

Aggregate Collateral Demand*

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

April 27, 2026

Abstract

We study aggregate collateral demand and document its effect on the Treasury convenience yield. Prices of safe assets reflect a variety of non-pecuniary benefits—safety, liquidity, and collateral services—whose relative importance is hard to pin down. We establish the causal pass-through from aggregate collateral demand shocks to the Treasury convenience yield by measuring *collateral sinks*, different ways in which safe assets are removed from circulation, even if temporarily. We show that the rapidly growing over \$1 trillion collateral-swap market exacerbates these effects in times of stress.

JEL Codes: E40, E51, G12, N20

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

*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 concurrence by other members of the research staff or the Board of Governors.

[†]New York University. Email: tlarits@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

Prices of safe assets reflect a 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 convenience. Some drivers—liquidity and safety—are partly observable with proxies like bid–ask spreads and credit default swap spreads. The collateral channel is harder to observe. In this paper, we isolate the convenience yield arising from Treasury use as collateral, untangling the ways that aggregate collateral demand can raise convenience yields, widen arbitrage spreads, and interact with dealer constraints to produce the spikes in safe-asset premia observed during stress.

We conceptualize aggregate collateral demand by introducing the notion of *collateral sinks*: uses of safe assets that temporarily withdraw them from circulation. While collateral can often be reused,¹ many Treasuries are posted or pledged in ways that leave them temporarily immobile. For example, margin posted to a central counterparty (CCP) against derivative positions cannot be rehypothecated. Some counterparties place legal or contractual constraints on what can be done with their collateral, sometimes prohibiting its reuse. Using regulatory data we estimate the total size of collateral sinks at \$3.4 trillion in 2025, following nearly a decade of steady growth. Such *sunk collateral* constitutes 15 percent of Treasuries outstanding—a striking magnitude—and varies substantially over our sample reaching back to 2016.

First, we show that the amount of sunk collateral is a key driver of convenience yields. This result is intuitive: by simple market-clearing logic, when collateral sinks grow, the supply of Treasuries available outside these sinks contracts, raising the convenience yield of the remaining supply. The collateral channel is important in explaining convenience yields above and beyond the two most commonly studied drivers: credit risk and liquidity. Adding our collateral sink measure to a specification controlling for both channels increases explanatory power by more than 60 percent at a daily frequency, comparable to the marginal contribution of bid-ask spreads and double the marginal contribution from credit risk.

We provide causal evidence for this channel using two identification strategies. The first approach relies on idiosyncrasies in the Chicago Mercantile Exchange’s (CME) margin-setting framework. The CME—the largest U.S.-based CCP and a major hub for trading in Treasury, interest-rate, and equity-related derivatives—adjusts margins in response to

¹See (Infante and Saravay, Forthcoming) on Treasury rehypothecation.

market volatility but also incorporates several other factors when setting margin requirements. These factors generate variation in margin requirements that is plausibly orthogonal to the market-wide forces that drive Treasury convenience yields. Instrumenting collateral sinks with shocks to margin demand, we show that a one-standard-deviation increase in sunk collateral relative to Treasuries outstanding raises convenience yields by about 3 basis points (cast in units comparable to the 3-month OIS–T-bill spread, which averages 15 bps with a standard deviation of 20 bps). Our second identification strategy uses a granular instrument constructed from banks’ trading desk notional exposures. This approach yields quantitatively similar estimates, reinforcing the causal pass-through from sunk collateral to convenience yields.

Collateral sinks help explain movements in collateral-intensive Treasury arbitrage spreads. A one-standard-deviation shock to collateral sinks widens Treasury swap spreads by about a basis point. The mechanism is straightforward: these trades require Treasuries as margin and financing collateral, so greater collateral scarcity raises funding costs and widens spreads. We do not find comparable effects for arbitrage dislocations with low Treasury-collateral requirements, such as cross-currency bases or the CDS-bond basis. This contrast suggests that collateral sinks identify a collateral-specific component of the dealer balance sheet costs that Du et al. (2023) show drive arbitrage spreads more broadly.

This substantial demand for collateral, and the resulting impact on Treasury prices, reflects both broad trends toward financialization as well as regulatory action. Dysfunction in opaque and largely bilateral collateral markets during the global financial crisis imprinted on 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 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. They also increased Treasury demand by requiring banks to hold additional safe assets under the liquidity coverage ratio.

