

The Imperfect Intermediation of Money-Like Assets*

Jeremy C. Stein
Harvard University and
NBER

Jonathan Wallen
Harvard Business School
Harvard University

November 2023

Abstract: We study supply-and-demand effects in the U.S. Treasury bill market by comparing the returns on T-bills to the administered policy rate on the Federal Reserve's reverse repurchase (RRP) facility. In spite of the arguably more money-like properties of an investment in RRP, we observe repeated episodes where one-month T-bill rates fall well below expected RRP rates. This gap frequently exceeds 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. In an effort to understand this phenomenon, we develop and test a simple model where the RRP-bill spread is policed by a group of heterogeneous money funds, who differ in their elasticity of substitution between the two assets. Our main finding is that when T-bills are scarce, and the spread is large, the marginal money fund is more inelastic, as the more elastic funds have already exhausted their holdings of T-bills. As a result, for a given shift in T-bill supply, the effect on rates is an order of magnitude larger when T-bills are scarce, and when more money funds are out of the market.

*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, and seminar participants at HBS and UMass Amherst for helpful feedback.

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, where money-market funds are among the most important players. However, unlike much of the work in this area, which focuses on the behavior of a single representative intermediary, we emphasize heterogeneities in the responsiveness of different money funds to supply and demand shocks in the T-bill market. We show that these heterogeneities play a crucial role in shaping the price impact of such shocks.

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 roughly 140 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 \$4.5 trillion in T-bills (excluding T-bills held by the Fed) and \$1.8 trillion in RRP outstanding as of July 31st, 2023.

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 July 2023.¹ 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.

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.

Yet, as Figure 1 shows, this hypothesis is rejected by the data, sometimes dramatically so. While T-bill and expected RRP rates are nearly identical from June 2021 through early 2022, after that they begin to diverge, with one-month T-bill rates falling well below expected RRP rates. This gap frequently exceeds 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.

In an effort to understand this puzzling behavior, we zero in on two types of frictions that can 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 14 percent over the period of September 2013 to July 2023.

Limited access implies that if, for example, a corporate treasurer 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 for our purposes 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.

This form of segmentation leaves the money funds themselves as the front-line players who are in the best position to police the RRP-bill spread. Here is where we encounter a second important friction: inelastic substitution on the part of these funds. While one might have expected 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—the evidence contradicts this presumption.² For reasons that

² 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.

we do not fully understand, money funds behave, at least in the short to medium run, as if they derive some non-pecuniary benefit from holding T-bills relative to RRP, leading them to have a lower elasticity of substitution than they otherwise might. While slow-moving capital is commonly attributed to information asymmetries or specialization in the pricing of risky assets (Duffie 2010), we find that capital is also slow to move between even the simplest and most transparent risk-free assets. Thus, this impediment to arbitrage may be more ubiquitous than is commonly believed.

In addition, and strikingly, 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 in a situation where the RRP-bill spread is 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 serve to 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.

Said more generally, in a world with heterogeneous money fund elasticities of substitution, an important state variable is the fraction of money funds that are at a corner, being completely out of T-bills and having substituted fully into RRP. Our model predicts that the impact of a given-sized T-bill supply shock on the RRP-bill spread will be larger when the fraction of money funds at a corner is higher. We find this prediction to be strongly supported by the data.

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.³ At this point, our model predicts that, in the absence of any intermediary who can connect the two markets, the RRP-bill spread becomes extremely sensitive to even modest shocks in T-bill 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.

This paper aims to make two contributions. First, and most concretely, we hope to shed some 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 quite disconnected from the rate paid by the Fed on its RRP facility would 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.”

One way to interpret our findings is that by limiting access to the RRP facility to money funds, the Fed’s “broad range of financial institutions” is, in fact, not broad enough to accomplish its objectives. Ex ante, it might have been plausible to think that money funds are a thin enough veil that giving them access to the RRP would effectively be just as good as giving direct access to e.g., those corporations and households that can in turn easily invest with the money funds. However, this seems not to be the case. Rather, as has been found in many other situations, frictions to financial intermediation can leave a noteworthy imprint on asset prices (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).⁴

The novelty here is that 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. Given this simplicity, as well as the magnitude of the stakes, the fact that frictions to financial intermediation play such an important role in this setting may be thought of as particularly surprising. In contemporaneous work, Doerr, Eren and Malamud (2023) also study the spread between RRP and T-bills, and the associated role

³ On May 1st, 2023, the Treasury projected in a letter to Congress that it would be unable to meet its obligations by June 1st.

⁴ For a more comprehensive review of the literature see He and Krishnamurthy (2018).

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 heterogeneous preferences among a set of perfectly competitive funds.

On a second, and more general level, we hope to contribute to the broader literature on intermediary asset pricing by highlighting the role of heterogeneity among intermediaries, and the implications this heterogeneity has for the responsiveness of asset prices to various shocks. The theoretical literature tends to model a representative financial intermediary 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).⁵ By contrast, we show that the fraction of intermediaries active in the market at any point in time 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.⁶ What makes our setting a particularly nice laboratory for examining this hypothesis is that we know precisely who 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.⁷

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 a heterogeneous group of money funds, who differ in their willingness to substitute between the two assets. In Section III, we examine a range of data that bear 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.

⁵ 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.

⁶ 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.

⁷ 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 Treasuries they have taken onto their balance sheets.

II. A Model with Heterogeneous Money Funds

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). The key distinction between the two groups is that money funds also have access to the Fed’s RRP facility, while treasurers do not, and can therefore only invest in T-bills. 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 key friction in the model: the hard segmentation between treasurers and money funds. Below, we will show that something close to this degree of segmentation effectively holds in the data. We assume that treasurers’ demand for T-bills is given by $X(r_b)$, where r_b is the endogenously-determined T-bill rate, and $X(\cdot)$ is an increasing function. The RRP facility rate is r_p and is exogenously set by the Federal Reserve.

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, whose ultimate source we are admittedly agnostic as to, is the second key friction in the model. As discussed above, it allows us to match the empirical fact that money funds tend 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:

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

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:

$$(r_p - r_b) = V'_i(T_i) = \frac{1-T_i}{b_i}, \quad (3)$$

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

$$T_i^* = 1 - b_i(r_p - r_b). \quad (4)$$

In order for equation (4) to apply, i.e., in order for fund i to in fact be at an interior solution, we require $T_i^* > 0$, 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 key 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:

$$X(r_b) = S - A. \quad (5)$$

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:

$$S \geq X(r_p) + A. \quad (6)$$

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:

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

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:

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

We can then use equation (13) to prove:

Proposition 1: In Region 2b, we have that $\frac{d^2 r_b}{dS db_U} < 0$, 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).

