Working Papers

FEDERAL RESERVE BANK OF BOSTON™

No. 24-7

The Effect of Primary Dealer Constraints on Intermediation in the Treasury Market

Falk Bräuning and Hillary Stein

Abstract:

Using confidential microdata, we show that shocks to primary dealers' risk-bearing constraints have significant effects on the US Treasury securities market. In response to tighter constraints, dealers reduce their Treasury positions, triggering a reduction in aggregate turnover and an increase in bid–ask spreads. These effects are more pronounced in securities that contribute more to the utilization of risk constraints. The impaired intermediation also affects Treasury yields, amplifying the yield response to net demand shifts. Moreover, tighter dealer constraints weaken Treasury auction outcomes: Bid-to-cover ratios decline, driven by dealers' less aggressive bidding, and the highest yield accepted by participants rises, thereby increasing the government's cost of issuing debt. Using our estimates, we back out key elasticities to show that the shadow cost of dealer constraints is as high as one-third of dealers' intermediation margin.

JEL Classifications: G10, G12, G18, G21

Keywords: Treasury market, primary dealers, intermediation, risk constraints

---

Falk Brauning (falk.brauning@bos.frb.org) is a senior economist and policy advisor in the Federal Reserve Bank of Boston Research Department. Hillary Stein (hillary.stein@bos.frb.org) is an economist in the Federal Reserve Bank of Boston Research Department.

The authors thank seminar participants at the Federal Reserve Bank of Boston for their feedback. They also thank Darrell Duffie, Joe Peek, and Matt Pritsker for valuable comments. Justin Sun provided excellent research assistance.

This version: July 2024 **https://doi.org/10.29412/res.wp.2024.07**

Introduction

The market for US Treasury securities is one of the most important financial markets in the world, with more than $33 trillion outstanding at the end of 2023. Treasuries play a crucial role in the global economy, serving as the primary means of financing the US federal government, a significant investment instrument and hedging vehicle for global investors, a risk-free benchmark for many other financial instruments, and an important tool for the Federal Reserve's implementation of monetary policy. However, despite its size and importance, the market is intermediated by a select group of market-making financial firms—typically affiliated with large banking organizations. These so-called primary dealers participate directly in the primary market for Treasuries, where the US government auctions new debt issuance; then, they resell the Treasuries to investors in the secondary markets. Primary dealers are also direct business counterparties to the Federal Reserve's monetary policy operations.

Against the backdrop of several recent disruptions in the Treasury market and increasing federal deficits, commentators have pointed out the potential vulnerability of this dealer-centric market to changes in primary dealers' constraints and their ability to efficiently intermediate the Treasury market (see, for example, Duffie, 2018, 2020).[1] Because banks comprise multiple organizational layers, these constraints might be at the bank level or individual business-unit level, and they may be driven by regulatory requirements or internal risk management. However, empirical evidence on the role of dealer constraints, regardless of the specific constraint, has been scarce due to the lack of available (micro) data needed for identification. Moreover, identification of the effect of dealer constraints on the Treasury market remains challenging because constraints respond endogenously to broader market conditions.

Our paper overcomes both challenges. We combine two detailed and rarely explored confidential sets of microdata, collected by the Federal Reserve for supervisory and monitoring purposes, that provide a unique window into primary dealers' risk limits, positions, turnover, and trading income.[2] Furthermore, we exploit two separate causal identification frameworks to analyze the effect of constraints on various business levels of the dealer-bank. First, we use a difference-in-differences strategy to study the impact of the supplementary leverage ratio (SLR), a regulatory constraint that targets the overall bank balance sheet. Second,

[1] For example, `https://www.wsj.com/finance/why-treasury-auctions-have-wall-street-on-edge-8385f15e`.

[2] As we explain in the data section, we use supervisory FR VV-1 data on risk limits of the trading desks of large banking organizations supervised by the Board of Governors of the Federal Reserve. We also use the confidential FR 2004 data that are used to monitor the performance of the primary dealers and the condition of the US government securities market. Although the FR 2004 does not constitute supervisory information, participation is required to obtain primary dealer status.

following Gabaix and Koijen (2020), we study the impact of internal risk-holding constraints (value-at-risk, or VaR, limits) at the trading-desk level by examining the aggregate impact of idiosyncratic shocks. By construction, these idiosyncratic shocks are exogenous to the broader market conditions that govern return dynamics and liquidity, yet they affect aggregate prices and quantities due to the granularity of the Treasury dealer market.

Our key results are as follows. Consistent with a model of a spread-charging intermediary with costly position holdings, we find that primary dealers reduce their Treasury holdings when they are faced with tighter constraints. The reduction in positions is accompanied by lower liquidity in the Treasury market because turnover decreases while bid–ask spreads increase. These core findings hold across the two constraints that we analyze, and they are economically sizable. First, when a bank's SLR decreases by 1 percentage point relative to other banks, the bank-affiliated primary dealer sees a 9.1 percentage point lower gross Treasury position, a 7.4 percentage point lower turnover, and a 9.8 percentage point stronger increase in margin. Second, when the primary dealers' trading desks, in aggregate, face a one-standard-deviation tighter risk-limit shock, they experience a 2.1 percent fall in net position and a 1.7 percent fall in turnover, while market-level bid–ask spreads increase by 2.4 percent.[3]

We verify the robustness of our results using differently constructed limit shocks, and we show that using raw limit changes leads to an attenuated and insignificant effect, underscoring the need to use idiosyncratic limit shocks that are orthogonal to general market developments. Moreover, when we study the risk-sensitive VaR constraints, we find that the reduction in exposure after tighter limit shocks is greater in securities with higher past return volatility—a key contributing factor to VaR-limit usage. Consistent with such a risk-sensitive adjustment, VaR-limit shocks have no significant effects on positions in Treasury bills (T-bills), which have a very low return variability; all results are concentrated in notes and bonds. On the other hand, we find that T-bills do react to changes in risk-insensitive SLR constraints.

How costly are the constraints imposed on primary dealers' exposure to Treasuries? The equilibrium condition of our model shows that these costs can be backed out from the bid–ask spread elasticities of turnover and positions. Using the estimates from the analysis of VaR limits and the SLR constraint, we obtain spread elasticities of turnover that are in the same ballpark: from –0.71 to –0.76. These estimates from our data are substantially lower (in absolute value) than the optimal elasticities that would prevail under profit maximization without constraints. We estimate larger spread elasticities of dealer positions, from –0.87 to –0.93; thus, positions respond more strongly to spread increases than to turnover. Taken together, our estimates of these spread elasticities imply a sizable economic cost of the position

[3]As we explain later, the SLR targets gross exposure, while VaR constrains net exposure.

constraints, with shadow costs accounting for about 26 to 33 percent of the intermediation spread, or about \$2.4 billion to \$3 billion per year.

Because the identification method used in our trading-desk-level analysis allows us to aggregate effects across banks, we also examine aggregate effects on Treasury prices and auctions. We show that changes in risk limits do affect yields in the secondary market. Specifically, we find that yields rise when dealers are more constrained against a backdrop of heightened net Treasury demand. We proxy Treasury demand by using declines in the effective duration of mortgage-backed securities (MBS), a connection explained by heightened refinancing expectations. As Hanson (2014) shows, when a large number of MBS refinance, bond investors will put the received funds into the Treasury market. This increase in net demand for Treasuries drives up prices, thereby lowering yields. We show that these effects are significantly stronger (about one-third greater) when dealers' intermediation capacity is constrained due to tighter risk constraints. Our results are robust to using foreign interest changes as demand shifters.

In addition to the liquidity and yield effects in the secondary market, we also find that primary dealers' constraints affect outcomes in the primary market. Specifically, we show that in response to a one-standard-deviation tighter risk limit, primary dealers bid less in Treasury auctions, as measured by their 5 percent lower bid-to-cover ratio. Non-primary dealers, on the other hand, experience no change in their bid-to-cover ratios, lending additional credibility to our results. Moreover, we find that when dealer risk limits are tight, the high yield of the auction—that is, the yield that clears the auction—increases significantly, by about 1 percent. This constitutes an increase in the government's cost of financing itself.

Related Literature Our findings provide direct support for theoretical contributions that highlight the role of constrained intermediaries in asset pricing (see, for example, He and Krishnamurthy, 2013; Brunnermeier and Sannikov, 2014). Like seminal work by Bernanke and Gertler (1989), these macro-finance models emphasize the general role of equity constraints of financial intermediaries in driving asset pricing cycles, thereby affecting the business cycle. They also follow in the tradition of Grossman and Miller (1988) in emphasizing how the risk-bearing capacity of intermediaries amplifies price responses to upswings in supply and demand.

A host of papers seek to provide empirical support for these theories. One strand of the literature examines return reversals and price surges, a phenomenon consistent with these theories (see, for example, Hendershott and Seasholes, 2007; Hendershott and Menkveld, 2014; Nagel, 2012).[4] Our approach allows us to more directly tie intermediary constraints to both

[4] See Duffie (2010) for a more complete review.

price dynamics and quantity responses because we employ detailed data on the specific actors in the Treasury market (primary dealers) and their constraints (for example, VaR constraints). There is also a segment of the literature that empirically ties a reduction in dealer balance sheet capacity to prices of risky assets. For example, Adrian, Moench and Shin (2010) and Adrian, Etula and Shin (2015) use publicly available bank balance sheet data to forecast asset prices. He, Kelly and Manela (2017) use a series of predictive time-series regressions to show that primary dealers' capital ratios have significant explanatory power for cross-sectional variation in expected asset returns. In recent work that also examines the Treasury market, Duffie et al. (2023) show that Treasury markets experience higher-than-expected illiquidity when dealers have higher inventories. Favara, Infante and Rezende (2022) study the effect of leverage constraints, and He, Nagel and Song (2022) incorporate an SLR constraint in a model of Treasury market liquidity. Our analysis has two distinct advantages over that segment of the literature. First, while those papers approximate intermediaries' limited inventory-carrying capacity with their inventory positions, we observe a direct measure of risk-bearing constraints. Second, our granular identification approach directly addresses the primary endogeneity concern in this literature: Capacity utilization and prices are both affected by market conditions. While much of the literature neglects to address this concern, Duffie et al. (2023) do seek to address it by orthogonalizing capacity utilization to yield volatility. Our granular identification approach is not subject to the concern that yield volatility might not be a sufficient control.

Our paper is also closely tied to the literature on government bond auctions and their effects on dealer inventory and debt prices. Fleming, Nguyen and Rosenberg (2023) highlight that Treasury issuance is the main driver of dealers' weekly inventory changes, finding that dealers are compensated for inventory risk by means of subsequent price appreciation of their holdings. Similarly, Lou, Yan and Zhang (2013) document that Treasury prices dip before auctions and recover afterward, and Beetsma et al. (2016) do the same for Italian debt. As in the return reversal literature, these papers all point to dealer balance sheet constraints as an underlying mechanism. We take this a step further by empirically showing the effects of dealer constraints using novel microdata.

The remainder of the paper is organized as follows. Section 1 introduces a conceptual framework to guide our empirical analysis. Section 2 discusses our data sources. Section 3 presents the empirical results and combines the insights into estimates of the shadow cost of the constraints. Section 4 concludes.

1 Conceptual Framework

To understand the implications of primary dealers' constraints in driving Treasury market liquidity, we present a basic model that formulates the key channel and from which we derive testable predictions. The model is an adapted version of the one in Barbiero et al. (2024).

Consider a representative dealer that intermediates demand and supply for Treasuries. The dealer charges a spread on its intermediation activities and can hold nonzero Treasury exposure at the end of the period. The model is set up in one period. The price of the Treasury is P, and we assume a symmetric bid–ask spread of $2s$ around the mid-price P. We work with the log price throughout: $p \equiv \log(P)$.

The bank takes as given the demand for Treasuries, $D(p+s)$, with $D' < 0$.[5] The supply of Treasuries is given by $S(p-s)$, with $S' > 0$. Let δ be the Treasury position that the bank holds at the end of the period, which is simply the difference between Treasuries sold and purchased. In this static model, the bank sells all this exposure in the future at $\mathbb{E}[p_1] = 0$. However, a generic marginal cost λ is imposed on the position, which can be motivated by a myriad of factors, including risk management and regulation, as we discuss later.[6]

The bank's expected profit is equal to the sum of the total margin from intermediation and the expected return of its exposure. The bank takes prices as given and chooses a spread and exposure to maximize expected profits:

$$\max_{s,\delta} \ \pi = s\left(D(e+s) + S(e-s)\right) - p\delta - \lambda\delta \tag{1}$$

$$\text{s.t. } \ \delta = D(e+s) - S(e-s), \tag{2}$$

where $\lambda > 0$ parameterizes the bank's marginal cost of holding nonzero net positions. A larger λ means that it is more costly for the bank to hold a nonzero exposure. Note that the minus sign before the term $p\delta$ comes from the fact that we normalize the expected future log price to zero, so the log return is $\mathbb{E}[p_1] - p = -p$. We also define turnover as the sum of supply and demand: $t = D(e+s) + S(e-s)$. Further, without loss of generality, fix $p = 0$, meaning that, in expectation, there are zero profits from holding positions, such that we focus only on the intermediation business of the bank. The first-order condition of the general problem is then given by

$$t + st' = \lambda\delta', \tag{3}$$

which says that the marginal intermediation income from increasing the spread needs to

[5]Note that $\tilde{D}(P) = \tilde{D}(e^{\log(P)})$, so any demand function $\tilde{D}(P)$ can be expressed as $D(p)$ with $D = \tilde{D} * p$. Also note that $\log(P(1+s)) = p + \log(1+s) \approx p + s$.

[6]In the appendix, we consider the case in which a convex cost is imposed.

be equal to the change in the position resulting from the spread increase multiplied by the marginal cost of holding the position. We use this key insight in the final section of this paper to compute key elasticities and to estimate the marginal cost of the constraint.

For now, to obtain a closed-form solution that guides our empirical analysis, assume linearity (in the log price) of the demand and supply functions of the form $D(p+s) = a-b(p+s)$ and $S(p-s) = c + d(p-s)$, with $a, b, c, d > 0$. Solving the first-order condition for the optimal spread and deriving the optimal exposure and turnover gives:

$$s^* = \frac{(a+c) + \lambda(b-d)}{2(b+d)} \tag{4}$$

$$\delta^* = \frac{2(a-c)(b+d) - (a+c)(b-d) - \lambda(b-d)^2}{2(b+d)} \tag{5}$$

$$t^* = \frac{(a+c) - \lambda(b-d)}{2}. \tag{6}$$

We build intuition for Equation 4 by recalling that the spread is chosen based on its dual impact on profit. First is its impact on intermediation income: A larger spread leads to a higher intermediation profit, all else being equal, though it distorts supply and demand. Second is its impact on holding cost: Because the spread distorts supply and demand, the bank may need to hold a larger position to clear the market, which is costly. An increase in supply and demand $(a + c)$ increases the size of the market, which increases the optimal spread due to the intermediation motive. However, an increase in the combined elasticities of supply and demand $(b + d)$ makes the spread increase more distortionary, decreasing the optimal spread. The position δ that the bank needs to take to clear the market decreases with the slope differences in supply and demand $(b - d)$. Thus, due to the holding cost motive, the optimal spread the bank charges rises with the slope differences for a given λ.