Without a concomitant increase in safe-asset issuance, these regulations have often proved binding (Caballero et al., 2017). Part of the increased demand for good collateral has been met by the private sector in the form of a large—over \$1 trillion USD—and growing *collateral*

swap market. Collateral swap contracts replace one type of collateral with another and allow investors to transform risky assets into safe assets. Such swaps are especially useful for meeting margin requirements by allowing investors to exchange high-haircut securities for lower-haircut safe assets, thereby lowering their capital needs. We estimate that these transactions reduce investors' margin capital by \$70 billion, allowing up to \$700 billion in additional levered exposures using average leverage ratios for hedge funds.

At the same time, the collateral swap market can amplify stress-period spikes in convenience yields. When banks are unconstrained, they can help investors meet margin requirements by providing lower-haircut collateral through collateral swaps. 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, while banks simultaneously supply less high-quality collateral via swaps, leaving investors unable to rely on the swap market to meet their rising margin requirements.

Our second main contribution is to show that the interaction of collateral sinks and collateral swaps can produce spikes in convenience yields when banks are constrained. Unconstrained banks attenuate the elasticity of the convenience yield with respect to collateral sinks by intermediating the collateral swap market. When banks grow more constrained, we show that the elasticity of the convenience yield with respect to collateral sinks increases by more than fourfold, helping explain the spikes in convenience yields that we observe in bad states.

We quantify this amplification using March 2020: under the actual path of dealer constraints, margin shocks cumulatively account for about three-quarters of the increase in the spike in the convenience yield, whereas a counterfactual in which banks maintained median intermediation capacity reduces the effect of margin shocks on the convenience yield to approximately zero.

We formalize the relationship between sunk collateral and convenience yields in a model where Treasuries convey various convenience services—liquidity, safety, collateral—and those services are temporarily lost once a Treasury is placed in a collateral sink. The convenience yield, then, equals the shadow value of the Treasury's pledgeability. The model shows that convenience yields are increasing in sunk collateral. The degree of pass-through from sunk collateral to convenience yields depends on dealer constraints; banks buffer the effect by intermediating collateral swaps when constraints are slack, but amplify it when constraints

are tight and banks hoard collateral.²

Combined, our analysis makes three novel contributions. First, we measure collateral sinks, and document their level, time variation, and composition. Second, we provide the first causal estimate of the elasticity of Treasury convenience yields with respect 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 cyclicity of dealer intermediation, to 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 affect 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 studies 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, and how markets transform collateral when safe assets become scarce. Infante and Saravay (Forthcoming) show a related mechanism, 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

²We will refer interchangeably to banks and dealers for brevity even though the broker-dealer subsidiary of a bank-holding company typically conducts collateral swaps.

Treasury convenience yields. Duffie et al. (2015) show how central clearing affects margin requirements in CDS markets.

Our paper contributes to the literature by isolating the collateral channel and its effect on convenience yields using collateral-demand shocks. 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.

1 Institutional Context

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

1.1 Margin

Financial market participants must post margin to reduce default risk 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 consists of two components: initial and variation. Traders owe initial margin when they initiate a derivative transaction, and variation margin is called as the contract's value changes.

Margin requirements change and reflect expectations about the derivative's risk. Figure 1 shows that the CME's margin requirements for the 10-year Treasury future range from 60 to 225 bps, and move closely—but not perfectly—with implied volatility; the same applies for S&P 500 index futures, although their margin requirements are a level shift higher, about 5 pp, reflecting their higher risk.³

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). 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.

³The Internet Appendix gives an example of the lifecycle of a derivative trade that clarifies the margin mechanics and provides additional institutional details.

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.⁴ 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.⁵

Table 1 provides summary statistics on our sample and shows that cash constitutes the majority of margin posted, averaging \$376 billion, or about two-thirds of all margin posted. Treasuries are the next largest collateral type, averaging \$70 billion or 12 percent of all margin, but exceeding 20 percent in some periods.

1.2 Liquidity Coverage Ratio

Margin is not the only regulatory source of collateral demand. The liquidity coverage ratio (LCR), introduced under Basel III and finalized in the U.S. in 2014, requires large banks to hold enough high-quality liquid assets (HQLA) to cover at least 100 percent of projected net cash outflows over a 30-day stress scenario. Different assets have different HQLA weights: Treasuries, for example, contribute to HQLA holding with no haircut or cap; agency MBS have a 15 percent haircut. The LCR creates structural Treasury demand. A bank whose HQLA falls short of its projected net cash outflows must hold additional assets to fill the gap. When banks choose to fill that gap with Treasuries, they are functionally sunk—held to meet a regulatory floor, unavailable for lending, pledging, or other uses that would keep them circulating.