The proposition forms the basis for our key 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 $r_b \leq r_p - (\frac{1}{b_L})$, 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. Again, the intuition is that as more money funds move to the corner and exit the T-bill market altogether, they provide less buffering capacity against subsequent shocks to T-bill supply.

III. Empirical Evidence

Data Sources

We collect data for a sample that spans from September 30th, 2013, to July 31st, 2023. 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. There are 140 such money funds and

they have in aggregate 2.985 trillion dollars 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 Treasury Direct to measure changes in the outstanding quantity of T-bills due to auctions and maturing T-bills. 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.

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 July 2023. 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.

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 RRP and T-bills is 22 percent over the full sample and 51 percent from June 2021 to July 2023. Thus, while our two-asset model 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 July 2023; 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^* in the model, i.e., the average portfolio weight in T-bills in the model's simplified two-asset setting.

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 T-bills yield weakly more than RRP (a negative RRP-bill spread) and 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 bps and the average T-bill portfolio ratio is 96 percent. By contrast, in Region 2, T-bills are meaningfully scarce, the RRP-bill spread is positive, and money funds hold some RRP. In the data, this corresponds to the interval from September 2013 to February 2018, as well as that from April 2021 to April 2023. Across these two periods the average RRP-bill spread is 9 bps and the average T-bill portfolio ratio is 43 percent. Within the latter part of Region 2, 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 bps

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:

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

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:

$$IFlow_{Bill,i,t} = w_{Bill,i,t-1}(A_{i,t} - A_{i,t-1}) \quad (17)$$

Thus the managerial rebalancing component is given by:

$$MFlow_{Bill,i,t} = \Delta D_{Bill,i,t} - IFlow_{Bill,i,t} = (w_{Bill,i,t} - w_{Bill,i,t-1})A_{i,t} \quad (18)$$

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:

$$\tilde{D}_{Bill,t} = D_{Bill,Sept\ 2013} + \sum_{\tau=1}^t IFlow_{Bill,\tau}, \quad (19)$$

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:

$$\widetilde{Ratio}_t = \frac{\tilde{D}_{Bill,t}}{\tilde{D}_{Bill,t} + \tilde{D}_{RRP,t}} \quad (20)$$

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 July 2023 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 6.38 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 6.38 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 2.36 percentage points more than a fund entirely invested in T-bills.⁸

Money Fund Elasticity of Substitution

The second key 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 3, we estimate aggregate regressions of the form:

$$\Delta w_{Bills,t} = \alpha - \beta \Delta Spread_t + \epsilon_t \quad (21)$$

where $\Delta w_{Bills,t}$ is the log percent change 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.⁹ 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 (21) 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

⁸ This 2.36 percentage point effect is the point estimate of 6.38 in column 4 of table 2 multiplied by the average RRP-bill spread of 37 bps ($2.36 = 6.38 \times 0.37$).

⁹ 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.

private supply of T-bills.¹⁰

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. If, to the contrary, the Treasury did tactically time the market in this way, we would confound supply and demand shocks, which would bias our estimate of β toward zero.

In the first column of Table 3, we estimate the first-stage regression for this instrumenting procedure, namely the effect of changes in T-bill supply on the RRP-bill spread (in units of basis points):

$$\Delta Spread_t = \alpha + \beta \Delta TBill Supply_t + \epsilon_t \quad (22)$$

where $\Delta TBill Supply_t$ is the log percent change in the private supply of T-bills. The point estimate in this column implies that for a 10 percent increase in the supply of T-bills, the RRP-bill spread decreases by 4.0 bps.

The second column of Table 3 estimates equation (21) 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 10 basis point increase in the RRP-bill spread is associated with money funds decreasing their T-bill portfolio weight by 9.4 percent. To be clear, this is not 9.4 percentage points, but rather 9.4 percent of the aggregate portfolio weight of money funds in T-bills. Money funds on average held 12.6 percent of their portfolio in T-bills from September 2013 to March 2021. Thus, in response to a 10 basis point increase in the RRP-bill spread, the regression implies that money funds would decrease their portfolio weight in T-bills by 1.2 percentage points ($1.2 = 12.6 \times 0.094$).

The third column of Table 3 again estimates equation (21), but this time by instrumental variables, using changes in T-bill 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 10 basis point increase in the RRP-bill spread due to a change in T-bill supply, money funds decrease their T-bill holdings sharply, by 80 percent, which amounts to 10.1 percentage points of portfolio weight. The fourth

¹⁰ The private supply of T-bills is the total notional of outstanding T-bills less what is held by the Federal Reserve.

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 illustrates this point, by plotting the actual decrease in money fund T-bill holdings and a counterfactual decrease that would have occurred if money funds had stopped buying any new T-bills after December 2021. 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.

As noted above, one potential problem with using changes in T-bill supply as an instrument for changes in the RRP-bill spread would be if the Treasury attempted to time the T-bill market, issuing a greater share of T-bills (as a share of total Treasury debt issuance) when demand for T-bills was unusually high. One way to address this problem is to instead use changes in the total supply of all U.S. government debt as an instrument, rather than changes in just T-bill supply. The identifying assumption now is the weaker condition that total changes in aggregate government borrowing do not respond to movements in the RRP-bill spread. As Appendix Table B1 shows, this approach yields estimates very similar to those in Table 3. This should not be too surprising, as changes in T-bill supply have a correlation of 0.81 with changes in the total supply of government debt.

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 3) 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. Figure 6 plots the kernel density of the estimated elasticities for our sample of funds. The median fund has an elasticity of 10.6 and the interquartile range of elasticities goes from 6.1 to 13.7.

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 7.4; Government funds and Prime funds are significantly more elastic with average elasticities of 10.4 and 11.2, respectively. This hints at the idea that funds with narrower mandates behave more inelastically. Treasury funds have the most restrictive mandate, and 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.¹¹

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 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.

The core economic mechanism underlying Proposition 1 of the model is that in response to a negative T-bill supply shock, more elastic funds sell their T-bills more aggressively, and therefore become constrained—i.e., hit a corner where they run out of T-bills—faster than do less elastic funds. To test this, we take the fund-level elasticities estimated in the earlier September 2013 to March 2021 period, and we check to see whether they predict the incidence of constraint in a later period, January 2022 to April 2023, which includes the interval of the most pronounced T-bill scarcity. Specifically, we group funds into two buckets, based on their early-period estimated elasticities: a low-elasticity bucket, composed of funds in the bottom quartile; and a high-elasticity bucket, composed of the remainder of funds.