To ensure that all equilibrium quantities (D^* and S^*) are positive, we need the following parameter restrictions:

$$\frac{ad - cd - 2bc}{d(b-d)} < \lambda < \frac{ab + 2ad - bc}{b(b-d)}.$$

Without loss of generality, we can simplify to the case in which $c = 0$. It is clear that because λ is positive, it must be the case that $b > d$.

We then derive the following derivatives that are at the core of our analysis:

$$\frac{\partial s}{\partial \lambda} = \frac{(b-d)}{2(b+d)} > 0 \tag{7}$$

$$\frac{\partial \delta}{\partial \lambda} = -\frac{(b-d)^2}{2(b+d)} < 0 \tag{8}$$

$$\frac{\partial t}{\partial \lambda} = -\frac{(b-d)}{2} < 0. \tag{9}$$

Thus, when it becomes more costly to hold Treasury exposure, the bank will increase the spread and reduce its exposure, and turnover will decline. In other words, an impairment in intermediation and Treasury market liquidity occurs. These three comparative statics serve as core testable predictions that guide our empirical analysis throughout the remainder of this paper.

What types of constraints do primary dealers face that make their Treasury exposure costly—in our highly stylized model captured by the parameter λ? As mentioned earlier, we focus on internal constraints imposed by risk-management practices as well as regulatory constraints, specifically those imposed by banking regulation. Banking regulations are relevant because the most important primary dealers are bank affiliated and therefore subject to banking regulation, which has been considerably expanded and tightened after the 2007–2008 financial crisis.

Because banks are composed of different businesses, different regulatory constraints may target the overall banking organization or individual business units. Furthermore, constraints may be of different natures—they may be risk sensitive or risk insensitive, and they may constrain different exposures, such as gross or net exposures. For example, capital regulation typically applies to the overall bank balance sheet, and there are both risk-weighted capital constraints and risk-insensitive leverage ratios. Other constraints target specific activity centered in certain business units, such as restrictions on trading desks. Generally, however, the constraints are intended to allocate resources (equity capital) across different bank businesses according to managerial or regulatory preferences.

In this paper, we focus on two key regulations—both enacted in response to the 2007–2008 financial crisis—to understand the effects of dealer constraints on the Treasury market. First, we study the effects of the supplementary leverage ratio (SLR), the US implementation of the leverage ratio outlined in the Basel III accords. The SLR is a risk-insensitive capital ratio that applies generally to large banks with more than $250 billion in total consolidated assets. Current regulation requires banks to hold a minimum ratio of 3 percent of Tier 1 capital measured against their total leverage exposure.[7] Total leverage exposure includes both on- and off-balance-sheet positions. Importantly, exposure also includes *gross* positions in securities held by bank-owned primary dealers. As a result, SLR-constrained banks may have to reduce or not be able to expand their trading assets, potentially impairing Treasury market making and liquidity.

Second, we focus on constraints on banks' trading activity and trading-desk exposure. Specifically, the Volcker rule set forth in the Dodd–Frank Act generally restricts banking entities from engaging in proprietary trading, except for certain purposes, such as market making.[8] Importantly, the rule generally also requires banking entities to establish internal mechanisms to ensure and monitor the risks from trading activity. For the largest banks, the rule includes reporting requirements of quantitative measurements to the supervisors. Such measurements include internal risk limits and usage. Although US Treasury securities are exempt from Volcker rule proprietary trading restrictions, in practice, all desks report these internal limits, even desks that solely trade Treasuries. Typically, such limits are formulated as VaR or sensitivity constraints that effectively restrict desks' *net* positions depending on the riskiness of the assets. Thus, while the Volcker rule does not directly restrict dealers' Treasury exposure, the associated risk-monitoring and reporting requirements provide a unique and detailed view into the internal risk limits of primary dealers.

Figure 1, Panel (a) shows the average SLR and VaR-limit usage, indicating substantial variation in both series as well as periods of co-movement.[9] Using bank-quarter data, the binned scatterplot in Panel (b) shows a negative correlation between changes in VaR-limit usage and changes in the SLR; that is, periods when trading desks' limits are heavily utilized are also periods when the SLR compresses. This could be because the Treasury holdings (positions) factor into both the utilization of VaR limits as well as the SLR through its exposure measure. Or it could be that when more equity capital is available at the bank level, VaR limits are eased. As we detail later, we exploit reasonably exogenous variation in these constraints—SLR and VaR—in two complementary analyses to assess the general relevance of dealer constraints on Treasury market functioning.[10]

[7]More stringent requirements for the largest and most systemic financial institutions add a capital buffer of 2 percent, for a total of 5 percent.

[8]Under this exemption, a market-making trading desk is allowed to hold inventory to meet but not exceed, on an ongoing basis, the reasonably expected near-term demands of customers. A market-making desk may also hedge the risks of its market-making activity under this exemption.

[9]Both SLR and VaR are constrained to exceed a certain threshold (imposed by regulation or management), but the optimal target SLR and VaR may well diverge from the constraint in practice. Therefore, the SLR- and VaR-utilization ratio is not necessarily informative about how constrained a given bank is.

[10]We are not arguing that one constraint is more important than the other, and we are not conducting an empirical horse race between the two constraints. In fact, given their different natures, how binding one constraint is versus the other likely depends on the individual bank as well as market conditions.

Figure 1: Dealer Constraint Utilization: Leverage Ratio and Value-at-Risk

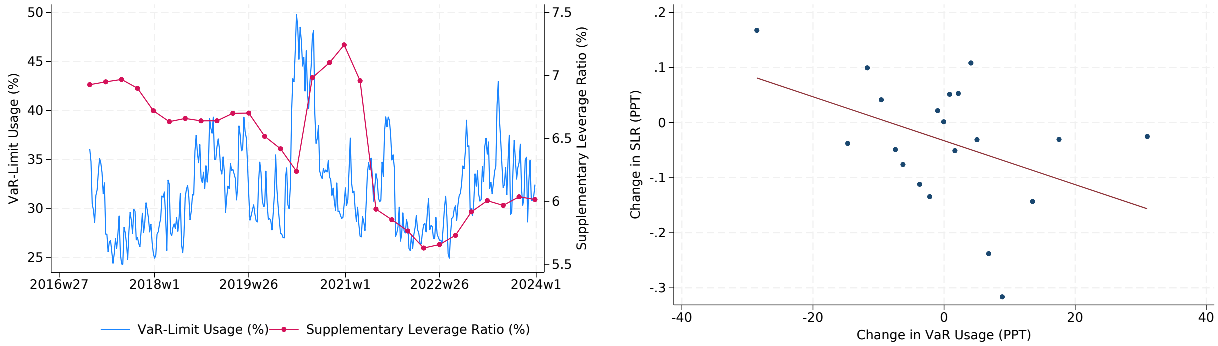

(a) Time Series of Means

(b) Binned Bank-Quarter Data

Notes: Panel (a) shows the average SLR and the average percentage VaR-limit usages of dealers in our sample based on a constant sample of dealers. A lower SLR means higher utilization of equity capital. For VaR-limit utilization, a higher percentage represents higher usage of the set limit. In Panel (b), using variation at the dealer-quarter level, we show a binned scatter plot and linear fit between quarterly SLR changes and quarterly VaR usage changes. *Sources*: FR VV-1, FR Y-9C, authors' calculations.

2 Data

Our analysis uses two main confidential data sets: the FR 2004 Government Securities Dealers Reports and the Regulation VV Quantitative Measurements (henceforth VV-1).[11]

The FR 2004 is at the core of our analysis, as it includes information on outright positions (FR 2004a) and turnover (FR 2004b)—important response variables to changes in constraints according to our theoretical predictions—of *all* primary dealers of US government securities.[12] Note that the FR 2004 does not constitute supervisory information, but the Federal Reserve uses these data to monitor the performance of its business counterparties (the primary dealers) and the condition of the US government securities market. This information enables the Federal Reserve to fulfill its responsibilities in open market operations and act as a fiscal agent for the US Department of the Treasury. The FR 2004 has been collected since the 1960s, and participation is required to obtain primary dealer status. Position and turnover data are collected weekly for various maturity buckets and security types; we focus on US Treasury securities, excluding Treasury Inflation-protected Securities (TIPS). Aggregate statistics of the FR 2004 data are available to the public on the Federal Reserve Bank of New York's website.

[11] The reported data are confidential, but the reporting form and instructions, including the list of variables collected by the Federal Reserve, are publicly available at `https://www.federalreserve.gov/apps/reportforms/Default.aspx`.

[12] A full list of the 24 primary dealers (as of March 20, 2024) is available at `https://www.newyorkfed.org/markets/primarydealers`.

Figure 4, panels (a) and (b) show the evolution of turnover and net positions (long minus short) during our sample period for relevant maturity buckets taken from the FR 2004 and aggregated across the banks in our sample.[13] Panel (a) shows that weekly turnover ranges from about \$100 billion to \$600 billion and tends to be highest for bills. Net positions are orders of magnitude smaller but still sizable, ranging from about \$5 billion to as much as \$70 billion for bills during the onset of the COVID-19 pandemic, when the Treasury Department issued large volumes of bills to fund the unprecedented fiscal stimulus. Appendix Figure A.1 breaks down the net position into long and short positions, showing that net positions are driven mostly by long positions, while short positions are generally small.

The second key data set is the supervisory FR VV-1, which provides desk-specific internal VaR limits and usage for the trading desks of banking organizations with average gross trading assets and liabilities over the preceding four calendar quarters equal to \$20 billion or more (see Barbiero et al., 2024). This data set has been collected since July 2014 to monitor compliance with the 2013 Volcker rule of the Dodd–Frank Act, which prohibits banks from engaging in proprietary trading. Banks are required to report their internal organization of trading desks, recording desk names and descriptions.[14] For each desk, a bank reports all internal risk limits as well as the usage of these limits. Risk limits fall primarily into two categories: sensitivity (for example, dollar duration) and VaR. We focus primarily on VaR limits. When desks report an upper and a lower limit, which may be relevant for two-sided sensitivity limits, we focus on the limit with the larger absolute value, which we call the "max limit." The VV-1 data also report desk-day-level profits, which we use to construct margin measures.

Fourteen of the banks in the VV-1 data are affiliated with primary dealers of Treasuries during our sample period. The primary dealers covered in our merged data sets are the most important ones affiliated with the largest bank holding companies (BHCs) in the United States. The primary dealers in our merged data account, on average, for about 75 percent of both total turnover and total net positions during our sample period. Hence, focusing on these institutions is relevant for estimating aggregate market effects. Using publicly available information, Appendix Table B.1 reports the primary dealers in our sample period and their coverage in the different data sets used in this study. Appendix Figure A.2 shows the number of primary dealers covered in the FR 2004 data during our sample period. The figure also shows the number of dealer-owning BHCs in the VV-1 data.

[13] We exclude from the graph securities with coupons due in more than 11 years but not more than 21 years and with coupons due in more that 21 years. These two maturity buckets generally have low volume.

[14] The Volcker rule defines a desk as a unit of organization that purchases or sells financial instruments for the bank's trading account. Desks are structured according to business strategies and to set and monitor trading limits, losses, and strategies. See 12 CFR § 248.3(e)(14)(ii).

Figure 2: Aggregate Turnover and Net Positions of Primary Dealers

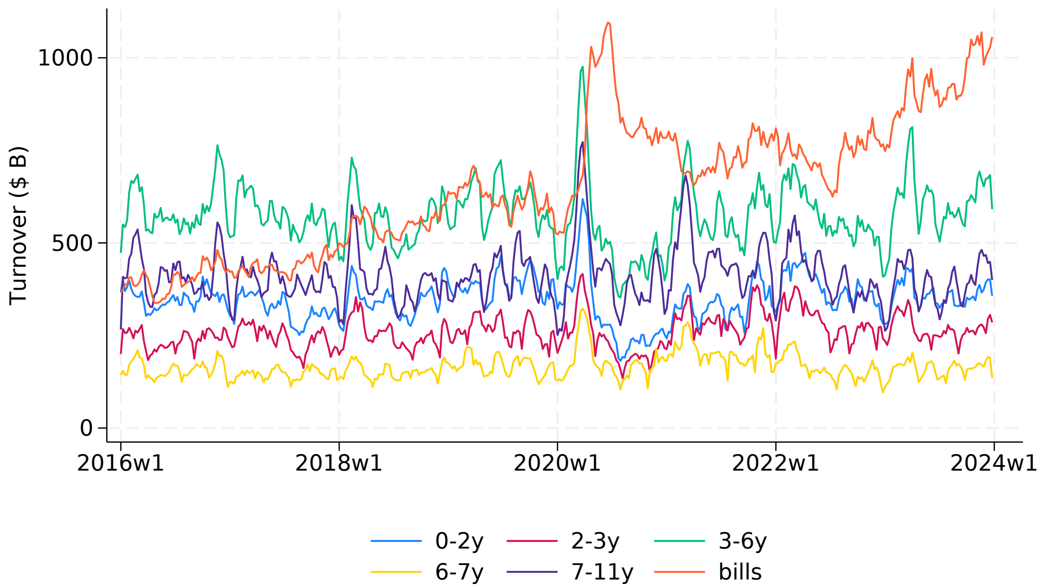

(a) Turnover $B

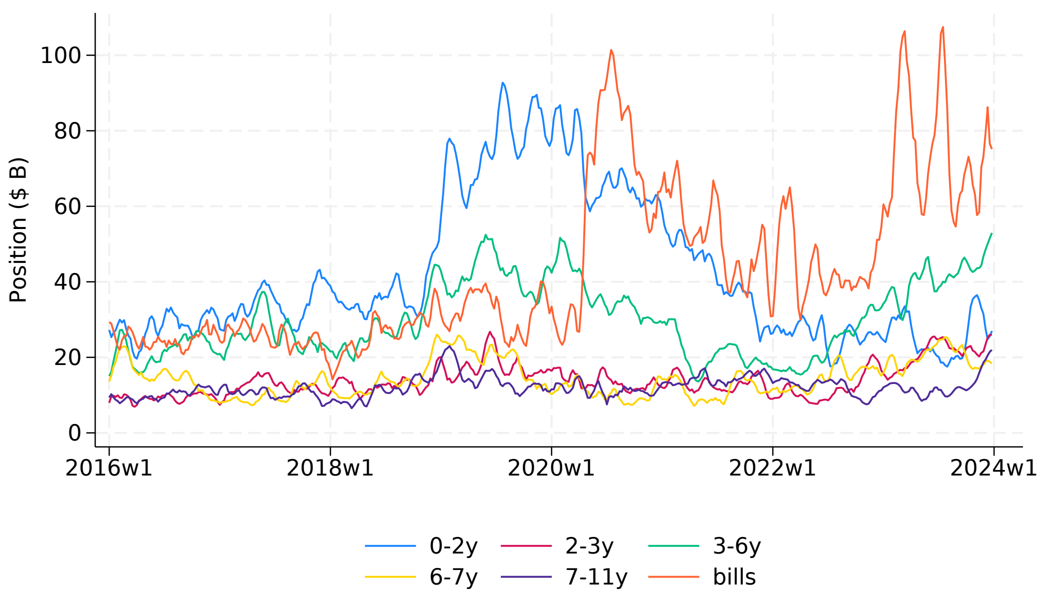

(b) Net position $B

Notes: Four-week moving averages based on weekly FR 2004 data. *Sources*: FR 2004, FR VV-1, authors' calculations.

