Aggregate Collateral Demand*

Toomas Laarits† Chase P. Ross‡ Sharon Y. Ross§

December 2, 2025

Abstract

We study aggregate collateral demand to understand its effect on the Treasury convenience yield. It is well documented that safe assets trade at a premium because they are convenient. That convenience depends on many non-pecuniary benefits—safety, liquidity, and collateral services—but their relative importance is hard to pin down. We show the causal pass-through from aggregate collateral demand shocks on the Treasury convenience yield by measuring *collateral sinks*, ways in which safe assets are removed from circulation, even if temporarily. We show how the rapidly growing $1 trillion collateral-swap market exacerbates these dynamics in stress.

JEL Codes: E40, E51, G12, N21

Keywords: margin requirement, safe asset, collateral swap, convenience yield, Treasury

*For comments and suggestions, we thank Andrew Metrick, Hannah Sheldon, and seminar participants at the U.S. Treasury. The analysis and conclusions set forth are those of the authors and do not indicate concurrency by members of the Board of Governors of the Federal Reserve System or its staffs.

†New York University. Email: tlaarits@stern.nyu.edu

‡Board of Governors of the Federal Reserve System. Email: chase.p.ross@frb.gov

§Board of Governors of the Federal Reserve System. Email: sharon.y.ross@frb.gov

Disclosure Statement

All three authors have nothing to disclose.

1 Introduction

Safe assets command an extra premium called the *convenience yield*. This premium arises from several sources, including deep liquidity, negligible credit risk, and use as collateral. U.S. Treasuries are the preeminent convenience-earning safe assets, but it is hard to disentangle the specific sources of their convenience yields. Some drivers—liquidity and safety—are partly observable, proxied by bid–ask spreads and credit default swap spreads. The collateral channel is harder to observe. In this paper, we isolate the collateral channel, untangling the ways that aggregate collateral demand can affect the convenience yield and how intermediaries relax, but at times exacerbate, those constraints.

We begin by conceptualizing aggregate collateral demand with *collateral sinks* that remove safe assets from circulation. It is difficult to pin down the elasticity of the convenience yield with respect to Treasuries outstanding because Treasuries can be reused and rehypothecated, sometimes many times over. Treasury velocity is well above one (Infante and Saravay, Forthcoming). Yet some Treasuries are posted or pledged in ways that leave them, at least temporarily, immobile. For example, margin posted to a central counterparty (CCP) against derivatives cannot be rehypothecated. Some counterparties place constraints—legal or contractual—on what can be done with their collateral, sometimes prohibiting its reuse. By simple market-clearing logic, when collateral sinks grow, the price of the remaining convenience rises.

Dysfunction in opaque and largely bilateral collateral markets during the global financial crisis imprinted on a generation of regulators and central banks the importance of liquidity. Regulations have been designed, implemented, and revised over the intervening years; collectively, they require the financial system to hold and consume substantially more safe assets than before. These reforms increased Treasury demand by requiring banks to hold additional safe assets under the liquidity coverage ratio. They also increased safe-asset demand with stricter margin requirements, pushing bilateral transactions (with lower margin requirements) toward CCPs (with higher margin requirements) and imposing higher capital charges on the remaining bilateral transactions. Without a concomitant increase in safe-asset issuance, these liquidity regulations have often proved binding (Caballero et al., 2017).

We estimate the total size of collateral sinks at $2.2 trillion in 2024, following nearly a decade of steady growth. Sunk collateral constitutes 9 percent of Treasuries outstanding,

a striking magnitude, and it varies substantially over our sample back to 2016. Measuring sunk collateral is nontrivial. We use supervisory data for the largest U.S. banks to make progress on this front. We identify three components of collateral sinks: repos, collateral swaps, and margin. Since much of the collateral in these markets is not sunk—e.g., the collateral in a repo can typically be reused—we focus on segments where we have strong priors that the collateral is sunk. This could be because certain counterparties never pledge collateral themselves or because the collateral is explicitly barred from being rehypothecated.

We provide novel data and analysis on a key and rapidly growing tool that the private sector uses to buffer the increased demand for safe assets: the $1 trillion *collateral swap* market. Collateral swaps directly transform non-safe assets into safe assets. Such swaps are especially useful for meeting margin requirements, and we estimate it reduces investor's margin capital by $70 billion by allowing investors to exchange high-haircut securities for lower-haircut safe assets to reduce their capital needs.

But the collateral swap market depends on banks. In normal times, when banks are unconstrained, they can help investors satisfy their margin requirements by providing collateral via collateral swaps. In bad times, when banks become constrained or need to source collateral for themselves, they pull back from the collateral swap market, and investors face a double whammy: volatility raises margin requirements, and simultaneously investors cannot rely on collateral swaps to meet their increasing margin requirements. Hence, a novel contribution of our paper is to show that the interaction of collateral sinks and collateral swaps produces spikes in convenience yields.

We formalize the mechanism in a stylized model where Treasuries supply different convenience services—liquidity, safety, collateral—and those services are lost once a Treasury is placed into a collateral sink. The convenience yield, then, equals the shadow value of the Treasury's pledgeability. The model shows that convenience yields are increasing on sunk collateral. The degree of pass-through depends on dealer constraints: banks buffer the effect by intermediating collateral swap when constraints are slack, but amplify it when constraints are tight and banks hoard collateral.[1] We use the model to generate predictions that we test in the data.

A simple time series regression of convenience yields on collateral sinks yields a strong

[1]We will refer interchangably to banks and dealers for brevity even though the broker-dealer subsidiary of a bank-holding company typically conduct collateral swaps.

relationship but is rife with endogeneity: reverse causality, omitted variable bias, and simultaneity bias. We take two approaches to better identify the effect of collateral sinks on convenience yields. The first approach relies on idiosyncrasies in the Chicago Mercantile Exchange's (CME) margin requirements. The CME is the largest U.S.-based CCP, and a major hub for trading in Treasury, interest rate, and equity-related derivatives. While the CME sets margins in response to heightened volatility, it also considers several other factors when setting margin requirements. These factors introduce idiosyncrasies in margin requirements that are plausibly orthogonal to the aggregate conditions that drive convenience yields. We use those idiosyncrasies as instruments for margin-driven variation in collateral sinks, and we use that variation to estimate the effect of collateral sinks on convenience yields. A one-standard-deviation increase in sunk collateral relative to Treasuries outstanding—equal to 0.07pp of Treasuries outstanding, or roughly $10.5 billion—-raises convenience yields by 2 basis points (cast to units comparable to the 3-month OIS–T-bill spread, which averages 15 bps with a standard deviation of 20 bps).

Our final empirical analysis examines the role of banks as pressure-release valves, attenuating the elasticity of the convenience yield with respect to collateral sinks when banks are unconstrained and intermediating the collateral swap market. When banks grow more constrained, we show that the elasticity increases by more than 8-fold, helping explain the spikes in convenience yields that we observe in bad states.

Combined, our analysis makes three novel contributions. First, we construct a new measure of non-rehypothecable collateral—collateral sinks—and document its level, time variation, and composition across margin, repos, and swaps. Second, we provide the first causal estimate of the elasticity of Treasury convenience yields to sunk collateral using two independent designs—idiosyncratic CME margin variation and a granular IV based on desk-level flows—and show that pass-through is state-dependent and rises when dealer balance-sheet constraints bind. Third, we document new facts about the collateral-swap market, including its scale and the cyclicality of dealer intermediation, and show that swap intermediation buffers the effect of sinks in normal times but contracts in stress, amplifying convenience yields. Together, these results isolate the role the collateral channel plays in safe-asset premia and show when dealer balance sheets mute or amplify it.

Related Literature A large literature has documented that Treasuries carry a convenience yield, meaning they carry a lower yield or higher price than simple fundamentals suggest (Duffie 1996, Longstaff 2004, Fleckenstein et al. 2014, Fleckenstein and Longstaff 2024). Several studies have shown how the supply and demand of safe assets affects the convenience yields, including the role that intermediary frictions play (Krishnamurthy and Vissing-Jorgensen 2012, Greenwood et al. 2015, He et al. 2022). Acharya and Laarits (2025) show that Treasuries' convenience yield is time-varying and must be earned by the asset's properties: it tends to compress when Treasuries are less effective as a hedge. Krishnamurthy and Ma (2024) give a comprehensive framework for these effects and explain how the equilibrium convenience yield is jointly determined by safe-asset demand and safe-asset supply from both public and private sources.

A closely related set of papers study the interplay of collateral markets and safe asset premia. Aggarwal et al. (2021) show how safe assets in European lending markets trade at a premium in collateral markets during stress, showing how markets transform collateral when safe assets become scarce. Infante and Saravay (Forthcoming) show a related mechanisms, the reuse of Treasuries in response to scarcity. Also related to our focus is work by Brunnermeier and Pedersen (2009) and Garleanu and Pedersen (2011) who show that binding margin requirements and funding constraints can push up the price of collateral and thereby increase Treasury convenience yields.

Our paper contributes to the literature by focusing on isolating the collateral channel on convenience yields by studying shocks to aggregate collateral demand that create causal shifts in the convenience yield. We also contribute to the literature by showing the role that the collateral swap market plays in directly amplifying or attenuating the pass-through of collateral demand shocks in safe asset prices.

2 Institutional Context

It is helpful to briefly describe how the financial system demands collateral for margin and financing.

2.1 Margin

Financial market participants must post margin to reduce the risk that they are unable to satisfy their obligations to counterparties if a position loses value. Margin requirements have long proved an invaluable risk management tool, and many flavors of derivatives, both bilateral and centrally-cleared, carry margin requirements. Margin typically constitutes two components: initial and variation. Traders owe initial margin when they initiate a derivative transaction, and variation margin is required when the value of the contract later fluctuates.

Margin requirements change and reflect expectations about the derivative's risk. Volatile derivatives, like Bitcoin futures, have margin requirements larger than 20 percent, while Treasuries have margin requirements in the single digits. Figure 1 shows that the CME's margin requirements 10-year Treasury future range from 60 to 225 bps, and moves closely—but not perfectly—with implied volatility; the same applies for S&P 500 index futures, although its margin requirements are a level shift higher, about 5pp, reflecting its higher risk.

An examples of the lifecycle of derivative trade makes the margin mechanics clear. Suppose an investor goes long $1 million, 10 contracts, of notional exposure to a 30-year U.S. Treasury future settling in September 2025. When the trader takes on the position, they post margin equal to the initial margin, which is $5,150 per contract for a total of $51,150. At initiation, the trader posts no variation margin because they start with no profit or loss. Suppose the price falls and the trader has a paper loss of $3,750. The equity in their account is now $51,150−$3,750=$47,750. Assuming that the initial margin requirement is equal to the maintenance margin requirement, the CCP issues a margin call for an additional $3,750 of collateral. Bilateral margin requirements conceptually mirror this process.

The financial system consumes substantial margin. The six largest U.S. banks alone posted nearly $800 billion in recent years (we describe the data sources in detail in the next section). The top panel of Figure 2 shows that margin volumes vary substantially over time, with a clear jump during the panicked stage of the Covid-19 pandemic, when margin jumped by $117 billion in the first week of March. Yet there are many other local maxima that do not obviously correspond to periods of market stresses.

Our $800 billion margin estimate is an undercount of margin requirements for the aggregate financial system because it does not include the full set of market participants. Assuming the ratio of centralized to total margin for the largest banks—26 percent—captures the true

market structure, the CFTC's estimate of $770 billion in early 2025 implies total margin posted of $2.8 trillion (see IA.1). As a point of comparison, total initial margin posted to CCPs globally was $1.3 trillion in 2024, a lower bound on total margin because it excludes bilateral margin requirements (CCP Global, 2024). Even more, margin requirements also include separate "default fund" contributions which are mutualized loss-sharing funds that CCPs require their members prefund. While this also requires large balances—CCPs required about $170 billion in 2024 globally—the number tends to be slower moving.

The trader can choose what type of collateral to post, including cash, Treasuries, or other securities, although variation margin is typically cash. At the CME, USD cash is accepted without limit or haircut, and Treasuries are accepted without limit but with haircuts ranging from 0.5 to 8 percent depending on tenor.[2] Traders can post a wide range of other securities as margin, subject to haircuts and caps; corporate bonds have a minimum 20 percent haircut and are capped at $2 billion of USD equivalent.[3]

Table 1 provides summary statistics on our sample and shows that cash constitutes the majority of margin posted, averaging $374 billion, or about two-thirds of all margin posted. Treasuries are the next largest collateral type, averaging $64 billion or 11 percent of all margin, but exceeding 20 percent in some periods. Banks pledge a wide set of high-quality liquid asset (HQLAs), although any given collateral type is small in level terms outside Treasuries, with agencies being the next largest. Non-HQLA assets constitute the remaining posted margin, about $100 billion, and spans a wide mix of asset classes—the largest of which include "other" and investment grade corporate debt, although both are magnitudes smaller than Treasuries. Table IA.1 shows the breakdown of margin between initial and variation margin; variation margin accounts for two-thirds of total margin, and Treasuries are more often used as initial margin.

Table 1 also shows the other side of margin requirements: notional exposures, defined as the sum of notional for derivatives carried as assets plus those carried as liabilities, based on daily data from 2014 through 2020 (we describe our treatment of post-2020 data below).

[2]Many CCPs also pay interest on USD cash balances. CCPs that are "designated financial market utilities" can maintain accounts at the Federal Reserves which are eligible to receive interest on reserve balances. CCPs can, in turn, pass though IORB subject to a spread on cash balances to their customers. These balances are included in the Federal Reserve's "other deposits' balance sheet item (Hull, 2024).

[3]See `https://www.cmegroup.com/solutions/clearing/financial-and-collateral-management/acceptable-collateral.html`.

Total notional exposure averages $219 trillion, mostly from desks trading rates ($120 trillion), FX ($44 trillion), and credit ($35 trillion). Figure 2 plots the time series of notional exposures. Banks post margin equal to 22 bps of their notional exposure on average, with cash accounting for 15 bps and Treasuries 3 bps. The bottom panel of Figure 2 shows that the margin ratio was perhaps falling until the Covid pandemic, then increased by 15 bps, and broadly stayed at higher average level following the pandemic, but with many local peaks and troughs. We provide additional details on how we estimate notional exposures in the Internet Appendix.

2.2 Secured Financing Transactions

Financial markets also consume collateral to raise funding through secured financing transactions (SFTs), which include repos, securities lending, collateral swaps, and other less common structures. Although participants use the different SFTs for different reasons, they all monetize the collateral services provided by the underlying security, often Treasuries. Much literature has documented the importance and behavior of these markets in recent years (Gorton and Metrick 2012, Copeland et al. 2012, Duffie 2018, Infante and Vardoulakis 2021, Chang et al. 2025). These collateral markets are extremely large, cumulatively amounting to $5 trillion, with the bulk coming from repos (see Figure IA.4).

An important complication is that a single CUSIP can back several SFTs simultaneously since banks reuse and rehypothecate collateral, sometimes several times over. Infante and Saravay (Forthcoming), for example, estimate that dealers pledge seven times more Treasuries than they own, made possibly by reusing and rehypothecating collateral the dealer receives itself. To isolate the contribution of collateral services to the convenience yield, it is important to control for these reuse dynamics, which we address next.

3 Model

We use a stylized model to analyze how collateral demands affect the convenience yield on Treasuries when they serve as collateral. The key friction is that posting Treasuries to meet collateral requirements—putting Treasuries into a collateral sink—sacrifices their money-like services, creating an endogenous opportunity cost that depends on the aggregate demand for collateral.

3.1 Environment

Consider a two-period economy with dates t and $t+1$. The economy consists of a representative household and a competitive set of dealers. Households trade three assets: a risky asset A with price p_A and gross payoff R_A, a Treasury security T with price p_T and gross payoff R_T, and a derivative contract D with per-unit gross payoff R_D where $\mathbb{E}_t[R_D] > 0$. The derivative has zero price at time t, reflecting its nature as a zero-cost position (such as a futures contract).

The derivative requires initial margin M per unit of derivative position D and accepts both Treasuries and risky assets as collateral, subject to haircuts h_T and h_A respectively, where $h_T < h_A < 1$. When an investor posts asset $i \in \{T, A\}$ as collateral, each unit delivers $(1 - h_i)p_i$ units of post-haircut margin.

A key friction in our model is that unencumbered Treasuries—those not posted as margin or otherwise committed—provide money-like convenience services to their holders. The household's preferences are given by:

$$U = \mathbb{E}_t[u(c_t) + \beta u(c_{t+1})] + \Omega(T^u), \tag{1}$$

where T^u denotes the quantity of unencumbered Treasuries, and $\Omega(\cdot)$ is an increasing and concave function capturing the non-pecuniary benefits of holding liquid, safe assets.

3.2 Household Problem

The representative household enters period t with an endowment y and chooses portfolio positions and collateral allocations. the household selects quantities of unencumbered and encumbered assets (A^u, T^u, A^m, T^m), where superscript u denotes unencumbered holdings and superscript m denotes assets posted as margin, with total holdings $A = A^u + A^m + s$ and $T = T^u + T^m$, where s denotes risky assets delivered to dealers for collateral swaps. The household also chooses its derivative position $d \geq 0$ and may engage in collateral swaps, obtaining $q \geq 0$ units of Treasuries via dealers to post as additional margin.

The household faces the following budget constraints:

$$\begin{aligned} c_t &= y - p_A A - p_T T - f(\chi)q, \\ c_{t+1} &= AR_A + TR_T + dR_D, \end{aligned}$$

where $f(\chi)q$ represents the fee paid for collateral swap services, with $f(\chi)$ denoting the per-unit fee that depends on the state of dealer balance sheets, parameterized by χ. Note that the household receives payoffs from all their asset holdings at $t+1$, including the s units temporarily posted to dealers, since swaps unwind at the end of period t.

The derivative's margin requirement takes the form:

$$(1 - h_T)p_T(T^m + q) + (1 - h_A)p_A A^m \geq Md. \tag{2}$$

When engaging in a collateral swap, the household delivers s units of the risky asset to a dealer and receives q units of Treasuries for posting, where the exchange ratio reflects dealer-specific haircuts (h_d^T, h_d^A):

$$q = \frac{(1 - h_d^A)p_A}{(1 - h_d^T)p_T}s.$$

Let m_{t+1} denote the household's stochastic discount factor, $\mu \geq 0$ the Lagrange multiplier on the margin constraint (2), and $\lambda_T \equiv \Omega'(T^u)$ the shadow value of holding an unencumbered Treasury. The household's optimization yields the following first-order conditions.

For unencumbered Treasuries, the standard Euler equation is modified by the convenience benefit:

$$1 = \mathbb{E}_t[m_{t+1}R_T] + \frac{\lambda_T}{p_T}. \tag{3}$$

This implies that the convenience yield on Treasuries, defined as $\mathrm{CY}_T \equiv \lambda_T/p_T$, equals the wedge between the required return on Treasuries and their pecuniary return: $\mathrm{CY}_T = 1 - \mathbb{E}_t[m_{t+1}R_T]$.

The decision to post Treasuries as margin reflects a trade-off between the margin value

and the forgone convenience services:

$$\mu(1-h_T)p_T = \begin{cases} \lambda_T, & \text{if } T^m > 0 \text{ and } q = 0, \\ \lambda_T + f(\chi), & \text{if } q > 0. \end{cases} \tag{4}$$

For risky collateral posted directly (when $A^m > 0$), which provides no convenience services:

$$\mu(1-h_A)p_A = 0. \tag{5}$$

The optimal derivative position satisfies:

$$\mathbb{E}_t[m_{t+1}R_D] = \mu M \quad \text{for } d > 0. \tag{6}$$

These conditions jointly determine the collateral selection rule: Treasuries are used for margin if their shadow posting cost (including any swap fees) does not exceed the shadow cost of posting risky assets.

3.3 Dealer Sector

A competitive dealer sector provides collateral transformation services through temporary securities lending. At the beginning of period t, each dealer chooses Treasury inventory ℓ to acquire at price p_T. The dealer then chooses swap volume $q \leq \ell$, temporarily lending q units of Treasuries to households who use them for margin posting. In exchange, the household temporarily delivers s units of the risky asset as collateral to the dealer. At the end of period t, the swap unwinds: the household returns the borrowed Treasuries to the dealer and receives back their risky assets. All payoffs at $t+1$ remain with the original owners—the dealer receives ℓR_T from their Treasury holdings, while the household receives the payoffs from their risky asset holdings.

The dealer's optimization problem is:

$$\max_{\ell,q} \quad \Pi = f(\chi)q + \mathbb{E}_t[m_{t+1}]\ell R_T - \ell p_T - \Psi(\ell;\chi) \quad \text{subject to} \quad 0 \leq q \leq \ell, \tag{7}$$

where the first term is swap fee revenue, the second term is the expected payoff from Treasury holdings (which the dealer retains regardless of swap activity), the third term is the cost of acquiring Treasury inventory, and $\Psi(\ell;\chi)$ represents convex balance-sheet costs with $\Psi_\ell > 0$ and $\Psi_{\ell\ell} > 0$.

Rearranging using the dealer's Euler equation for Treasuries, $p_T = \mathbb{E}t[mt + 1R_T]$ (since dealers do not derive convenience yields), the problem simplifies to:

$$\max_{\ell,q} \quad \Pi = f(\chi)q - \Psi(\ell;\chi) \quad \text{subject to} \quad 0 \le q \le \ell. \tag{8}$$

The first-order conditions are:

$$\text{w.r.t. } \ell: \quad \Psi_\ell(\ell;\chi) = \nu, \ \text{w.r.t. } q: \qquad\qquad f(\chi) = \nu, \tag{9}$$

where $\nu \ge 0$ is the multiplier on the constraint $q \le \ell$. When swaps are active ($q > 0$) and the constraint binds ($q = \ell$), we have:

$$f(\chi) = \Psi_\ell(\ell;\chi). \tag{10}$$

Competition drives the swap fee to the marginal balance-sheet cost. The state variable χ captures factors that tighten dealer balance sheets, with $f_\chi(\chi) > 0$ reflecting that tighter balance-sheet constraints increase the marginal cost of intermediation.