Figure 7 then plots, for this later period, the fraction of funds in a given month that are constrained, where we define a constrained fund as one that has a T-bill portfolio weight of less

¹¹ In a separate cut, we find that bank-affiliated money funds are significantly more elastic than non-affiliated funds, with elasticities of 12.5 and 8.9, respectively.

than 5 percent and a portfolio weight of at least 10 percent in either RRP or T-bills.¹² As can be seen in the figure, the more elastic funds (the blue line) become constrained significantly faster than do the less elastic funds (red line). For example, by November 2022, 70 percent of elastic funds are constrained, while only 28 percent of inelastic funds are constrained. Again, this statement is based on fund-level estimates of elasticity from a non-overlapping period, which suggests that these estimates do in fact carry meaningful information about persistent fund-level differences. This validates a central premise of our model, namely that, in the face of rising T-bill scarcity, more elastic funds will hit the corner first.

On a separate note, Figure 7 also 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.

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 April 28th to the end of July of 2023, T-bill supply increased by \$848 billion, the RRP facility experienced outflows of \$504 billion, and the RRP-bill spread dropped to nearly zero. Consistent with these developments, the aggregate fraction of constrained money funds fell back down to 34 percent by the end of July.

As emphasized in our 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, illustrated in Figure 7, 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 9 depicts this latter effect, plotting the average elasticity of the remaining unconstrained money funds weighted by their T-bill holdings, using the same methodology and sample period as in Figure 7. As the more elastic funds become constrained and leave the market, the average elasticity of the remaining marginal players in T-bills falls by more than half, from an elasticity of 8.6 in January 2022 to an elasticity of 3.0 by April 2023. Not surprisingly, this pattern also reverses itself after the raising of the debt

¹² This portfolio weight requirement of at least 10 percent in RRP or T-bills precludes funds that are not invested in either RRP or T-bills from being labeled as constrained. This requirement is important for data prior to the money fund reform in October 2016, when many funds held nearly zero RRP or T-bills. Figure 2 shows that during this early sample period, the money fund sector held more commercial paper and certificates of deposit and less RRP and T-bills.

ceiling and the subsequent large issuance of T-bills. By the end of July 2023, with a substantial fraction of money funds re-establishing long positions in T-bills, the average elasticity of the unconstrained fund has increased by to 7.0.

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 our model, these two effects combine to drive our Proposition 1, which states 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 4 displays our tests of this hypothesis. Using a sample period that runs from September 2013 to April 2023, we estimate the following regression in monthly data:

$$\Delta Spread_t = \alpha + \beta_0 \Delta TBill Supply_t + \beta_1 Scarce TBills_t + \gamma \Delta TBill Supply_t \times Scarce TBills_t + \epsilon_t \quad (23)$$

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.

The first column of Table 4 shows that, unconditionally over the full sample from September 2013 to April 2023, a one percent decrease in the supply of T-bills leads a 0.77 bps increase in the RRP-bill spread.¹³ In the second column of Table 4, we show the interactive specification in equation (23). 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 T-bill supply causes the RRP-bill spread to increase by 5.3 bps, which is 15 times as large as the impact that arises when T-bills are not scarce. Moreover, while the regression with just changes in T-bill supply in the first column explains just 6 percent of the variation in the RRP-bill spread, conditioning on the period of T-bill scarcity in the second column raises the R-squared of the regression to 30 percent.

Figure 10 illustrates this result graphically, with a scatter plot of the changes in the RRP-

¹³ Note that this effect is larger than the 0.44 bps figure that we estimated in column 1 of Table 3, due to the difference in sample; recall that in Table 3 we focused on an early-period sample to estimate price impact so that we could then make out-of-sample predictions.

bill spread and 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 2 of Table 4.

Finally, in the third column of Table 4, we consider the following alternative specification:

$$\begin{aligned} \Delta Spread_t = & \alpha + \beta_0 \Delta TBill Supply_t + \beta_1 Constrained Share_t \\ & + \gamma \Delta TBill Supply_t \times Constrained Share_t + \epsilon_t, \end{aligned} \tag{24}$$

where $Constrained Share_t$ 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 Proposition 1 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 25 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 3 imply that when 25 percent of money funds are constrained, as was the case in May of 2022, a 1 percent decrease in T-bill supply causes the RRP-bill spread to increase by 1.7 bps. 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 bps. 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 1st, 2023, the Treasury stated

that it expected that it would be unable to meet its obligations beginning on June 1st.¹⁴ 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 11 shows the average yield on T-bills maturing from May to August of 2023. The average is taken over the interval from May 2nd, 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 30th yields an average of 153 bps less over the course of the month than the T-bill maturing on June 1st. The average RRP-bill spread for T-bills maturing prior to June 1st is 81 bps. 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—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. 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

¹⁴ This projection was later revised to June 6th on May 26th, 2023.

create time-varying market dynamics. This is because we know exactly who the relevant intermediaries are—i.e., the money funds—and because we can cleanly observe when each one is individually 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.¹⁵ 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.

¹⁵ 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.
- 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., & Vandeweyer, Q. (2023). Treasury bill shortages and the pricing of short-term assets. *Journal of Finance*, forthcoming.
- 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 us 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 July 2023