Because desk delineations in the VV-1 data vary according to internal bank organization, rather than product, we select for the analysis desks that, based on their names and descriptions, are likely to be the primary dealer desks. A few desk descriptions actually say, "This desk is a primary dealer in US Treasuries." For banks that do not indicate which desks are primary dealers, we select the desks that trade US Treasury securities. In general, most banks have only one desk that trades US Treasury securities at a time, though there are a couple exceptions. Thus, we aggregate desks to the bank level.

We employ a few additional data sets in our analysis. First, to collect information on the capitalization of the primary dealers' BHCs, we employ the publicly available FR Y-9C data. These data provide quarterly balance sheet and income statements for the largest BHCs in the United States, including intermediate holding companies of foreign banking organizations. Crucially, the data also include regulatory capital ratios, such as the SLR. We also collect information on the BHCs' total Treasury and reserves holdings to assess their contribution to risk-insensitive exposure. For simplicity, we refer to BHCs as *banks* in the rest of this article.

Second, we use information from individual Treasury auctions. These publicly available data include security information, such as initial maturity; auction characteristics, such as the auction date and type; and details on the auction outcome, such as the high-yield and the bid-to-cover ratio as well bids submitted and accepted by primary dealers and other auction participants.[15]

Third, we collect comprehensive data on all Treasury yields and bid–ask spreads at the CUSIP–day level from the Center for Research in Security Prices (CRSP). In most of our analysis, we aggregate these data into maturity buckets to match the buckets in the FR 2004 data. We aggregate the spreads and yields into maturity buckets according to weighted means using the publicly outstanding amount of each security as weights.

Finally, in our analysis on Treasury yields, we proxy Treasury demand using two sources: the index of modified adjusted duration of US fixed-rate MBS from Bloomberg and the Euro Interbank Offered Rate (EURIBOR) changes from Haver.

3 Results

In this section, we estimate the effects of dealer constraints on the Treasury market. Guided by the model insights, we focus on the effects on dealers' positions and their intermediation ability as measured by turnover and intermediation spread. We also look at the impact on yields and study Treasury auctions. We then use identified shifts in quantities and prices to compute key elasticities that characterize the cost of the constraints.

A key empirical identification challenge is that risk constraints may respond endogenously to broad market conditions. Neglecting the endogeneity of risk limits generally leads to

[15]See `https://fiscaldata.treasury.gov/datasets/treasury-securities-auctions-data/`.

biased estimates of their effects. We identify the limit shock and its transmission using two different identification schemes that, in turn, exploit different constraints—regulatory capital constraints (SLR) and internal risk limits (VaR)—and empirical methods. Yet, as we show, they produce similar results, as we estimate similar spread elasticities of turnover and position and the related shadow cost of the constraints.

For identification of the effects of the SLR, we exploit a surprise change in the SLR regulation that excludes Treasury holdings from the SLR exposure measures. This change allows us to study the effect of the SLR using a standard difference-in-differences framework. In our second approach, we estimate the effect of VaR-limit changes on the Treasury market. For shock identification, we exploit the supervisory microdata on daily internal risk limit to construct shocks to dealers' aggregate risk-taking capacity that are driven by idiosyncratic factors—that is, they are exogenous to the aggregate dynamics in the Treasury market by construction—following the granular shock identification of Gabaix and Koijen (2020).

We discuss the identification and estimation results of each approach in the next two subsections.

3.1 Results on SLR Constraint

Identification On April 1, 2020, amid severe strains in the Treasury market at the onset of the COVID-19 outbreak, the Board of Governors of the Federal Reserve System announced that US Treasury securities would be excluded from the SLR-relevant exposure and that the change would take effect immediately.[16] As a result, primary dealers could increase their position in Treasury securities without impairing their banks' SLRs. Figure 3 shows that the exemption of Treasuries and reserves from the SLR computation boosted banks' capital ratios. Compared with a hypothetical SLR that would prevail without the exemptions, the actual SLR is about 1 percentage point higher. The figure also reveals an improvement in Treasury market liquidity that was concurrent with the SLR relief, as shown by a decline in the Merrill Lynch Option Volatility Estimate (MOVE) index.

We estimate the effects of the SLR policy change using a standard difference-in-differences approach. We expect the effect of the temporary exemption of Treasuries from the SLR to boost the holdings of lower SLR dealers more than those of high SLR dealers, as lower SLR dealers were presumably more constrained by the rule and should have had a higher marginal

[16] Deposits at Federal Reserve banks were excluded from the exposure. The stated goal of the temporary change was "to ease strains in the Treasury market resulting from the coronavirus and increase banking organizations' ability to provide credit to households and businesses." At the time of the announcement, the Fed estimated that the change would temporarily decrease Tier 1 capital requirements for bank holding companies by approximately 2 percent in aggregate. The change was in effect until March 31, 2021. See `https://www.federalreserve.gov/newsevents/pressreleases/bcreg20200401a.htm`.

Figure 3: SLR Change and Treasury Market Liquidity

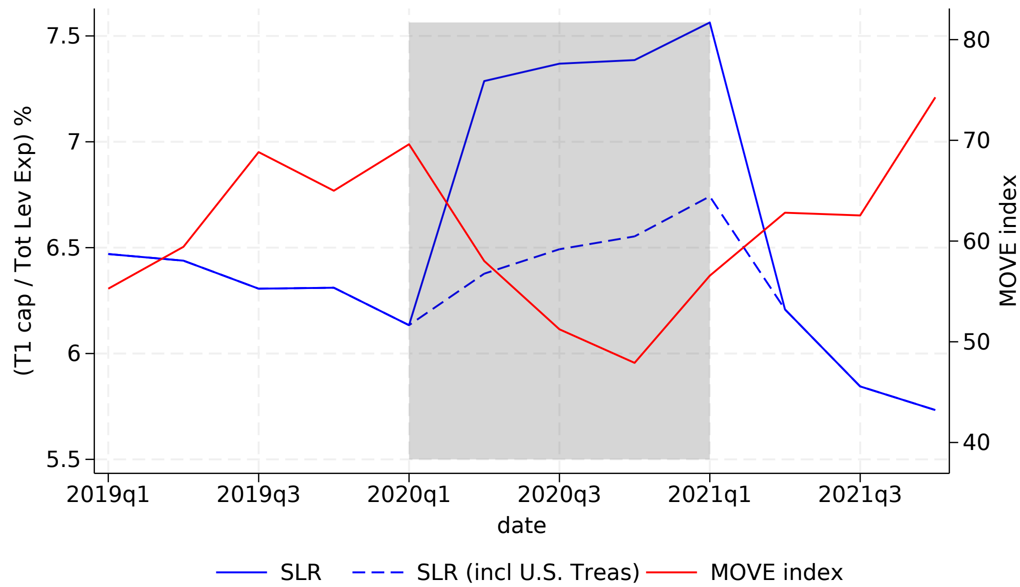

Notes: This figure shows the effect of the exemption of exposure in Treasury securities and reserves on the SLR. Reported are the average SLR (blue line) and the hypothetical SLR that would prevail without the exemption. The sample is restricted to BHCs that own a primary dealer. The MOVE index, reported on the right-hand side scale, is a market-implied measure of bond market volatility. The shaded area indicates the period when the exemption was in effect. *Sources*: FR Y-9C, Haver, authors' calculations.

(shadow) cost of adding exposure in Treasuries before the regulatory change. Our formal regression model is:

$$\text{Log(Position)}_{i,t} = \beta \text{Post-Change}_t \times \text{SLR}_i^{2019q4} + \alpha_i + \alpha_t + \epsilon_{i,t}, \tag{10}$$

where the dependent variable is the gross position in Treasuries of dealer i at the end of week t. The dummy variable Post-Change equals one as of week 14, 2020, that is, when the policy change was announced and became effective. We measure the continuous treatment variable, SLR, of dealer i as of 2019:Q4, that is, before the COVID-19 outbreak, although results are similar using 2020:Q1. Our regression also includes time fixed effects to account for common time trends in the Treasury market, such as issuance by the US Department of Treasury. Dealer fixed effects account for structural cross-sectional differences in holdings. Note that those fixed effects also absorb the level (uninteracted) effects of SLR and Post-Change.

The key parameter of interest is β, which measures the *differential* effect of the policy change depending on the dealer's SLR after partialling out the fixed effects.[17] Given the difference-in-differences setup, we cannot identify the level effect of the policy change (the aggregate effect) that is absorbed by the time fixed effects. For statistical inference, we compute standard errors clustered at the week level. Given the small numbers of dealers in the sample, we do not (additionally) cluster by dealer. Our baseline sample includes 2020:W8 through 2020:W28, a narrow window around the change, but later, we discuss the robustness of our results using different windows. Table 1 shows the summary statistics for our baseline regression sample.

Table 1: Summary Statistics of Bank-Week-Level Data

	mean	sd	p25	p50	p75	count
Turnover ($B)	209.45	122.68	131.64	203.46	275.62	171
Net Position ($B)	16.31	18.27	3.70	11.88	18.35	171
Gross Position ($B)	34.27	27.62	14.36	27.65	45.93	171
Margin (P&L over Gross Position, %)	0.08	0.12	0.00	0.04	0.12	171
SLR (%)	7.13	1.85	6.21	6.36	7.07	171
SLR Diff (%)	0.57	0.25	0.47	0.66	0.73	171

Notes: Data are at the bank-week frequency, and summary statistics are reported for the baseline sample from 2020:W8 through 2020:W28. The margin is computed as primary dealers' trading-desk profits and losses reported in the VV-1 data as a percentage of their gross position reported in the 2004 data. SLR is the 2019:Q4 SLR ratio, and SLR Diff is the hypothetical change in the 2019:Q4 SLR if Treasuries and reserves were exempt from the exposure. *Sources*: FR 2004, FR Y-9C, authors' calculations.

One key assumption for causal identification in the difference-in-differences approach is the parallel trends assumption, which holds in our application, as we show below. Another assumption requires that the policy change is not concurrent with other events that could affect the treated group differentially from the control group. This means that the change in the SLR rule must not be concurrent with other changes that differentially affect banks' Treasury trading positions depending on their SLR.[18] The high frequency of the dealers' positions and turnover observed in the FR 2004 data allows us to convincingly mitigate such concerns.[19] Specifically, we can rule out that, during the week of the SLR policy change,

[17]A comparison between BHC-affiliated dealers—that are subject to SLR regulation—and dealers that are unaffiliated with BHCs shows that the two groups were on different trends before the policy change, invalidating an important identification assumption (see below). We therefore cannot use this control group.

[18]For example, lower SLR banks may have higher exposure to commercial credit risk. If concurrent with the change in the SLR rule, the default risk of firms declined—for example, because another public policy intended to foster firm health was enacted—this could lead to a reduction in corporate-loan-loss provisions by lower SLR banks. The freed-up capital could then be redeployed to their affiliated dealers for trading purposes. In such a case, we would spuriously attribute a change in trading by lower SLR banks to the change in the SLR rule, when, in reality, it could come from changes in corporate credit risk.

[19]Most financial statements available to the public, such as those for public firms (Compustat) or commercial banks (Call reports), are at the quarterly frequency.

other relevant policy changes were announced or implemented, or that any relevant financial market developments or broader economic events occurred. For example, the potentially relevant Primary Dealer Credit Facility (PDCF) was announced on March 17, 2020, and began operating on March 20, which was two weeks before the change in the SLR rule. Appendix Figure A.3 reports key responses of the Federal Reserve to the COVID-19 crisis, showing that at the time of the SLR change, no other policy changes occurred.

Table 2: Effect of SLR Change on Dealers' Positions

	Log(Position)		
	Bank Level	Bank-Maturity Level	
	(1)	(2)	(3)
Post-Change × SLR^{2019q4}	-0.102***	-0.091***	
	(0.009)	(0.007)	
Post-Change × ΔSLR			0.672***
			(0.096)
Week FE	Yes		
Bank FE	Yes		
Week*Maturity FE		Yes	Yes
Bank*Maturity FE		Yes	Yes
R^2	0.969	0.940	0.940
R^2 within	0.224	0.064	0.063
N	171	1,026	1,026

Notes: The table reports the differential effects of the SLR policy on primary dealers' net positions, depending on their SLR. In column (1), the data are at the bank-week level, and the dependent variable is the logarithm of total gross position in Treasury securities of dealer i in week t. In columns (2) and (3), the data are at the bank-maturity-week level, and the dependent variable is the logarithm of gross position in Treasury securities in maturity bucket m of dealer i in week t. Post-SLR is an indicator that equals one as of week 14, 2020, and zero otherwise. SLR is the 2019:Q4 SLR ratio as a percentage. ΔSLR is the (hypothetical) increase in the 2019:Q4 SLR if the policy enacted on week 14 of 2020 applied. The sample period covers 2020:W8 through 2020:W28. Fixed effects are included as indicated in the bottom of the table. Robust standard errors are clustered at the week level and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1. *Sources*: FR 2004, FR Y-9C, authors' calculations.

Results on Positions We now present our baseline results for how dealer positions responded to the SLR policy change. Table 2, column (1) reports the estimation results of Equation (10). The coefficient estimate of β is highly significant and negative, indicating that, after the SLR policy change, dealers of lower SLR banks significantly increased their gross position in Treasuries. The estimate indicates a relative increase of 10.2 percentage points for a 1 percentage point increase in the SLR. Note again that in this difference-in-differences setup, the level effect of the policy change is absorbed by the week fixed effects. Similarly, the SLR differential effect before the policy change is absorbed by the bank fixed effects.

We next zoom in on the dynamics of the effect by estimating SLR differentials for each week; that is, instead of interacting the SLR variable with a post-change dummy, we estimate different β coefficients for each week. This analysis is important for two reasons. First, it allows us to assess how fast the differential effect materializes and how persistent it is—information that is not contained in the average differential effect estimated in Table 1, column (1). Second, estimating week-by-week differential effects allows us to test whether both the treated group and the control group were on parallel trends before the policy change—a key assumption in the difference-in-differences identification setting. The results reported in Figure 4 show that the parallel trends assumption clearly holds for the period before the policy change. Once the policy change is announced and becomes effective (week 14), the differential effect immediately becomes negative and highly statistically significant, with values smaller than –0.1. The differential effect remains significant for several weeks, until it decreases in size and becomes indistinguishable from zero at the end of 2020 (Appendix Figure A.4).

In columns (2) and (3) of Table 1, we further substantiate our baseline difference-in-differences estimate. First, in column (2), we exploit the maturity dimension of the FR 2004 data. Specifically, we estimate regression models similarly to equation (10) but without aggregating dealers' positions across maturity buckets. Thus, our unit of observation is at the bank*maturity*week level, which allows us to control for maturity*week fixed effects and dealer*maturity fixed effects.[20] Inclusion of such multi-way fixed effects allows us to further control for heterogeneity that may affect our estimates or precision. Column (2), however, shows that such additional controls do not materially affect our estimate, which is now –0.09, close to our baseline estimate in column (1). Moreover, the standard errors do not change substantially, lending additional support that our estimate is picking up a causal relationship between the policy change and the differential positions depending on SLR.