3.4 Equilibrium

An equilibrium consists of quantities $(A^u, T^u, A^m, T^m, d, q, s, \ell)$, shadow prices (μ, λ_T), and fee $f(\chi)$ such that: (i) the household's first-order conditions (3)–(6) and constraints are satisfied, (ii) dealer optimality conditions (7) hold, (iii) the collateral swap market clears with $q = \frac{(1-h_d^A)p_A}{(1-h_d^T)p_T}s$, and (iv) asset markets clear. The government's Treasury supply is fixed at Θ, with market clearing in the Treasury market requiring:

$$\Theta = T^u + T^m + \ell,$$

where ℓ represents the dealers' total Treasury holdings. The risky asset market clears with the household holding $A = A^u + A^m + s$ units total.

3.5 Testable Implications

The model generates four main empirical predictions relating margin demands to Treasury convenience yields and collateral composition. All comparative statics are local, assuming interior solutions and positive convenience yields.

Prediction 1. *Convenience Yields Increase with Margin Demands.*

An increase in the aggregate margin sink Md raises the shadow value of margin capacity μ through equation (5). Via equation (3), this reduces unencumbered Treasury holdings T^u and increases their shadow value λ_T, leading to:

$$\frac{\partial \mathrm{CY}_T}{\partial (Md)} > 0, \quad \frac{\partial \mathrm{CY}_T}{\partial h_T} > 0, \quad \frac{\partial \mathrm{CY}_T}{\partial h_A} < 0. \tag{11}$$

Higher Treasury haircuts increase convenience yields by reducing the margin capacity per unit of Treasury posted, while higher haircuts on risky assets have the opposite effect by making Treasuries relatively more attractive as collateral.

Prediction 2. *Pass-Through Intensifies with Dealer Balance-Sheet Tightness.*

When collateral swaps are active, the sensitivity of convenience yields to margin demands increases with dealer balance-sheet costs:

$$\frac{\partial \mathrm{CY}_T}{\partial S} = (1 - h_T)\frac{\partial \mu}{\partial S}, \quad \text{where } \frac{\partial \mu}{\partial S} \text{ increases in } f(\chi). \tag{12}$$

Tighter dealer balance sheets reduce swap activity, causing unencumbered Treasury holdings to decline more rapidly per unit increase in margin demands, thereby steepening the relationship between margin sinks and convenience yields.

Prediction 3. *Collateral Composition Responds to Convenience Yields.*

Define $s_T \equiv \frac{(1-h_T)p_T(T^m+q)}{Md}$ as the share of margin covered by Treasuries. Holding the aggregate margin demand Md fixed:

$$\frac{\partial s_T}{\partial \mathrm{CY}_T} < 0, \quad \frac{\partial s_T}{\partial h_T} < 0. \tag{13}$$

Higher convenience yields reduce Treasury usage for margin as their opportunity cost rises, while higher Treasury haircuts reduce their attractiveness as collateral.

Prediction 4. *Shadow Posting Costs Drive Collateral Substitution.*

The total cost of delivering one unit of margin using Treasuries equals:

$$\frac{\lambda_T}{(1-h_T)p_T} + \frac{f(\chi)}{(1-h_T)p_T}, \tag{14}$$

where the first term represents the opportunity cost of forgone convenience services and the second term captures swap fees when Treasuries are obtained through dealers. Even when Treasuries are not actively posted, higher convenience yields shift the optimal collateral mix toward eligible assets with lower opportunity costs.

4 Data

The paper relies primarily on three datasets: trading desk data (FR VV-1), high frequency balance sheet data (FR2052a), and a wide set of Treasury convenience yield estimates calculated by Acharya and Laarits (2025).

We collect notional derivative exposure data from FR VV-1, which the Fed collects to implement the Volcker Rule. Banks with at least $20 billion in average gross trading assets and liabilities—excluding trading in U.S. Treasuries and agencies—are required to report daily information on the value-at-risk, profitability, and exposures to the Fed. Our main panel is at the date by bank by trading desk level, running from 2014 through 2024. The data has comprehensive coverage of trading activity by banks. While fewer than 20 firms are required to report metrics, these firms span 93 percent of U.S-based trading activity. Recent research has used this data to study banks' trading activities (Falato et al. 2025, Lu and Wallen 2024). We limit our analysis to the set of banks that consistently provide data in both datasets, which includes U.S. GSIBs.

The FR2052a data provides high frequency, granular data on the largest banks' consolidated balance sheets, which the Fed uses to implement liquidity regulations. We focus on data at the bank-holding company; for brevity, we will refer to BHCs as banks. The data includes quantities on exposures but not prices or rates. We filter the data in two substantive

ways: first, we focus on the set of banks that consistently file at a daily frequency, which include the six largest U.S. banks: Bank of America, Citigroup, Goldman Sachs, JP Morgan, Morgan Stanley, and Wells Fargo. We provide additional details on the data in the Internet Appendix.

The literature has studied several proxies of the Treasury convenience yield. We aggregate across several proxies by focusing on the first principal component, PC1, of the proxies following Acharya and Laarits (2025) using data from 2005–2024:

1. General collateral finance repo – Treasury bill spread
2. OIS – Treasury bill spread
3. Fed funds – Treasury bill spread
4. Negative Z-spread from Greenwood et al. (2015)
5. 10-year TIPS–Treasury premium from Fleckenstein et al. (2014)
6. 30-year OIS swap spread from Feldhütter and Lando (2008) and Du et al. (2023)

We select these series because they have data for the full period 2005 to 2024. For readability, we rescale the first principal component into the units of the 3-month OIS/Treasury spread by regressing the spread on the component and using the fitted values as the rescaled principal component. The Internet Appendix provides details. We provide additional data on these measures and other control variables in the Internet Appendix.

5 Collateral Sinks

We now estimate the volume of collateral *consumed* by the financial system. We say that collateral has been consumed if it is held in a *collateral sink*. Because collateral can be reused, repledged, and rehypothecated, pledging $1 of collateral does not imply that the pledgeable collateral stock has fallen by $1. We expect that convenience yields will respond to the supply of collateral that is pledgeable, regardless of whether that collateral has already been pledged, perhaps even many times over. Collateral sinks, by contrast, represents locations where collateral cannot be reused, and hence is removed from the stock of circulating collateral,

even if briefly. This section shows that collateral sinks are economically large—standing at $2.2 trillion—which provides motivation that changes in sunk collateral corresponds to changes in aggregate collateral and its attendant effect on convenience yields.

5.1 Construction

We measure collateral sinks as the sum of collateral posted in two separate markets: margin and SFTs. We limit ourselves to the subsets of these markets where we have evidence or strong priors that the collateral stop circulating.

SFT Collateral Sinks Intuitively, many types of SFT counterparties constitute collateral sinks. For example, when a bank pledges collateral to a money fund as part of a repo, the money fund does not rehypothecate the collateral, hence money funds are collateral sinks. While the data is not sufficiently granular to identify the SFT counterparty's identity, it does provide counterparty buckets.[4] We estimate SFT collateral sinks by comparing the collateral flows between banks and that counterparty type; we identify SFT collateral sinks as counterparties that receive collateral from banks but do not post collateral to banks.

For each counterparty type c and settlement method s (bilateral, FICC, triparty, or other), we calculate the collateral flow imbalance after aggregating across all banks with

$$\text{Imbalance}_{t,c,s} = \frac{\text{Pledged}_{t,c,s} - \text{Received}_{t,c,s}}{\text{Pledged}_{t,c,s} + \text{Received}_{t,c,s}},$$

where pledged and received are from the bank's perspective. We include settlement type because some settlement types are more likely to be used in markets in which counterparties aggressively reuse collateral (bilateral) while others less so (triparty). If $\text{Imbalance}_{t,c,s} = 1$ then all collateral flows between the banks and that counterparty type are collateral pledged from the bank to the counterparty; the flows are strictly one way. By contrast, if the bank is pledging and receiving the same amount, $\text{Imbalance}_{t,c,s} = 0$.

We define collateral sinks as cells with $\text{Imbalance}_{t,c,s} > 0.5$, however we test several

[4]There are 19 counterparty types: bank, broker dealer, Central bank, debt issuing special purpose entity, financial market utility, GSE, investment company or advisor, multilateral development bank, non-financial corporate, non-regulated fund, other, other financial entity, other supervised non-bank financial entity, other supranational, pension fund, public sector entity, retail, sovereign, and supervised non-bank financial entity.

other thresholds and show that our main estimate is closely correlated with looser or tighter restrictions (e.g., a positive imbalance, an imbalance of at least 0.75 or 0.9) given the large share of cells clustered near 1. Empirically, most counterparty types only receive collateral from banks, as shown in the histogram of Figure 3, indicating that many counterparty types are clearly collateral sinks. For example, the median imbalance for both non-financial corporate and GSEs in the triparty market is 1. By contrast, collateral flows with dealers are two-sided, and the median imbalance is about 0 for "supervised non-bank financial entities" in FICC, which includes broker dealers, which represents a sizeable share of interdealer trading. The Internet Appendix provides medians by counterparty and settlement type (Figure IA.5).

Margin Collateral Sinks We estimate margin-induced collateral sinks as the sum of initial and variation margin received by or posted to the house account that satisfies any of the following conditions:

- it is not rehypothecable;
- it is held as segregated cash;
- it uses centralized settlement (including principal and agent transactions); or
- it is exchange traded (including principal and agent transactions).

We assume that margin posted to a CCP is not rehypothecated, even if in some cases the CCP may retain the contractual right to do so.[5] We could also include other CCP contributions—like clearing members' default fund contributions—but this data is available beginning only in 2022, it is many magnitudes smaller than the other categories, and it is relatively slower moving, so its dynamics will marginally affect the level but not the daily variation in sunk collateral.

A basic question is why we include cash in our *collateral* sink measure. The answer is because collateral is interchangeable in almost all of the settings we define as collateral sinks, as described in prediction 4.[6] So even if a bank satisfies the collateral demand with $1 of

[5]Regulations require that certain types of derivatives are centrally cleared even though they are negotiated OTC. Derivatives exchanges all use a CCP.

[6]While variation is typically cash, it is not always required to be cash. Empirically, about a quarter of variation margin in the sample is posted in non-cash form; the data does not allow us to distinguish when cash is the only permissible way to post variation margin, although some CCPs require this.

cash, it could have also done it with a Treasury. Higher aggregate demand in collateral sinks will make the collateral services component of Treasuries more valuable because opportunity costs increase even if Treasuries are not actually placed in the sink.

Measuring collateral sinks is hard because we observe their market value, not separate quantities and prices. A sink can grow either because more collateral is locked in or because prices rise on a fixed quantity. We address this in three ways. First, our main regressors scale the sink by the market value of publicly held Treasuries outstanding, which we define with S_t.[7] This controls for trend growth in the system and for price-level movements in Treasuries. Second, margin-driven sinks are less sensitive to pure price moves because margin is set against the market value of posted collateral; if prices rise enough, existing collateral can satisfy higher requirements without new posting. Third, for SFTs we focus on the maturity (cash) value of the trade rather than the market value of the pledged security. The haircut equals the difference between the collateral's market value and the maturity value; once an SFT starts, its maturity value is fixed, and different collateral types can secure comparable post-haircut funding. Using maturity value therefore strips out interim price moves and lets us compare collateral across types on a like-for-like, post-haircut basis.

5.2 Collateral Sink Summary Statistics

Collateral sinks stand at $2.2 trillion in 2024, equal to 10 percent of Treasuries outstanding, as shown in Figure 4. SFTs account for more than half of the total, with repos/sec lending averaging $700 billion (49 percent of sunk collateral) and collateral swaps $210 billion (15 percent). Margin averages $510 billion (36 percent of collateral sinks), although its share has approached 50 percent of the total at times.

Is $2.2 trillion of sunk collateral a big number? We argue yes for two reasons. First, it is equal to 9 percent of the total stock of privately held Treasuries and about 30 percent more than the total Treasury holdings at private depository institutions. To be sure, most of the sunk collateral is not Treasuries, yet much of it is swappable to Treasuries and hence they represent a source of incremental collateral demand for Treasuries. Second, it is a lower bound estimate because it does not include data from any banks outside the six in our sample, it

[7] We describe how we calculate the market value of Treasuries in the Internet Appendix. In particular, we use a five-day moving average ($t-4$ to t) to remove mechanical variation stemming from bill issuance and SOMA weekly Wednesday portfolio disclosures.

does not include collateral received via SFTs which is not rehypothecable, nor does it include margin posted by banks' clients (due to data availability) or banks' other CCP contributions. A key assumption for our empirics is that collateral sink dynamics for the largest banks are representative of the broader financial system, which is reasonable because the largest U.S. banks are large components of both the SFT market and post a disproportionate share of margin to CCPs.

Table 2 provides the composition of sunk collateral. Treasuries account for 24 percent, other HQLA 36 percent, cash 16 percent, and other non-HQLA 24 percent. Sunk Treasuries are the most volatile, both in level terms ($142 billion) and relative to its mean (44 percent). This is suggestive evidence that banks use Treasuries to satisfy volatile aggregate collateral demands on the margin, which we study in detail in the next section.

5.3 Collateral Sinks and Convenience Yields

Prediction 1 shows that convenience yields rise when more collateral is sunk. Table 3 formalizes this prediction by regressing PC1 on S_t along with a battery of controls. In levels, a 1 pp increase in S_t ($\approx$ 1.25 SD) is associated with a 8 bps higher PC1 ($\approx$ 1 SD). The relationship between the suggests that collateral sinks track aggregate collateral demand insofar as it varies with the convenience yield.

The levels regression is not causal, and does not reject the possibility that more collateral is sunk simultaneously with periods with higher convenience yields, meaning they both happen to be high at the same time rather than reflecting some underlying relationship. If volatility increases, margin requirements can increase while demand for safety pushes the convenience yield up without a causal relationship between them.

We can improve the regression by focusing on daily changes, which removes time trends and better estimates the relationship by looking at high frequency changes in the two variables. Table 3 also compare changes on changes; the full specification in column (6) shows that a 1pp increase in ΔS_t raises $\Delta PC1_t$ by 4.5 bps after controls, equal to over one standard deviation. While a more reliable estimate of the relationship between aggregate collateral demand and convenience yields, it is still imperfect; it is not causal. In the next section we provide better identified results of this regression.

Which component of collateral sinks drive the result? Table 4 breaks the aggregate

measure into its separate components—repo/sec lending, collateral swaps, and margin. All three variables have positive and roughly similar magnitude effects, and both the margin and collateral swap components are significantly different from zero, while standard error for collateral swaps is somewhat larger even though its magnitude is similar to the others. Because the margin component does not rely on the imbalance filter, it is perhaps unsurprising that the margin component is estimated most precisely whereas the other two components are measured with noise that might attenuate their coefficients.

We also provide suggestive evidence for the portfolio-substitution prediction that Treasuries' share of sunk collateral falls as the convenience yield rises, since posting a high-convenience Treasury is more costly at the margin. Indeed, the September 2025 Senior Credit Officer Survey's special questions report a broad shift toward using securities, namely Treasuries, for variation margin, with net increases across dealer clients and one-third of accepting dealers expecting further growth over the next year.[8] In our data, the Treasury share of margin posted reaches a new high in 2025 (Figure IA.6), and shows a strong negative relationship between that share and PC1 as illustrated in Figure 5, confirming that lower convenience yields coincides with more use of Treasuries as collateral.[9]

6 Identifying the Effects of Aggregate Collateral Demand Shocks

The challenge to studying the effect of aggregate collateral demand on the convenience yield is that much variation in the two is due to other forces like general risk appetite and financial market conditions. Although we have seen that the two are closely correlated in Table 3, we now use two separate approaches to isolate plausibly exogenous shocks to collateral dynamics to trace out its effects on convenience yields: a IV based on idiosyncratic margin choices by the CME and a granular IV (GIV) using trading desk flows.

[8] See `https://www.federalreserve.gov/data/scoos/files/scoos_202509.pdf`.
[9] We give the correlation of these variables in Table IA.2.

6.1 Idiosyncratic Margin Requirements

Margin requirements principally vary with the derivative's underlying risk, yet margin-setters still have idiosyncratic preferences that may drive a wedge between the actuarial "fair" margin requirement and the actual requirement. We focus on the idiosyncratic margin choices of CME, possibly the largest single margin-setter. While CME's margin requirements closely follow volatility measures, there are nontrivial and time-varying residuals unexplained by a wide set of aggregate factors. Our basic identifying assumption is that these residuals are not systematically correlated with unobservable aggregate factors but instead correlated with the CME's distinct risk management preferences.

To be sure, CME's margin requirements reflect a complex set of house rules. The calculation depends on volatility and the range of expected price ranges, look-back windows and stress add-ons, floors and buffers and assumptions about how quickly positions can unwind under stress. There are also product-specific charges related to delivery and concentration, and adjustments stemming risks faced by similar products. The factors are also updated on a discrete schedule that introduces step changes through time. The combination of these factors means that CME margins depend on more than simple actuarial estimates of the product's volatility or other standard risk measures—measures that we think could drive the convenience yield. Therefore we have good reason to believe that the residual from a regression of margin requirements on standard risk measures is not hidden macro risk, but an artifact of CME's risk-management preferences for tail hedging, smoothing through the cycle, and microstructure frictions.

We can quickly confirm this by regressing the change in margin requirements for E-mini S&P500 futures on changes in the VIX, which yields an R^2 of 40 percent. But splitting the regression to reflect the sign of the change in the VIX we see a clear a asymmetric response in margin requirements: a 1 unit increase in the VIX is associated with a 3.6pp increase in margin requirements, while a 1 unit decrease in the VIX decreases margin requirements by 3pp. Margin requirements increase 21 percent faster than they decrease for identically-sized moves in the VIX. See Table IA.3.

We exploit this fact in a two step process. First, we estimate the residuals for several benchmark CME products' margin requirements after controlling for a wide set of aggregate

risk measures:

$$\Delta(\text{Margin Requirement})_{c,t} = \alpha + \beta' X_t + \varepsilon_{c,t}$$

where X_t is a vector of controls and c denotes a contract. When margin requirements increase by more than the factors would predict, then $\varepsilon_{c,t} > 0$. We focus on margin requirements for six top-tier futures contracts across several asset classes: commodities (WTI Crude), equities (E-mini S&P500, E-mini Dow), FX (Euro/USD), and rates (2 and 10-year Treasury notes). For each contract, we regress the daily changes in its margin requirements on the following controls (all in daily changes): VIX, Treasury implied volatility, FX implied volatility, crude oil ETF implied volatility, S&P500 return, effective funds rate, Baa/Aaa spread, U.S. 5-year CDS spread, 5-year Treasury bid-ask spread, and the 5-year/5-year inflation breakeven rate. We describe the variables and their construction in IA.A. We winsorize the residuals at the 5th and 95th percentile to reduce the influence of outliers. We aggregate the residuals to a single daily time series of margin shocks by either taking an average of the z-scored residuals, Z_t^{Avg}, or by calculating the first principal component across them on a given day, Z_t^{PC}.

We then use the shocks as instruments for the change in the collateral sink ratio in a two-step process. First, we regress ΔS_t, the change in the collateral sink ratio, on the margin shocks:

$$\Delta S_t = \alpha_1 + \pi Z_t + \delta' X_t + u_t.$$

Then, we regress the first principal component of the convenience yield on the estimated $\widehat{\Delta S}_t$

$$\Delta \text{PC1}_t = \alpha_2 + \beta \widehat{\Delta S}_t + \kappa' X_t + \varepsilon_t.$$

To make interpretation easier, we standardize both Z_t and the collateral sink ratio to have mean 0 and unit standard deviation.

The top panel of Table 5 shows the second stage results, with the first three columns showing results using Z_t^{Avg}, the last three using Z_t^{PC}, and the columns vary in the set of second-stage controls. Using the estimates in column (3), a 1 standard deviation increase in the instrumented change in the collateral sink ratio (6 bps) raises $\Delta PC1$ by 2 basis points, roughly half of its standard deviation. The coefficient is similar across the specifications and

statistically significant, ranging from 2 to 2.5.

The second panel of Table 5 shows the first stage results. The relevance condition requires that the instrument is sufficiently related to the instrumented variable which is confirmed by the F-statistics shown in the bottom row. Moreover, the table confirms the economic intuition that tighter-than-expected margin should increase sunk collateral rather than decrease it. We expect this: the direct reason is that sunk collateral includes margin posted directly by the banks, so it should mechanically increase when they are required to post more margin. Further, our measure of sunk collateral includes collateral swaps and other transactions through which the banks' counterparties themselves obtain collateral to post as margin. As robustness, the Internet Appendix regresses both Z_t variables on different measures of posted margin and finds a strong direct relationship (Table IA.4).

Is it plausible that CME margin requirements alone can drive aggregate collateral dynamics? We believe so for two reasons. First, the CME is consistently the first or second largest CCP globally with respect to total required initial margin, hence its decisions have material effects on aggregate collateral demand. Second, even though the CME only sets margin requirements for its contract, CCPs compete with one another to attract volume and hence changes in the CME's requirements are likely tightly correlated with margin requirements at other CCPs (Park and Abruzzo, 2016). Anecdotal evidence also suggests that bilateral exposures also, in part, set margin in reference to information implicitly provided in CME's margin requirements, so the innovations to the CME's margin requirements are also likely closely correlated with margin conditions for bilateral contracts.

The instruments are constructed from contract-level residuals conditional on a wide set of aggregate risk controls, so the residuals reflect CME-specific risk management choices rather than economy-wide risk. Therefore, the residuals shift collateral demand mechanically through margin requirements but, conditional on X_t, should not affect convenience yields except through its effect on sunk collateral, as necessary for the exclusion restriction. As a check, we include a placebo regression where we swap the instrument Z_t with its 30-day lead, Z_{t+30}, and find no effects in either stage (Table IA.5).

