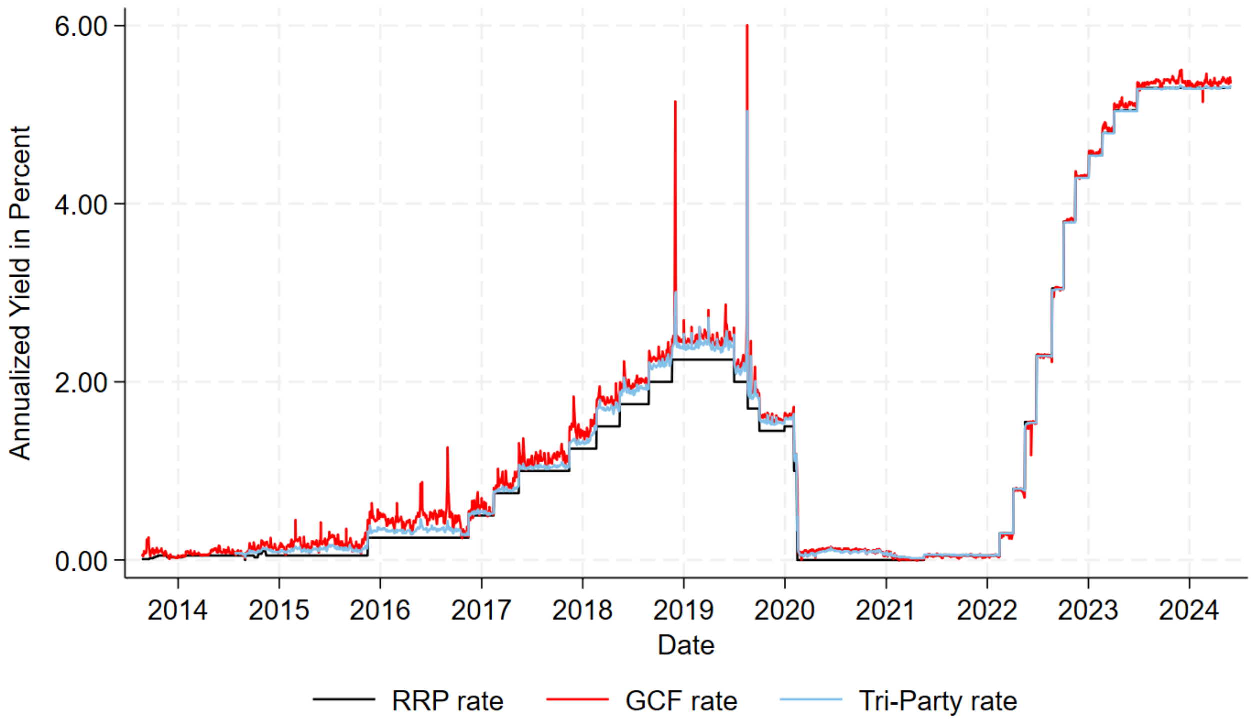

Figure A1 shows the RRP rate (black), general collateral financing (GCF) rate (red) and the Tri-Party repo rate (light blue). All three repo rates are overnight rates. The GCF rate is a weighted average rate for general collateral repo contracts as reported by the Depository Trust & Clearing Corporation (DTCC). The Tri-Party repo rate is the rate at which large dealer banks borrow from cash lenders, such as money funds, and is reported by the Office of Financial Research (OFR). The RRP and GCF rates span from September 30, 2014, to June 30, 2024, while the Tri-Party rate spans from August 22, 2014 to June 30, 2024.