In column (3), we use a slightly different treatment variable. Instead, of the 2019:Q4–SLR ratio, we compute the (hypothetical) increase in the 2019:Q4 SLR that each dealer bank would face if Treasuries and reserves were exempt from the SLR exposure related to the policy change. The idea behind this variable is that one may argue that dealer banks that gain more from the policy change (in terms of easing, that is, increasing, their SLR) may

[20]We weight observations with the within-dealer-week share of positions in each maturity to mimic the bank-level results in column (1).

Figure 4: Differential SLR Effect on Dealer Positions

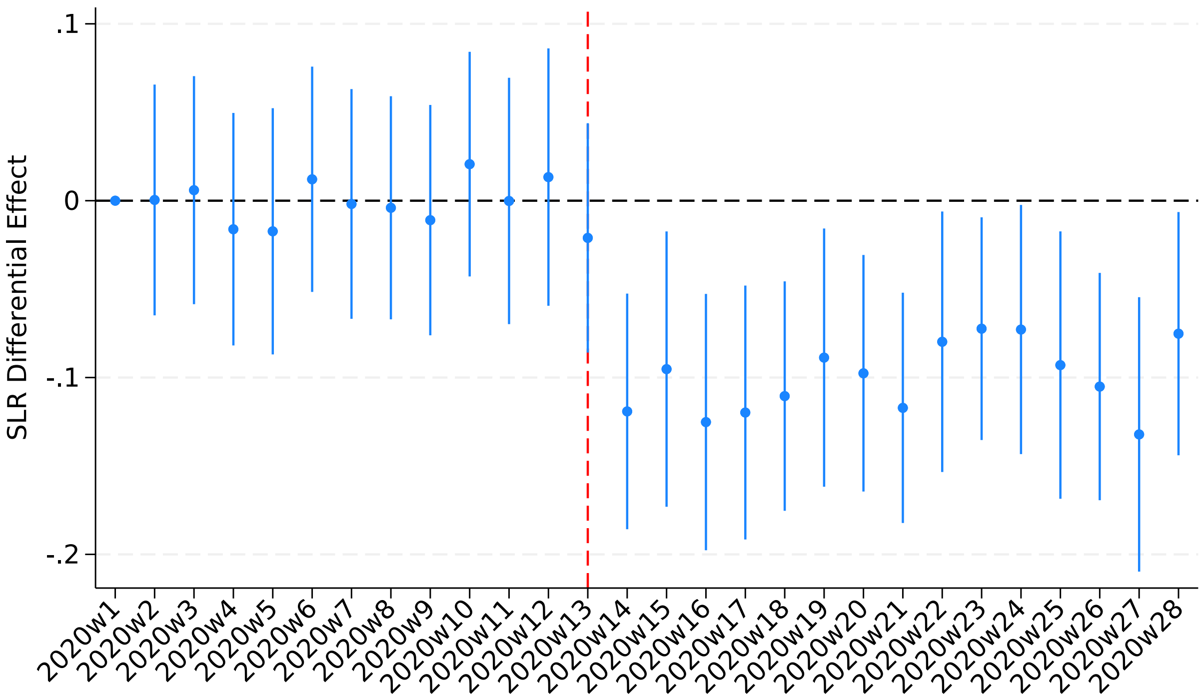

Notes: The figure shows coefficient estimates $\{\hat{\beta}_t\}$ from the following regression: $\log(\text{Position})_{i,t} = \beta_t \text{Post-SLR}_t \times \text{SLR}_i^{2019q4} + \alpha_t + \alpha_i + \epsilon_{i,t}$. Post-SLR is a dummy equal to one after 2020:W14, and zero otherwise. Positions are gross positions, in line with gross notionals entering the SLR exposure measure. The sample period runs from 2020:W1 through 2020:W26. The SLR change was announced on April 1, 2020 (week 14), and became effective immediately. The vertical red line indicates the last week before the change. The bars represent 90 percent confidence intervals based on robust standard errors. *Sources*: FR 2004, FR Y-9C, authors' calculations.

more strongly increase their Treasury positions, as opposed to banks with lower SLR ratios.[21] The significantly positive point estimate of 0.67 in column (3) supports the view that banks also (relatively) increased their dealers' Treasury positions when their SLR increased more from the policy change. Overall, this analysis confirms our key point: The SLR change had strong effects on dealers' Treasury positions.

Results on Liquidity Understanding the effects of the SLR on dealer positions is the first step in understanding the effects on market liquidity. The channel involves dealers that are constrained by their inventory holdings, thereby impairing their ability to make markets, which has adverse effects on liquidity. As a next step, we investigate how dealers' turnover in

[21]The level of the SLR and its increase under the exemptions are highly correlated in our set of relevant BHCs, with a correlation coefficient of –0.8259. That said, we are not arguing that one or the other variable is a better measure of the constraint, nor are we trying to horse-race the two variables.

Treasuries responded to the SLR policy change. Table 3, column (1) shows the estimation results of regression model (10), replacing the left-hand side variable with the logarithm of dealers' weekly turnover in Treasuries. The estimate shows that after the policy change, lower SLR banks significantly increased their turnover relative to higher SLR banks. The differential effect is about 5 percentage points per 1 percentage point increase in the SLR. Column (2) reports analogous results when we look at the maturity-bank-level data, which allows us to control for additional fixed effects, as discussed earlier. Compared with column (1), we find a somewhat larger differential effect (in absolute value) of about 7.4 percentage points.

Table 3: Effect of SLR Change on Liquidity

	Log(Turnover)		Margin	Log(Margin)
	Bank Level	Bank-Mat Level	Bank Level	Bank Level
	(1)	(2)	(3)	(4)
Post-Change × SLR^{2019q4}	-0.050**	-0.074***	0.081**	0.098*
	(0.020)	(0.025)	(0.039)	(0.050)
Week FE	Yes		Yes	Yes
Bank FE	Yes		Yes	Yes
Week*Maturity FE		Yes		
Bank*Maturity FE		Yes		
R^2	0.977	0.944	0.667	0.833
R^2 within	0.073	0.041	0.013	0.011
N	171	1,026	171	147

Notes: The table reports the differential effects of the SLR policy on primary dealers' turnover and margin, depending on their SLR. In column (1), the data are at the bank-week level, and the dependent variable is the logarithm of total turnover in Treasury securities of dealer i in week t. In columns (2) and (3), the data are at the bank-maturity-week level. In column (2), the dependent variable is the logarithm of turnover in Treasury securities in maturity bucket m of dealer i in week t. In column (3), the dependent variable is the total weekly profits of dealers' trading desks (retrieved from VV-1 data) relative to the preceding week's position (margin) and expressed as a ratio of its standard deviation. In column (4), similar results are obtained with the dependent variable being the logarithm of the intermediation margin. Post-SLR is an indicator that equals one as of week 14 of 2020 and zero otherwise. SLR is the 2019:Q4 SLR ratio as a percentage. The sample period covers 2020:W8 through 2020:W28. Fixed effects are included as indicated in the bottom of the table. Robust standard errors are clustered at the week level and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1. *Sources*: FR 2004, FR Y-9C, FR VV-1, authors' calculations.

So far, we have shown that lifting dealer constraints led to larger positions and higher turnover for more constrained dealers. Given that a binding constraint implies a higher (marginal) shadow cost of maintaining positions and intermediating Treasuries, we also expect that, once the constraints are lifted, the affected dealers' intermediation margins would decline because the marginal cost declines. In column (3), we therefore look at the response of weekly total trading profits of primary dealers' trading desks relative to the preceding week's gross position. This variable is a natural margin measure because it captures the income from trading relative to the position held. We again use the same regression model (10) but swap out the dependent variable against our margin measure. The significant positive coefficient estimate confirms our hypothesis: After the policy change, lower SLR dealers decreased their profit margins relative to higher SLR banks. Similar results hold when we look at the logarithm of margin (column 4).[22] The effects are sizable, with a differential effect of about 8 to 10 percentage points per 1 percentage point increase in the SLR. Overall, the evidence on the responses of positions, turnover, and margins is consistent with the SLR change easing dealer constraints and improving Treasury market liquidity.

3.2 Result on VaR Limits

Identification For our VaR-limit shock identification, we combine two insights from Gabaix and Koijen (2020). First, we identify dealer-level idiosyncratic risk VaR-limit changes, that is, changes in limits that are exogenous to the overall evolution of the Treasury market. Second, the granularity of the dealer-centric market—a few large primary dealers account for a large share of Treasury market intermediation—allows idiosyncratic limit changes to affect aggregate quantities. Figure 5 shows the distribution of nonzero limit changes at the bank-day level that we exploit in our analysis.

Our granular approach to isolating risk-limit shocks is based on two steps. First, we recover dealer-day-level idiosyncratic limit changes as the residuals from a saturated regression of dealer-day-level limits, controlling for time trends and dealer-specific effects. Specifically, in our baseline specification, we model the log limits changes at the dealer-day (d,τ) level using a saturated fixed-effects model:

$$\Delta \log \text{Limit}_{d,\tau} = \sum_{\substack{h=-20 \\ h\neq 0}}^{20} \gamma_h \log \text{Limit}_{d,\tau-h} + \sum_{\substack{h=-20 \\ h\neq 0}}^{20} \omega_h \text{Usage}_{d,\tau-h} + \alpha_d + \alpha_\tau + e_{d,\tau}. \tag{11}$$

To help us understand the variation in changes in (log) limits, Table 4 shows the R^2 of the different sets of fixed effects included in our baseline model, as well as the explanatory power of the full model shown in column (7). Columns (1) and (3) show that time fixed effects alone can explain only about 11.5 percent of the variation. Bank fixed effects have no explanatory power (column 2). On the other hand, column (4) shows that adding lagged limits (20 lags) to the specification explains about 24 percent of the variation (net of fixed

[22] Due to some negative profits, the sample size is smaller when we look at the log margin as the response variable, which is not defined for negative profits.

Figure 5: Distribution of VaR-Limit Changes

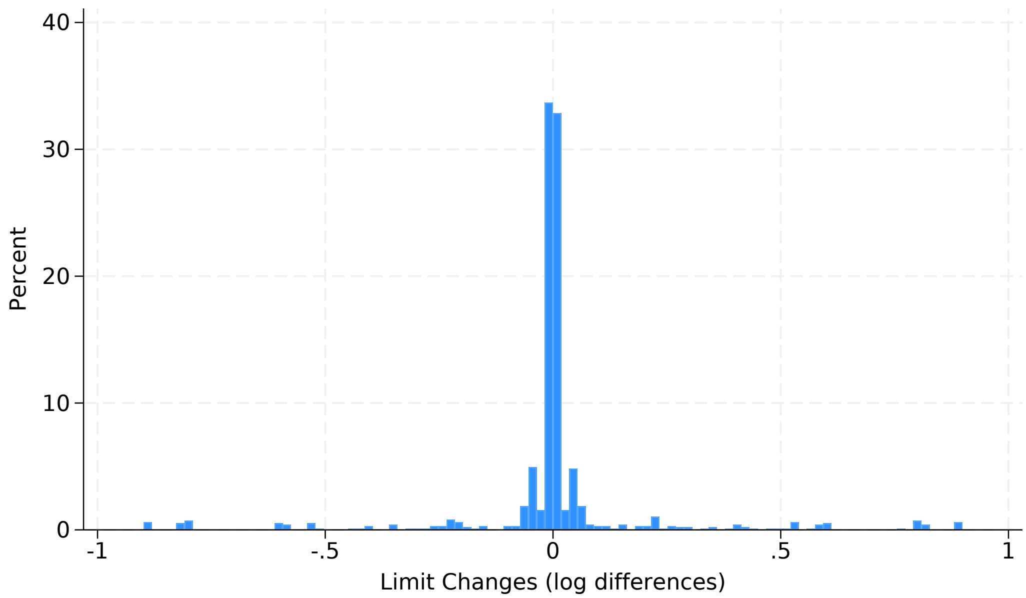

Notes: This figure shows the distribution of the log difference of VaR limits at the bank-day level. Only nonzero changes are included, and they are trimmed at top and bottom 2.5 percent. *Sources*: FR VV-1, authors' calculations.

Table 4: Variation in Log Limit Changes

	Δ Log Limit						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Time FE	Yes	No	Yes	Yes	Yes	Yes	Yes
Bank FE	No	Yes	Yes	Yes	Yes	Yes	Yes
Lagged Limits	No	No	No	Yes	No	Yes	Yes
Lagged Usage	No	No	No	No	Yes	Yes	Yes
Leads Limits+Usage	No	No	No	No	No	No	Yes
R^2	0.115	0.000	0.115	0.328	0.117	0.329	0.466
R^2 within	0.000	0.000	0.000	0.241	0.003	0.242	0.397
N	11,547	11,547	11,547	11,547	11,547	11,547	11,547

Notes: The dependent variable is the daily change in log limits. Lagged Limits and Usage include 20 lags of either variable. Leads refers to 20 leads. The table shows that, net of fixed effects, past limit changes and usage explain about a quarter of the variation in limit changes (within R^2 shown in column 6). The sample is held constant across specifications. *Sources*: FR VV-1, authors' calculations.

effects). The most saturated model, in column (7), also adds lagged usage and 20 leads of usage and limits and explains about 46 percent of the variation in limit changes.

In robustness tests, we allow for a common factor structure in the residuals of Equation

(11) that may not be captured by our extensive set of fixed effects and controls:

$$e_{d,\tau} = \lambda_d' f_\tau + \epsilon_{d,\tau}, \tag{12}$$

where f_τ is a vector of common factors, and λ_d are the associated loadings of dealer d. The residual $\epsilon_{d,\tau}$ represents the truly idiosyncratic components. We estimate the factor structure using principal component analysis of the residuals obtained from a least squares regression of Equation (11). In our baseline robustness test, we remove two common factors, but the results are robust to removing more common factors.

These idiosyncratic limit shocks may result from a variety of factors that cause desks to change their existing limits, impose additional limits, or remove existing limits. Such factors may include changes to management or trading personnel, changes to trading strategy, or changes to bank strategy that cause a reallocation of risk across divisions. In our data, we find that many of the largest shocks are driven by the effect of foreign holidays on banks with large Treasury desks in foreign countries. When these desks are closed due to a foreign holiday, it constitutes a significant reduction in intermediation capacity in the Treasury market because other desks are not dynamic enough to increase their limits in response. Despite the importance of foreign holidays, we will show that our results do not depend on them.

The residual of Equation (12) is our baseline measure of idiosyncratic (dealer-level) limit innovations. To obtain our final maturity-level aggregate limit shock, we aggregate the dealer-level limit innovations as exposure-weighted means:

$$\text{Limit Shock}_{m,t} = -\sum_b w_{b,t-1}^m \sum_{\tau \in t} \hat{e}_{d,\tau}, \tag{13}$$

where the inner sum aggregates from a daily to weekly frequency, and the weight $w_{b,t-1}^m = \text{abs}(\delta_{b,m,t-1})/\sum_b \text{abs}(\delta_{b,m,t-1})$ is, for each maturity m, the share of the net position held by primary dealer b in the total net position of all primary dealers. Thus, maturity-level shocks are constructed as an exposure-weighted mean of bank-level shocks reflecting the importance of a dealer's exposure in a given maturity bucket. The idea here is that limit shocks to primary dealers that have larger positions relative to other banks in a given maturity should matter more to that maturity.[23] Appendix Figure A.6 depicts our baseline limit shocks at the maturity-quarter level. For robust inference, all shocks are winsorized at the 2.5th and

[23]Appendix Figure A.5 shows the concentration of weights (Herfindhal index) by maturity bucket. Appendix Table B.2 shows that within bank maturity, the weights are highly stable over time. That said, note how our identification hinges on the random shocks, while exposure shares are allowed to be endogenous (Borusyak, Hull and Jaravel, 2021).