Our IV estimates of the effect of collateral sinks on the convenience yield is approximately an order of magnitude larger than the corresponding OLS estimate. Several forms of endogeneity could in principle bias the OLS coefficient. Measurement error in our proxy for aggregate collateral demand would attenuate the OLS estimate toward zero. At the same

time, simultaneity or reverse causality would likely push in the opposite direction: periods of higher convenience yields could coincide with greater margin posting, biasing OLS upward. We interpret $\beta_{IV} > \beta_{OLS}$ as evidence that attenuation bias is the dominating effect. Our daily collateral sink measure is necessarily a proxy for the true unobserved and market-wide aggregate collateral sink; we observe only a handful of large U.S. banks that, while capturing a large share of collateral dynamics, does not include foreign banks or smaller U.S. banks. Our IVs—either based on the CME margin changes or using trading desk flows in the next section—better isolate the exogenous component of the variation in collateral demand and minimizes the noisy error-driven variation in the collateral sink proxy, thereby correcting for the attenuation bias inherent in the OLS estimate.

6.2 Idiosyncratic Shocks from Trading Desks

We now provide an alternative approach to estimating the effect of collateral demand on convenience yields using trading desk data to form a GIV. Insofar as different trading desks receive idiosyncratic derivative trading flows, and so long as some of the trading desks account for a disproportionately large share of total flows, then the idiosyncratic flows will not be cancelled out in aggregate (Gabaix and Koijen, 2024). Suppose there are 100 trading desks across banks—each desk specializing in in a different asset class or product—and each desks' trading flows depend on a combination of aggregate and idiosyncratic factors. If the desks were equally sized, then idiosyncratic shocks to the trading flows on each would, on average, cancel out. However, if a handful of those desks represent a disproportionately large share of total trading flows, then the idiosyncratic shocks to those desks will not cancel out in aggregate. Since we capture flows using derivative gross notional, larger notional amounts require larger posted margin. In this case, we can trace out the effect of idiosyncratic shocks from trading flow—equal to an idiosyncratic shock to collateral demand via posted margin—on convenience yields.

We filter the trading data in several ways to create a comparable panel of trading desks. In particular, we limit the sample to years 2016 through 2020 when the data consistently provides the level of notional derivative exposures across the full set of banks, however there are some days when data is missing for one or more desks, so the panel is not strictly balanced. We require the desk has a non-zero derivatives position at some point in the sample, that the

desk reports data consistently through the sample, and that the desk averages at least $5 million in value-at-risk. The resulting sample spans 93 desks across eight asset classes and six banks, implying an average of 16 desks per bank. Each resulting desk by date observation includes the desk's total value-at-risk (VaR), the notional of its derivatives carried as assets and liabilities, and the value of long and short securities positions. We calculate the change in (1+ the log of) each of these variables, and we winsorize the daily changes at the 5th and 95th percentile.

Derivative trading data are measured in gross notional, but margin requirements scale with risk and vary depending on the derivative's risk; $1 notional at a relatively low risk desk does not increase margin demand the same way a $1 exposure does at a higher risk desk. Treating the notional flows across all desks equally risks biasing the results towards those with larger notional flows, not necessarily those that require more margin. We therefore convert each desk's notional exposure into a margin-equivalent measure based on how that desk's overall VaR varies with its derivatives positions, holding other variables fixed. For each desk i within an asset class a we estimate the contribution of derivatives and securities positions on VaR using

$$\begin{aligned}\Delta \ln(\text{VaR}_{i,t}) = {} & \beta^a_{\text{sec},A} \Delta \ln(\text{SecAsset}_{i,t}) + \beta^a_{\text{sec},L} \Delta \ln(\text{SecLiab}_{i,t}) \\ & + \beta^a_{\text{der},A} \Delta \ln(\text{DerivAsset}_{i,t}) + \beta^a_{\text{der},L} \Delta \ln(\text{DerivLiab}_{i,t}) \\ & + \mu_i + \lambda_{b(i),t} + \eta_{a,t} + \varepsilon_{i,t}.\end{aligned} \tag{15}$$

We then construct the derivative-driven change in VaR using the estimated relationship for each asset class

$$\Delta \ln(\text{VaR}^{\text{deriv}}_{i,t}) = \hat{\beta}^a_{\text{der},A} \Delta \ln(\text{DerivAsset}_{i,t}) + \hat{\beta}^a_{\text{der},L} \Delta \ln(\text{DerivLiab}_{i,t}). \tag{16}$$

This construction isolates the risk stemming from derivative trading activity, holding fixed the effect of the desks' changing securities portfolios. The key assumption in this construction is that the VaR attributable to derivatives is proportional to margin demand, since we do not observe desk-specific margin.

We then extract the idiosyncratic component of derivatives-induced VaR using the date

by desk panel with

$$\Delta \ln(\text{VaR}_{i,t}^{deriv}) = \alpha + \mu_i + \lambda_{b,t} + \eta_{a,t} + u_{i,t},$$

where μ_i are desk fixed effects, $\lambda_{b,t}$ are bank by date fixed effects, and $\eta_{a,t}$ are asset class by date fixed effects. The fixed effects play an important role in stripping out non-idiosyncratic effects, including persistent differences across desks, shocks common to all desks within a bank on a given date like funding shocks or bank-wide risk adjustments, and shocks common to all desks within an asset class across banks. After stripping out these effects, $u_{i,t}$ is designed to capture the variation, which we expect reflects idiosyncratic variation. We winsorize $u_{i,t}$ at the 5th and 95th levels to reduce the influence of outliers.

We calculate the GIV by comparing value-weighted residuals and equal-weighted residuals

$$GIV_t = \sum_i w_{i,t-1} u_{i,t} - \frac{1}{N} \sum_i u_{i,t}$$

where the weights are estimated using the desks' average risk-adjusted notional in the previous week ($s_{i,t-1}$) to ensure pre-determination:

$$s_{i,t-1} = \left|\hat{\beta}_{\text{der},A}^{a(i)}\right| \text{DerivAssetNotional}_{i,t-1} + \left|\hat{\beta}_{\text{der},L}^{a(i)}\right| \text{DerivLiabNotional}_{i,t-1}, \tag{17}$$

$$w_{i,t} = \frac{s_{i,t-1}}{\sum_j s_{j,t-1}}. \tag{18}$$

The GIV increases when large desks receive idiosyncratic shocks not offset by average shocks captured in the equal-weighted term. We use absolute values of the betas to prevent negative weights, however excluding observations with a negative weight does not meaningfully change our results. We provide the beta estimates by asset class (Figure IA.7) and a plot of GIV (Figure IA.8) in the Internet Appendix.

We then proceed with the two-stage least squares: the first stage shows how sunk collateral varies as a result of idiosyncratic derivative trading flows:

$$\Delta S_t = \alpha_1 + \pi GIV_t + \delta' X_t + e_t.$$

The second stage regresses the change in convenience yields on the instrumented collateral

sink measure

$$\Delta PC1 = \alpha_2 + \gamma \widehat{\Delta S_t} + \kappa' X_t + \varepsilon_t.$$

For additional robustness, we also include a set of results that swaps the collateral sink ratio with $\Delta \ln(\text{Sunk Margin}_t)$. To make interpretation easier, we standardize GIV, the collateral sink ratio, and the change in log margin posted to have mean 0 and unit standard deviation.

Panel A of Table 6 reports the second stage estimates.[10] Because we cast $\Delta PC1$ to basis points (scaled to the 3-month OIS-Tbill spread) and standardize the endogenous regression, a one standard deviation increase in the instrumented collateral sink ratio increases the convenience yield by 1.9 to 2.6 basis points, depending on the control specification. The standard deviation of $\Delta PC1$ is 4.4 basis points, so the effect of a one-standard deviation increase in the instrumented ratio raises convenience yields by 0.4 to 0.6 standard deviations. The first stage implies that a one standard deviation increase in GIV lifts the sink ratio by 0.14 standard deviations, so in reduced form units the same 1σ shock raises $\Delta PC1$ by 0.3 basis points (e.g., 0.14×1.9). Using sunk margin yields similar magnitudes. Notably, the estimated magnitudes are especially close to those estimated in the CME setting, providing further support for the validity of the estimates.

Panel B of Table 6 reports the first stage estimates. The first three columns use the change in the collateral sink ratio as the endogenous regressor, while the last three repeat the design using for robustness. The first stage shows a strong positive relationship between the collateral sink ratio and the GIV across control specifications: a 1 standard deviation increase in GIV corresponds to an increase of about 0.12 standard deviations in both the endogenous regressors with large F-statistics, ranging from 15 to 34, with the lowest values in the columns without controls. Given that the GIV is designed to identify collateral demand shocks stemming from margin requirements, it is not surprising that the F-statistics for the more direct endogenous regressor—sunk margin—is consistently larger by a factor of about three.

Panel C of Table 6 reports the OLS estimates. We again see that the instrumented coefficients are a magnitude larger than the OLS estimates, and the coefficients are positive but not significant. As discussed in the previous setting with CME margin requirements, we

[10]We provide tables reporting all controls for the GIV in the Additional Exhibits section.

believe this likely reflects measurement error in our collateral sink measure that attenuates the OLS relationship: all collateral in sinks should affect convenience yields, yet we only observe those that are reported by six banks. Moreover, in Panel C, it is further limited to the same sample when the GIV is available and therefore loses more than 4 years of daily observations relative to the CME margin approach.

The validity of the GIV relies on three key identifying assumptions: its granularity, the residualization of aggregate factors, and our estimate of margin requirements by desk. First, the approach requires that trading desks vary substantially in size. This is empirically true: on average, the largest 20 desks account for roughly 50 percent of total VaR and 90 percent of our risk-adjusted size measure $w_{i,t-1}$.

Second, the approach requires that the residualization accurately strips out aggregate factors. Suppose it did not. In that case, the GIV would be correlated with some aggregate factor—say, risk sentiment—so that our regressions would only tell us that increased risk sentiment coincides with both more sunk collateral and larger convenience yields; the result would be biased by omitted variable and simultaneity bias. We implement a simple placebo: if the GIV simply reflects aggregate factors, then the weighting is immaterial. We construct a placebo GIV that reverses the order of the weights, assigning the largest desk's weight to the smallest desk, the second-largest desk's weight to the second-smallest desk, and so on. We then calculate the placebo GIV in the same way. Table IA.6 shows that the placebo GIV does not generate similar results and Table IA.7 shows that the two instruments are not correlated, both indicating that is does a good job of stripping out aggregate factors and isolating truly idiosyncratic factors.

Third, we assume that desks with more VaR attributable to derivative positions have larger margin requirements. Were the universe of the derivatives trades on centrally cleared platforms, this would be trivially true. Instead, trading desks face several types of sophisticated counterparties, some of which may bilaterally negotiate lower margin requirements. Ideally, we could simply look at margin posted by trading desks, but that data is unavailable. In that case, the concern is that larger gross notional positions for the bank's trading desks could still coincide with lower margin requirements if the composition of their counterparties dramatically shifts to those that require less margin. We empirically reject this by regressing total margin posted by the banks—which, unlike sunk margin, includes only margin posted by banks rather than margin posted to or by banks—on our risk-adjusted notional $s_{i,t}$ and

find a strong positive relationship (Table IA.8).

7 Dealer Dynamics and Collateral Swaps

We now show how dealers use *collateral swaps* to alleviate the effect of collateral sinks on convenience yields. When banks are unconstrained, they can intermediate collateral markets, attenuating the relationship between sunk collateral and convenience yields. But when they grow constrained, banks stop buffering collateral scarcity, driving convenience yields up, as described in Prediction 2.

7.1 Institutional Background

Collateral swaps replaces one type of collateral with another, either explicitly using a collateral swap or implicitly by using a combination of repos. Investors can engage in collateral swaps to obtain collateral for margin. The market has grown dramatically as post-crisis reforms increased collateral demand without an attendant increase in collateral issuance, with more than $1 trillion of swaps outstanding in recent years.

We trace out the logic of a collateral swap in Figure 6. An investor facing a margin call—say an equity investor—can post either its own equities at a 30 percent haircut or, using a collateral swap, Treasuries at a 2 percent haircut. So long as the dealer's haircut on equities plus its fee combine to create a lower effective haircut than the CCP's haircut, the investor can use the collateral swap to more efficiently satisfy its margin call.

The dealer providing a collateral swap can source the collateral from a supplier—perhaps an insurance company or a pension fund—or from its own inventory; the figure shows the former example. The dealer can also use the collateral flows and fees to finance its own inventories. Importantly, the collateral swap has two legs: an upgrade and a downgrade. The dealer engages in a collateral *downgrade* when it receives equities and pledges Treasuries, shown in the left green box. By contrast, the dealer engages in a collateral *upgrade* when it pledges the equities and receives Treasuries, shown on right green box. Some investors likely also prefer to post Treasuries as collateral even when they have cash to avoid counterparty risk exposures to the CCP itself.

Investors can engage in collateral swaps for several reasons, including to obtain collateral

for margin, obtain collateral for delivery, and short sell. We show that margin demands are an important driver of this market. Collateral swaps share many features with repos, and the collateral swap shown in Figure 6 can be implemented in simultaneous, matched book repos, but requires the use of cash in intermediate steps, introducing more operational complexity to the transactions. Notably, repos and collateral swaps differ because repos transform assets—typically safe assets like Treasuries—into cash; collateral swaps directly swap one type of asset for another.

We plot the total size of the collateral swap market in Figure 7, showing the market has grown substantially over the past decade, comparable to the total quantity of margin posted by banks and their customers.[11] The figure shows the amount of collateral pledged and received by banks are roughly equal, indicating that banks largely intermediate the market. But the fact that collateral pledged consistently slightly exceeds collateral received indicates that banks pledge some of their own inventory to swaps. Table 7 shows that sovereign bonds constitute the majority of total collateral pledged by banks in swaps, consistent with their counterparty using the swaps to acquire high quality collateral. Equities are also pledged in large volumes, likely for short sales, although over the past decade the pledging of high quality government bonds has grown nearly double the pace of equities, showing that collateral swaps are an increasingly important tool to source high quality collateral rather than short sales. On average, banks do not use the swap market to source collateral for themselves, since the net pledged values on the bottom of the table are positive.[12]

The top half of Figure 8 shows that banks largely pledge collateral through the swap to banks ($351 bn) and investment companies ($113 bn). The bottom half of Figure 8 shows that the swaps transfer high quality collateral from insurance companies and public sector firms to intermediaries and non-regulated funds.

7.2 Haircut Arbitrage

The figure makes clear that the key feature of collateral swaps as a collateral transformation service is haircut arbitrage in two dimensions: (1) it is only valuable when the dealers can

[11] Figure IA.11 compares the size of the repo market, collateral swap market, client long and short positions, and posted margin.

[12] The Internet Appendix also provides a timeseries of collateral upgrades and downgrades from the banks' perspectives, and figures that show the net pledged position by banks for all collateral types and Treasuries.

charge lower haircuts than the CCP, which they typically do; and (2) dealers are willing to offer lower haircuts only when they, themselves, can borrow at even lower haircuts.

Collateral swaps are economically viable so long as they reduce a derivative investor's capital requirements by arbitraging the difference between haircuts offered by the CCP and the dealer. We can price a collateral swap by no arbitrage by comparing the two options the investor faces when facing a margin requirement of M:

1. Post low quality collateral to the CCP with value V_ℓ at haircut h^ℓ_{ccp}, fund the shortfall at the investor's cost of capital r_c less the CCP's rate paid on posted cash r_m

2. Engage in a collateral swap with a dealer: posting low quality collateral to the dealer at h^ℓ_d in exchange for Treasuries; post the Treasuries to the CCP at haircut h^t_{ccp}, funding the haircut at the investor's cost of capital r_c less the CCP's interest paid on cash r_m and paying the fee quoted on the notional of Treasuries received[13]

For simplicity, we assume the dealers sizes the Treasuries delivered in the collateral swap to be equal across the two collateral types, meaning there is no cash exchanged for haircut differentials between the investor and the dealer at the initial leg. By no arbitrage, we can write the fee at which the investor would be indifferent between posting to the CCP and using the collateral swap by equating the returns from the two trades:

$$\underbrace{\left(M - (1-h^\ell_{ccp})V_\ell\right)(r_c - r_m)}_{\text{option (1)}} = \underbrace{\left(M - (1-h^t_{ccp})(1-h^\ell_d)V_\ell\right)(r_c - r_m) \;+\; f\,(1-h^\ell_d)V_\ell}_{\text{option (2)}}.$$

The difference between these two options is the savings per dollar of low quality collateral posted to the swap:

$$\frac{\text{Savings}}{V_\ell} = (r_c - r_m)\left[(1-h^t_{ccp})\frac{1-h^\ell_d}{1-h^t_d} - (1-h^\ell_{ccp})\right] - f\frac{1-h^\ell_d}{1-h^t_d}. \tag{19}$$

[13]We do not separately include the repo rate differentials that the dealer charges or receives on the two legs, instead bundling these together into the fee; the intuition is unchanged if the fee is not separately quoted and instead embedded into haircuts or repo rates. We also ignore the fee that the CCP charges on non-cash posted, since it has a small effect and typically ranges 10bps on the post-haircut collateral value.

$$\frac{\text{Savings}}{V_\ell} = (r_c - r_m)\left[(1 - h^t_{ccp})(1 - h^\ell_d) - (1 - h^\ell_{ccp})\right] - f(1 - h^\ell_d). \tag{20}$$

The expression indicates that the savings are larger—and the incentive to engage in collateral swaps—depending on two factors: (1) investor cost of capital relative to r_m and (2) the difference in low-quality collateral haircuts at the dealer and CCP. A larger spread between the haircuts make collateral swap more profitable; we plot the tradeoff between these two forces in Figure 9.

A simple numerical example makes clear. Suppose the lower quality collateral are 10-year corporate bonds, which are swapped for 10-year Treasuries, and $h^t_{ccp} = 4.5\%$, $h^\ell_{ccp} = 25\%$, $h^t_d = 1\%$, $h^\ell_d = 10\%$, and $V_\ell = \$100$. Two parameters are less easily observed: the cost of capital and the collateral swap fee. We assume their cost of capital is equal to the short rate (IORB) plus 8 percentage points; on average, this gives a cost of capital around 10 percent, roughly in line with the return on equity for insurance companies and slightly lower than the average annualized return for hedge funds in recent years.[14] There is little data on collateral swap fees, but market commentary suggests it is in the range of 50 to 100 bps, so we assume $f = 75$bps. Plugging in the values, the savings provided by the swap are 50 basis points on an annualized basis. Figure 9, which uses the same parameters, shows that the collateral swap would be profitable so long as the $r_c - r_m$ was at least 6 percent.

Since it's hard to know the cost of capital for such a diverse set of investors, it is perhaps easier to quantify the captial reduction. We define capital reduction for the investor as the difference in the amount of cash the investor has to get externally as

$$\left[(1 - h^t_{ccp})(1 - h^\ell_d) - f(1 - h^\ell_d) - (1 - h^\ell_{ccp})\right] V_\ell. \tag{21}$$

Figure 10 estimates collateral swaps reduce capital required by about \$70 billion in 2024 (estimated using (21)), \$60 billion of which stems from swapping equity collateral. The savings are harder to estimate since it requires assumptions about investors' cost of capital, but using the assumptions in the above example yields an average savings rate (as measured in (20)) of 1 percent, amounting to about \$5.5 in 2024. The bulk of the savings stems from

[14] See the OFR "Hedge Fund Monitor" at `https://www.financialresearch.gov/hedge-fund-monitor/categories/risk-management/chart-69/`.

the difference in banks' equity repo haircuts of 10 percent compared to the CME's 30 percent. Importantly, these estimates are likely upper bounds because a material share of collateral swaps are likely related to short-selling or non-margin related motives. Still, even if only half of collateral swaps are used for margin capital reduction, it would still constitute a $30 billion capital reduction. The Internet Appendix provides additional details.

7.3 Dealer Collateral Swaps Capacity

Dealers' ability to alleviate collateral scarcity through collateral swaps varies with their own constraints. These constraints could stem from balance sheet limits, the bank's own haircuts, and changes in the bank's liquidity and risk management. Hence the ability of bank-intermediated collateral swaps to offset margin-induced collateral shortages falls in bad times, even though margin-induced collateral demand rises in bad states.

We test state-dependent pass-through to convenience yields by estimating the sensitivity of $\Delta PC1$ to the margin shocks Z_t estimated in section 6.1. Each Z_t is the idiosyncratic component of margin-requirement changes after controlling for several standard risk proxies. We proxy for banks' collateral swap intermediation capacity two ways. First, we calculate a constraint ratio that compares the collateral supplied through collateral swaps to total margin posted to or by banks. Lower values indicate that banks intermediate less collateral per unit of margin demand. Second, we aggregate total value-at-risk across banks' trading desks, then multiply by -1 so lower values indicate tighter constraints. Higher VaR tightens market-based risk limits and consumes bank capital, raising the internal cost of funding and reducing the profitability of low-margin intermediation. Both measures therefore directly capture states when it is more expensive for the banks to intermediate collateral swaps. The Internet Appendix includes plots for both series.

We then sort the two measures into quintiles, with the lowest quintile reflecting periods of tight constraints.[15]

[15]Most directly, collateral swaps depend not only on the haircut difference between CCPs and the bank but also on a second haircut arbitrage: banks must be able to borrow collateral at lower haircuts and lend it at higher haircuts. Hence the collateral swap market critically depends on the banks' haircut spread. Because banks borrow at higher haircuts in bad states, the profitability of collateral swaps falls in bad states, pushing down swap volumes relative to margin demand.

We estimate the state-dependent sensitivity of convenience yields using

$$\Delta PC1_t = \alpha + \beta Z_t + \sum_{q=2}^{5} \mu_q \mathbb{I}_{q,t} + \sum_{q=2}^{5} \delta_q (Z_t \mathbb{I}_{q,t}) + \gamma' X_t + \varepsilon_t$$

where $\mathbb{I}_{q,t}$ is an indicator equal to one when date t falls into a specific quintile of constraint, X_t is a vector of controls, and $Z_t \in \{Z_t^{PC}, Z_t^{Avg}\}$ are the standardized CME margin shocks. The coefficients μ_q capture level differences in $\Delta PC1_t$ across quintiles, and δ_q capture how the slope with respect to Z_t in quintile q differs from the most constrained quintile. Quintile specific slopes are $s_1 = \beta$ for the most constrained quintile and $s_q = \beta + \delta_q$ for the others.