1.3 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

⁴Many CCPs also pay interest on USD cash balances. CCPs that are “designated financial market utilities” can maintain accounts at the Federal Reserve which are eligible to receive interest on reserve balances. CCPs can, in turn, pass through 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).

⁵See <https://www.cmegroup.com/solutions/clearing/financial-and-collateral-management/acceptable-collateral.html>.

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.1).

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 possible by reusing and rehypothecating collateral they receive. To isolate the contribution of collateral services to the convenience yield, it is important to control for these reuse dynamics, which we address next.

2 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.

2.1 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 return R_A , a Treasury security T with price p_T and gross return R_T , and a derivative contract D with per-unit gross return 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:

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

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.

2.2 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:

$$\begin{aligned} c_t &= y - p_A A - p_T T - f(\chi)q - \kappa(A^m), \\ c_{t+1} &= p_A A R_A + p_T T R_T + d R_D, \end{aligned}$$

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 . $\kappa(A^m)$ represents a small convex cost to posting collateral. 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:

$$(1 - h_T)p_T(T^m + q) + (1 - h_A)p_A A^m \geq M d. \quad (2)$$

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 $\lambda_T \equiv \Omega'(T^u)/u'(c_t)$ the shadow value of holding an unencumbered Treasury in consumption units. The household's optimization yields the following first-order conditions.

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

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

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

The decision to post Treasuries as margin reflects a trade-off between the margin value and the forgone convenience services:

$$\frac{\mu(1 - h_T)p_T}{u'(c_t)} = \begin{cases} \lambda_T \leq f(\chi), & \text{if } T^m > 0 \text{ and } q = 0, \\ \lambda_T = f(\chi), & \text{if } T^m > 0 \text{ and } q > 0. \end{cases} \quad (4)$$

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

$$\frac{\mu(1 - h_A)p_A}{u'(c_t)} = 1 - \mathbb{E}_t[m_{t+1}R_A]. \quad (5)$$

When risky assets are posted as collateral, the shadow value of margin capacity μ must offset the risk premium on A , $1 - \mathbb{E}_t[m_{t+1}R_A]$, scaled by the haircut. When the margin constraint binds ($\mu > 0$), any risky asset holdings are fully encumbered.

The optimal derivative position satisfies:

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

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.

2.3 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.

The dealer's optimization problem is:

$$\max_{\ell, q} \quad \Pi = f(\chi)q - \Psi(\ell; \chi) \quad \text{subject to} \quad 0 \leq q \leq \ell. \quad (7)$$

The first-order conditions are:

$$\text{w.r.t. } \ell: \quad \Psi_\ell(\ell; \chi) = \nu, \quad \text{w.r.t. } q: \quad f(\chi) = \nu, \quad (8)$$

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

$$f(\chi) = \Psi_\ell(\ell; \chi). \quad (9)$$

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.

2.4 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.

2.5 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:

$$\frac{\partial \text{CY}_T}{\partial (Md)} > 0, \quad \frac{\partial \text{CY}_T}{\partial h_T} > 0, \quad \frac{\partial \text{CY}_T}{\partial h_A} > 0. \quad (10)$$

Higher Treasury haircuts increase convenience yields by reducing the margin capacity per unit of Treasury posted. Higher haircuts on risky assets also increase convenience yields, but through a different channel: by reducing the margin capacity of risky collateral, they shift margin demand toward Treasuries, lowering unencumbered holdings.

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:

$$\frac{\partial \text{CY}_T}{\partial S} = (1 - h_T) \frac{\partial \mu}{\partial S}, \quad \text{where } \frac{\partial \mu}{\partial S} \text{ increases in } f(\chi). \quad (11)$$

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:

$$\frac{\partial s_T}{\partial \text{CY}_T} < 0, \quad \frac{\partial s_T}{\partial h_T} < 0. \quad (12)$$

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:

$$\frac{\lambda_T}{(1 - h_T)p_T} + \frac{f(\chi)}{(1 - h_T)p_T}, \quad (13)$$

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.

3 Data

The paper relies primarily on three datasets: trading desk data (FR VV-1), high frequency balance sheet data (FR2052a), and a 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 a date bank-desk panel, running from 2014 through 2024. The data have 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.

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; we will refer to BHCs as banks for brevity. The data include quantities on exposures but not prices or rates. We focus on the set of banks that consistently file at a daily frequency. 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 cover 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.

4 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, represent 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 \$3.4 trillion—which motivates that changes in sunk collateral correspond to changes in aggregate collateral and its attendant effect on convenience yields.

4.1 Construction

We measure collateral sinks as the sum of collateral used in three ways: collateral posted as margin or to back SFTs, and collateral sunk to satisfy bank liquidity regulations. We limit ourselves to the subsets of these markets where we have evidence or strong priors that the collateral stops 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.⁶ 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),

⁶There 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.