97.5th percentiles. Finally, we standardize the limit shocks in each estimation sample to have mean zero and unit variance.

We use these granularly identified limit shocks to assess the Treasury market response. Specifically, we regress three key dependent variables—net positions, turnover, and bid–ask spread—on the limit shock using our merged data set at the week-maturity level. Formally, we use the following linear regression model:

$$\Delta y_{m,t} = \beta_1 \text{Limit Shock}_{m,t} + \gamma' X_{m,t} + \alpha_m + \alpha_t + u_{m,t}, \tag{14}$$

where $y_{m,t}$ is the outcome variable, m indexes the maturity bucket, and t indexes weeks. In our baseline analysis, $y_{m,t}$ is either (1) the log of the aggregate net positions in maturity bucket m, (2) the log aggregate turnover in maturity bucket m, or (3) the log bid–ask spread associated with maturity bucket m.[24] The specification also includes week fixed effects (α_t) and maturity bucket fixed effects (α_m) to account for common time effects and structural differences in liquidity across different maturity buckets. The lagged values of the dependent variables are included in the vector of controls. For statistical inference, we compute standard errors clustered at the week level.[25]

Our baseline estimation sample covers weekly data from 2016:W1 through 2023:W52. However, given potential confounders during the COVID-19 period—including changes in bank capital regulation that we discussed earlier—we show robustness to excluding the 2020:Q1–2021:Q1 period. As we will explain later, in our baseline analysis, we focus on the maturity buckets of bonds and notes as reported in the FR 2004 data (that is, we exclude bills from our analysis). Table 5 reports summary statistics of key variables in our baseline maturity-week sample.

Positions Table 6 presents our baseline results regarding how VaR-limit shocks affect dealers' net positions. The estimated coefficients show that, in response to a tightening limit shock, primary dealers decrease their net position. The parameter estimate of –0.021 in column (1) suggests that, in response to a one-standard-deviation tightening limit shock, primary dealers reduce their net position by about 2.1 percent. Considering that a one-standard-deviation limit shock corresponds to a 0.52 percent (surprise) limit reduction (see Table 5), this effect is sizable. Column (2) shows that the effect is quantitatively similar

[24]Net positions are aggregated as the sum of absolute values of net positions across all dealers, thus measuring the total amount of dealer net position. Aggregate turnover is simply the sum of value of all transactions reported by the dealers in a given maturity. Bid–ask spreads are aggregated as weighted means across all Treasuries within the same maturity bucket.

[25]We do not also cluster by maturity bucket due to the small number of buckets, although the standard errors would be similar if we did so.

Table 5: Summary Statistics of Baseline Maturity-Week-Level Data

	mean	sd	p25	p50	p75	count
Turnover ($B)	351.39	171.40	218.15	336.19	454.52	2080
Net Position ($B)	22.53	15.75	11.80	16.88	28.08	2080
Gross Position ($B)	65.79	37.42	36.05	45.91	93.74	2080
Bid-Ask Spread (% Mid Point)	0.05	0.03	0.03	0.04	0.08	2073
Margin (Income over Net Position, %)	1.23	1.31	0.41	0.71	1.45	2073
Margin (Income over Gross Position, %)	0.38	0.36	0.12	0.24	0.50	2073
Limit Shock (%)	-0.14	0.52	-0.04	-0.01	-0.01	2080
Raw Limit Change (%)	-0.07	0.51	-0.04	0.00	0.03	2080
Yield-to-Maturity (%)	2.02	1.27	1.07	1.81	2.78	2080
Adj. MBS Duration	4.18	1.28	3.07	4.51	5.20	1935
EURIBOR Change (bps)	1.96	5.98	-0.42	0.00	0.66	2070

Notes: Positions are aggregate net positions computed as the sum of net positions across all dealers by maturity and week. Income = Spread*Turnover. MBS duration in years. EURIBOR Changes represent two-week changes. *Sources*: FR 2004, FR VV-1, CRSP, Bloomberg, authors' calculations.

when we exclude the COVID-19 period from the sample. Column (3) shows that, when using raw limit changes that are endogenous to market conditions, we find an attenuated and insignificant effect of –0.008. Thus, using limit shocks that are exogenous to broader market conditions is important for identification.

Appendix Table B.3 shows that our results are robust to differently constructed limit shocks. First, column (1) reports results when we control for the primary dealer's *past* positions in Treasuries as reported in the FR 2004 data, in addition to the past limit utilization that we include in our baseline specification. Second, in column (2), we report results when we construct the limit shocks incorporating a factor analysis of the desk-day-level limit innovations to ensure that no remaining factor structure is driving our results. Third, because part of our identification comes from foreign holidays, we also remove foreign holidays from the estimation sample of the idiosyncratic desk innovations (column 3) to show that our identification does not hinge on foreign holidays. Fourth, in an alternative approach, we remove primary dealers with foreign desks altogether from the sample (column 4). While these different ways to construct limit shocks lead to the same conclusions, as Appendix Table B.4 shows, our effects are driven mainly by the period before 2020, when our limit shocks exhibit higher variation.

In our baseline analysis, we exclude Treasury bills from the sample. In Table 7, column (1), we report results when we include bills in our analysis, but we allow for a differential limit shock effect. The estimates show that limit shocks affect only coupon bonds (non-bills), not bills, with the effect on bills about 10 times smaller in magnitude. This result is intuitive given

Table 6: Limit Shocks and Primary Dealers' Net Positions

	Log(Position)		
	(1)	(2)	(3)
Limit Shock	-0.021***	-0.021**	
	(0.007)	(0.008)	
Raw Limit Change			-0.008
			(0.007)
Lagged LHS	Yes	Yes	Yes
Maturity FE	Yes	Yes	Yes
Week FE	Yes	Yes	Yes
Excl. COVID	No	Yes	No
R^2	0.945	0.942	0.945
R^2 within	0.733	0.725	0.732
N	2,075	1,815	2,075

Notes: The table reports the effects of a one-standard-deviation tightening limit shock on primary dealers' net positions. The dependent variable is the logarithm of the absolute value of dealers' net position in Treasury securities. The data are at the maturity-bucket and week levels. The sample period covers 2016:W1 through 2023:W52. The COVID-19 period is excluded or not, as indicated in the fixed effects panel. Robust standard errors are clustered at the week level and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1. *Sources*: FR 2004, FR VV-1, authors' calculations.

that we focus on VaR limits. Indeed, Figure A.7 shows that bills, as opposed to Treasury securities in other maturity buckets, exhibit strikingly lower return variability. Therefore, bills contribute substantially less to the use of VaR constraints. As a result, it is economically rational that dealers, in response to a tighter risk limit, adjust their net positions in other maturities that contribute more to the tighter constraint because doing so will lead to the highest marginal reduction in the constraint utilization.

We can test this hypothesis directly by estimating heterogeneous effects of limit shocks on net positions depending on the variability of returns. To do so, we interact our limit-shock variable with the standard deviation of log price differences computed at the security level based on a 30-day rolling window, which we then aggregate to the maturity-bucket level, the unit of observation in our data on net positions. Table 7, columns (2) and (3) clearly show that exposure adjustments are more sensitive in maturity buckets that exhibit higher return volatility. This observation holds both when we include Treasury bills in our sample and when we exclude them. Thus, the result that dealers cut back more on exposure in securities

Table 7: Heterogeneous Effects Depending on Return Variability

	Log(Position)		
	(1)	(2)	(3)
Non-Bills × Limit Shock	-0.0154**		
	(0.008)		
Bills × Limit Shock	0.00153		
	(0.010)		
Limit Shock		-0.00987	-0.0201***
		(0.006)	(0.007)
Return Variability		0.00989**	0.00875**
		(0.004)	(0.004)
Return Variability × Limit Shock		-0.0184**	-0.0171**
		(0.008)	(0.009)
Lagged LHS	Yes	Yes	Yes
Maturity FE	Yes	Yes	Yes
Week FE	Yes	Yes	Yes
Excl. Bills	No	No	Yes
R^2	0.944	0.944	0.950
R^2 within	0.749	0.749	0.740
N	2,696	2,696	2,281

Notes: This table extends the baseline analysis of Table 6 to include T-bills (column 1). In columns (2) and (3), the table shows the differential effect of standardized limit shocks on dealer positions depending on Treasury return variability. We measure return variability by the standardized coefficient of variation of the daily log price change computed for each maturity bucket based on a 30-day trailing rolling window. The data are at the maturity-bucket and week levels. The sample period covers 2016:W1 through 2023:W52. Robust standard errors are clustered at the week level and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1. *Sources*: FR 2004, FR VV-1, CRSP, authors' calculations.

with higher return variability also holds within the set of coupon bonds.[26] Overall, this analysis lends additional support for the economic channel driving our identified empirical relationships.

Liquidity We next study the effects of limit shocks on Treasury market liquidity, as tighter limits should induce dealers to hold fewer positions, leading to lower turnover and an increase in bid–ask spreads. Table 8 reports the baseline liquidity results that confirm these model predictions. Column (1) shows that a one-standard-deviation limit shock leads to a decrease in turnover of about 1.7 percent. Given the average weekly turnover of about $351 billion,

[26]Appendix Figure A.8 reports our baseline effects for each maturity bucket separately.

this corresponds to a decline in turnover of about \$6 billion. The effect is robust to excluding the COVID-19 period (column 2). Tighter limits also increase bid–ask spreads. Column (3) shows that, in response to a one-standard-deviation limit shock, bid–ask spreads increase by about 2.4 percent. The results are robust to excluding the COVID-19 period (column 4). We find qualitatively similar results when looking at measures of Treasury market liquidity other than bid–ask spreads and turnover.[27]

Table 8: Limit Shocks and Liquidity: Turnover and Bid–Ask Spreads

	Log(Turnover)		Log(Bid-Ask)	
	(1)	(2)	(3)	(4)
Limit Shock	-0.017*	-0.019**	0.024***	0.027***
	(0.009)	(0.009)	(0.007)	(0.008)
Lagged LHS	Yes	Yes	Yes	Yes
Maturity FE	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes
Excl. COVID	No	Yes	No	Yes
R^2	0.866	0.872	0.943	0.938
R^2 within	0.017	0.013	0.564	0.571
N	2,075	1,815	2,061	1,803

Notes: The table reports the effects of a one-standard-deviation tightening limit shock on turnover and bid–ask spreads. In columns (1) and (2), the dependent variable is the logarithm of primary dealers' total turnover in Treasury securities. In columns (3) and (4), the dependent variable is the logarithm of the bid–ask spread (relative to the midpoint) of Treasury securities. The data are at the maturity-bucket and week levels. The sample period covers 2016:W1 through 2023:W52. The COVID-19 period is excluded or not, as indicated in the fixed effects panel. Robust standard errors are clustered at the week level and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1. *Sources*: FR 2004, FR VV-1, CRSP, authors' calculations.

Yields We have shown that limit shocks affect primary dealers' net positions and impair intermediation in the secondary market, as bid–ask spreads increase and turnover decreases. What about the effects on yields? Our framework in Section 1 abstracts away from endogenous price movements by fixing the price and focusing on partial equilibrium results. However, as the general equilibrium model in Barbiero et al. (2024) shows, limit shocks should amplify price movements in the direction predicted by net demand. The intuition here is that when dealers are more constrained, they must receive greater compensation for intermediating the market, and this comes in the form of a higher return. Since dealers must take a larger

[27]For example, Table B.5 shows that liquidity declines after limit shocks when we measure liquidity with the logarithm of the absolute value of the yield error.

position when net demand is higher, an increase in the cost of this position would cause a larger price increase or a larger yield decrease (in the case of bonds) than otherwise would occur. In our regression analysis, this would manifest as a significant negative coefficient on the interaction term in a regression of yield on our limit shock variable and a proxy of net demand.

In our analysis, we proxy for net Treasury demand in two different ways. First, we use MBS duration. As Hanson (2014) explains, a decline in MBS duration indicates an increase in refinancing expectations. In turn, MBS refinancing prompts bond investors to put their received funds into Treasuries, driving up net demand for Treasuries and reducing Treasury yields. Table 5 shows that the average adjusted MBS duration (of fixed rate mortgages) in our sample period is about 4.8 years. In our regressions, we use the lagged value of the *negative* log change in the MBS duration as a proxy for Treasury net demand shifts. An increase in this variable means that MBS duration declined during the last week and Treasury demand increased.

Second, we use euro-area monetary policy rate changes. If the foreign policy as proxied by EURIBOR falls, investing in Treasuries becomes more attractive, all else being equal. Therefore, we expect that a decline in foreign interest rates leads to an upward demand shift for Treasuries, especially by international investors, which would manifest as a higher price and lower yield of US Treasuries. We compute the EURIBOR changes as two-week changes because capital may need time to adjust. In our regressions, we use the *negative* change such that an increase in this variable means that the EURIBOR fell and Treasury demand increased. We also looked at interest rate changes in other major foreign economies, and the results, which we do not report, were qualitatively similar to our other results.

Table 9 reports the results, with columns (1) through (3) focusing on the demand-shifter proxy derived from MBS duration, and columns (4) through (6) focusing on the foreign interest rate change. Column (1) shows that when the MBS duration decreases, Treasury yields fall, consistent with this variable capturing an increase in Treasury demand. Because we have standardized the demand-shifter variable and the dependent variable is in basis points, the coefficient of –1.325 means that yields drop by about 1.3 basis points in response to a one-standard-deviation change. Importantly, in columns (2) and especially (3), where we add time fixed effects, we find that when the Treasury demand shift is accompanied by a tighter-limit shock, yields do decrease more than otherwise. Relative to the average effect of the demand shifter (column 1), the differential effect per standard-deviation tighter limit shock is about 26 percent (–0.345/–1.325). Columns (4) through (6) provide additional evidence for the amplification of the effect of net demand shifts on yields. When the EURIBOR rate falls, US yields decline, but more so when primary dealers experience a limit shock that tightens their constraints. Relative to the average effect of the demand shifter (column 4), the differential effect per standard-deviation tighter limit shock is about 29 percent (–0.748/–2.525).

Table 9: Effect on Yields: Interaction with Treasury Demand Shifters

	Dep. Var.: Yield, bps					
	Shifter: Δ MBS duration			Shifter: Δ EURIBOR		
	(1)	(2)	(3)	(4)	(5)	(6)
Treasury Demand Shifter	-1.325***	-1.324***		-2.525***	-2.553***	
	(0.450)	(0.451)		(0.640)	(0.641)	
Limit Shock	-0.0534	-0.0689	-0.0609	-0.0342	0.286	0.213
	(0.093)	(0.092)	(0.065)	(0.098)	(0.184)	(0.175)
Treasury Demand Shifter × Limit Shock		-0.231	-0.345**		-0.969**	-0.748*
		(0.190)	(0.165)		(0.428)	(0.452)
Lagged LHS	Yes	Yes	Yes	Yes	Yes	Yes
Maturity FE	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	No	No	Yes	No	No	Yes
R^2	0.994	0.994	0.999	0.995	0.995	1.000
R^2 within	0.994	0.994	0.987	0.995	0.995	0.989
N	1,930	1,930	1,930	2,070	2,070	2,070

Notes: The dependent variable is the traded yield-to-maturity, as a percentage. For each maturity bucket, the yield is computed as the outstanding-volume-weighted mean yield of all outstanding Treasuries in that maturity bucket. In columns (1) through (3), the Treasury Demand Shifter is the lagged value of the *negative* log change in the MBS duration; in columns (4) through (6), it is the *negative value* of the two-week change in the EURIBOR. In each case, the demand shifter has been standardized to have unit variance and mean zero in each regression sample, and an increase in the variable means higher demand for Treasuries. Limit Shock is our baseline standardized VaR-limit shock. Robust standard errors are clustered at the week level and reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1. *Sources*: FR VV-1, FR 2004, Bloomberg, CRSP, Haver, authors' calculations.