Figure 11 reports the quintile-specific marginal effects s_q by quintile. The most constrained quintile shows a slope of 0.8 bps for a 1 σ shock, versus -0.01 in the least constrained quintile. The sensitivity is largely monotonic, decreasing as constraints fall. The results are similar using all combinations of Z_t measures and constraint measures, including the ratio or VaR. The Internet Appendix provides the full regression table. The figure shows that the elasticity in the most constrained state is 8 times the average slope in the other quintiles.

8 Conclusion

Safe assets earn their convenience yields through several channels. We provide new evidence on a typically unobserved facet: aggregate collateral demand. We use detailed administrative data to show that aggregate collateral demand is an important driver of convenience yields. Moreover, we provide novel evidence on how the collateral swap market helps, but that banks' ability to intermediate the swaps covaries with financial conditions. Intermediaries hold more liquidity post-crisis to meet regulations, contributing to substantially higher aggregate collateral demand. Yet we show that a nontrivial share of that aggregate collateral is satisfied only with the help of intermediaries. When those intermediaries are unwilling or unable to shuffle collateral around the financial system from those who have collateral to those who need it, collateral and liquidity buffers may prove less stable than expected.

References

Viral V Acharya and Toomas Laarits. When do treasuries earn the convenience yield?: A hedging perspective. *Journal of Finance*, 2025.

Reena Aggarwal, Jennie Bai, and Luc Laeven. Safe-asset shortages: Evidence from the european government bond lending market. *Journal of Financial and Quantitative Analysis*, 56(8):2689–2719, 2021.

Markus K Brunnermeier and Lasse Heje Pedersen. Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6):2201–2238, 2009.

Ricardo J Caballero, Emmanuel Farhi, and Pierre-Olivier Gourinchas. The safe assets shortage conundrum. *Journal of Economic Perspectives*, 31(3):29–46, 2017.

CCP Global. The year in clearing: 2024. Report, The Global Association of Central Counterparties, 2024.

Jin-Wook Chang, Elizabeth Klee, and Vladimir Yankov. Rewiring repo. 2025.

Adam Copeland, Darrell Duffie, Antoine Martin, and Susan McLaughlin. Key mechanics of the us tri-party repo market. 2012.

Wenxin Du, Benjamin Hebert, and Amy Wang Huber. Are intermediary constraints priced? *Review of Financial Studies*, 36(4):1464–1507, 2023.

Darrell Duffie. Special repo rates. *The Journal of Finance*, 51(2):493–526, 1996.

Darrell Duffie. Financial regulatory reform after the crisis: An assessment. *Management Science*, 64(10):4835–4857, 2018.

Antonio Falato, Diana Iercosan, and Filip Zikes. Banks as regulated traders. *Journal of Financial Economics*, 170:104080, 2025. ISSN 0304-405X. doi: https://doi.org/10.1016/j.jfineco.2025.104080.

Peter Feldhütter and David Lando. Decomposing swap spreads. *Journal of Financial Economics*, 88(2):375–405, 2008.

Matthias Fleckenstein and Francis A Longstaff. Treasury richness. *The Journal of Finance*, 79(4):2797–2844, 2024.

Matthias Fleckenstein, Francis A Longstaff, and Hanno Lustig. The tips-treasury bond puzzle. *the Journal of Finance*, 69(5):2151–2197, 2014.

Xavier Gabaix and Ralph SJ Koijen. Granular instrumental variables. *Journal of Political Economy*, 132(7):2274–2303, 2024.

Nicolae Garleanu and Lasse Heje Pedersen. Margin-based asset pricing and deviations from the law of one price. *Review of Financial Studies*, 24(6):1980–2022, 2011.

Gary Gorton and Andrew Metrick. Securitized banking and the run on repo. *Journal of Financial economics*, 104(3):425–451, 2012.

Robin Greenwood, Samuel G Hanson, and Jeremy C Stein. A comparative-advantage approach to government debt maturity. *The Journal of Finance*, 70(4):1683–1722, 2015.

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

Cindy E Hull. Are the federal reserve's "other deposits" sensitive to changes in interest rates? *Chicago Fed Letter*, 504:1–8, 2024.

Sebastian Infante and Zack Saravay. Treasury re-use and the demand for safe assets. *Review of Financial Studies*, Forthcoming.

Sebastian Infante and Alexandros P Vardoulakis. Collateral runs. *The Review of Financial Studies*, 34(6):2949–2992, 2021.

Arvind Krishnamurthy and Yiming Ma. The demand and supply of convenience assets. *Annual Review of Financial Economics*, 17, 2024.

Arvind Krishnamurthy and Annette Vissing-Jorgensen. The aggregate demand for treasury debt. *Journal of Political Economy*, 120(2):233–267, 2012.

Francis A Longstaff. The flight-to-liquidity premium in us treasury bond prices. *Journal of Business*, 77(3), 2004.

Lina Lu and Jonathan Wallen. What do bank trading desks do? Technical report, Harvard Working Paper, 2024.

Yang-Ho Park and Nicole Abruzzo. An empirical analysis of futures margin changes: Determinants and policy implications. *Journal of Financial Services Research*, 49(1):65–100, 2016.

9 Figures

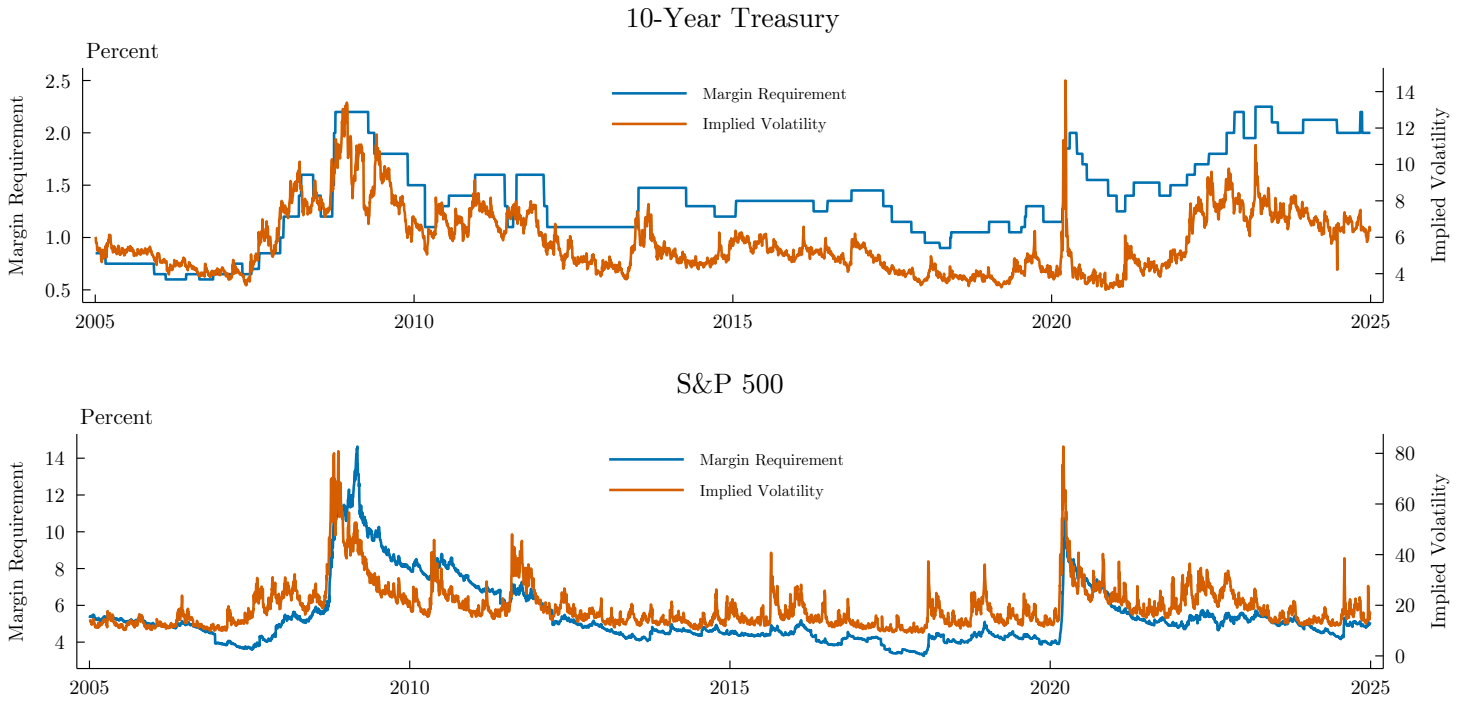


Figure 1: CME Margin Requirements. Red line plots the CME's margin requirement for 10-year Treasury (TY) and E-mini S&P 500 (ES) futures contracts expressed as percent of the contract notional. Blue line plots the implied volatility for the same contracts.

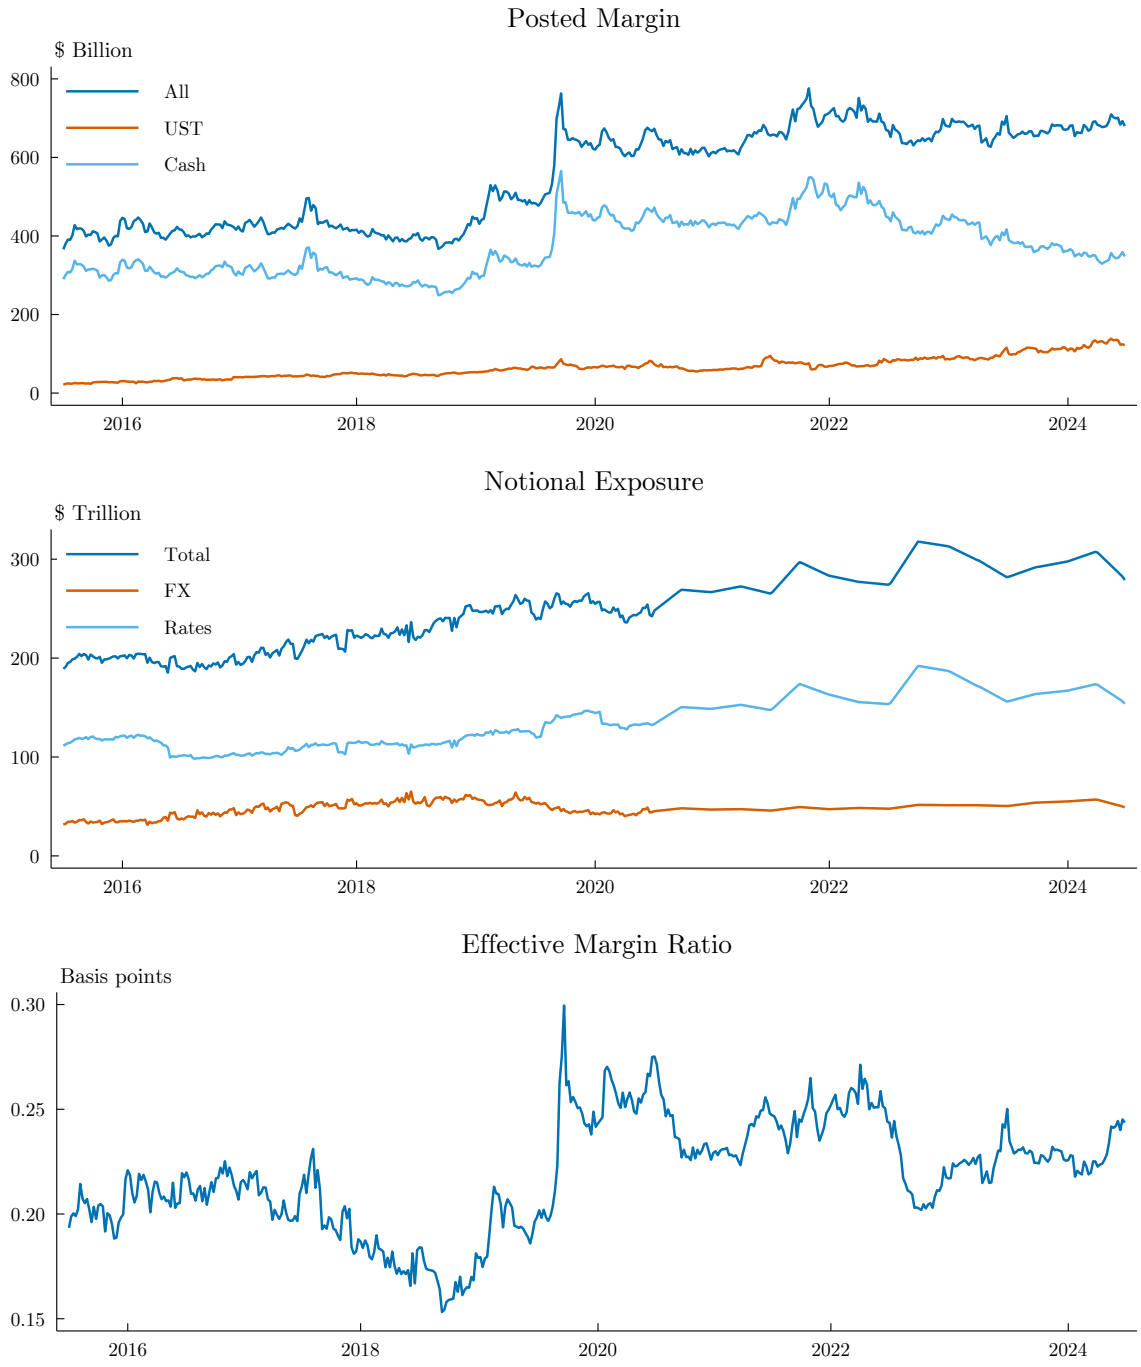


Figure 2: Margin Posted, Notional, and the *Effective Margin Ratio*. Top panel plots total margin posted, cash posted as margin, and Treasuries posted as margin. Middle panel plots total notional exposure, along with exposures for specifically for rates and FX desks. Bottom panel plots the Effective Margin Ratio, which is the ratio of posted margin to notional exposures. Plot uses week-end values.

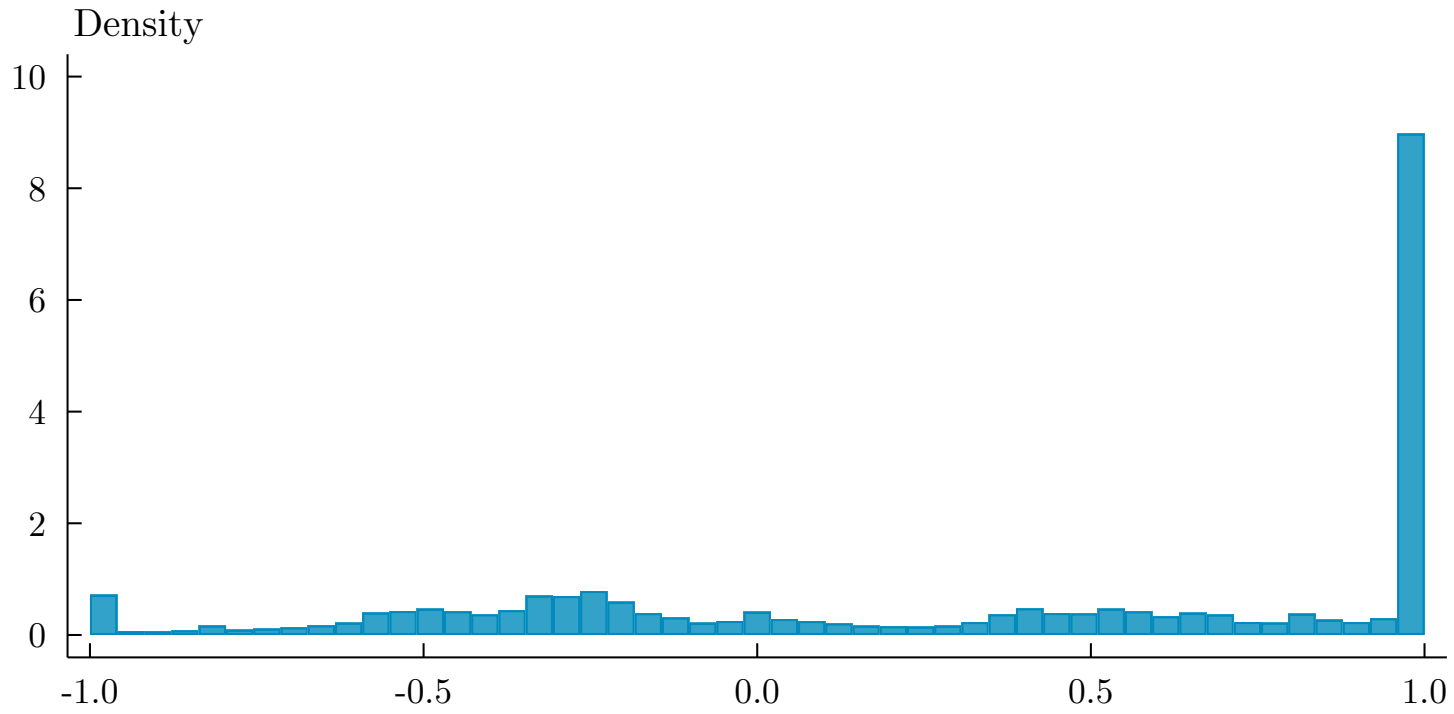


Figure 3: Distribution of Collateral Flow Imbalances. Figure is a histogram of the collateral flow imbalance variables described in section 5. Data aggregates across all banks in our sample, and the imbalance measure is at the date by counterparty by settlement type level.

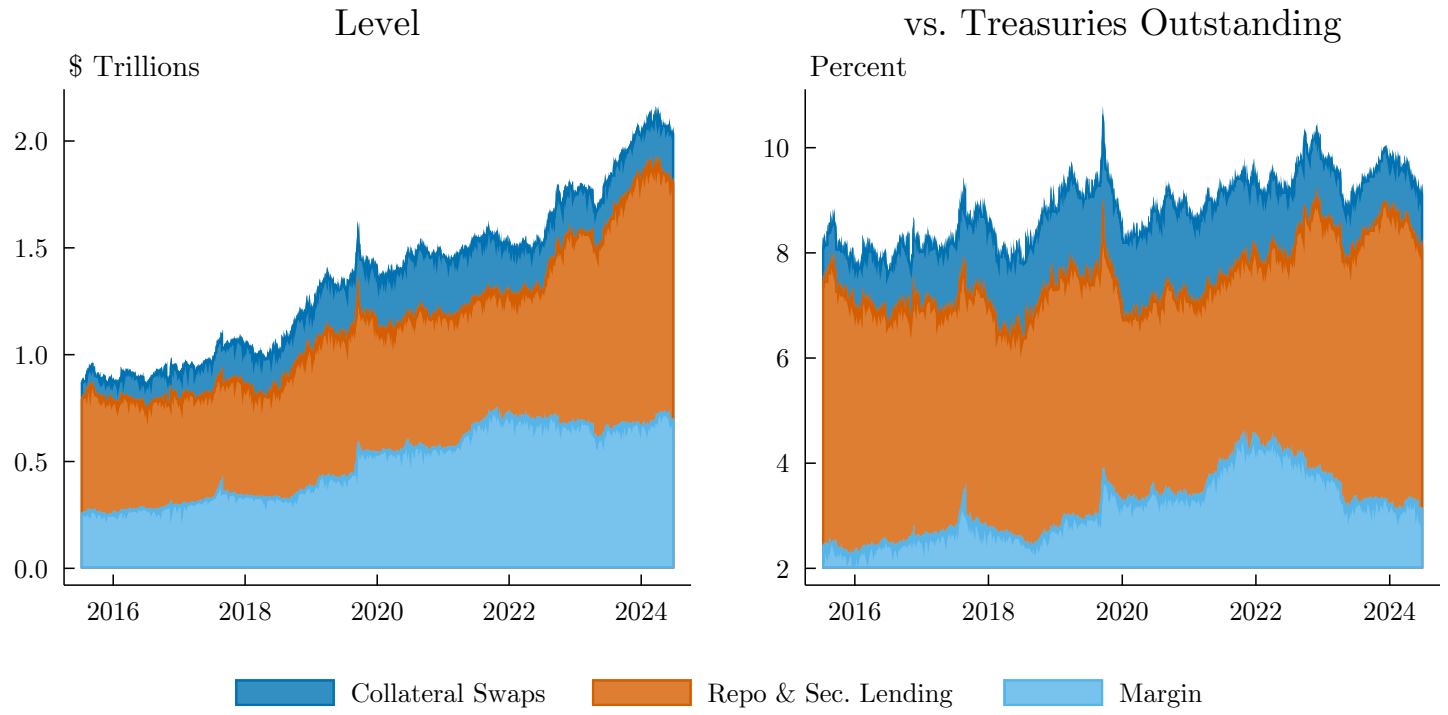


Figure 4: Collateral Sinks. Figure plots the total amount of sunk collateral using the methodology described in Section 5, aggregating across all collateral types and all banks in our sample. Left panel shows the level and right panel compares to the market value of publicly held Treasuries outstanding.