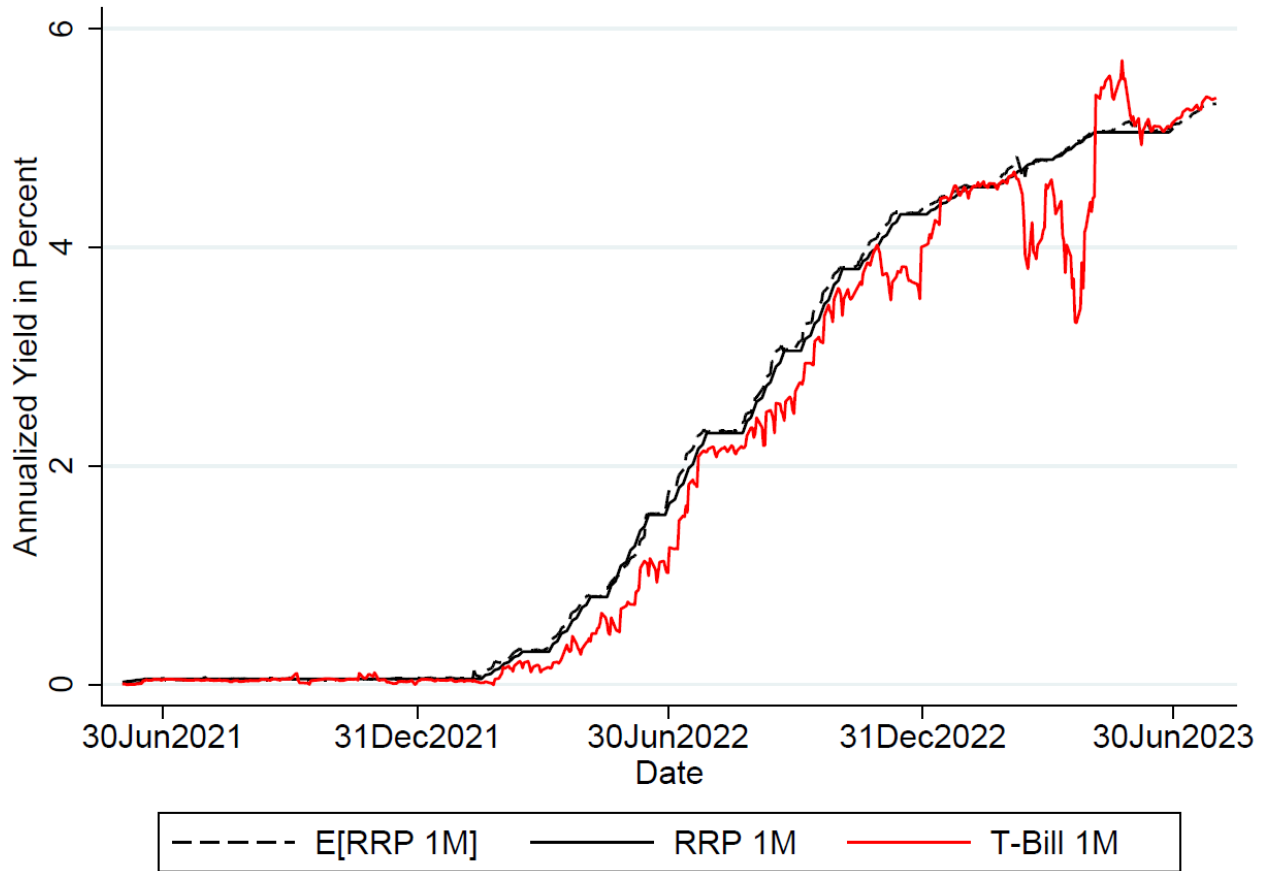


Figure 1 shows the one-month expected and realized returns on the Fed’s RRP facility, as well as the one-month T-bill rate. The expected one-month RRP rate (dashed black line) is measured as the overnight RRP rate, plus the one-month OIS rate, minus the overnight Fed Funds rate. The realized 1-month 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 July 31, 2023 (the realized 1-month RRP rate ends one month earlier, June 30th, 2023).

Figure 2
Money Fund Portfolio Composition
September 2013-July 2023

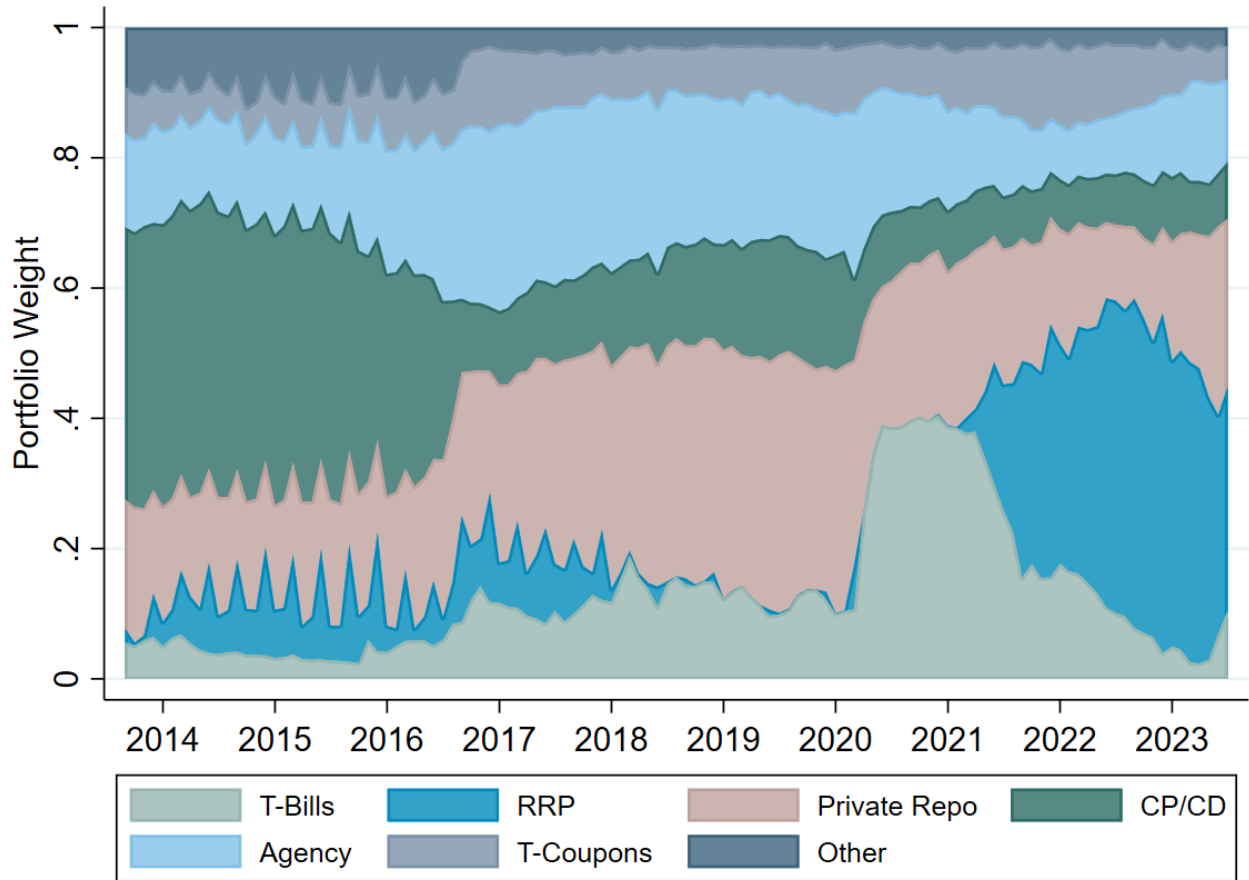


Figure 2 shows the portfolio weights by asset class for the sum of all money funds in our sample. The asset classes include Treasury Bills, RRP, reverse repo with the private sector, commercial paper and certificates of deposits, agency bonds, Treasury coupons, and other. The sample spans from the beginning of the RRP program, Sept 2013, to July 2023.