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 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 mass of cells clustered near 1. Empirically, most counterparty types only receive collateral from banks, as shown in the histogram of Figure IA.2, 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.3).

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.⁷ We could also include other CCP contributions—like clearing members’ default fund contributions—but those data are available beginning only in 2022, are orders of magnitude smaller than the other categories, and are relatively slower moving, so they 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 4.⁸ So even if a bank satisfies the collateral demand with \$1 of 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, we work with the sink scaled by the market value of publicly held Treasuries, which we define as S_t .⁹ 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-proceeds) 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.

LCR-Induced Collateral Sinks Banks must also hold HQLA to satisfy the Liquidity Coverage Ratio, which requires that high-quality liquid assets cover at least 100 percent of

⁷Regulations require that certain types of derivatives are centrally cleared even though they are negotiated OTC. Derivative exchanges all use a CCP.

⁸While 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.

⁹We 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.

projected net cash outflows over 30 days. Some of these required holdings would be held absent the regulation, but to the extent that the LCR binds, it forces banks to hold Treasuries and other HQLA that cannot be deployed elsewhere. We estimate the LCR-induced Treasury sink as the volume of Treasuries that banks must hold to satisfy the requirement, defined as

$$\text{LCR Sunk Treasuries} = \max\{0, \text{Net Cash Outflows} - \text{HQLA ex Treasuries}\}.$$

When non-Treasury HQLAs—reserves, agency MBS, and other qualifying assets—are insufficient to cover projected outflows, the shortfall must be met with Treasuries, which are then effectively sunk.¹⁰ We compute net cash outflows, the LCR ratio denominator, using the Federal Reserve’s FR 2052a reporting data mapped to the provisions of Regulation WW. Given high-frequency seasonality in net cash outflows, we smooth the series with a five-day moving average, from $t - 4$ to t .

Ideally, we would observe the Treasuries each bank would hold absent the LCR and compare that counterfactual to actual holdings; the difference would be the true LCR-induced sink. That comparison is impossible, so we instead use the lower bound defined above. This measure is conservative: it assumes banks exhaust all non-Treasury HQLA before turning to Treasuries, so it understates the true LCR-induced sink whenever banks choose to hold Treasuries in place of other eligible assets. The key concern is that banks’ portfolio composition across HQLA categories is itself endogenous—for instance, a bank expecting higher Treasury convenience yields might tilt its HQLA buffer away from Treasuries and toward reserves, generating a correlation between our sink measure and convenience yields that does not reflect a causal scarcity channel.

We view this concern as strengthening rather than undermining our results, for two reasons. First, endogenous rebalancing works against our finding: when Treasury convenience yields are high, banks have an incentive to substitute toward alternatives in lieu of Treasuries in their HQLA buffers, which reduces our measured LCR sink precisely when convenience

¹⁰Our required Treasury calculation assumes that removing Treasuries from the HQLA pool does not alter the binding status of the LCR’s composition constraints, which limit Level 2 assets to 40 percent of total HQLA and Level 2B assets to 15 percent. Because the banks in our sample hold predominantly Level 1 assets, these caps do not bind in practice, and recomputing the caps under the counterfactual without Treasuries yields negligibly different required Treasury amounts. The results are also robust to including a management buffer, for example, by calculating the 5th percentile of observed bank by day LCR ratios, and imposing that as the effective lower bound.

yields are elevated. Any such rebalancing therefore attenuates our estimates rather than generating a spurious positive relationship between the sink and convenience yields. Second, our instrumental variable strategy, which exploits variation in derivatives margin requirements, is orthogonal to banks' HQLA portfolio decisions, so any bias from endogenous composition would attenuate the OLS estimates without contaminating the IV.

In the Internet Appendix, we compare four constructions of the collateral sink: the benchmark measure that includes the LCR lower bound, an alternative that uses the full LCR numerator (total HQLA holdings), and a version that excludes all LCR-related holdings, and a version of our benchmark that incorporates estimated management buffers since banks typically leave some buffer between their preferred effective lower bound LCR ratio and the regulatory minimum. Daily changes in the three series are closely correlated, and the regression estimates are quantitatively similar across specifications (Figure IA.4 and Table IA.1).

4.2 Collateral Sink Summary Statistics

Collateral sinks stand at \$3.4 trillion in 2025, equal to 15 percent of Treasuries outstanding, as shown in Figure 3. SFTs account for almost half of the total, with repos/sec lending averaging \$732 billion (41 percent of sunk collateral) and collateral swaps \$215 billion (12 percent). Margin averages \$526 billion (29 percent of collateral sinks), although its share has approached 50 percent of the total at times. LCR-related sunk Treasuries stand at \$351 billion, or 19 percent of total sunk collateral.