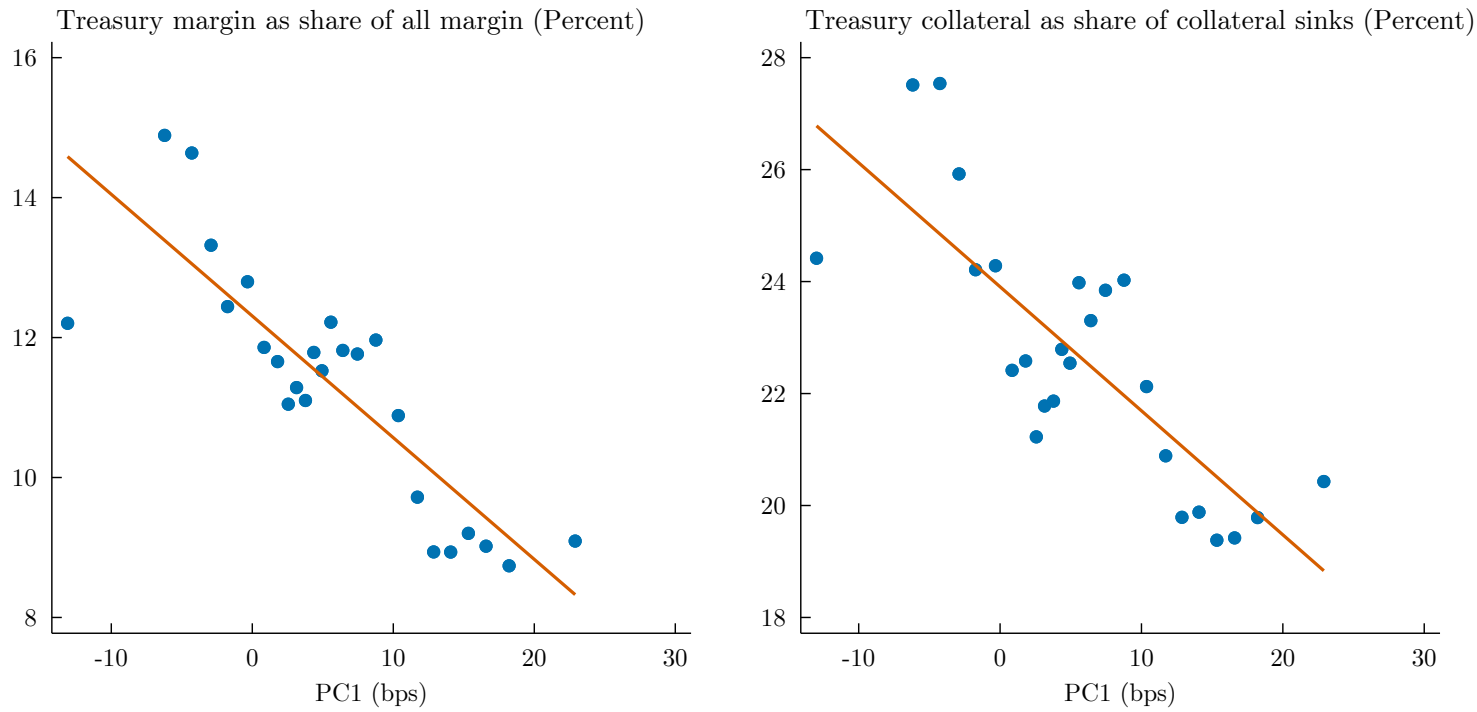


Figure 5: Treasury Share of Margin and Collateral Sinks vs. PC1. Left panel provides binscatter plot of Treasury margin as a share of all margin against PC1; right panel plots the Treasury collateral share of collateral sinks.

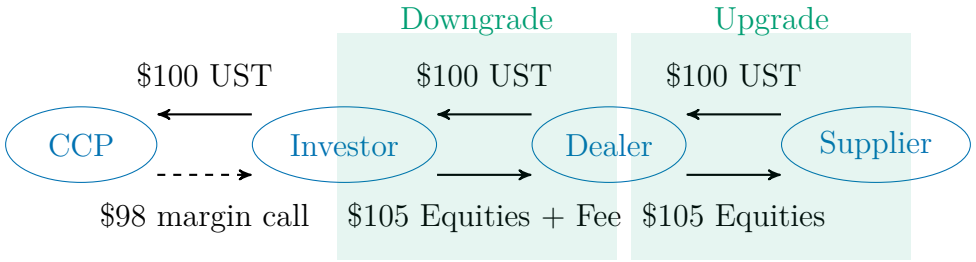


Figure 6: Collateral Swap Example. Figure provides an example collateral swap in which an invest transforms equities into Treasuries (UST).

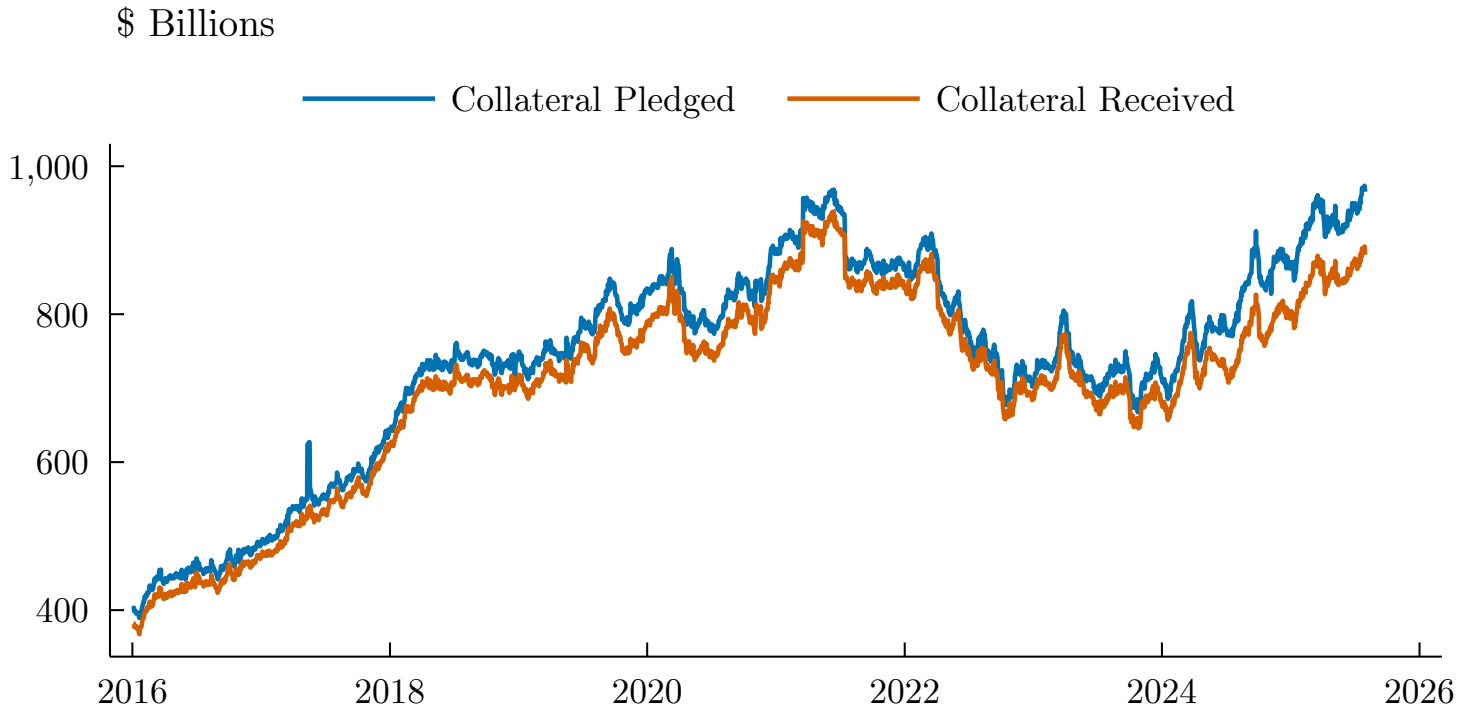


Figure 7: Collateral Swaps Outstanding. Figure plots the total market value of collateral pledged or received by banks through collateral swaps.

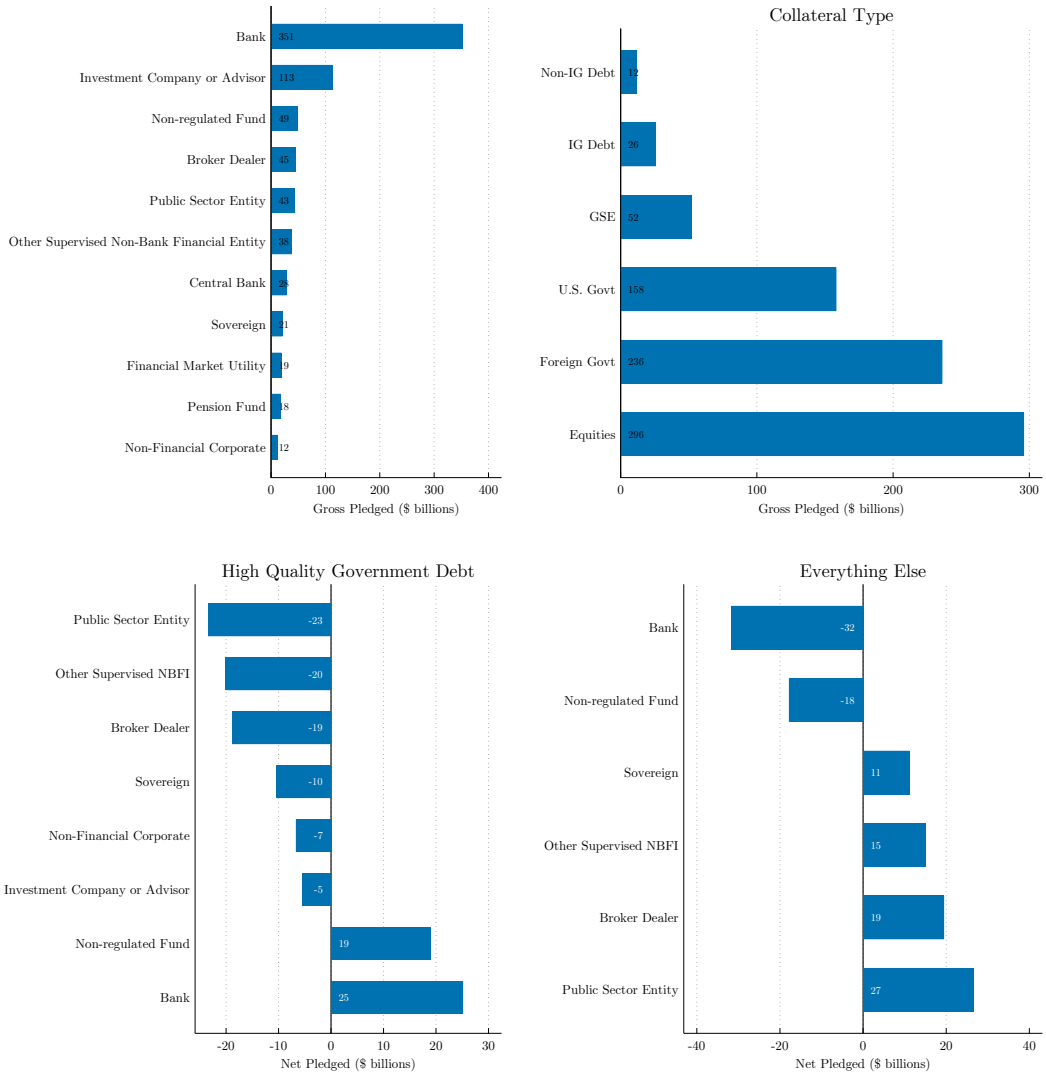


Figure 8: Collateral Swap Market Structure. Figure plots flows in collateral swaps averaged from May 2022 to December 2024, matching the period we have detailed counterparty data.

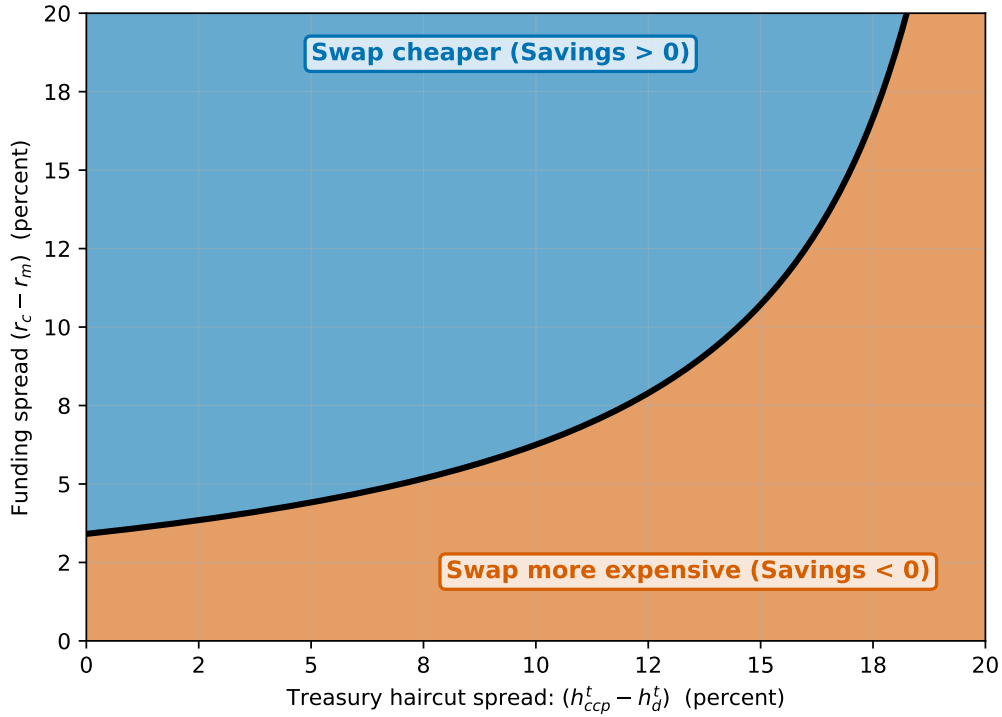


Figure 9: Regions of Collateral Swap Viability. Figure plots the regions where collateral swaps are economically profitable for an investor needing to post margin, as a function of the spread between the investor's cost of capital and the rate paid by the CCP on cash $(r_c - r_m)$ against the difference in Treasury haircuts. Figure assumes

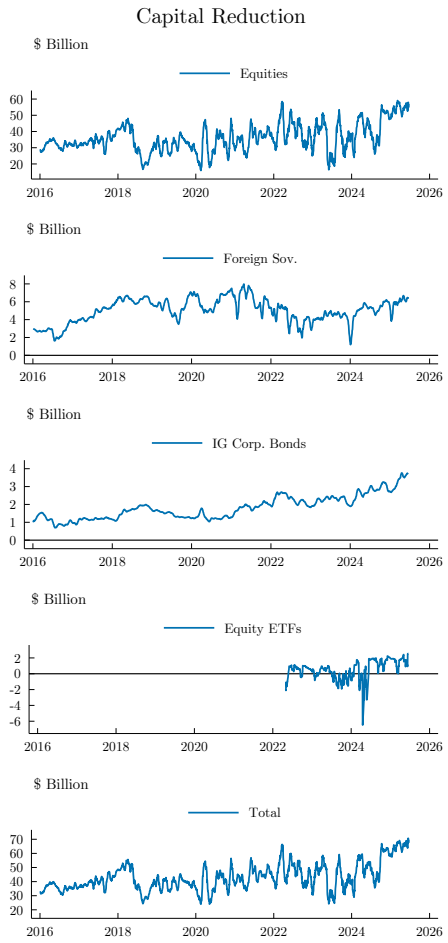


Figure 10: Estimated Margin Capital Reduction provided by Collateral Swaps. Figure uses CME haircuts, time-varying dealer reverse repo haircuts, and the volume of risky collateral posted to banks through collateral swaps and calculates the capital reduction using (21). See section IA.B.1 for methodology details.

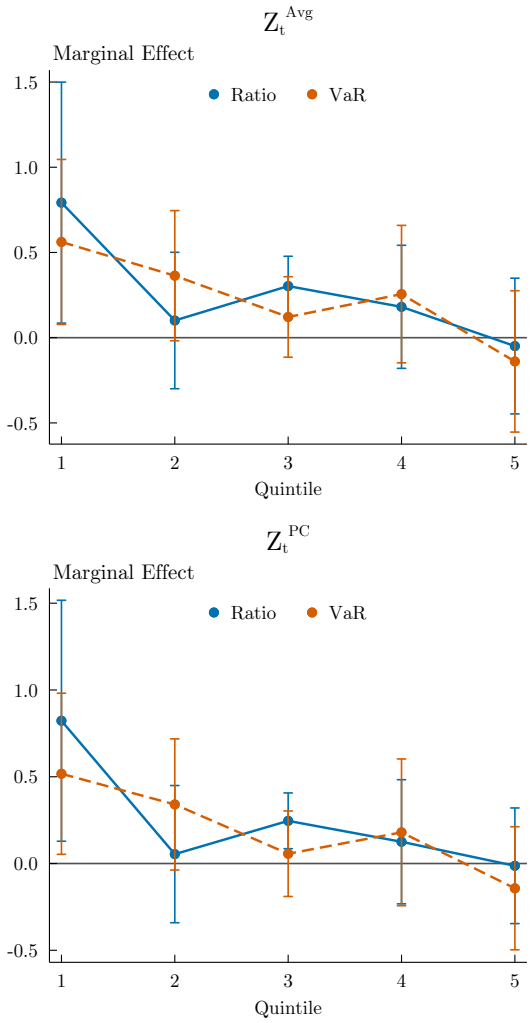


Figure 11: State-dependent Passthrough. Figure plots the total marginal effect of Z_t on $\Delta PC1_t$ by bank constraint state, which are calculated either using the ratio of total collateral pledged by banks in collateral swaps to total margin posted by or to banks or alternatively by the total value-at-risk across banks' trading desks. Top panel shows the effects when using Z_t^{Avg} and bottom panel uses Z_t^{PC}.

10 Tables

Panel A: Margin Posted ($ billions)	Obs.	Mean	SD	Min	Max
Cash	2,217	374	74	247	571
USTs	2,217	64	26	21	142
HQLA	2,217	447	155	231	723
Non-HQLA	2,217	106	44	26	178
Total	2,217	553	124	359	776

Panel B: ΔMargin Posted ($ billions)	Obs.	Mean	SD	Min	Max
Cash	2,187	0.0	5.0	-43.2	50.8
USTs	2,187	0.0	1.4	-15.8	9.1
HQLA	2,187	0.2	5.4	-48.5	62.7
Non-HQLA	2,187	0.0	3.2	-109.3	16.8
Total	2,187	0.2	5.8	-46.6	59.5

Panel C: Margin Composition (Percent)	Obs.	Mean	SD	Min	Max
Cash	2,217	69	6	49	80
USTs	2,217	11	3	6	20
HQLA	2,217	79	11	64	96
Non-HQLA	2,217	21	11	4	36

Panel D: Notional Exposure ($ billions)	Obs.	Mean	SD	Min	Max
Rates	1,639	119,674	13,940	89,443	155,211
FX	1,639	44,085	9,186	23,561	73,994
Credit	1,639	35,337	5,087	21,419	44,683
Equity	1,639	7,212	3,580	3,259	24,096
Securitized Products	1,639	3,479	2,302	627	9,436
Commodity	1,639	2,955	1,431	1,744	16,057
XVA	1,639	2,832	1,104	780	5,258
Unclassified	1,639	2,883	1,865	956	7,849
Total	1,639	219,236	23,739	149,012	269,815

Panel E: Margin Ratio (Basis Points)	Obs.	Mean	SD	Min	Max
Cash Margin Ratio	2,217	15.0	2.2	10.3	22.2
UST Margin Ratio	2,217	2.5	0.7	1.1	4.9
Total Margin Ratio	2,217	21.9	2.7	15.2	30.0

Table 1: Margin Summary Statistics. Margin data is daily market value of margin posted by large banks, excluding that posted on behalf of customers, from 2016 to 2024; notional exposure trading data is for the same set of banks from 2014 to 2021. UST is US Treasuries, HQLA are high-quality liquid assets. Panel B presents daily changes in margin posted. Panel C shows the share of total margin that each subset constitutes. Panel D provides notional exposure by desk type, excluding desk types that have average daily notional less than $1 billion, although these desks are included in the total. Margin ratio is the ratio of margin posted relative to notional exposures on the same day. All values are aggregated across all banks in the sample before calculating moments.

Repo and Securities Lending ($ billions)	Obs	Mean	SD	Min	Max
Cash	2,350	0	0	0	0
Treasuries	2,226	201	104	77	542
Other HQLA	2,226	265	63	191	472
Other Non-HQLA	2,226	201	36	136	307
Total	2,226	666	189	468	1,244

Collateral Swaps ($ billions)	Obs	Mean	SD	Min	Max
Cash	2,350	0	0	0	0
Treasuries	2,226	25	10	5	58
Other HQLA	2,226	122	37	44	192
Other Non-HQLA	2,226	57	15	20	89
Total	2,226	204	56	72	299

Margin ($ billions)	Obs	Mean	SD	Min	Max
Cash	2,226	224	61	150	407
Treasuries	2,226	97	41	34	193
Other HQLA	2,226	109	65	20	248
Other Non-HQLA	2,226	73	29	37	148
Total	2,226	503	169	251	761

Aggregated ($ billions)	Obs	Mean	SD	Min	Max
Cash	2,226	224	61	150	407
Treasuries	2,226	323	142	140	771
Other HQLA	2,226	496	133	289	796
Other Non-HQLA	2,226	330	74	233	534
Total	2,226	1,373	356	865	2,164

vs. Treasuries Outstanding (percent)	Obs	Mean	SD	Min	Max
Repo & Sec. Lending	2,213	4.4	0.6	3.4	5.8
Collateral Swaps	2,213	1.3	0.3	0.7	1.9
Margin	2,213	3.2	0.6	2.3	4.6
Total (S_t)	2,213	8.9	0.6	7.7	10.8

Table 2: Collateral Sink Summary Statistics. Table provides summary statistics for the collateral sink data described in Section 5. Aggregated reflects the sum of sunk collateral across repo/securities lending, collateral swaps, and margin. HQLA is high quality liquid assets. All values are aggregated across all banks in the sample before calculating moments. Bottom panel calculated relative to the total market value of publicly held Treasuries outstanding averaged over the previous five days; See Internet Appendix for construction.