Figure 3
RRP-Bill Spread and Portfolio Shares
September 2013 to July 2023

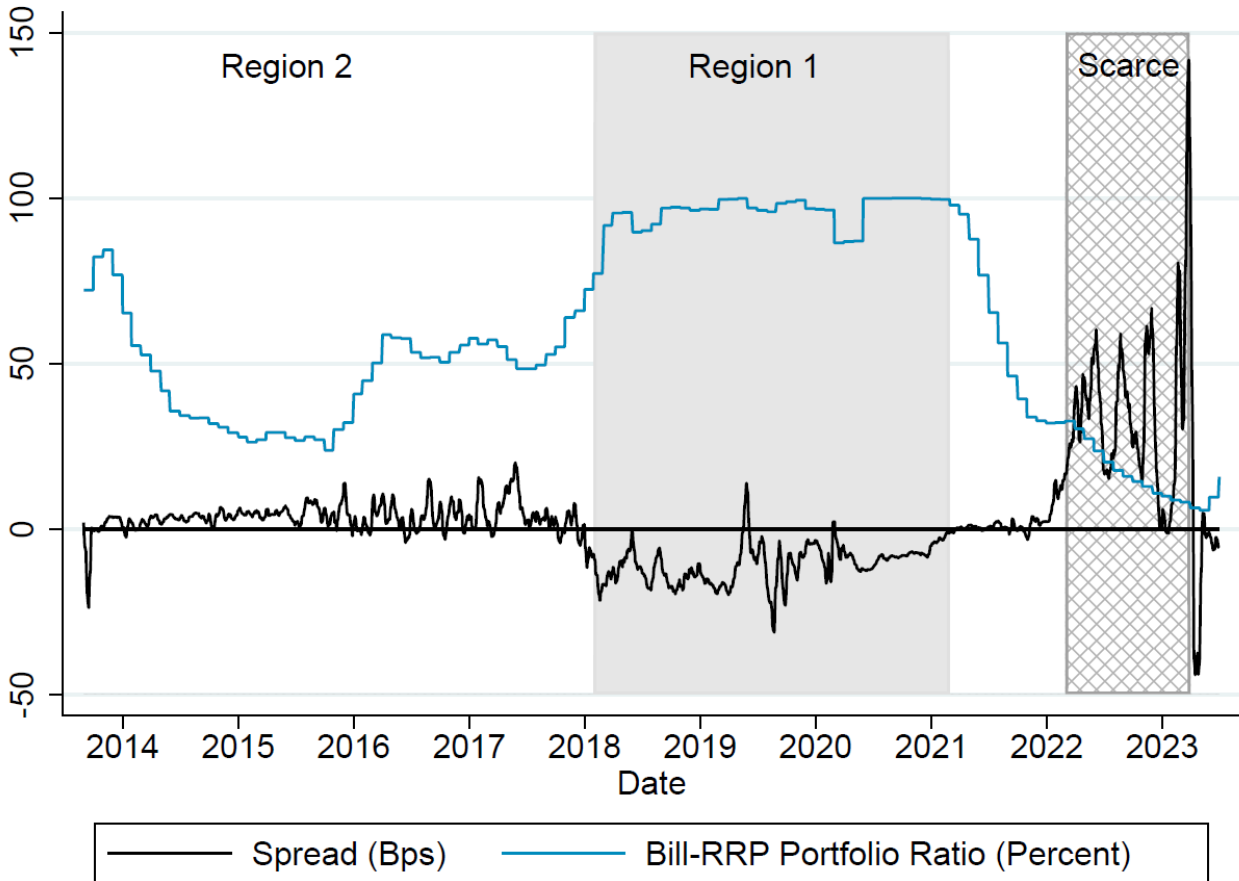


Figure 3 highlights the data corresponding to regions 1 and 2 of the model. Spread is the one-month RRP-bill spread in units of basis points, which we show as a 1-week moving average of daily data. T-bill portfolio share is the portfolio weight in T-bills divided by the sum of the portfolio weight in 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 spans from March 2018 to March 2021 (shaded area). During this period, the average RRP-bill spread is -11 bps and the average T-bill portfolio ratio is 96 percent. Region 2 spans from September 2013 to February 2018 and again from April 2021 to April 2023. During this period, the average RRP-bill spread is 9 bps 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 (hatched area) During this period the average RRP-bill spread is 37 bps and the average T-bill portfolio share is 15 percent. The debt ceiling crisis causes the sharply negative spread that troughs on May 15th, 2023.

Figure 4
Variation in T-Bill Portfolio Share Due to Investor Flows
September 2013 to July 2023