Effects on Auctions We also study the effects of limit shocks on the primary market for Treasuries. The Treasury Department auctions off Treasuries by maturity at monthly intervals. Among the main participants in these auctions are primary dealers, and they are required by their primary-dealer status to bid at auctions. Naturally, we would expect that shifts in primary dealers' constraints will affect their bidding behavior.

During our sample period, auctions were conducted as single-price auctions. These are "Dutch" auctions, meaning that all successful bidders pay the same price; that is, they obtain the same yield. All prospective buyers place competitive bids, submitting to the Treasury Department a price-quantity schedule that represents the payout (yield) they want to receive and the amount they want to buy. The Treasury Department first accepts the bids with the lowest yield—that is, the ones that will cost the government the least—then the next-lowest, scaling the ladder toward its predetermined borrowing goal. All bidders receive the highest rate needed to clear the auction—the "high yield"—which in turn determines the cost of the debt issuance for the Treasury Department.

Table 10: Effect of Limit Shocks on Treasury Auctions

	Log(Bid-to-Cover Ratio)				Log(High Yield)	
	Dealers		Non-Dealers			
	(1)	(2)	(3)	(4)	(5)	(6)
Limit Shock	-0.048***	-0.051***	-0.001	-0.000	0.006	0.009**
	(0.014)	(0.015)	(0.004)	(0.004)	(0.005)	(0.005)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Maturity FE	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes	Yes	Yes
Excl. COVID	No	Yes	No	Yes	No	Yes
R^2	0.780	0.777	0.811	0.806	0.916	0.899
R^2 within	0.436	0.432	0.471	0.446	0.029	0.022
N	558	496	558	496	558	496

Notes: The table reports the effect of standardized limit shocks on auction outcomes. In columns (1) through (4), the logarithm of the bid-to-cover ratio is broken down by primary dealers and non-dealer participants. In columns (5) and (6), the dependent variable is the logarithm of the high yield. Each security auction is one observation. Controls include the bid-to-cover ratios and log total amount accepted at the last auction in the same tenor, as well as 14 daily lags of yield to maturity with the coefficients allowed to vary by tenor. The sample excludes T-bill auctions but includes all notes and bonds of all tenors. *Sources*: FR 2004, FR VV-1, Treasury Securities Auctions Data, CRSP, authors' calculations.

A crucial metric for the demand at an auction is the "bid-to-cover" ratio. It is calculated by dividing the total dollar amount of bids submitted in the auction by the amount of securities offered for sale. Typically, the higher the ratio, the more demand there is for Treasuries. Summary statistics at the auction level are reported in Appendix Table B.6.

We analyze the effects of limit shocks on two auction outcomes, the bid-to-cover ratio and the high yield.[28] The data breakdown also allows us to study the bidding behavior (bid-to-cover) of primary dealers versus other auction participants. Our regressions are similar to the ones we described earlier, but our data dimension is now the auction level, and we match the weekly limit shocks to the week preceding the auction date. To estimate the effect on the high yield, we control for 14 daily lags of the maturity-specific yield and include the bid-to-cover ratio and (log) amount in the preceding auction of the same maturity.

[28] Summary statistics for auction variables are given in Table B.6 in the Appendix.

Table 10 shows that primary dealers reduce their bidding in the auctions when they face a tightening-limit shock in the week before the auction (columns 1 and 2). The coefficient estimates indicate a 5 percent drop in the bid-to-cover ratio in response to a one-standard-deviation limit shock. On the other hand, in columns (3) and (4), we find a precisely estimated zero effect on the bidding behavior of non-dealers, which is what we would expect for a limit shock constructed from idiosyncratic dealer behavior. Columns (5) and (6) show that the high yield, the yield/price that clears the auction, increases along with the inward shift in the dealers' Treasury demand. The estimates indicate a 1 percent increase in response to a one-standard-deviation limit shock. Thus, shocks to dealers' constraints have direct effects on the financing cost to the US government.

3.3 Economic Implications

Our core findings show that position constraints matter for primary dealers' intermediation capacity in the Treasury market, as they strongly affect position, turnover, and spread, consistent with our model predictions. This holds for two separate identification schemes and independent periods. Specifically, in the previous section, using granularly identified shocks, we focused on estimating aggregate limit-shock elasticities of positions, turnover, and spread (Tables 6 and 8). The estimates are sizable, and we next want to more closely assess the economic effects (costs) of dealer constraints. While we do not aim to provide a full welfare analysis, our estimates nevertheless allow us to shed light on the values of key elasticities and the shadow cost of dealer constraints.

To do so, we revisit the first-order condition of the generic dealer problem, which states that, for profit maximization, the marginal return from charging a spread—intermediation income—has to equal the marginal cost of holding a position. This holds irrespective of the linearity of the supply-and-demand functions. The first-order condition can be re-written as:

$$t + st' = \lambda\delta' \quad \Rightarrow \quad \mu(1 + \epsilon_t) = \lambda\epsilon_\delta, \tag{15}$$

where $\mu \equiv \frac{s \cdot t}{\delta}$ is a margin measure equal to spread income relative to position, and ϵ_t is the spread elasticity of turnover. This term measures the percentage change in turnover in response to a percentage change in the spread. On the right-hand side, the marginal cost is the shadow cost of the constraint (λ), which is generally nonzero if the constraint binds, multiplied by the spread elasticity of the net position, ϵ_δ, measuring the percentage change in the net positions to a percentage change in the spread.

We can back out these crucial spread elasticities from the limit-shock elasticities of turnover, position, and the spread that we estimated in Tables 6 and 8. We define these

Table 11: Estimated Elasticities and Shadow Cost of Constraints

	Bid-Ask Spread Elasticity of		Shadow Cost (% of Margin)
	Turnover, ϵ_t	Position, ϵ_δ	$\frac{1+\epsilon_t}{\epsilon_\delta} \times 100$
	(1)	(2)	(3)
VaR estimates	-0.71	-0.87	33
SLR estimates	-0.76	-0.93	26

Note: The table reports the spread elasticity of turnover and position as implied by our estimates from the analysis of either VaR limits or the SLR policy change. $\frac{1+\epsilon_t}{\epsilon_\delta} = \frac{\lambda}{\mu}$ is the shadow cost relative to the margin. Note that the numbers in the quotients are semi-elasticities with respect to a one-standard-deviation VaR-limit shock or a percentage point SLR increase. The corresponding elasticities can be backed out by rescaling the estimates.

limit-shock elasticities as $e_x \equiv \frac{dx}{d\lambda}\frac{\lambda}{x}$ for $x = t, \delta, s$. The spread elasticity of turnover and position, respectively, can be derived as

$$\epsilon_t \equiv \frac{dt}{ds}\frac{s}{t} = e_t/e_s \quad \text{and} \quad \epsilon_\delta \equiv \frac{d\delta}{ds}\frac{s}{\delta} = e_\delta/e_s.$$

In Tables 6 and 8, to facilitate the interpretation of the magnitudes, we present the estimates using standardized limit shocks (unit variance). Note that such a rescaling of coefficients is irrelevant to the ratio of the effects that determines the spread elasticity.

Table 11 presents our estimates of the implied spread elasticities. Specifically, from our VaR-limit analysis, we estimate a spread elasticity of turnover of about $\hat{\epsilon}_t^{VaR} = -0.71 = -0.017/0.024$; that is, when the spread increases by 1 percent, turnover falls by 0.71 percent. We find a larger spread elasticity of position (in absolute value) of $\hat{\epsilon}_t^{VaR} = -0.87 = -0.021/0.024$. We can obtain similar estimates of the spread elasticities from our difference-in-differences analysis of the SLR change. In this case, we compute the spread elasticities as the ratios of differential growth in quantities (that is, turnover and position) and spread to the policy change. Our baseline estimates of $\hat{\epsilon}_t^{SLR} = -0.76 = -0.074/0.098$ and $\hat{\epsilon}_\delta^{SLR} = -0.93 = -0.091/0.098$ are similar in magnitude to those we obtain from our analysis of VaR shocks, although the elasticities are now obtained from growth rates.[29]

The estimated elasticities are smaller than unity (in absolute value), consistent with binding constraints. Without a (binding) position constraint (that is, $\lambda = 0$), the marginal intermediation profit needs to be zero (see Equation (15)). This is achieved at the point where the spread elasticity of turnover is –1. At this point, the gain from an increasing spread is

[29]For SLR elasticities, we use the coefficients in Table 2, column 2 and Table 3, columns 2 and 4. Thus, for position and turnover, we use conservatively identified estimates from our maturity-level analysis.

offset by the drop in volume. Hence, elasticity estimates of less than one (in absolute values) suggest that the bank operates on a part of the turnover curve, where it should expand business to decrease the marginal return, but constraints on both gross and net positions prohibit such an expansion of trading.

What do these elasticities say about the economic cost of these constraints? Rearranging the first-order condition further shows that

$$\frac{1+\epsilon_t}{\epsilon_\delta} = \lambda/\mu.$$

Thus, we can use a ratio of our identified spread elasticities to compute the marginal cost of the constraints as a share of the intermediation margin. Table 11 shows that our elasticities imply a sizable marginal cost from position constraints. Those costs are as high as 26 percent (SLR estimate) and 33 percent (VaR estimate) of the intermediation margin. Notice that the ratio of the elasticities identifies only the relative shadow cost of the constraint (relative to the margin).

We can use the estimated spread elasticities to shed additional light on the economic importance of primary dealers' constraints. Specifically, we can estimate the marginal dollar cost by evaluating the first-order condition using the elasticities and computing the margin based on the average spread, turnover, and position in our data. For the margin measure, we obtain $\hat{\mu} = 0.007798 = \frac{0.0005 \cdot 351.39}{22.53}$. Using the spread elasticities based on the VaR analysis, we then estimate that the marginal return is about $0.0023 = 0.007798(1 - 0.71)$, implying an estimated shadow cost of $\hat{\lambda} = 0.0026 = 0.0023/0.87$, so 0.26 percent. To back out the dollar amount, we multiply the shadow cost by a \$22.53 billion average weekly position, which gives a marginal foregone dollar profit of \$58.57 million per week, or \$3.05 billion per year.[30]

4 Conclusion

In this paper, we study the effects of primary dealers' constraints on the Treasury securities market. We use detailed confidential microdata on their risk limits, positions, turnover, and profits as well as two separate identification schemes to show that constraints on exposure have sizable effects on the market. Tighter constraints impair market liquidity: Dealers hold smaller positions, turnover falls, and intermediation spreads increase. This is the case across multiple settings involving different types of dealer limits. Tighter constraints also lead to amplified yield effects from demand shifts, and they dampen dealers' bidding behavior in

[30]If we use the spread elasticities obtained from the SLR analysis, we obtain a slightly lower estimate implying a shadow cost of \$45.3 million per week, or \$2.36 billion annualized.

Treasury auctions, thereby increasing the US government's financing cost.

Our intuitive findings are consistent with a stylized model of a spread-charging intermediary (market maker) that is subject to holding constraints. The equilibrium conditions show that the spread and turnover are crucially affected by position constraints and depend on the shadow cost of the constraint. This shadow cost is equal to the ratio of the spread elasticities of turnover and position, capturing the optimal trade-off between intermediation gains and the position-holding (shadow) cost under the constraint. From our identified estimates of shock responses, we back out the bid–ask spread elasticities of dealer turnover and dealers' positions to estimate a sizable shadow cost of the constraint of about 26 percent to 33 percent of their margin.

Our results have several policy implications, as Treasury markets are at the heart of the financial system, both in the United States and globally. From a regulatory perspective, it is important to understand that regulatory constraints that target broad bank-level exposure can impair bank-affiliated dealers' intermediation capacity, which is crucial for Treasury market liquidity. Perhaps more broadly, our results highlight how aggregate intermediation capacity depends on a granular set of key dealers and is therefore subject to idiosyncratic dealer shocks. An impairment of liquidity in the Treasury market in turn has implications for both monetary policy and fiscal policy. Monetary policy transmission may be affected in situations where yield movements are amplified due to demand shocks. Meanwhile, constraints on primary dealers may prevent them from absorbing additional federal government debt, which may be especially relevant in the current environment of rising federal debt issuance.

References

Adrian, Tobias, Emanuel Moench, and Hyun Song Shin. 2010. "Financial Intermediation, Asset Prices and Macroeconomic Dynamics." *Federal Reserve Bank of New York Staff Reports*.

Adrian, Tobias, Erkko Etula, and Hyun Song Shin. 2015. "Risk Appetite and Exchange Rates." *Federal Reserve Bank of New York Staff Reports*.

Barbiero, Omar, Falk Bräuning, Gustavo Joaquim, and Hillary Stein. 2024. "Dealer Risk Limits and Currency Returns."

Beetsma, Roel, Massimo Giuliodori, Frank de Jong, and Daniel Widijanto. 2016. "Price effects of sovereign debt auctions in the euro-zone: The role of the crisis." *Journal of Financial Intermediation*, 25: 30–53.

Bernanke, Ben, and Mark Gertler. 1989. "Agency Costs, Net Worth, and Business Fluctuations." *American Economic Review*, 79(1): 14–31.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2021. "Quasi-Experimental Shift-Share Research Designs." *The Review of Economic Studies*, 89(1): 181–213.

Brunnermeier, Markus K., and Yuliy Sannikov. 2014. "A Macroeconomic Model with a Financial Sector." *American Economic Review*, 104(2): 379–421.

Duffie, Darrell. 2010. "Presidential Address: Asset Price Dynamics with Slow-Moving Capital." *The Journal of Finance*, 65(4): 1237–1267.

Duffie, Darrell. 2018. "Financial Regulatory Reform After the Crisis: An Assessment." *Management Science*, 64(10): 4835–4857.

Duffie, Darrell. 2020. "Still the World's Safe Haven? Redesigning the U.S. Treasury Market After the COVID-19 Crisis." Brookings Hutchins Center Working Paper.

Duffie, Darrell, Michael J. Fleming, Frank M. Keane, Claire Nelson, Orr Shachar, and Peter Van Tassel. 2023. "Dealer Capacity and U.S. Treasury Market Functionality." *Federal Reserve Bank of New York Staff Reports*.

Favara, Giovanni, Sebastian Infante, and Marcelo Rezende. 2022. "Leverage Regulations and Treasury Market Participation: Evidence from Credit Line Drawdowns."

Fleming, Michael, Giang Nguyen, and Joshua Rosenberg. 2023. "How do treasury dealers manage their positions?" *Federal Reserve Bank of New York Staff Reports*.