	$PC1_t$			$\Delta PC1_t$		
	(1)	(2)	(3)	(4)	(5)	(6)
Collateral Sink Measures						
Collateral Sinks/Treasuries Outstanding$_t$	6.699***	6.894***	7.660***			
	(2.71)	(2.87)	(2.82)			
ΔCollateral Sinks/Treasuries Outstanding$_t$				4.538**	4.693***	4.546**
				(2.54)	(2.69)	(2.54)
Equity Implied Vol.$_t$		0.042	0.100			
		(0.57)	(0.93)			
Treasury Implied Vol.$_t$		−0.559	−0.682			
		(−1.05)	(−1.24)			
FX Implied Vol.$_t$		0.263	0.105			
		(0.40)	(0.18)			
Oil ETF Implied Vol.$_t$			0.006			
			(0.53)			
Effective Fed Funds Rate$_t$			25.939***			
			(5.08)			
BAA-AAA Spread$_t$			−0.579			
			(−0.16)			
US CDS$_t$			18.144			
			(1.40)			
5yr Treasury Bid-Ask Spread$_t$			1103.257			
			(1.41)			
5y5y Breakeven Inflation$_t$			−1.503			
			(−0.34)			
ΔEquity Implied Vol.$_t$					0.005	−0.089
					(0.12)	(−1.05)
ΔTreasury Implied Vol.$_t$					0.240	−0.078
					(0.66)	(−0.18)
ΔFX Implied Vol.$_t$						−0.175
						(−0.50)
ΔOil ETF Implied Vol.$_t$						0.002
						(0.29)
S&P 500 Return$_t$						−18.733
						(−1.25)
ΔEffective Fed Funds Rate$_t$						20.039***
						(4.76)
ΔBAA-AAA Spread$_t$						−3.650
						(−0.96)
ΔUS CDS$_t$						25.488**
						(2.29)
Δ5yr Treasury Bid-Ask Spread$_t$						1118.436*
						(1.91)
Δ5y5y Breakeven Inflation$_t$						−12.771***
						(−3.97)
N	2,212	2,159	2,153	2,180	2,122	2,108
R^2	0.02	0.02	0.23	0.01	0.01	0.10
Controls	No	Vol	All	No	Vol	All
Time FE	Yes	Yes	Yes	No	No	No

Table 3: Treasury Convenience Yield and Collateral Sinks. Table shows the regression of the first principal component (PC1) or its changes on sunk collateral relative to total Treasuries outstanding. Levels regression include weekly fixed effects. Constant omitted. Columns labeled "Vol" include volatility controls: changes in the VIX and Treasury implied volatility. Columns labeled "All" include the full control set: changes in the VIX, Treasury, FX, and oil implied volatilities; the S&P 500 return; the change in the effective federal funds rate; and changes in the Baa–AAA spread, the U.S. 5-year CDS spread, the 5-year Treasury bid–ask spread, and the 5-year–5-year breakeven rate. t-statistics shown using heterosketastic and autocorrelation consistent standard errors using the Newey and West (1994) automatic lag selection procedure where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	$\Delta PC1_t$					
	(1)	(2)	(3)	(4)	(5)	(6)
Collateral Sink Components						
ΔSunk Margin/Treasuries Outstanding$_t$	6.940***	6.714***				
	(2.82)	(2.69)				
ΔSunk Repo Collateral/Treasuries Outstanding$_t$			3.448	3.548		
			(1.49)	(1.50)		
ΔSunk Collateral Swap Collateral/Treasuries Outstanding$_t$					10.498*	9.992*
					(1.65)	(1.91)
Controls						
ΔEquity Implied Vol.$_t$		−0.090		−0.089		−0.093
		(−1.13)		(−1.14)		(−1.17)
ΔTreasury Implied Vol.$_t$		−0.107		−0.049		−0.076
		(−0.25)		(−0.11)		(−0.18)
ΔFX Implied Vol.$_t$		−0.156		−0.176		−0.094
		(−0.43)		(−0.43)		(−0.26)
ΔOil ETF Implied Vol.$_t$		0.000		0.003		0.001
		(0.05)		(0.27)		(0.06)
S&P 500 Return$_t$		−19.302		−16.997		−17.760
		(−1.31)		(−1.20)		(−1.21)
ΔEffective Fed Funds Rate$_t$		20.043***		20.134***		20.094***
		(4.75)		(4.85)		(4.79)
ΔBAA-AAA Spread$_t$		−3.645		−3.278		−3.422
		(−0.98)		(−0.92)		(−0.93)
ΔUS CDS$_t$		24.740**		25.272**		24.843**
		(2.27)		(2.21)		(2.33)
Δ5yr Treasury Bid-Ask Spread$_t$		1136.088*		1154.950*		1193.928**
		(1.91)		(1.91)		(1.97)
Δ5y5y Breakeven Inflation$_t$		−13.160***		−12.947***		−12.861***
		(−4.07)		(−3.91)		(−3.99)
N	2,180	2,108	2,180	2,108	2,180	2,108
R^2	0.00	0.09	0.00	0.09	0.00	0.09
Controls	No	All	No	All	No	All
Time FE	No	No	No	No	No	No

Table 4: Treasury Convenience Yield and Collateral Sink Components. Table shows the regression of the changes in the first principal component (PC1) changes in sunk collateral by type. Controls include equity implied volatility is the VIX, Treasury implied volatility is the implied volatility of 10-year Treasury futures, Treasuries outstanding is market value, US CDS is the senior unsecured 5-year credit default swap spread for US government. Includes monthly fixed effects. t-statistics shown using heterosketastic and autocorrelation consistent standard errors using the Newey and West (1994) automatic lag selection procedure where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Second Stage

	Dependent Var.: $\Delta PC1_t$					
	(1)	(2)	(3)	(4)	(5)	(6)
ΔS_t	2.30***	2.37***	2.01***	2.43**	2.53**	2.01**
	(2.69)	(2.85)	(2.85)	(2.44)	(2.52)	(2.54)
N	2,108	2,108	2,108	2,108	2,108	2,108
Controls	No	Vol.	All	No	Vol.	All

Panel B: First Stage

	ΔS_t					
	(1)	(2)	(3)	(4)	(5)	(6)
Z_t^{Avg}	0.14***	0.15***	0.14***			
	(6.44)	(6.36)	(5.31)			
Z_t^{PC}				0.12***	0.13***	0.13***
				(5.35)	(5.18)	(4.42)
N	2,108	2,108	2,108	2,108	2,108	2,108
$F-stat$	41.43	40.43	28.20	28.59	26.80	19.56
Controls	No	Vol.	All	No	Vol.	All

Panel C: OLS

	Dependent Var.: $\Delta PC1_t$					
	(1)	(2)	(3)	(4)	(5)	(6)
ΔS_t	0.32***	0.34***	0.32***	0.32***	0.34***	0.32***
	(2.73)	(2.84)	(2.65)	(2.73)	(2.84)	(2.65)
N	2,180	2,118	2,108	2,180	2,118	2,108
R^2	0.01	0.01	0.10	0.01	0.01	0.10
Controls	No	Vol.	All	No	Vol.	All

Table 5: CME Residual Regression. Table presents the first-stage, second-stage, and OLS estimates described in Section 6.1. Columns labeled "Vol" include volatility controls: changes in the VIX and Treasury implied volatility. Columns labeled "All" include the full control set: changes in the VIX, Treasury, FX, and oil implied volatilities; the S&P 500 return; the change in the effective federal funds rate; and changes in the Baa–AAA spread, the U.S. 5-year CDS spread, the 5-year Treasury bid–ask spread, and the 5-year–5-year breakeven rate. Constant omitted. t-statistics shown using heterosketastic and autocorrelation consistent standard errors using the Newey and West (1994) automatic lag selection procedure where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Second Stage						
	Dependent Var.: ΔPC1_t					
	(1)	(2)	(3)	(4)	(5)	(6)
$\widehat{\Delta S_t}$	2.19***	2.02***	1.85***			
	(2.84)	(2.82)	(2.63)			
$\Delta \ln(\widehat{\text{Sunk Margin}_t})$				2.64***	2.37***	2.16***
				(3.09)	(2.98)	(2.83)
N	1,218	1,169	1,163	1,218	1,169	1,163
Controls	No	Vol.	All	No	Vol.	All

Panel B: First Stage						
	ΔS_t			$\Delta \ln(\text{Sunk Margin}_t)$		
	(1)	(2)	(3)	(4)	(5)	(6)
GIV_t	0.13***	0.14***	0.14***	0.11***	0.12***	0.12***
	(3.95)	(4.46)	(4.44)	(4.97)	(5.70)	(5.81)
N	1,218	1,169	1,163	1,218	1,169	1,163
$F-stat$	15.63	19.90	19.76	24.73	32.50	33.74
Controls	No	Vol.	All	No	Vol.	All

Panel C: OLS						
	Dependent Var.: ΔPC1_t					
	(1)	(2)	(3)	(4)	(5)	(6)
ΔS_t	0.22*	0.22	0.20			
	(1.81)	(1.58)	(1.39)			
$\Delta \ln(\text{Sunk Margin}_t)$				0.02	0.01	0.00
				(0.27)	(0.14)	(0.03)
N	1,218	1,169	1,163	1,218	1,169	1,163
R^2	0.01	0.01	0.05	0.00	0.00	0.04
Controls	No	Vol.	All	No	Vol.	All

Table 6: Granular IV. Table presents the first-stage, second-stage, and OLS estimates described in Section 6.2. Columns (1)–(3) use ΔS_t as the endogenous regressor, and columns (4)–(6) use $\Delta \ln \text{Sunk Margin}_t$. Columns labeled "Vol" include volatility controls: changes in the VIX and Treasury implied volatility. Columns labeled "All" include the full control set: changes in the VIX, Treasury, FX, and oil implied volatilities; the S&P 500 return; the change in the effective federal funds rate; and changes in the Baa–AAA spread, the U.S. 5-year CDS spread, the 5-year Treasury bid–ask spread, and the 5-year–5-year breakeven rate. Constant omitted. t-statistics shown using heterosketastic and autocorrelation consistent standard errors using the Newey and West (1994) automatic lag selection procedure where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

$ billions	**U.S. Treasuries**	**Non-U.S. Gov't Debt**	**Equities**	**Other**
Panel A: Collateral Pledged				
Mean	160	190	327	128
Std. Dev. (monthly)	21	9	31	12
Avg. Maturity (Days)	14	15	80	101
2016 to 2024 Growth	166%	43%	125%	13%
Panel B: Collateral Received				
Mean	277	188	275	58
Std. Dev. (monthly)	28	12	20	6
Avg. Maturity (Days)	84	55	11	187
2016 to 2024 Growth	97%	66%	73%	154%
Panel C: Net Collateral Pledged				
Mean	-118	2	52	69
Std. Dev. (monthly)	23	15	30	8

Table 7: Collateral Swap Summary Statistics. Table provides summary statistics on collateral swaps on average in 2024.

Internet Appendix to

Aggregate Collateral Demand

Toomas Laarits, Chase P. Ross, Sharon Y. Ross

The Internet Appendix consists of four sections. Section IA.A provides additional details about the data. Section IA.B provides additional discussion and results to supplement the main text. Section IA.C presents additional tables and figures.

IA.A Data

IA.A.1 FR 2052a Data Details

We focus on data provided by the largest banks that consistently provide daily balance sheet data through the full sample, from 2016 through 2024. We limit to the top-level consolidated bank holding company and we limit to dates when all banks report data and exclude internal transactions. We clean the data in several ways to reduce outliers, namely around a reporting requirement change in May 2022 that changed several dimensions of the data collection, namely the classification of posted margin (which provided more granular data on where the margin is posted) and counterparty types (which provided more granular definitions of counterparties). We manually adjust the series in these cases to create a consistent and comparable data over the sample. We manually clean the data in a handful of other cases and dates when reported data is likely an outlier or when the reported data is incomplete.

The data provides the market value for some assets and provides maturity values for other assets. Notably, posted margin is reported at the market value of the underlying collateral, while SFTs are reported based on the maturity value of the SFT itself, not the market value of the underlying collateral, with the difference reflecting a haircut. We focus on the maturity value of the SFT, which reflects the post-haircut market value of the collateral.

IA.A.2 FR VV-1 Data Details

We use the FR VV-1 data to estimate total gross notional exposures for banks' trading desks. The data reporting requirements changed in 2021, and we filter the data by limiting to the set of banks that report consistently through the full sample. We exclude a handful of desks with outlier positions, focusing principally on those that report consistently through the full sample and we winsorize level variables—including PnL, market values, notional values, and value-at-risk levels at the 0.1 and 99.9 level by bank to reduce the influence of outliers. We

assign desks to asset classes based on the text of the desk name or its description, and desks that do not match any of our asset class keywords are denoted unclassified, although these constitute a small share of the sample.

One challenge to estimating notional exposures for banks, and their realized margin requirement ratios, is that banks stopped reporting trading desk notional exposures in FR VV-1 beginning in 2021, instead reporting daily gross notional flows. We estimate post-2020 trading desk notional exposures using the bank-holding companies' public quarter Y9C reports.

FR VV-1 data provides notional exposures for trading desk until 2020. Following a change in reporting requirements, banks switched to report daily gross notional flows in 2021. It is not possible to reconstruct the notional level from gross flows without making several assumptions about the maturity structure of the existing contract base. Instead, we estimate the daily notional levels using public quarterly filings that large bank-holding companies report in the FR Y-9C. In particular, we use data from Schedule HR-C which provides the notional principal amount of derivative contracts by clearing (OTC vs. centrally clear), asset class (interest rate, FX, equity, etc) and maturity (1 year or less, 1 to 5 years, over 5 year). The variables correspond to BHCKS582 to BHCKS623.

The two datasets provide related by different measures of notional exposures. The VV-1 data spans derivatives booked on covered trading desks, while the Y9-C data spans all derivative positions held by the holding company, which includes both the dealer (where the trading desk sits) along with non-dealer subsidiaries, namely their depository institutions. On average, aggregate VV-1 data spans about 80 percent of total notional exposures reported in Y9-C, with the difference likely stemming from asset-liability management and hedging activities in the BHC's bank subsidiary.

We estimate the daily notional level in several steps. First, we collapse both the Y-9C and VV-1 data to comparable broad asset classes: credit, equity, foreign exchange, rates, and unclassified. We calculate total notional from VV-1 by summing the absolute value of notional exposures of derivatives carried as assets and those carried as liabilities. To reduce step changes introduced by quarterly data, we calculate a moving average of the Y-9C that assigns to date t the weighted average of the Y9-C value from the previous quarter and the next quarter, where the weight is the number of days since the previous quarter end. We merge the two panels together and calculate the average ratio of VV-1 notional to the Y-9C notional by bank, weighted by the asset class's share of a bank's total notional in a given quarter to reduce the influence of outliers. We estimate daily levels by multiplying the Y-9C notional by the estimated ratio, and we add the level difference between the estimated 2020 Q4 value and the actual VV-1 level. We plot the actual trading notional, estimated trading

notional, and the Y-9C notional in Online Appendix Figure IA.3.

For our set of banks, we calculate *effective margin requirement* using

$$\text{Effective Margin Requirement}_t = \frac{\text{Margin Posted}_t}{\text{Notional Exposures}_t}.$$

We term this the *effective* margin requirement because it represents the amount of margin the bank must post given its portfolio of derivatives, rather than a contract-by-contract weighted average. Banks post margin equal to 27 bps of their notional exposure on average, with cash accounting for 20 bps and Treasuries 3 bps. The bottom panel of Figure 2 shows that the margin ratio was perhaps falling until the Covid pandemic, then increased by 15 bps, and broadly stayed at higher average level following the pandemic, but with many local peaks and troughs.

An immediate concern is that our estimate of the margin ratio is biased by our use of estimated daily notional exposures beginning in 2021. As robustness, we note that variation in the numerator—margin posted—explains roughly 3 times more of the variation in the ratio than the denominator in the period we observe the numerator and denominator perfect. We decompose the log variance of the ratio, such that Var(ln(Margin Ratio)) = Var(ln(Margin Posted)) + Var(ln(Notional exposures)) - 2 Cov(ln(Margin Posted), ln(Notional exposures)). Var(ln(Margin Posted)) explains 172% and Var(ln(Notional exposures) explains 63 % of the variance of the log ratio, and the covariance term explains -135%.

IA.A.3 Convenience Yield Measures

1. **GCF repo – Treasury bill spread:** three-month general collateral financing repo rate minus the three-month T-bill yield.

2. **OIS – Treasury bill spread:** three-month overnight-indexed swap rate minus the three-month T-bill yield.

3. **Fed funds – Treasury bill spread:** spread of effective Fed funds rate and three-month T-bill yield; maturity-matched using OIS curve.

4. **Negative Z-spread:** the average T-bill yield (4 to 26 weeks to maturity) minus the fitted Treasury curve; value multiplied by negative one so that larger values indicate higher convenience. Greenwood et al. (2015).

5. **10-year TIPS – Treasury spread:** yield on a synthetic nominal 10-Year Treasury equal to the nominal TIPS yield plus a matched-maturity inflation swap minus the actual nominal Treasury yield; Fleckenstein et al. (2014).

6. **30-year OIS swap - Treasury spread:** 30-year overnight-indexed swap quoted rate minus maturity-matched nominal Treasury yield estimated from fitted yield curve. Feldhütter and Lando (2008); Du et al. (2023).

IA.A.4 Other Variables Description

- Market value of Treasuries outstanding. We collect Treasury auction data from TreasuryDirect to calculate the face value of Treasuries outstanding at a daily frequency; we exclude SOMA purchases from the face value by either using the most recent SOMA portfolio holding reported by the New York Fed (which is reported weekly on Wednesdays) or, if that number is unavailable, SOMA's purchase at issuance as reported by TreasuryDirect. We then merge the resulting face value of publicly held Treasuries with the CRSP daily Treasury file, and calculate the market value of Treasuries outstanding by multiplying the face value by the dirty price (the sum of the clean price and accrued interest). We require a CUSIP has data from both TreasuryDirect and price data from CRSP to be included. We then calculate a five day moving average, from $t-4$ to t, to use as the denominator in the ratios so remove mechanical variation over the week stemming from bill issuance patterns and the weekly Wednesday SOMA portfolio disclosures.

- Equity Implied Volatility. VIXCLS from FRED.

- Treasury Implied Volatility. Implied volatility for second 10-year Treasury note from Morgan Markets.

- FX Implied Volatility. Implied volatility for G10 currencies three-month at-the-money forward options, weighted by turnover from Morgan Markets.

- Oil Implied Volatility. Implied volatility of CBOE crude oil ETF from FRED.

- U.S. Credit Default Swap spread. We use the 5-year senior unsecured credit default swap for U.S. sovereign debt denominated in EUR with document clause of CR (before September 2014) or CR14 (after September 2014). Data from Markit.

- 5-year/5-year breakeven inflation rate. Calculated using hot-run yields from Morgan Markets.

IA.B Additional Results

IA.B.1 Collateral Swap Savings

We estimate collateral swap savings by using collateral swap volumes and bank haircut data from FR2052a. We estimate haircuts by focusing on reverse repos and securities borrowing transactions with no more than 91 days to maturity. Within each cell, defined at the collateral class by SFT type (reverse repo or securities borrowing) by maturity by bank by date, we calculate the haircut as $1 -$ SFT loan value/SFT collateral value. We exclude reverse repos and securities borrowing transactions that are likely driven by short-sale motives by requiring the SFT has a positive haircut. We then take the median haircut within the collateral type and date across the remaining cells.

We focus on collateral types that are comparable to CME's acceptable collateral. For equities, this includes U.S- and foreign-listed common equity securities (classes E-1 through E-4 in FR2052a, including the LCR qualifying classes). For ETFs, which are available only beginning in May 2022, it includes both U.S- and foreign-listed ETFs (E-5 and E-6). For foreign sovereign bonds, we include debt issued by non-U.S. Sovereigns (excluding central banks) with a 0 percent risk weight (S-1-Q). For IG corporate bonds, we include both LCR and non-LCR qualifying investment grade corporate debt (IG-1 and IG-1-Q). We also collect collateral swap data for the same collateral classes to estimate the value of risky collateral pledged to banks. We assume CME haircuts are 4.5 percent for Treasuries, 25 percent for ETFs and investment grade corporate bonds, and 30 percent for equities. We assume a collateral swap fee of 0.75 percent, and $r_c - r_m$ of 8 percent.

IA.C Appendix Figures and Tables

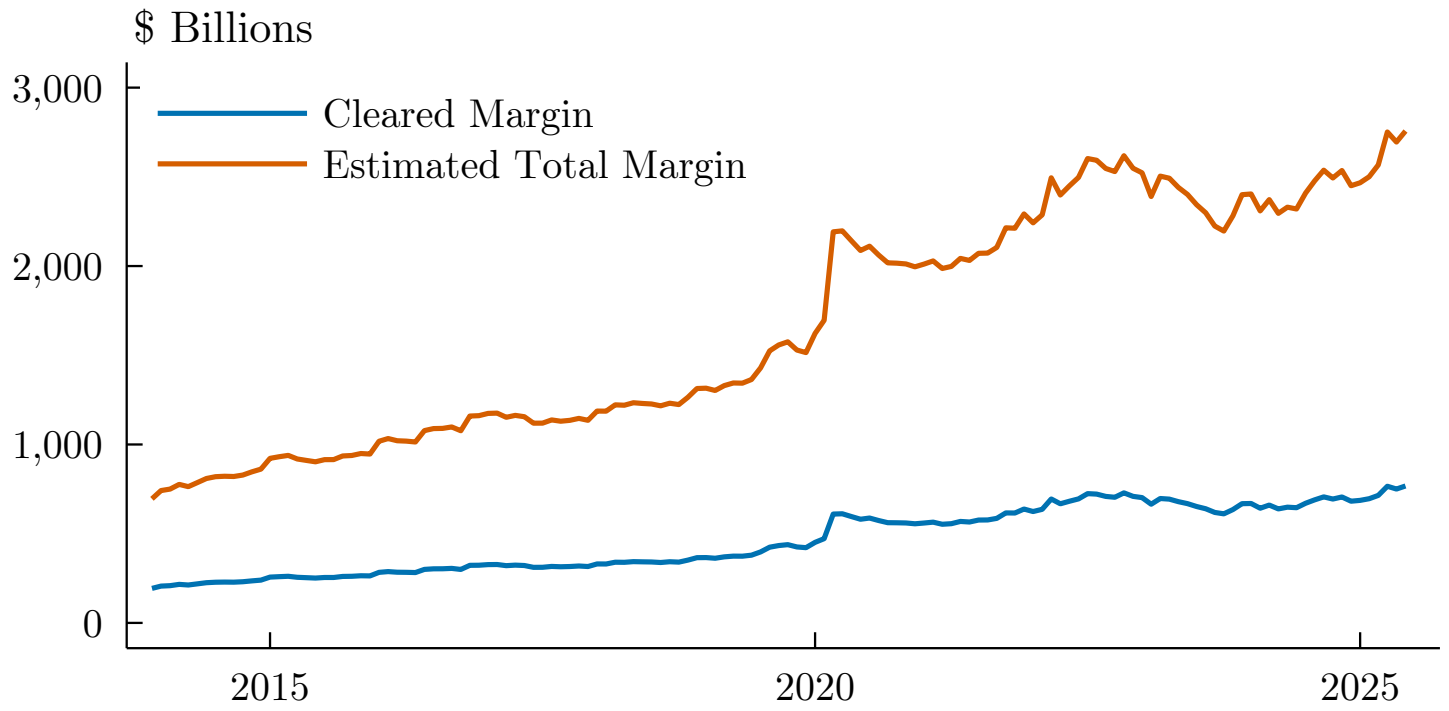


Figure IA.1: Margin Magnitudes. Figure plots the total cleared margin reported by the CFTC as well as an estimate of total margin—including bilateral—that adjusts the CFTC number using the average of the largest banks' bilateral margin posted as a share of total posted margin.