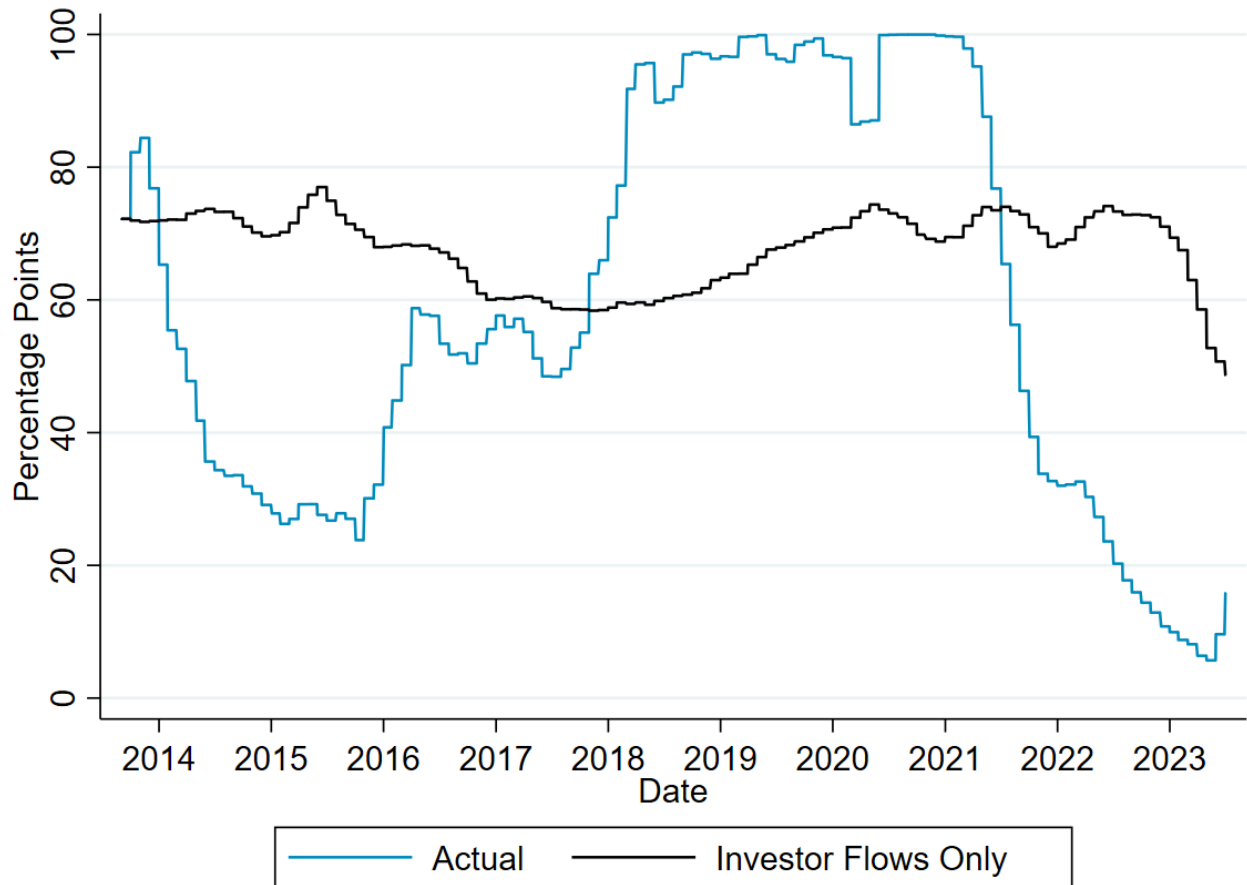


Figure 4 shows the actual T-bill portfolio share (also shown in Figure 3) and a counterfactual T-bill portfolio share driven only by investor flows, and assuming no active rebalancing by money fund managers. Both series shown are 3-month moving averages of monthly data so as to smooth the quarter-end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

Figure 5
Counterfactual T-bill Holdings with No New T-bill Purchases
June 2013 to December 2022




Figure 5 shows the actual T-bill holdings of our sample of money funds and the counterfactual holdings if money funds stopped buying T-bills and let their portfolio holdings roll off after December 2021 (vertical black line). 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 a 33 percent decrease in actual T-bill holdings.

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 10.6 and the interquartile range spans from 6.1 to 13.7. We label all funds with an elasticity in the bottom quartile as low elasticity funds and all other funds as high elasticity funds.

Figure 7
Predictable Cross-Sectional Variation in Fund Constraints
January 2022 to July 2023

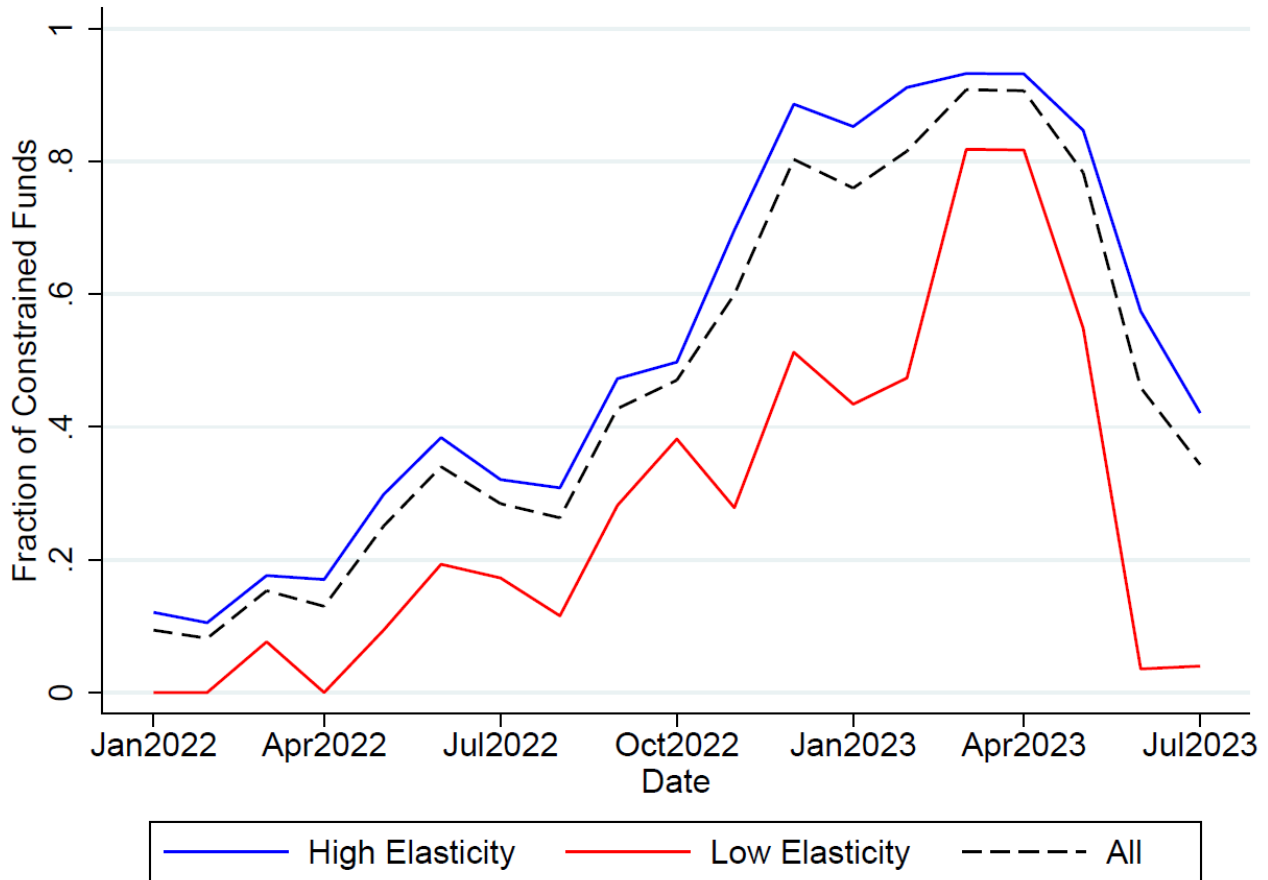


Figure 7 shows the different rates at which high elasticity and low elasticity funds become constrained. The sample covers January 2022 to July 2023, which includes the period of extreme T-bill scarcity from April 2022 to April 2023. We define a fund as being constrained if it has less than 5 percent of its portfolio in T-bills, and at least 10 percent of its portfolio in either T-bills or RRP. We measure fund-level elasticity as described in Figure 6 using non-overlapping data from the earlier period September 2013 to March 2021 (Region 1 and the early Region 2 periods). Low elasticity funds are in the bottom quartile of fund elasticity in this period and have a median elasticity of 2.5. High elasticity funds are in the second to fourth quartile of fund elasticity in this period and have a median elasticity of 11.3. The blue line shows the AUM-weighted fraction of high elasticity funds that are constrained. The red line shows the AUM-weighted fraction of low elasticity funds that are constrained. Finally, the black line shows the AUM-weighted fraction of all funds that are constrained; note that it is mechanically closer to the blue line because the high-elasticity group comprises 75 percent of all fund assets. The figure shows that funds that were previously more elastic in the earlier period became constrained faster in the later period than funds that were previously less elastic. Note that by April 2023, over 90 percent of funds (weighted by AUM) are deemed constrained. After the resolution of the debt ceiling crisis, the supply of T-bills increased by 473 billion and the RRP-Bill spread shrank to nearly zero (April 28th, 2023, to July 31st, 2023). In response to this increase in the supply of T-bills, the fraction of constrained funds with an inelastic demand for T-bills fell sharply compared to that of more elastic funds.