Is \$3.4 trillion of sunk collateral a big number? Yes, for two reasons. First, it is equal to 15 percent of the total stock of privately held Treasuries and about 80 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 can be swapped into 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 nor does it include margin posted by banks' clients (due to data availability) as well as 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 given that 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 39 percent, other HQLA 29 percent, cash 12 percent, and other non-HQLA 20 percent. Sunk Treasuries are the most volatile, both in level terms (\$457 billion) and relative to their mean. 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.¹¹

4.3 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 (≈ 0.6 SD) is associated with a 4 bps higher PC1 (≈ 1 SD). The relationship between the two suggests that collateral sinks track aggregate collateral demand insofar as it varies with the convenience yield.

The levels regression is not causal and cannot 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 compares changes on changes; the full specification in column (6) shows that a 1pp increase in ΔS_t raises $\Delta PC1_t$ by about 3 bps after controls, equal to about one standard deviation in daily changes of convenience yields. 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.¹²

While the coefficient on ΔS_t is statistically significant and economically large, a natural

¹¹In the Internet Appendix, we show that collateral sinks show some quarter-end seasonality, falling somewhat in the few days of each quarter, likely consistent with regulatory window dressing.

¹²In the Internet Appendix, we provide several robustness tests: we regress individual convenience yield measures on the collateral sink measure, we split the convenience yield measures into short-term and long-term and show that short-term convenience yields tend to load more on the collateral sink measure, and we show the relationships are robust to removing Treasuries from the numerator and instead comparing with the change in log of collateral sinks, including controlling for change in Treasuries outstanding.

question is how its explanatory power compares to the channels the literature has traditionally emphasized. We address this with an incremental R^2 exercise. For a given specification, we estimate the full model including all regressors, then re-estimate the model dropping each regressor in turn. The percent increase in R^2 from restoring a given variable measures its marginal contribution to explained variation. Because daily first-difference regressions have low R^2 in absolute terms—as is standard at daily frequency—we express each variable’s contribution as the percentage increase in R^2 relative to the specification that excludes it.

Figure 4 presents the results across two specifications. Panel A includes proxies for standard convenience yield explanatory variables: sovereign credit risk (Δ US CDS) and Treasury market liquidity (Δ 5yr Bid-Ask Spread), alongside collateral sinks. Both the collateral sink and liquidity measures increase R^2 by 60 percent, while credit risk increases it by about 30 percent. Panel B adds changes in the VIX, the BAA-AAA spread, breakeven inflation, the effective federal funds rate, implied volatility measures, and equity returns. Collateral sinks remain among the largest contributors across the kitchen sink specification, with a marginal contribution of 6 percent even in the most saturated specification, larger than credit risk (3 percent) but smaller than the contribution from the Fed funds rate (nearly 230 percent), inflation (14 percent) and liquidity (8 percent). These results show that collateral sinks are not merely statistically significant but quantitatively important relative to the credit risk and liquidity channels that Krishnamurthy and Vissing-Jorgensen (2012) and others have highlighted as first-order drivers of convenience yields. Table IA.2 in the appendix reports the underlying regressions.

Which component of collateral sinks drive the result? Table IA.3 breaks the aggregate measure into its separate components—repo/sec lending, collateral swaps, margin, and LCR-related holdings. The variables have positive coefficients, and the margin and collateral swap components are significantly different from zero, while the standard error for collateral swaps is somewhat larger even though its magnitude is similar to the others. The LCR-related component is weakly positive and close to zero. Because the margin component does not rely on the imbalance filter or rely on the assumptions used to estimate LCR-related holdings, it is perhaps unsurprising that the margin component is estimated most precisely whereas the other 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.¹³ In our data, the Treasury share of margin posted reaches a new high in 2024 (Figure IA.5), and shows a strong negative relationship between that share and PC1 as illustrated in Figure 5, confirming that lower convenience yields coincide with more use of Treasuries as collateral.¹⁴

4.4 Collateral Sinks and Arbitrage Spreads

If collateral sinks capture real variation in aggregate collateral demand, their pricing effects should extend beyond the convenience yield measures studied above. We test this using arbitrage spreads—mispricings between near-identical assets that persist because exploiting them requires collateral-intensive balance sheet capacity. Swap spreads are a common example: a dealer who trades against a negative swap spread must finance a long Treasury position, tying up collateral (Boyarchenko et al., 2018). When collateral is scarce, that trade costs more, and the spread widens. More broadly, Du et al. (2023) show that dealer balance sheet costs jointly rationalize swap spreads and CIP deviations post-GFC; our test asks whether collateral sinks identify a specific component of those costs.