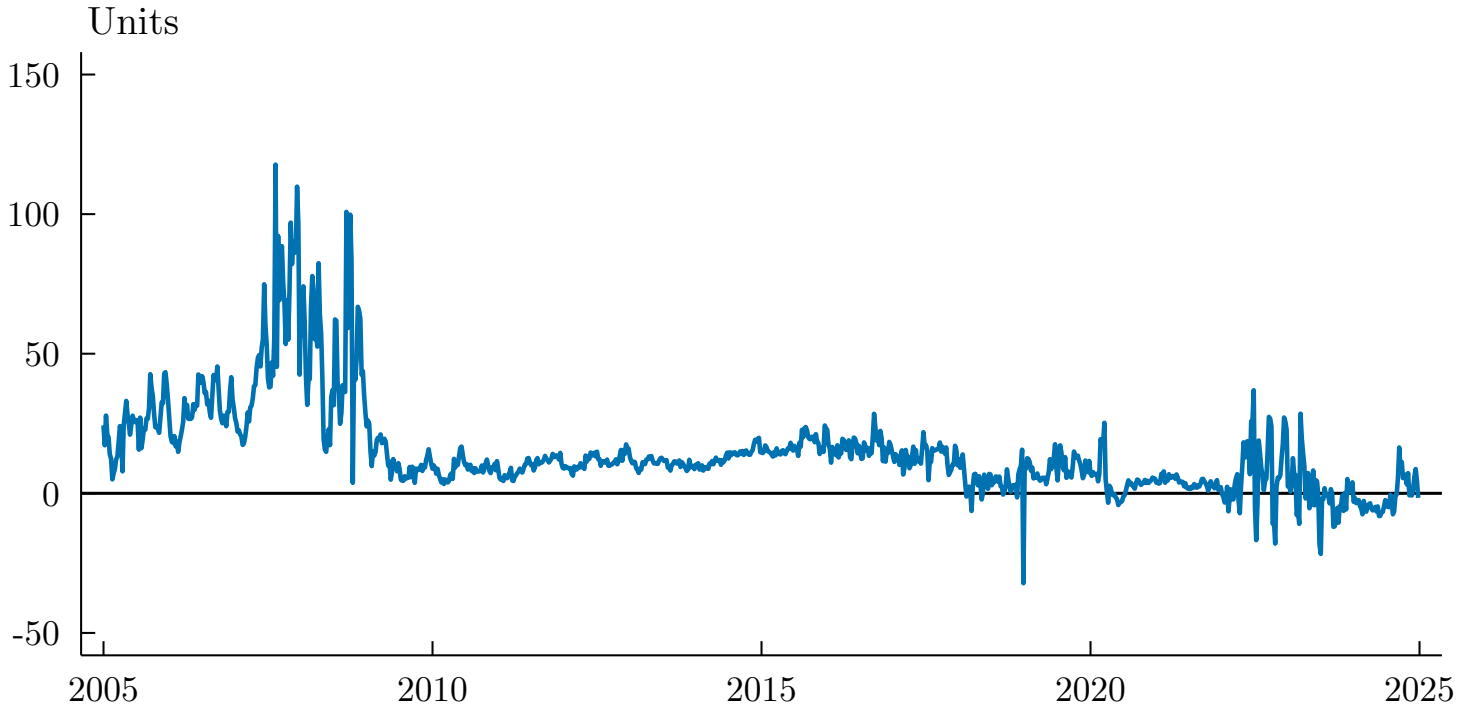


Figure IA.2: Convenience Yield PC1. Plots the first principal component of the six convenience yield proxies enumerated in the Internet Appendix. Estimated from data from 2005 to 2024, and standardized to have zero mean and unit variance. Plot is weekly frequency.

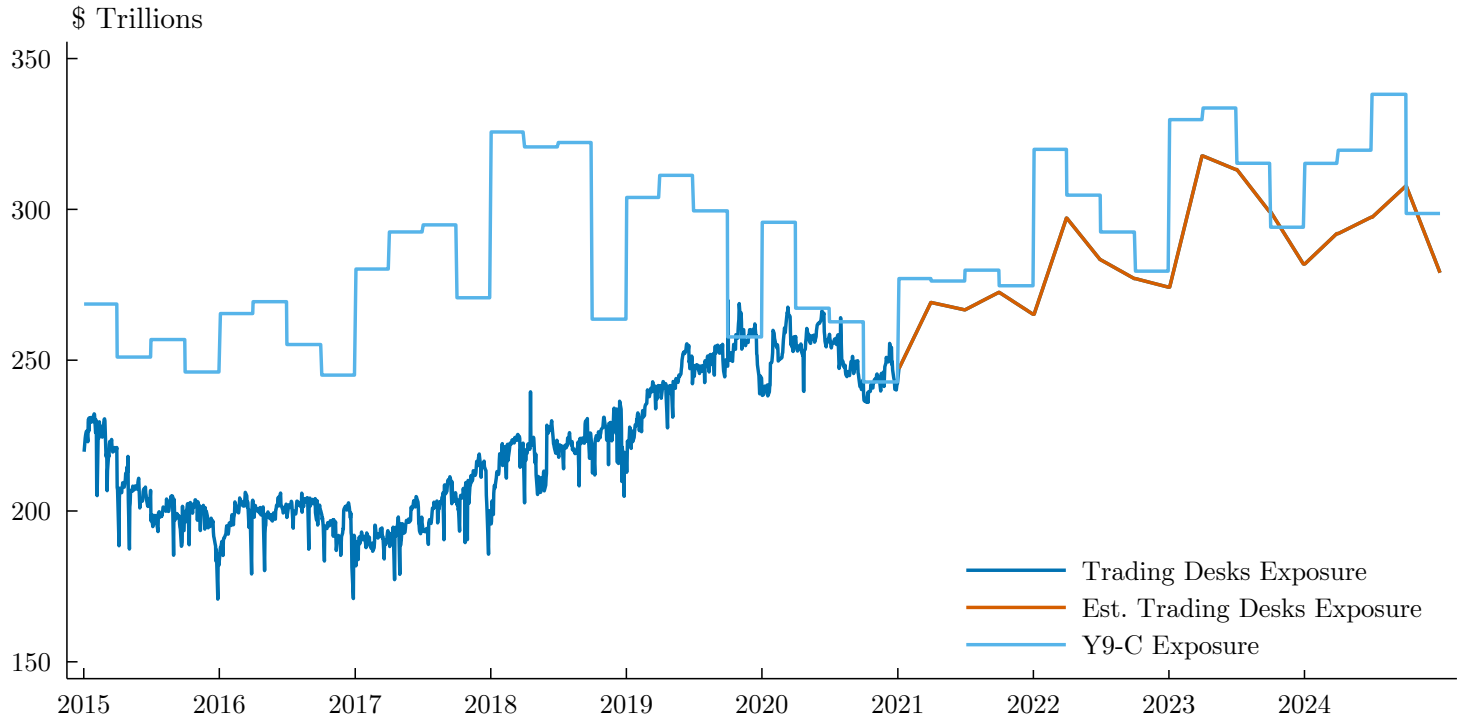


Figure IA.3: Estimated Notional Post 2020. Plots the actual VV-1 notional, estimated VV-1 notional, and actual Y9-C notional. See Online Appendix for discussion of the estimation.

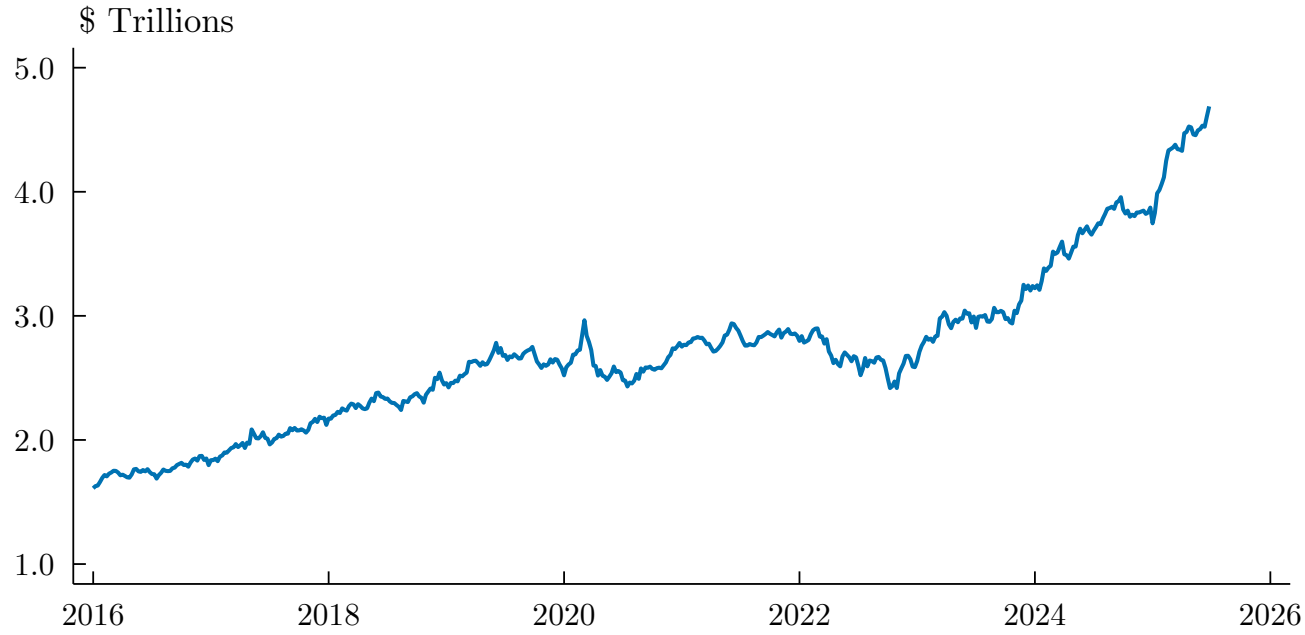


Figure IA.4: Secured Financing Transactions. Figure provides the pre-haircut value of all collateral pledged by banks in our sample in repos, securities lending transaction, and collateral swaps. Plot uses week-end values.

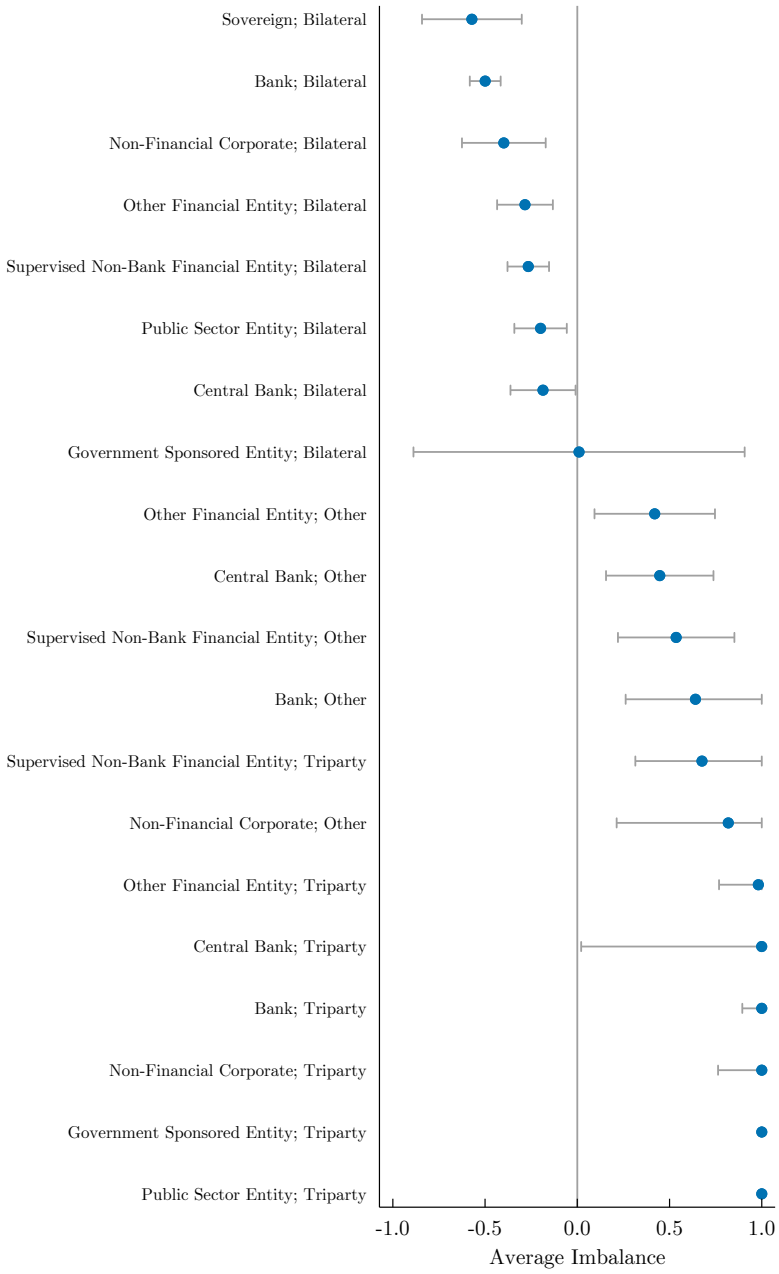


Figure IA.5: Median Collateral Flow Imbalances by Counterparty and Settlement. Figure provides the median collateral imbalance by counterparty and settlement type; imbalance variable described in section 5. Data aggregates across all banks in our sample; bars denote two standard deviations and are truncated at -1 and 1. To preserve confidentiality, we plot only the settlement $\times$ counterparty pairs that appear across all banks in the sample.

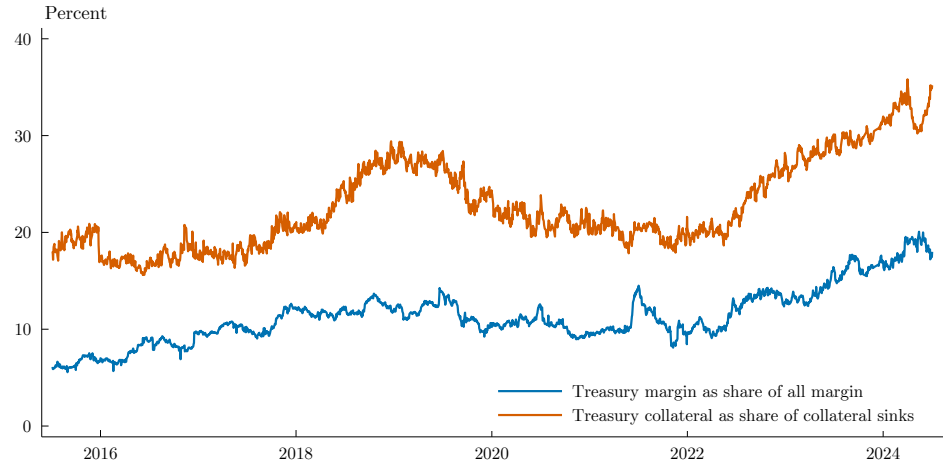


Figure IA.6: Treasuries in Margin and Collateral Sinks. Figure plots Treasury margin as a share of all margin and Treasury collateral as a share of all collateral sinks.

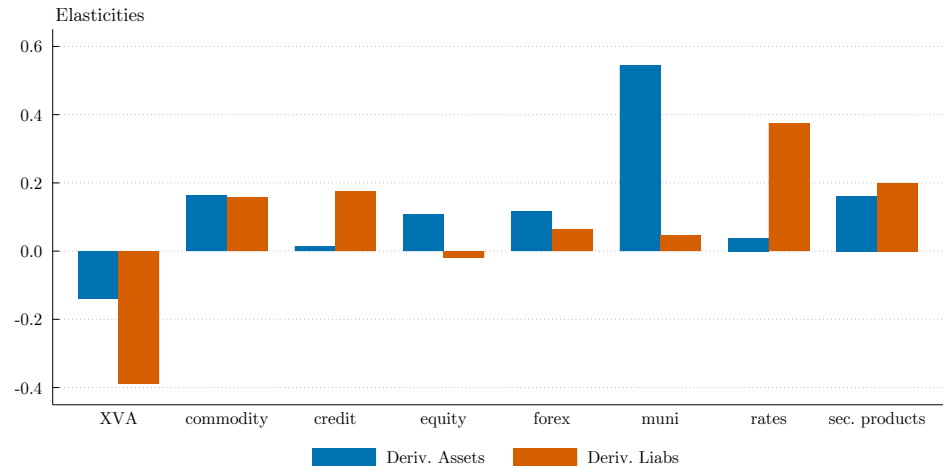


Figure IA.7: Derivative Beta Estimates by Asset Class. Figure plots $\beta^a_{der,A}$ and $\beta^a_{der,L}$ in (15) by asset class.

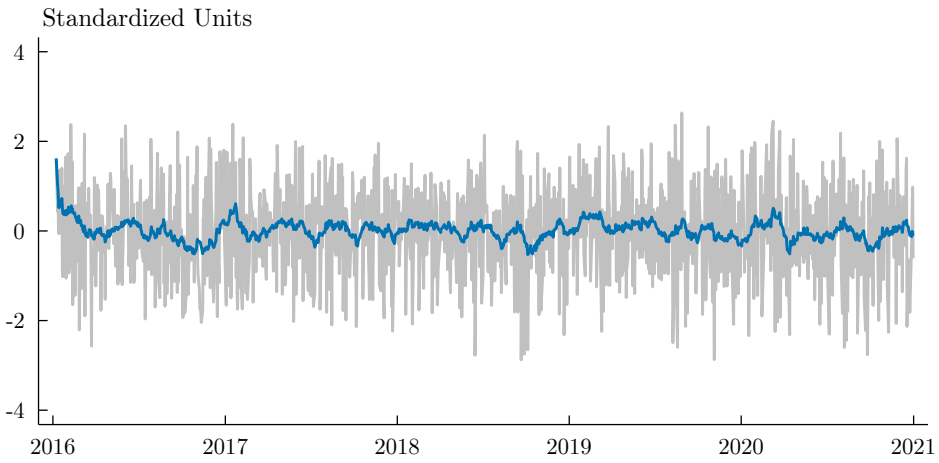


Figure IA.8: GIV. Figure plots the one-month moving average of GIV in blue along with its daily values in grey. We standardized GIV to have zero mean and unit standard deviation.

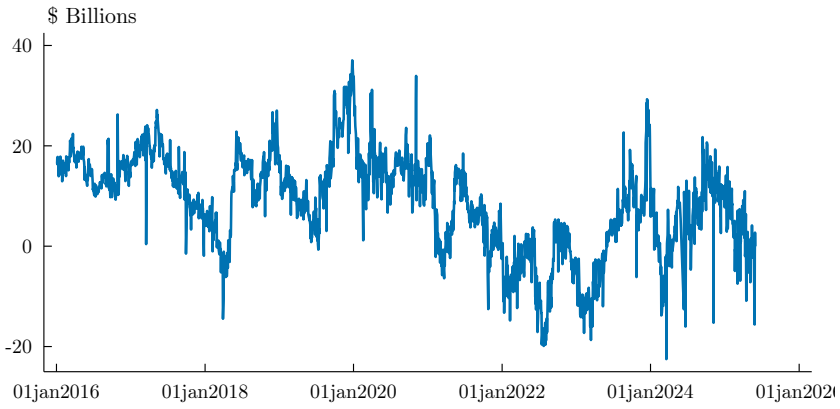


Figure IA.9: Net collateral pledged through collateral swaps. Figure plots the net market value of collateral pledged by banks using collateral swaps.

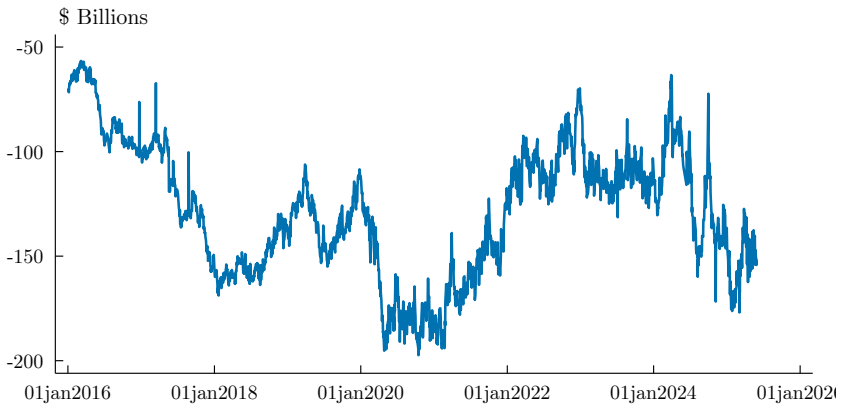


Figure IA.10: Net Treasuries pledged through collateral swaps. Figure plots the net market value of Treasuries pledged by banks using collateral swaps.