Gabaix, Xavier, and Ralph S. J. Koijen. 2020. "Granular Instrumental Variables."

Grossman, Sanford J., and Merton H. Miller. 1988. "Liquidity and Market Structure." *The Journal of Finance*, 43(3): 617–633.

Hanson, Samuel G. 2014. "Mortgage convexity." *Journal of Financial Economics*, 113(2): 270–299.

Hendershott, Terrence, and Albert J. Menkveld. 2014. "Price pressures." *Journal of Financial Economics*, 114(3): 405–423.

Hendershott, Terrence, and Mark S. Seasholes. 2007. "Market Maker Inventories and Stock Prices." *American Economic Review*, 97(2): 210–214.

He, Zhiguo, and Arvind Krishnamurthy. 2013. "Intermediary Asset Pricing." *American Economic Review*, 103(2): 732–70.

He, Zhiguo, Bryan Kelly, and Asaf Manela. 2017. "Intermediary asset pricing: New evidence from many asset classes." *Journal of Financial Economics*, 126(1): 1–35.

He, Zhiguo, Stefan Nagel, and Zhaogang Song. 2022. "Treasury inconvenience yields during the COVID-19 crisis." *Journal of Financial Economics*, 143(1): 57–79.

Lou, Dong, Hongjun Yan, and Jinfan Zhang. 2013. "Anticipated and Repeated Shocks in Liquid Markets." *The Review of Financial Studies*, 26(8): 1891–1912.

Nagel, Stefan. 2012. "Evaporating Liquidity." *The Review of Financial Studies*, 25(7): 2005–2039.

A Additional Figures

Figure A.1: Long, Short, and Net Positions by Mat Bucket

Source: FR 2004, authors' calculations.

Figure A.2: Primary Dealer Coverage in FR 2004 and FR VV-1 Data

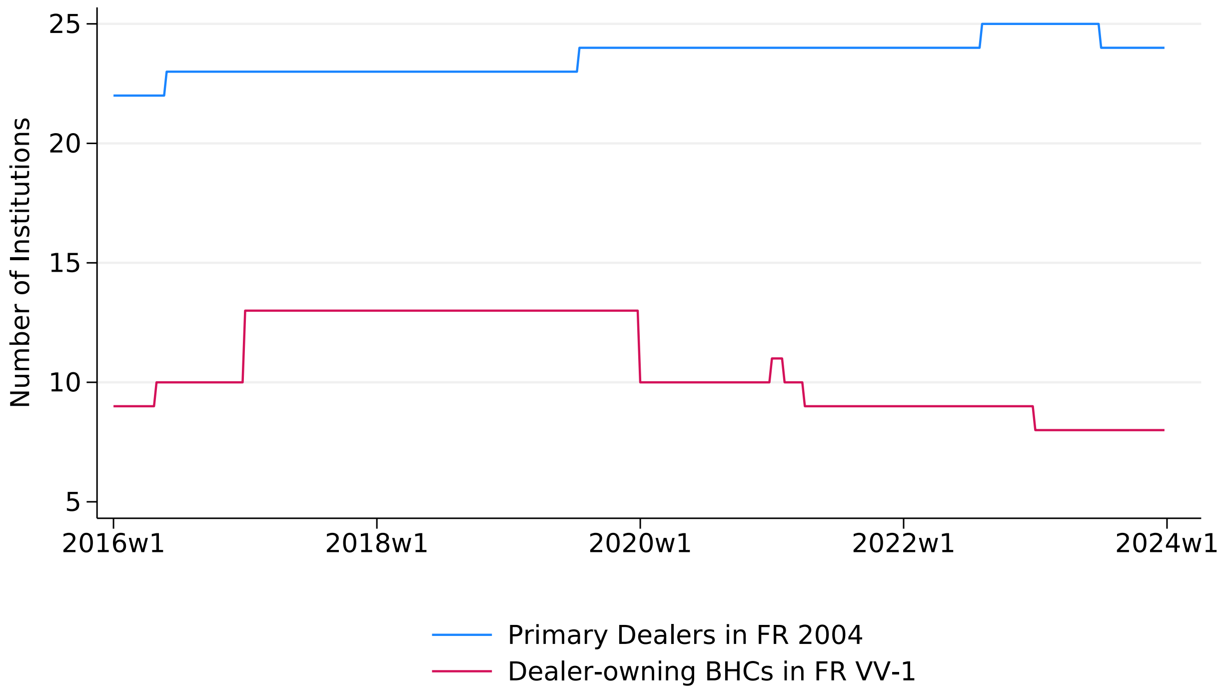

Notes: This figure shows the number of primary dealers in the FR 2004 data during our sample period from 2016:W1 through 2023:W52. The figure also shows the coverage of BHCs that own a primary dealer in the FR VV-1 data. More details on the two data sets are discussed in the data section. *Sources*: FR 2004, FR VV-1, authors' calculations.

Figure A.3: Federal Reserve Responses to COVID-19 Crisis

Facility	Sector targeted	Funding	Date announced	Date opened	Date closed	Maximum capacity ($ bil.)	Peak assets ($ bil.)	Assets as of 12/8/21 ($ bil.)	Treasury backstop ($ bil.)
Commercial Paper Funding Facility (CPFF)	Commercial paper market	Fed, Treasury (ESF)	3/17/20	4/14/20	3/31/21	Unlimited	4.2	0.0	10.0
Main Street Lending Program (MSLP)*	Small and mid-sized businesses, non-profits	Fed, Treasury (CARES Act)	4/9/20	7/6/20†	1/8/21	600.0	16.6	13.4	75.0
Money Market Mutual Fund Liquidity Facility (MMLF)	Money market mutual funds	Fed, Treasury (ESF)	3/18/20	3/23/20	3/31/21	Unlimited	53.2	0.0	10.0
Municipal Liquidity Facility (MLF)*	State and local governments	Fed, Treasury (CARES Act)	4/9/20	5/26/20	12/31/20	500.0	6.4	4.2	35.0
Paycheck Protection Program Liquidity Facility (PPPLF)*	Small businesses	Fed	4/9/20	4/16/20	7/30/21	953.0‡	90.6	39.9	-
Primary Dealer Credit Facility (PDCF)	Broker-dealers	Fed	3/17/20	3/20/20	3/31/21	Unlimited	33.4	0.0	-
Primary Market Corporate Credit Facility (PMCCF)*	Large businesses	Fed, Treasury (CARES Act)	3/23/20	6/29/20	12/31/20	750.0	0.0	0.0	50.0
Secondary Market Corporate Credit Facility (SMCCF)*	Large businesses, exchange-traded funds	Fed, Treasury (CARES Act)	3/23/20	5/12/20†	12/31/20	Combined with PMCCF	14.3	0.0	25.0
Term Asset-Backed Securities Loan Facility (TALF)	Securities markets (e.g. student, auto, & credit card loans)	Fed, Treasury (CARES Act)	3/23/20	6/17/20	12/31/20	100.0	4.1	1.4	10.0

Notes: Key Federal Reserve programs to support the economy at the onset of the COVID-19 crisis. *Source*: https://www.brookings.edu/articles/fed-response-to-covid19/.

Figure A.4: Effect of SLR Change on Dealers' Positions—Long Sample

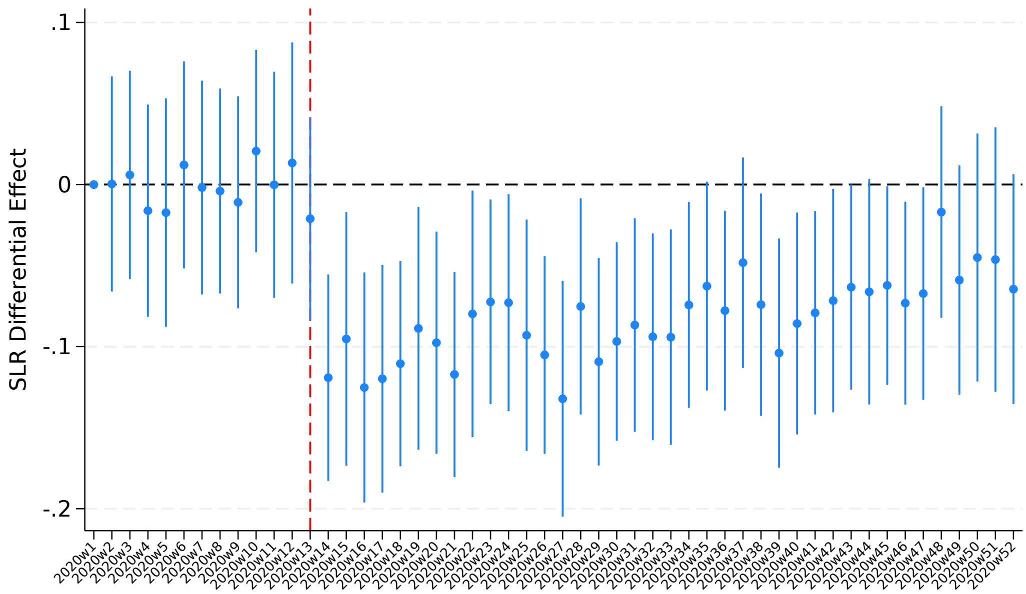

Notes: The figure shows coefficient estimates $\{\hat{\beta}_t\}$ from the following regression: $\log(\text{Position})_{i,t} = \beta_t \text{SLR Change}_t \times \text{SLR}_i^{2019q4} + \alpha_t + \alpha_i + \epsilon_{i,t}$. SLR Change is a dummy equal to one after 2020:W14 and zero otherwise. Positions are gross positions, in line with gross notionals entering the SLR exposure measure. The sample period runs from 2020:W1 through 2020:W52. The SLR change was announced on April 1, 2020 (week 14), and became effective immediately. The vertical red line indicates the last week before the change. The bars represent 90 percent confidence intervals based on robust standard errors. *Sources*: FR 2004, FR Y-9C, authors' calculations.

Figure A.5: Concentration of Weights by Maturity

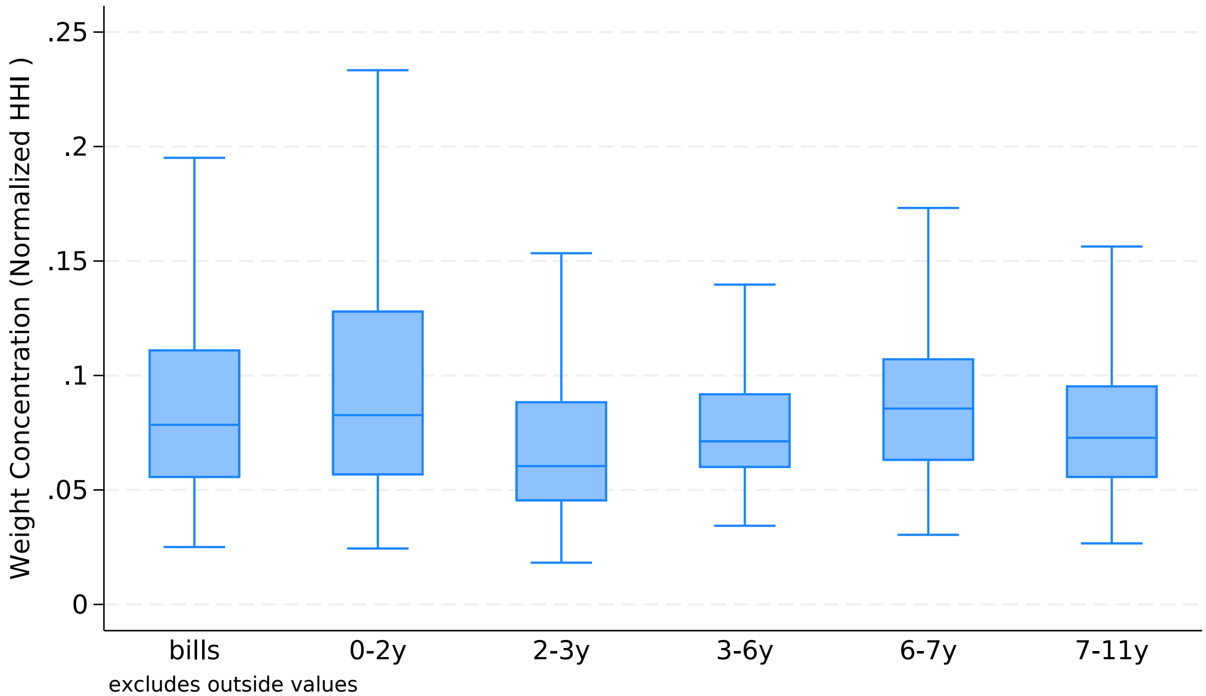

Notes: HHI (Herfindhal index) is computed as the sum of the squared weights summed for each maturity and week. Normalized HHI is (HHI-1/N)/(1-1/N). *Sources*: FR 2004, authors' calculations.

Figure A.6: Limit Shocks by Maturity Bucket

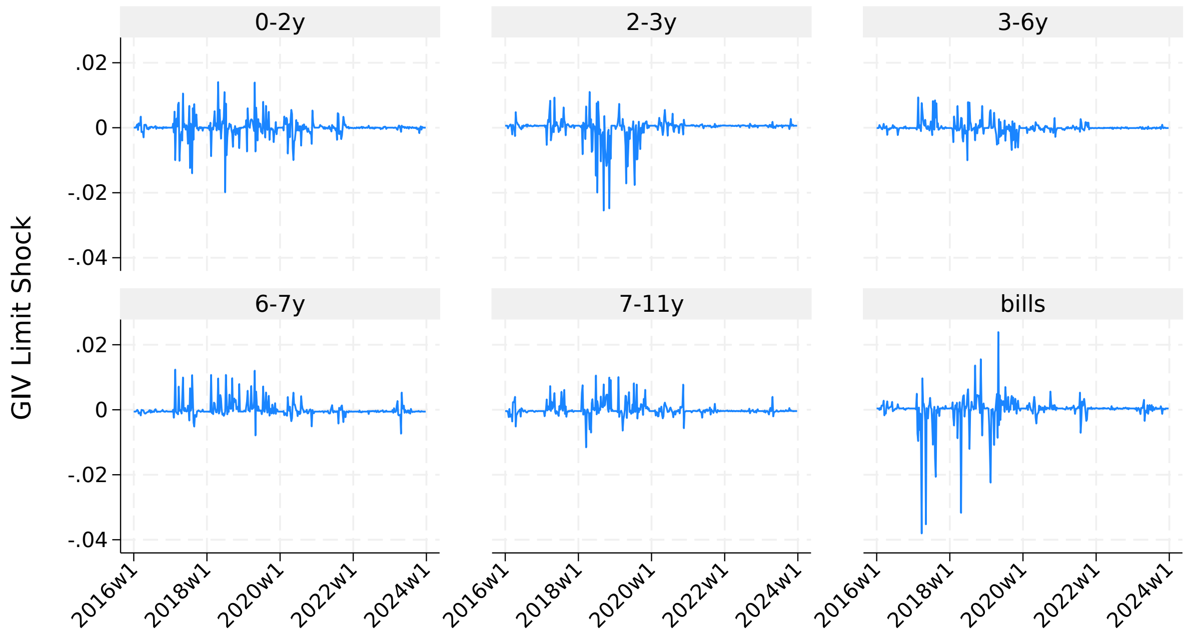

Notes: Limit shocks are orthogonalized with respect to the fixed-effects structure used in the main regressions.
Sources: FR 2004, FR VV-1, authors' calculations.