Figure 8
RRP and T-bill Supply
April 2023 to July 2023

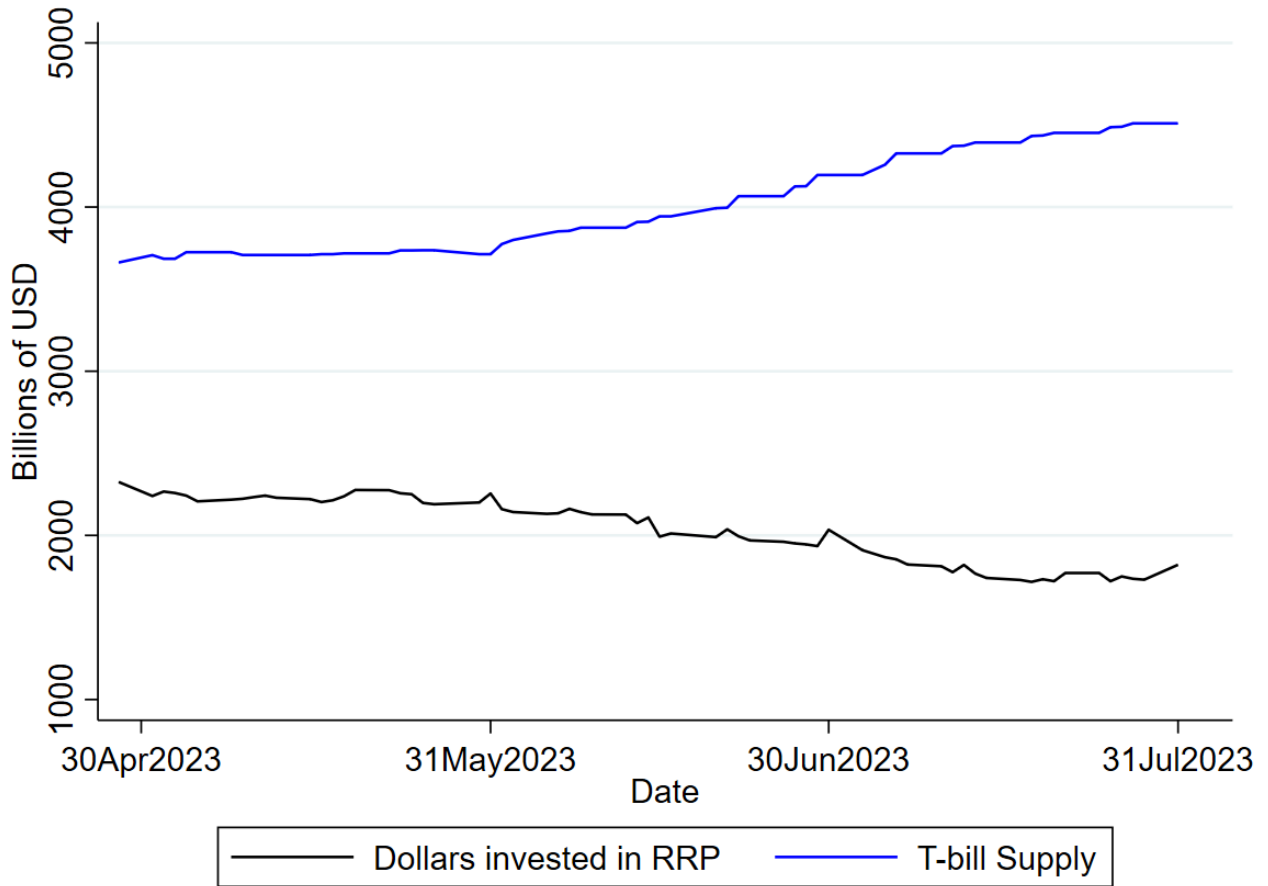


Figure 8 shows the supply of T-bills and dollars invested in the RRP facility from April 2023 to July 2023. Over this period, T-bill supply increased by 848 billion dollars and the RRP facility experienced outflows of 504 billion dollars, which is 59 cents for every dollar of increased T-bill supply.

Figure 9
Average Elasticity of Unconstrained Money Funds
January 2022 to July 2023