Table 4 regresses daily changes in arbitrage spreads on daily changes in S_t . Panel A reports results for Treasury swap spreads at the 2-, 5-, 10-, and 30-year maturities as well as their first principal component. A 1 pp increase in ΔS_t is associated with a 1 bp widening in 10-year swap spreads, and the relationship is positive and significant across all maturities. The effect on the first principal component is similarly strong, confirming that collateral sinks load onto the common component of swap spread variation rather than any single maturity.

Panel B turns to arbitrage spreads that are less Treasury-collateral-intensive: the EUR, JPY, and GBP cross-currency bases and the investment-grade and high-yield CDS-bond bases. We do not find comparable effects for any of these spreads. The contrast between panels is informative. CIP arbitrage requires balance sheet capacity but not specifically Treasury collateral, and the CDS-bond basis similarly does not require a dealer to finance a Treasury

¹³See https://www.federalreserve.gov/data/scoos/files/scoos_202509.pdf.

¹⁴We give the correlation of these variables in Table IA.4.

position. Swap spread arbitrage, by contrast, is collateral-intensive: exploiting it requires holding and financing Treasuries. Collateral sinks thus operate through a collateral-specific channel that is distinct from—but complementary to—the general balance sheet cost channel emphasized by Du et al. (2023).

Like the levels regression in the previous subsection, this OLS specification is not causal—common shocks could drive both collateral sinks and arbitrage spreads. We view these results as establishing a robust association and rely on the instrumental-variable estimates in Section 5.1 for causal identification.

5 Identifying the Effects of Aggregate Collateral Demand Shocks

The challenge in studying the effect of aggregate collateral demand on the convenience yield is that much of the 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 their effects on convenience yields: an IV based on idiosyncratic margin choices by the CME and a granular IV (GIV) using trading desk flows.

5.1 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 are 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 movements, 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 from 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 rather 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 asymmetric response in margin requirements: a one-unit increase in the VIX is associated with a 3.7pp increase in margin requirements, while a one-unit decrease in the VIX decreases margin requirements by 2.7pp. Margin requirements increase 33 percent faster than they decrease for identically-sized moves in the VIX. See Table IA.5.

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&P 500, 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 taking an average of the z -scored residuals, Z_t .¹⁵

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 $\widehat{\Delta S}_t$

$$\Delta PC1_t = \alpha_2 + \beta \widehat{\Delta S}_t + \kappa' X_t + \varepsilon_t.$$

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. The first two columns instrument for sunk margin relative to Treasuries outstanding, the middle two for collateral sinks excluding LCR-related components, and the last two for the full collateral sink measure; odd-numbered columns include no controls and even-numbered columns include the full control set. The second-stage coefficient is positive and statistically significant across all specifications, ranging from 1.9 to 3.6 bps. Using the estimates in column (6), a one-standard deviation increase in the instrumented change in the collateral sink ratio (8 bps) raises $\Delta PC1$ by 2.4 basis points, roughly 60 percent of its standard deviation. The coefficient is stable as the endogenous variable broadens from sunk margin to the full sink measure, indicating that the margin-driven variation alone is sufficient to identify the effect.

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. The first-stage F -statistics are strongest in the first four columns, where the endogenous variable is sunk margin and non-LCR-related components, and decline as the endogenous variable broadens to include LCR-related holdings that the margin instrument does not directly move. 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

¹⁵Alternative aggregations, like using the first principal component across the residuals, yield similar results.

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 Z_t on different measures of posted margin and finds a strong direct relationship (Table IA.6).

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 sets margin requirements only for its contracts, 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 margins with reference to information implicitly provided in CME's margin requirements, so innovations to CME's margin requirements are 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 their 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 one-month lead and find no effects in either stage (Table IA.7).

Our IV estimates of the effect of collateral sinks on the convenience yield are approximately an order of magnitude larger than the corresponding OLS estimates. 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 dominant 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, do 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 minimize the noisy error-driven variation in the collateral sink proxy, thereby correcting for the attenuation bias inherent in the OLS estimate.

5.2 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 a different asset class or product—and each desk’s 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 flows—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 provide the level of notional derivative exposures across the full set of banks, however there are some days when data are missing for one or more desks, so the panel is not strictly balanced. We require that each 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 $(1 + \text{the log of})$ each of these variables, and we winsorize the daily

changes at the 5th and 95th percentiles.