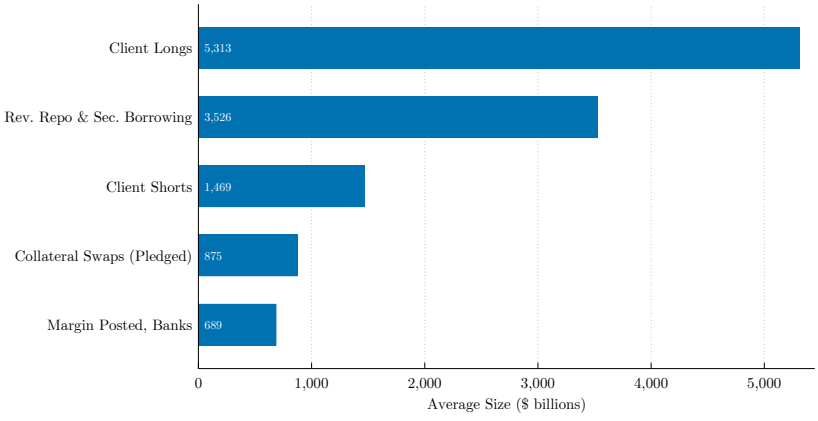


Figure IA.11: Magnitude Comparison. Figure plots the average size of several markets in 2024.

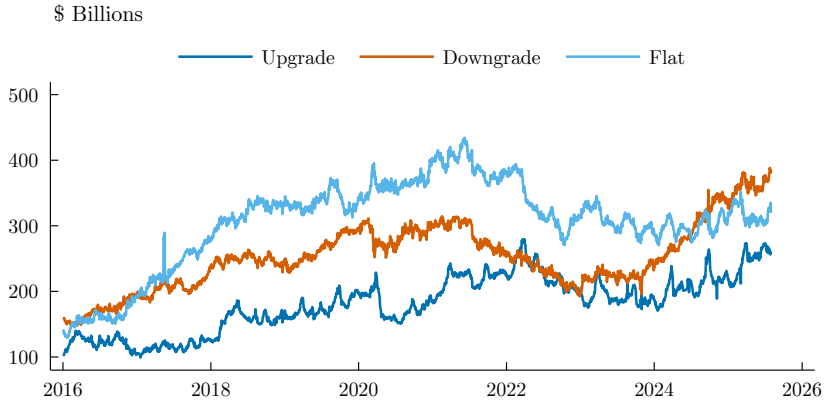


Figure IA.12: Upgrades and Downgrades through Collateral Swaps. Figure plots the total notional of collateral swaps that are upgrades, downgrades, or flat for the bank.

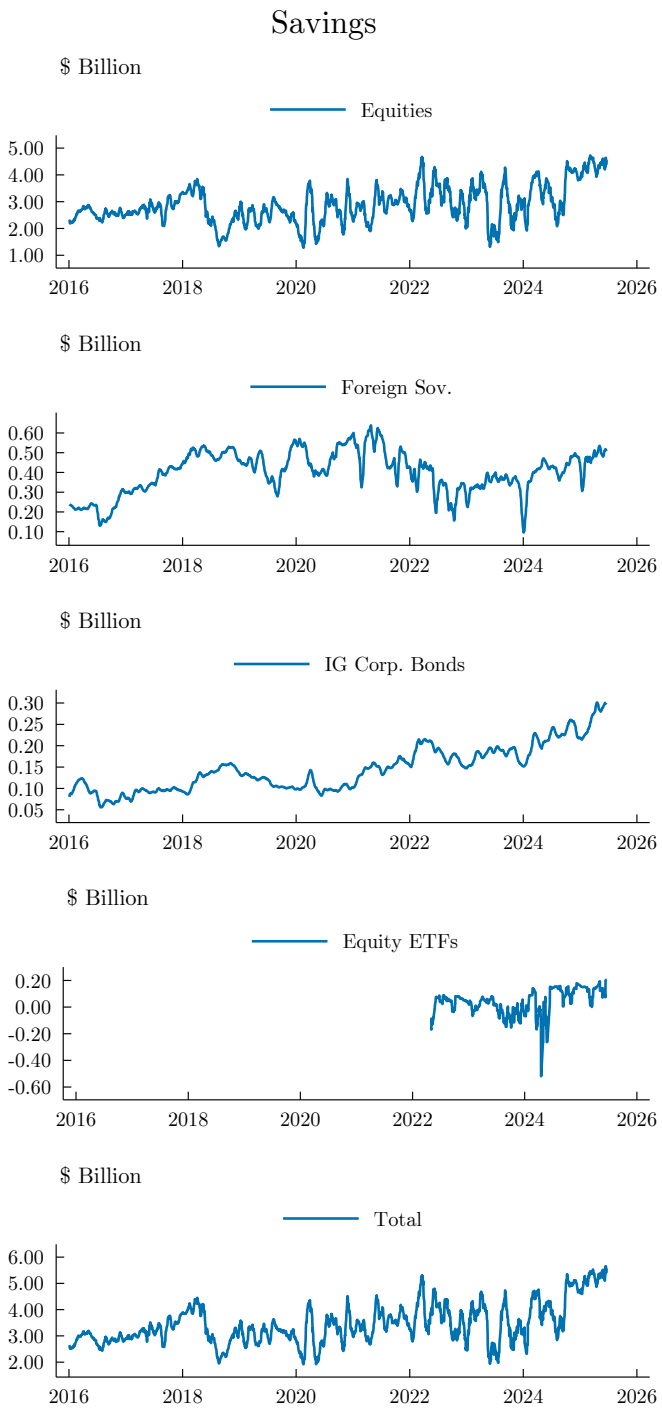


Figure IA.13: Estimated Savings provided by Collateral Swaps. Figure plots the estimated savings using (20) multiplied by the volume of risky collateral posted to banks through collateral swaps. Figure uses CME haircuts, time-varying dealer reverse repo haircuts, and the volume of risky collateral posted to banks through collateral swaps. See section IA.B.1 for methodology details.

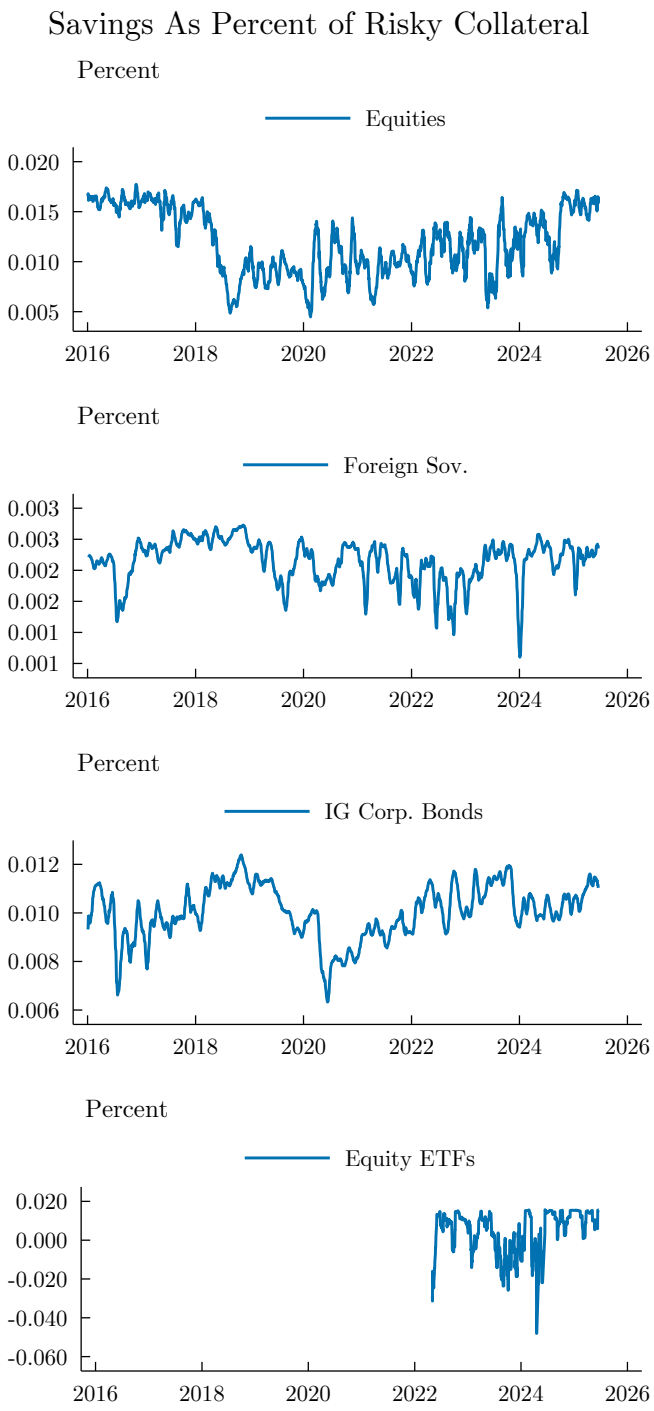


Figure IA.14: Estimated Savings provided by Collateral Swaps. Figure plots the estimated savings using (20). Figure uses CME haircuts, time-varying dealer reverse repo haircuts, and the volume of risky collateral posted to banks through collateral swaps. See section IA.B.1 for methodology details.

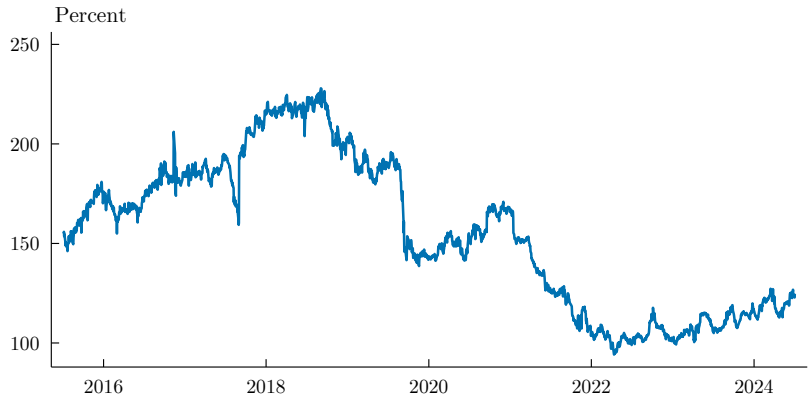


Figure IA.15: Ratio of Collateral Swaps to Margin. Figure plots the ratio of collateral pledged by banks to margin posted to or by banks.

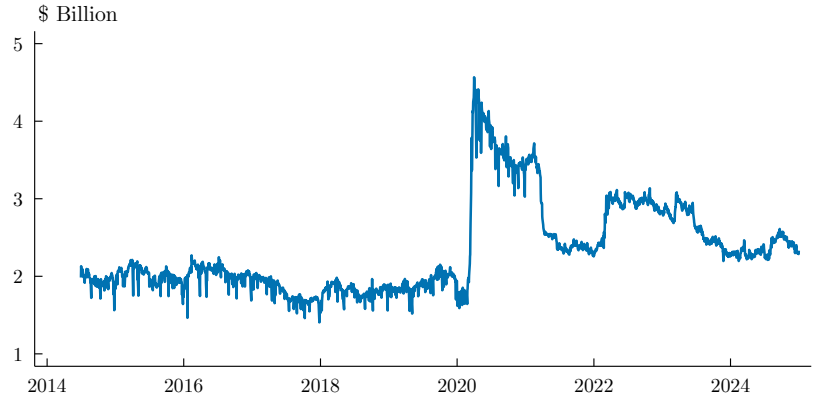


Figure IA.16: VaR. Figure plots the total VaR across the banks in our sample.

Panel A: Initial Margin ($ billions)	Obs.	Mean	SD	Min	Max
Cash	2,217	59	23	20	121
USTs	2,217	48	21	9	100
HQLA	2,217	158	73	37	286
Non-HQLA	2,217	31	9	10	53
Total	2,217	190	81	51	331

Panel B: Variation Margin ($ billions)	Obs.	Mean	SD	Min	Max
Cash	2,217	316	56	202	470
USTs	2,217	17	6	8	45
HQLA	2,217	288	88	147	497
Non-HQLA	2,217	75	48	1	148
Total	2,217	363	56	239	536

Table IA.1: Initial and Variation Margin Summary Statistics. Margin data is daily market value of margin posted by large banks, excluding that posted on behalf of customers, from 2016 to 2024. All values are aggregated across all banks in the sample before calculating moments.

Correlation of PC1 with:		
Treasury margin as share of all margin	-0.49^{***}	(p=0.00; $N = 2,217$)
Treasury collateral as share of collateral sinks	-0.39^{***}	(p=0.00; $N = 2,214$)

Table IA.2: Correlation of PC1 and Treasury Portfolio Choice. Table presents the correlation of PC1 with Treasury margin as a share of all margin and Treasury collateral as a share of collateral sinks. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	ΔMargin Requirement$_t$			
	(1)	(2)	(3)	(4)
ΔVIX_t	3.348***			
	(12.58)			
$\Delta\text{VIX}_t \times \mathbf{1}_{\Delta\text{VIX}_t<0}$		2.974***		
		(5.93)		
$\Delta\text{VIX}_t \times \mathbf{1}_{\Delta\text{VIX}_t>0}$		3.599***		
		(14.86)		
Constant	−0.010	−0.352	5.463***	−4.624***
	(−0.08)	(−1.53)	(14.78)	(−19.02)
N	5,030	5,030	2,301	2,727
R^2	0.40	0.40	0.00	0.00
Sample	All	All	ΔMargin > 0	ΔMargin < 0

Table IA.3: Margin Requirements and VIX. Table presents the relationship between changes in CME margin requirements for E-mini S&P 500 Futures and the VIX. t-statistics shown using heterosketastic and autocorrelation consistent standard errors where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Z_t^{Avg}				Z_t^{PC}			
	(1) $\Delta \ln$ (Central. Margin)	(2) ΔCentral. Ratio	(3) $\Delta \ln$ (Margin)	(4) ΔMargin Ratio	(5) $\Delta \ln$ (Central. Margin)	(6) ΔCentral. Ratio	(7) $\Delta \ln$ (Margin)	(8) ΔMargin Ratio
Z_t^{Avg}	0.001**	0.003**	0.001***	0.003***				
	(2.44)	(2.14)	(4.41)	(4.12)				
Z_t^{PC}					0.001**	0.003**	0.001***	0.003***
					(2.40)	(2.11)	(3.98)	(3.71)
N	2,237	2,236	2,231	2,230	2,237	2,236	2,231	2,230
R^2	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01

Table IA.4: Margin Requirements and CME Shock Measures. Table presents the relationship between different measures of margin and the Z_t measures described in section 6.1. $\Delta \ln$ (Central. Margin) is the change in the log of sunk margin that is posted to a CCP or exchange-traded. ΔCentral. Ratio is the level of centralized margin normalized by the market value of Treasuries outstanding. The other margin variables include all sunk margin, either the difference in the logs or after standardizing against Treasuries outstanding. Z_t variables standardized to have mean zero and unit standard deviation. Constant omitted. t-statistics shown using heterosketastic and autocorrelation consistent standard errors where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Second Stage

	Dependent Var.: $\Delta PC1_t$					
	(1)	(2)	(3)	(4)	(5)	(6)
ΔS_t	53.23	7.52	22.39	24.33	8.35	16.42
	(0.02)	(0.15)	(0.13)	(0.07)	(0.21)	(0.19)
N	2,062	2,000	1,990	2,062	2,000	1,990
Controls	No	Vol.	All	No	Vol.	All

Panel B: First Stage

	ΔS_t					
	(1)	(2)	(3)	(4)	(5)	(6)
Z^{Avg}_{t+1m}	0.00	0.00	0.00			
	(0.02)	(0.17)	(0.13)			
Z^{PC}_{t+1m}				0.00	0.01	0.00
				(0.07)	(0.23)	(0.20)
N	2,062	2,000	1,990	2,062	2,000	1,990
$F-stat$	0.00	0.03	0.02	0.01	0.05	0.04
Controls	No	Vol.	All	No	Vol.	All

Panel C: OLS

	Dependent Var.: $\Delta PC1_t$					
	(1)	(2)	(3)	(4)	(5)	(6)
ΔS_{t+1m}	0.08	0.09	0.08	0.08	0.09	0.08
	(0.61)	(0.68)	(0.61)	(0.61)	(0.68)	(0.61)
N	2,152	2,081	2,071	2,152	2,081	2,071
R^2	0.00	0.00	0.09	0.00	0.00	0.09
Controls	No	Vol.	All	No	Vol.	All

Table IA.5: CME Residual Regression Placebo. Table presents the first-stage, second-stage, and OLS estimates described in Section 6.1 using the placebo shock, which shifts the Z_t shocks forward to one month to Z_{t+1m}. Columns labeled "Vol" include volatility controls: changes in the VIX and Treasury implied volatility. Columns labeled "All" include the full control set: changes in the VIX, Treasury, FX, and oil implied volatilities; the S&P 500 return; the change in the effective federal funds rate; and changes in the Baa–AAA spread, the U.S. 5-year CDS spread, the 5-year Treasury bid–ask spread, and the 5-year–5-year breakeven rate. t-statistics shown using heterosketastic and autocorrelation consistent standard errors using the Newey and West (1994) automatic lag selection procedure where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Second Stage						
	Dependent Var.: ΔPC1_t					
	(1)	(2)	(3)	(4)	(5)	(6)
$\widehat{\Delta S_t}$	−0.38	−0.53	−0.18			
	(−0.09)	(−0.16)	(−0.05)			
$\Delta \ln(\widehat{\text{Sunk Margin}_t})$				−1.16	−2.28	−0.31
				(−0.08)	(−0.12)	(−0.05)
N	1,218	1,169	1,163	1,218	1,169	1,163
Controls	No	Vol.	All	No	Vol.	All

Panel B: First Stage						
	ΔS_t			$\Delta \ln(\text{Sunk Margin}_t)$		
	(1)	(2)	(3)	(4)	(5)	(6)
GIV Placebo_t	0.02	0.03	0.03	0.01	0.01	0.02
	(0.72)	(0.89)	(0.90)	(0.21)	(0.18)	(0.46)
N	1,218	1,169	1,163	1,218	1,169	1,163
$F-stat$	0.52	0.79	0.80	0.04	0.03	0.21
Controls	No	Vol.	All	No	Vol.	All

Table IA.6: Granular IV Placebo. Table presents the first-stage, second-stage, and OLS estimates described in Section 6.2. The regressions are identical to those in Table 6 with the exception that the weights used to estimate the GIV are reversed, as described in section 6.2. Columns (1)–(3) use ΔS_t as the endogenous regressor, and columns (4)–(6) use $\Delta \ln \text{Sunk Margin}_t$. Columns labeled "Vol" include volatility controls: changes in the VIX and Treasury implied volatility. Columns labeled "All" include the full control set: changes in the VIX, Treasury, FX, and oil implied volatilities; the S&P 500 return; the change in the effective federal funds rate; and changes in the Baa–AAA spread, the U.S. 5-year CDS spread, the 5-year Treasury bid–ask spread, and the 5-year–5-year breakeven rate. t-statistics shown using heterosketastic and autocorrelation consistent standard errors using the Newey and West (1994) automatic lag selection procedure where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Dependent Var.: GIV_t
	(1)
GIV Placebo$_t$	−0.0130
	(−0.33)
N	1,255
R^2	0.00

Table IA.7: GIV vs. Placebo GIV. Table presents the first stage regression describe in section 6.2. t-statistics shown using heterosketastic and autocorrelation consistent standard errors using the Newey and West (1994) automatic lag selection procedure where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Constant omitted.

	Dependent Var.: ln(Margin Posted by Banks$_t$)		
	(1)	(2)	(3)
$\ln(s_{i,t})$	0.378***	0.204***	0.194***
	(67.65)	(34.51)	(11.92)
$\ln(\text{VaR}_{i,t})$		0.466***	
		(43.68)	
N	7,100	7,100	7,100
R^2	0.55	0.61	0.83
FE	No	No	Yes

Table IA.8: Risk-Adjusted Notional vs. Margin Posted. Table regresses the risk-adjusted derivative notional $s_{i,t}$ described in section 6.2 on posted margin as reported in FR2052a and cumulative VaR across the sample of desks used in the GIV analysis. Panel is at the bank by date level. Fixed effects row includes date and bank fixed effects. Note that cumulative VaR across desks is not necessarily equal to the bank's aggregate VaR given possible hedging benefits across desks. t-statistics shown using heterosketastic and autocorrelation consistent standard errors using the Newey and West (1994) automatic lag selection procedure where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Constant omitted.

	Z_t^{PC}		Z_t^{Avg}	
	(1) Ratio	(2) VaR	(3) Ratio	(4) VaR
Z_t	0.823**	0.517**	0.793**	0.562**
	(2.32)	(2.18)	(2.20)	(2.27)
Quintile 2	0.109	−0.033	0.110	−0.029
	(0.73)	(−0.20)	(0.72)	(−0.18)
Quintile 3	0.043	0.078	0.042	0.076
	(0.31)	(0.49)	(0.29)	(0.50)
Quintile 4	0.076	0.019	0.077	0.019
	(0.55)	(0.11)	(0.54)	(0.11)
Quintile 5	0.040	−0.075	0.042	−0.072
	(0.30)	(−0.47)	(0.30)	(−0.48)
Quintile 2 × Z_t	−0.769*	−0.177	−0.692	−0.198
	(−1.78)	(−0.59)	(−1.60)	(−0.65)
Quintile 3 × Z_t	−0.577	−0.461	−0.490	−0.440
	(−1.61)	(−1.60)	(−1.33)	(−1.50)
Quintile 4 × Z_t	−0.697*	−0.338	−0.612	−0.306
	(−1.78)	(−1.06)	(−1.53)	(−0.96)
Quintile 5 × Z_t	−0.836**	−0.660**	−0.842**	−0.701**
	(−2.11)	(−2.08)	(−2.01)	(−2.04)
N	2,120	2,159	2,120	2,159
R^2	0.10	0.09	0.10	0.10

Table IA.9: Passthrough Regression. Table presents the first stage regression describe in section 6.2. All columns include the following controls: changes in the VIX, Treasury, FX, and oil implied volatilities; the S&P 500 return; the change in the effective federal funds rate; and changes in the Baa–AAA spread, the U.S. 5-year CDS spread, the 5-year Treasury bid–ask spread, and the 5-year–5-year breakeven rate. Quintile 5 reflects lowest constraint quintile, quintile 1 highest. Constant omitted. t-statistics shown using heterosketastic and autocorrelation consistent standard errors using the Newey and West (1994) automatic lag selection procedure where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Constant omitted.