Figure A.7: Return Variability by Maturity Bucket

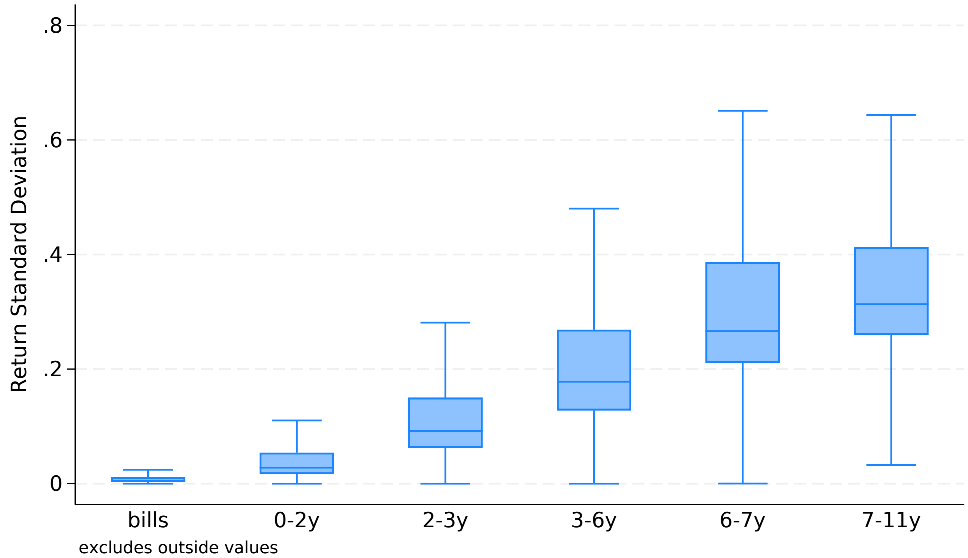

Notes: The figure shows, for each maturity bucket, the distribution of return variability, computed as the log price change (multiplied by 100, so it is approximate percentage returns) at the CUSIP-day level during a 45-(trading-)day rolling window. *Sources*: CRSP, authors' calculations.

Figure A.8: Primary Dealers' Risk Limit Changes and Treasury Position by Maturity

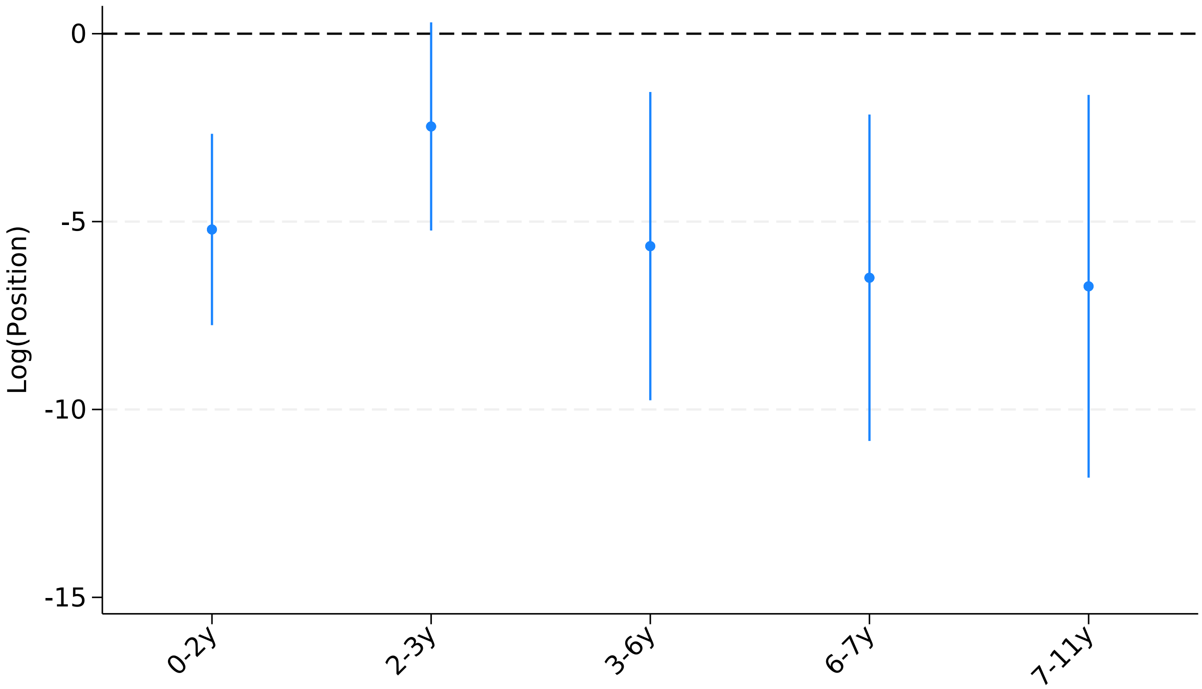

Notes: This table shows the results of Equation 14 for specific maturity buckets, with the net position outcome variable; specifically, it shows the interaction effect of the limit shock with the categorical maturity bucket variable. *Sources*: FR 2004, FR VV-1, authors' calculations.

B Additional Tables

Table B.1: List of Primary Dealers

Primary Dealer	FR2004 Reporting	BHC owned	SLR	Volcker
ASL Capital Markets Inc.	2022w32 - 2023w52			
BMO Capital Markets Corp.	2015w1 - 2021w49	✓		
BNP Paribas Securities Corp.	2015w1 - 2023w52	✓		
Bank of Montreal, Chicago Branch	2021w50 - 2023w52	✓		
Bank of Nova Scotia, New York Agency	2015w1 - 2023w52	✓		✓
Barclays Capital Inc.	2015w1 - 2023w52	✓	✓	✓
BofA Securities, Inc.	2019w20 - 2023w52	✓	✓	✓
Cantor Fitzgerald & Co.	2015w1 - 2023w52			
Citigroup Global Markets Inc.	2015w1 - 2023w52	✓	✓	✓
Credit Suisse Securities (USA) LLC	2015w1 - 2017w45	✓	✓	✓
Credit Suisse, New York Branch	2017w46 - 2023w26	✓	✓	✓
Daiwa Capital Markets America Inc.	2015w1 - 2023w52			
Deutsche Bank Securities Inc.	2015w1 - 2023w52	✓		✓
Goldman Sachs & Co. LLC	2015w1 - 2023w52		✓	✓
HBSC Securities (USA) INc.	2015w1 - 2023w52	✓	✓	✓
J.P. Morgan Securities LLC	2015w1 - 2023w52	✓	✓	✓
Jefferies LLC	2015w1 - 2023w52			
Merrill Lynch, Pierce, Fenner & Smith Inc.	2015w1 - 2019w19	✓		
Mizuho Securities USA LLC	2015w1 - 2023w52	✓		✓
Morgan Stanley & Co. LLC	2015w1 - 2023w52	✓	✓	✓
NatWest Markets Securities Inc.	2015w1 - 2023w52			
Nomura Securities International, Inc.	2015w1 - 2023w52			
RBC Capital Markets, LLC	2015w1 - 2023w52	✓		✓
SG Americas Securities, LLC	2015w1 - 2015w48	✓		
Santander US Capital Markets LLC	2019w29 - 2023w52	✓		
Societe Generale, New York Branch	2015w49 - 2023w52	✓		
TD Securities (USA) LLC	2015w1 - 2023w52	✓		
UBS Securities LLC	2015w1 - 2023w52	✓	✓	✓
Wells Fargo Securities, LLC	2016w22 - 2023w52	✓	✓	✓

Notes: The table reports the current and historical primary dealers in our sample period. BHC ownership, SLR reporting, and Volcker rule reporting are based on at least one occurrence during the sample period. Note that all primary dealers are required to file FR 2004. SLR and Volcker rule reporting requirements are based on publicly observable criteria. This table uses only publicly available information. *Sources*: FR Y-9C, Federal Reserve Bank of New York website.

Table B.2: Autoregressive Model of Weights Used in Shock Aggregation

	Weights		
	(1)	(2)	(3)
Lagged Weight	0.876***	0.818***	
	(0.002)	(0.002)	
Maturity Bucket 1 (Bills) × Lagged Weight			0.780***
			(0.005)
Maturity Bucket 2 (0-2y) × Lagged Weight			0.894***
			(0.005)
Maturity Bucket 3 (2-3y) × Lagged Weight			0.768***
			(0.005)
Maturity Bucket 4 (3-6y) × Lagged Weight			0.855***
			(0.005)
Maturity Bucket 5 (6-7y) × Lagged Weight			0.829***
			(0.005)
Maturity Bucket 6 (7-11y) × Lagged Weight			0.762***
			(0.005)
Maturity FE	No	Yes	Yes
Week FE	No	Yes	Yes
Bank FE	No	Yes	Yes
R^2	0.767	0.774	0.776
R^2 within	0.767	0.669	0.673
N	58,794	58,794	58,794

Notes: The table reports autoregressive coefficients on the weights $w^m_{b,t}$ used in the aggregation of idiosyncratic limit changes to construct the maturity-level limit shock; see (13). The most saturated estimated model of column (3) is $w^m_{b,t} = \phi^m w^m_{b,t-1} + \alpha_b + \alpha_t + \alpha_m + \epsilon^m_{b,t}$. *Sources*: FR 2004, authors' calculations.

Table B.3: Baseline Results Using Alternative Shocks

	Log(Position)			
	(1)	(2)	(3)	(4)
Limit Shock (w/ FR 2004 controls)	-0.022***			
	(0.008)			
Limit Shock (w/ PCA)		-0.012**		
		(0.006)		
Limit Shock (excl. Holidays)			-0.010*	
			(0.005)	
Limit Shock (excl. dealers w/ foreign desks)				-0.009**
				(0.004)
Lagged LHS	Yes	Yes	Yes	Yes
Maturity FE	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes
R^2	0.945	0.945	0.945	0.945
R^2 within	0.733	0.732	0.732	0.732
N	2,075	2,075	2,075	2,075

Notes: Robustness of baseline result to differently constructed Limit Shocks. *Sources*: FR VV-1, FR 2004, authors' calculations.

Table B.4: Baseline Results Using Alternative Sample Periods

	Log(Position)			
	$t < 2019$	$t \geq 2019$	$t \geq 2020$	All
	(1)	(2)	(3)	(4)
Limit Shock	-0.028**	-0.014*	-0.005	
	(0.012)	(0.008)	(0.005)	
Year 2016 × Limit Shock				-0.041
				(0.068)
Year 2017 × Limit Shock				-0.020**
				(0.009)
Year 2018 × Limit Shock				-0.021*
				(0.012)
Year 2019 × Limit Shock				-0.020*
				(0.012)
Year 2020 × Limit Shock				-0.041*
				(0.022)
Year 2021 × Limit Shock				0.040
				(0.051)
Year 2022 × Limit Shock				0.307
				(0.233)
Year 2023 × Limit Shock				0.011
				(0.045)
Lagged LHS	Yes	Yes	Yes	Yes
Maturity FE	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes
R^2	0.933	0.950	0.941	0.945
R^2 within	0.505	0.794	0.799	0.733
N	775	1,300	1,040	2,075

Notes: The table reports our baseline estimates for different sample periods and interacted with year dummies. As this table shows, our baseline effects are driven by the first half of the sample, where we have considerable variation in limit shocks; see Appendix Figure A.6. *Sources*: FR VV-1, FR 2004, authors' calculations.

Table B.5: Log Liquidity Index on LHS

	Liquidity Index	
	(1)	(2)
Limit Shock	-0.089*	-0.096
	(0.051)	(0.058)
Lagged LHS	Yes	Yes
Maturity FE	Yes	Yes
Week FE	Yes	Yes
Excl. COVID	No	Yes
R^2	0.653	0.648
R^2 within	0.314	0.333
N	2,040	1,780

Notes: The dependent variable is the logarithm of the absolute value of the yield error. The data are at the maturity-bucket-by-week level. *Sources*: FR VV-1, FR 2004, Morgan Markets, authors' calculations.

Table B.6: Summary Statistics on Treasury Auctions

	2-Year	3-Year	5-Year	7-Year	10-Year	20-Year	30-Year
Total Accepted ($ B)	44.68	44.41	48.06	41.54	32.81	20.20	20.91
Primary Dealer Accepted (%)	27.87	29.30	22.00	17.91	21.62	16.33	20.13
Bid-to-Cover	2.66	2.62	2.44	2.47	2.47	2.52	2.33
High Yield (%)	1.87	1.93	2.06	2.21	2.31	2.79	2.76
Number of Auctions	95.00	96.00	95.00	95.00	97.00	44.00	96.00

Notes: The sample excludes T-bills auctions. For the 10-, 20-, and 30-year buckets, we include the few securities with one or two months shorter maturity in those buckets. For example, we consider a security with a maturity of nine years and 10 months equivalent to a 10-year Treasury. *Sources*: Treasury Securities Auctions Data, authors' calculations.

C Conceptual Framework with Convex Holding Costs

Here, we consider a bank that faces a convex (quadratic) cost of holding a position leading to an increasing marginal cost, as opposed to the linear cost function in our baseline model that implies a constant marginal cost. Rewrite their optimization function as

$$\max_{s,\delta} \ \pi = s\left(D(e+s) + S(e-s)\right) - p\delta - \frac{\gamma}{2}\delta^2 \tag{16}$$

$$\text{s.t. } \ \delta = D(e+s) - S(e-s), \tag{17}$$

where $\gamma > 0$ parameterizes the bank's cost of holding nonzero net positions. The first-order-condition of the general problem is now given by

$$t + st' = \underbrace{\gamma\delta}_{\text{MC}}\,\delta'. \tag{18}$$

Again, assume linearity (in the log price) of the demand and supply functions of the form $D(p+s) = a - b(p+s)$ and $S(p-s) = c + d(p-s)$, with $a, b, c, d > 0$. Solving the first-order condition for the optimal spread and deriving the optimal exposure and turnover gives:

$$s^* = \frac{1}{\nu}\left[(a+c) + \gamma(a-c)(b-d)\right] \tag{19}$$

$$\delta^* = \frac{1}{\nu}\left[2(a-c)(b+d) - (a+c)(b-d)\right] \tag{20}$$

$$t^* = \frac{1}{\nu}\left[(a+c)(b+d) + 2\gamma(b-d)(bc-ad)\right], \tag{21}$$

where $\nu = 2(b+d) + \gamma(b-d)^2$.

To ensure that all equilibrium quantities (D^* and S^*) are positive, we need the following parameter restrictions:

$$\frac{ab + 2ad - bc}{(b-d)(ad-bc)} < \gamma < \frac{ad - cd - 2bc}{(b-d)(ad-bc)}.$$

Without loss of generality, we can again simplify to the case where $c = 0$, and because γ is positive, it must be the case that $b > d$.

We then derive the following derivatives, which have the same sign as in Section 1:

$$\frac{\partial \delta}{\partial \gamma} = -(b-d)\zeta < 0 \tag{22}$$

$$\frac{\partial t}{\partial \gamma} = -(b+d)\zeta < 0 \tag{23}$$

$$\frac{\partial s}{\partial \gamma} = \zeta > 0, \tag{24}$$

where $\zeta = \nu^{-2}a(b-d)(b+3d) > 0$.