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

$$\begin{aligned}\Delta \ln(\text{VaR}_{i,t}) &= \beta_{\text{sec},A}^a \Delta \ln(\text{SecAsset}_{i,t}) + \beta_{\text{sec},L}^a \Delta \ln(\text{SecLiab}_{i,t}) \\ &+ \beta_{\text{der},A}^a \Delta \ln(\text{DerivAsset}_{i,t}) + \beta_{\text{der},L}^a \Delta \ln(\text{DerivLiab}_{i,t}) \\ &+ \mu_i + \lambda_{b(i),t} + \eta_{a,t} + \varepsilon_{i,t}.\end{aligned}\tag{14}$$

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

$$\Delta \ln(\text{VaR}_{i,t}^{\text{deriv}}) = \hat{\beta}_{\text{der},A}^a \Delta \ln(\text{DerivAsset}_{i,t}) + \hat{\beta}_{\text{der},L}^a \Delta \ln(\text{DerivLiab}_{i,t}).\tag{15}$$

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 data by desk panel with

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

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:

$$s_{i,t-1} = \left| \hat{\beta}_{\text{der},A}^{a(i)} \right| \text{DerivAssetNotional}_{i,t-1} + \left| \hat{\beta}_{\text{der},L}^{a(i)} \right| \text{DerivLiabNotional}_{i,t-1}, \quad (16)$$

$$w_{i,t} = \frac{s_{i,t-1}}{\sum_j s_{j,t-1}}. \quad (17)$$

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.6) and a plot of GIV (Figure IA.7) 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

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

As with the CME approach, we report results for three endogenous variables of increasing breadth: sunk margin relative to Treasuries outstanding, collateral sinks excluding LCR-related components, and the full collateral sink measure. To make interpretation easier, we standardize GIV and each endogenous variable to have mean 0 and unit standard deviation.

Panel A of Table 6 reports the second stage estimates. Because we cast $\Delta PC1$ to basis points (scaled to the 3-month OIS-Tbill spread) and standardize the endogenous regressors, a one standard deviation increase in the instrumented collateral sink ratio increases the convenience yield by 1.8 to 3.6 basis points, depending on the specification. The standard deviation of $\Delta PC1$ is 3.8 basis points, so the effect of a one-standard deviation increase in the instrumented ratio raises convenience yields by 0.47 to 0.95 standard deviations. The second-stage point estimates are stable across all six columns, indicating that the result is not sensitive to which component of the sink is instrumented. 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 two columns use sunk margin as the endogenous regressor, the middle two use collateral sinks excluding LCR-related components, and the last two use the full collateral sink measure. The first stage shows a strong positive relationship between the GIV and each endogenous variable: a 1 standard deviation increase in GIV corresponds to an increase of about 0.1 standard deviations in the endogenous regressor. As with the CME approach, first-stage F -statistics are strongest in the first two columns, where the endogenous variable is sunk margin, and decline as the endogenous variable broadens to include components the instrument does not directly move. The stability of the second-stage estimates across specifications indicates that the margin-driven variation alone is sufficient to identify the effect of collateral scarcity on convenience yields.

Panel C of Table 6 reports the OLS estimates. We again see that the instrumented coefficients are approximately an order of magnitude larger than the OLS estimates, and the OLS coefficients are positive but not significant. As discussed in the previous setting with CME margin requirements, we 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, the sample is limited to dates 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.8 shows that the placebo GIV does not generate similar results and Table IA.9 shows that the two instruments are not correlated, both indicating that it 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.10).

6 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.

6.1 The Collateral Swap Market

Collateral swaps replace 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. 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.

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.¹⁶ The figure shows the amount of collateral pledged and received by banks is 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. On average, banks do not use the swap market to source collateral for themselves, since the net pledged values are positive.¹⁷

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 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 price collateral swaps

¹⁶Figure IA.8 compares the size of the repo market, collateral swap market, client long and short positions, and posted margin.

¹⁷The Internet Appendix provides additional details on the collateral swap market.

using a no-arbitrage argument in Section IA.C and use it to estimate the capital savings they provide.

The left panel of Figure 8 estimates collateral swaps reduce capital required by more than \$70 billion in 2025, \$60 billion of which stems from swapping equity collateral.¹⁸ The bulk of the savings stems from 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 \$35 billion capital reduction. The Internet Appendix provides additional details.