Figure A2
Growth in Private Supply of Treasuries
September 2013 to June 2024

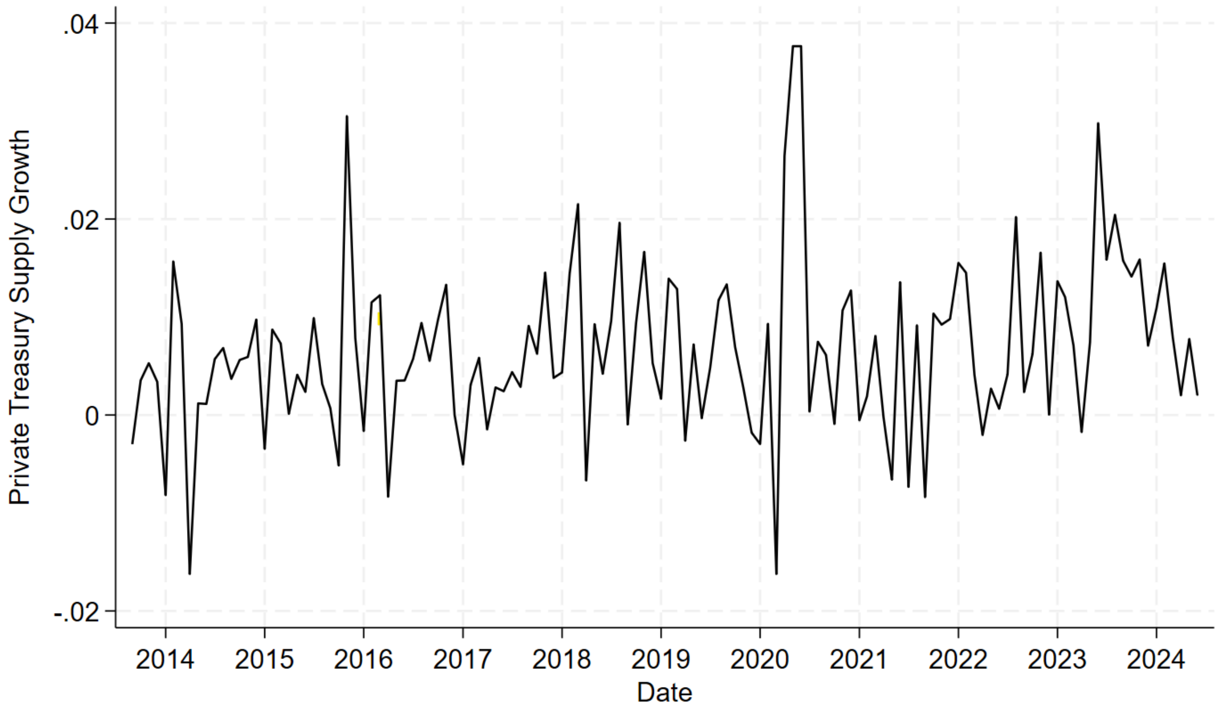

Figure A2 shows the time series of the growth in private supply of Treasuries, which we measure in monthly log differences, over our sample period from September 2013 to June 2024.

Table A1
Corporate Holdings of Corporate Debt and Derivatives Exposure
2004 to 2021

	Dep Variable: Corp Debt Share				
	(1)	(2)	(3)	(4)	(5)
Derivatives Exposure	1.818	1.401	1.517	1.409	0.182
	(1.36)	(1.04)	(1.12)	(1.99)	(1.17)
Firm Size	0.021	0.030***	0.031***	0.060**	0.020**
	(0.01)	(0.01)	(0.01)	(0.03)	(0.01)
Treasury Share					0.530***
					(0.07)
Time FE	Y	Y	Y	Y	Y
Industry FE	N	Y	Y	Y	Y
Time × Industry FE	N	N	Y	Y	Y
Subsample Pos Treasury	N	N	N	Y	N
Adjusted R^2	0.01	0.23	0.20	0.27	0.46
N	2,730	2,730	2,730	757	2,730

Table A1 shows the association between corporate holdings of corporate debt and derivatives exposure. The structure of the table mirrors that of Table 5, except that the dependent variable is Corp Debt Share, which we measure as corporate holdings of US corporate debt divided by the sum of cash and cash equivalents and US corporate debt. Furthermore, column 5 controls for lagged Treasury Share. Standard errors are clustered by firm.

Table A2
Corporate Holdings of Treasuries and Derivatives P&L
2004 to 2021

	Dep Variable: Treasury Share				
	(1)	(2)	(3)	(4)	(5)
Derivatives P&L	0.081	-0.440	-0.078	-0.261	-0.330
	(1.13)	(1.06)	(1.13)	(2.17)	(0.98)
Firm Size	0.012	0.020	0.021	0.034	0.005
	(0.02)	(0.01)	(0.02)	(0.03)	(0.01)
Treasury Share					0.580***
					(0.08)
Time FE	Y	Y	Y	Y	Y
Industry FE	N	Y	Y	Y	Y
Time × Industry FE	N	N	Y	Y	Y
Subsample Pos Treasury	N	N	N	Y	N
Adjusted R^2	0.00	0.17	0.14	0.10	0.41
N	2,730	2,730	2,730	757	2,730

Table A2 shows the association between corporate holdings of Treasuries and derivatives P&L. The structure of the table mirrors that of Table 5, except that the primary independent variable is derivatives P&L, rather than derivative exposure, which is given by the absolute value of derivatives P&L. Derivatives P&L is derivatives profits and loss divided by the firm's total assets. Standard errors are clustered by firm.