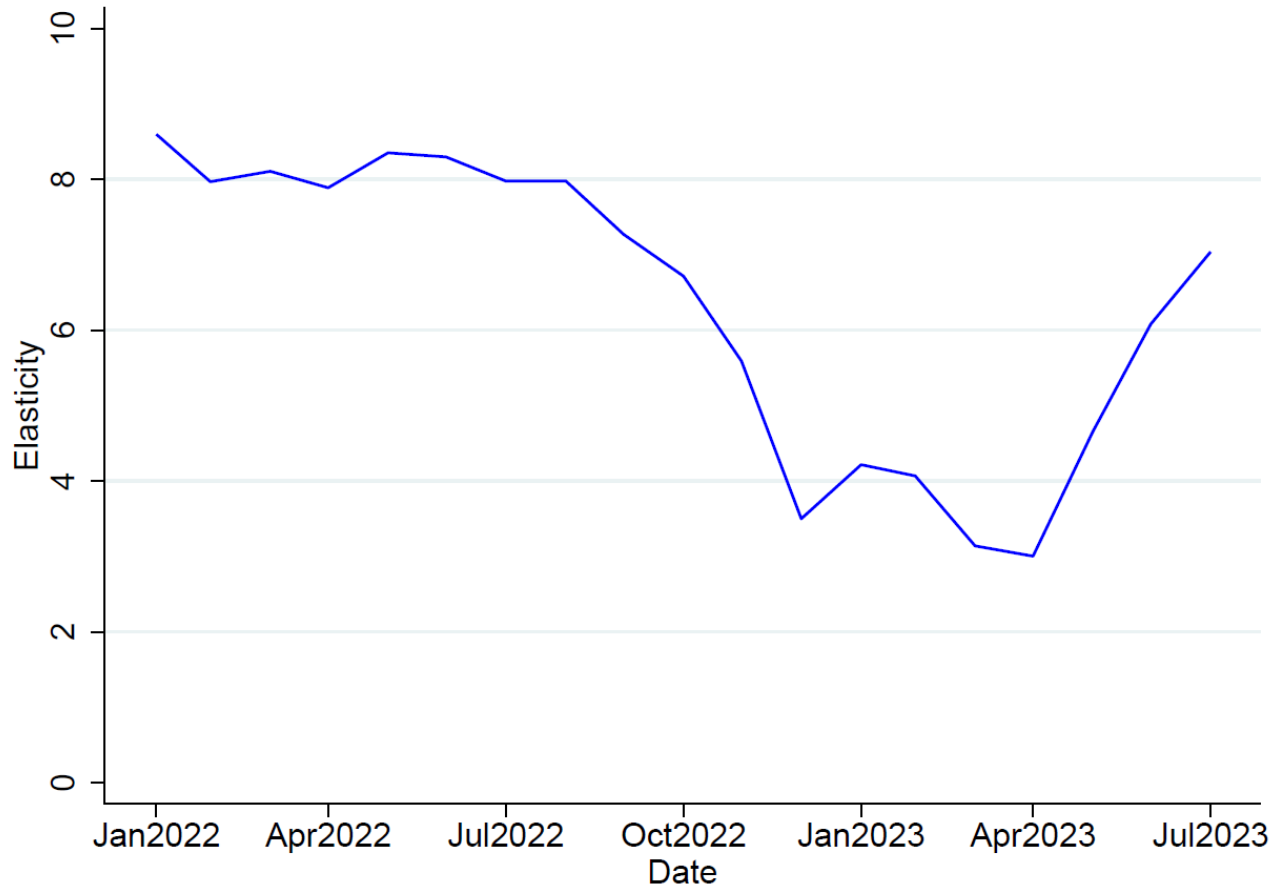


Figure 9 plots the average elasticity of unconstrained money funds weighted by their T-bill holdings. We estimate fund-level elasticity of demand for T-bills from September 2013 to March 2021 at a monthly horizon as described in the text and shown in Figure 6. The average elasticity of unconstrained money decreases sharply as the more elastic funds sell their T-bill holdings faster and become constrained.

Figure 10
Change in RRP-Bill Spread and T-bill Supply
September 2013 to April 2023




Figure 10 shows a scatter plot for monthly changes in the RRP-bill spread and the supply of T-bills. The points are grouped by whether T-bills are not scarce (September 2013 to March 2022) or scarce (April 2022 to April 2023). The black line shows a linear fit of the data for when T-bills are not scarce (black dots) and the blue line similarly shows this for when T-bills are scarce (blue squares).

Figure 11
T-Bill and RRP Yields by Maturity
May 2023




Figure 11 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 July 2023

	Mean	p50	Std. Dev	Min	Max	<i>N</i>
E[RRP Rate] 1M in Bps	104	27	135	-13	531	2,374
T-Bill Rate 1M in Bps	101	26	132	-11	571	2,374
RRP-Bill Spread in Bps	3	1	18	-61	166	2,374
AUM in Millions	28,448	14,242	40,884	16	281,933	12,487
Days to Maturity	33	34	14	1	103	12,487
Number of Funds	105	104	5	95	114	119
Weight T-Bills in Pct	11	3	18	0	100	12,487
Weight RRP in Pct	10	0	17	0	99	12,487

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 July 31st, 2023. 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 118 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 July 2023

	Dep Variable: Investor Flow Pct			
	(1)	(2)	(3)	(4)
Fund Return	3.216*** (0.70)	6.809*** (1.12)	5.044*** (1.08)	6.384*** (1.34)
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%	11%
N	12,342	12,342	12,342	4,054

Table 2 estimates the sensitivity of investor flows to money funds returns for our sample from September 2013 to July 2023. 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 6.38 percentage-point increase in AUM due to investor flows over a quarter. Standard errors are clustered by fund.

Table 3
Aggregate Elasticity of Substitution
September 2013 to March 2021

	Δ Spread	Δ T-Bill Portfolio Weight		
	First Stage	OLS	IV	IV Quarterly
Δ T-bill Supply	-0.40*** (0.12)			
Δ Spread		-0.94*** (0.28)	-8.00*** (2.65)	-6.13*** (1.62)
N	90	90	90	30

Table 3 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. The T-bill supply is the privately available supply of T-bills (outstanding notional minus T-bills held by the Federal Reserve). The first column shows the first stage regression: a 1 percent increase in T-bill supply decreases the RRP-Bill spread by 0.4 bps. The second column presents an OLS regression: a 10 bps increase in the spread is associated with a 9.4 percent decrease in money fund T-bill holdings, which is 1.2 percent 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 T-bills to estimate a causal effect of changes in spread on T-bill portfolio weight. We find that a 10 bps increase in the spread due to a decrease in T-bill supply causes money funds to decrease their T-bill portfolio weight by 80 percent, which is 10.1 percentage points of portfolio weight. The fourth column shows the estimate at a quarterly frequency, where we sample the data at quarter-ends (note that the observation count decreases by a factor of 3). Standard errors are robust.

Table 4
Effect of Changes to T-bill Supply on RRP-Bill Spread
September 2013 to April 2023

	Dep Variable: Δ Spread		
	(1)	(2)	(3)
Δ T-bill Supply	-0.770*** (0.25)	-0.351*** (0.11)	0.078 (0.27)
Scarce T-bills		0.079 (0.08)	
Scarce T-bills \times Δ T-bill Supply		-4.99*** (2.00)	
Constrained Share			0.141 (0.10)
Constrained Share \times Δ T-bill Supply			-6.67** (2.87)
Adjusted R^2	6%	30%	27%
N	115	115	115

Table 4 shows the effect of changes in T-bill supply on the RRP-bill spread. This is a monthly time-series regression. T-bill supply is the privately available supply of T-bills. Scarce T-bills is an indicator variable equal to 1 for April 2022 to April 2023 and 0 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 1 shows that an unconditional 1 percent decrease in T-bill supply causes the RRP-bill spread to increase by 0.77 bps. Column 2 shows that the effect of changes in T-bill supply is much larger when T-bills are scarce: a 1 percent decrease in T-bill supply causes a 5.3 bps increase in the spread. Column 3 uses a continuous version of the scarcity measure: when half of the money funds are constrained, a 1 percent decrease in T-bill supply causes the RRP-bill spread to increase by 3.3 bps.

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

Table B1
Aggregate Elasticity of Substitution:
Alternative Instrument for RRP-Bill Spread
September 2013 to March 2021

	Δ Spread	Δ T-Bill Portfolio Weight		
	First Stage	OLS	IV	IV Quarterly
Δ Total Debt	-1.33*** (0.12)			
Δ Spread		-0.94*** (0.28)	-9.91*** (2.83)	-6.38*** (1.65)
<i>N</i>	90	90	90	30

Table B1 estimates the aggregate elasticity of substitution as in Table 3 but now uses changes in the total marketable debt of the U.S. government, rather than changes in T-bill supply, as an instrument for changes in the RRP-bill spread. The first stage coefficient is larger because changes in total debt are less volatile than changes in T-bills. However, the IV estimates are very similar to those in Table 3.