The capital reduction represents equity that investors would otherwise need to source externally, and for levered investors, freed equity supports a multiple of its face value in levered exposures. The right panel of Figure 8 reports that large qualifying hedge funds have operated at average gross leverage ratios of roughly 10:1 in recent years. Applying this ratio implies that the \$70 billion capital reduction could support approximately \$700 billion in additional notional exposure capacity—a large figure given the hedge fund industry likely has on the order of \$3 trillion in total assets.¹⁹

6.2 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 5.1. 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

¹⁸Figure IA.9 provides the savings by asset class.

¹⁹Gross leverage ratios are from the OFR Hedge Fund Monitor, computed as the ratio of gross assets to net assets from SEC Form PF filings, for the 11th through 50th largest hedge funds by gross assets. See the Federal Reserve’s Flow of Funds series “Hedge Funds: Total Financial Assets.”

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.²⁰

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 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 9 reports the quintile-specific marginal effects s_q by quintile. The most constrained quintile shows a slope of 0.7 bps for a 1 σ shock, versus -0.1 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 4 times the average slope in the other quintiles.

We illustrate this amplification mechanism using the March 2020 episode. Figure 10 plots banks' collateral swap intermediation alongside margin demand and convenience yields over the January–June 2020 window. Panel B shows that margin demand surged beginning in late February as volatility spiked and CCPs raised requirements. Collateral pledged through swaps

²⁰Most 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.

did not keep pace—if anything, swap volumes fell as banks pulled back from intermediation. Panel A shows the resulting collapse in the swap-to-margin ratio, which fell from its January 2020 average of 191 percent to 141 percent at the trough, a decline of 50 percentage points. The pattern is consistent with banks being less able to intermediate collateral across clients. Panel C shows that convenience yields spiked beginning in early March, remained elevated through most of the month, and did not return to pre-crisis levels until after the Fed’s interventions took hold, eventually falling below their January starting point.

The disconnect between margin demand and swap supply is the mechanism in Prediction 2: when banks grow constrained, they stop buffering collateral scarcity through swaps, leaving investors to face margin calls without recourse to the collateral transformation that swaps provide. Rising margin requirements and falling swap intermediation together amplify the pass-through from collateral demand to convenience yields.

We quantify the amplification using a counterfactual exercise. For each trading day in March 2020, we assign the constraint ratio to its expanding-window quintile using data available as of that date, then apply the corresponding quintile-specific slope estimated analogously to those in Figure 9 to that day’s margin shock Z_t . The cumulative contribution over the month measures the total implied effect of margin shocks on convenience yields through the collateral sink channel, accounting for the time-varying state of dealer intermediation.

Figure 11 plots the cumulative contributions under two scenarios. Under the actual path of dealer constraints—which placed most of March 2020 in the most constrained quintile—the cumulative contribution of margin shocks reaches 15 bps by late March, accounting for roughly three-quarters of the observed 20 bp increase in the convenience yield PC1. Under the counterfactual in which banks maintain median intermediation capacity throughout the episode, the same margin shocks contribute less than 1 bp. The difference—approximately 14 bps—represents the amplification cost of dealer withdrawal from the collateral swap market.

These estimates are complementary to other accounts of the March 2020 Treasury market stress. He et al. (2022) show that selling pressure from large holders of Treasuries, absorbed by dealers subject to SLR constraints, drove Treasury yields above OIS rates. Our channel operates through a distinct margin-induced shock to collateral demand, amplified by the same underlying dealer constraints that limited balance sheet capacity for absorbing sales. The two mechanisms reinforce each other: selling pressure and margin calls both increase dealer balance sheet usage, and the resulting tightening of constraints simultaneously reduces dealers’

willingness to absorb Treasury supply and to intermediate collateral swaps. That the collateral sink channel alone can account for a substantial share of the observed convenience yield spike suggests that margin-driven collateral demand is quantitatively important alongside the selling-pressure channel.

7 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 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.
- Nina Boyarchenko, Pooja Gupta, Nick Steele, Jacqueline Yen, et al. Negative swap spreads. *Federal Reserve Bank of New York Economic Policy Review*, 24(2):1–14, 2018.
- 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.
- Darrell Duffie, Martin Scheicher, and Guillaume Vuillemeys. Central clearing and collateral demand. *Journal of Financial Economics*, 116(2):237–256, 2015.
- 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.
- Whitney K Newey and Kenneth D West. Automatic lag selection in covariance matrix estimation. *The review of economic studies*, 61(4):631–653, 1994.

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.

8 Figures




Figure 1: CME Margin Requirements. Blue 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. Red line plots the implied volatility for the same contracts.




Figure 2: Margin Posted. Figure plots total margin posted, cash posted as margin, and Treasuries posted as margin. Plot uses week-end values.

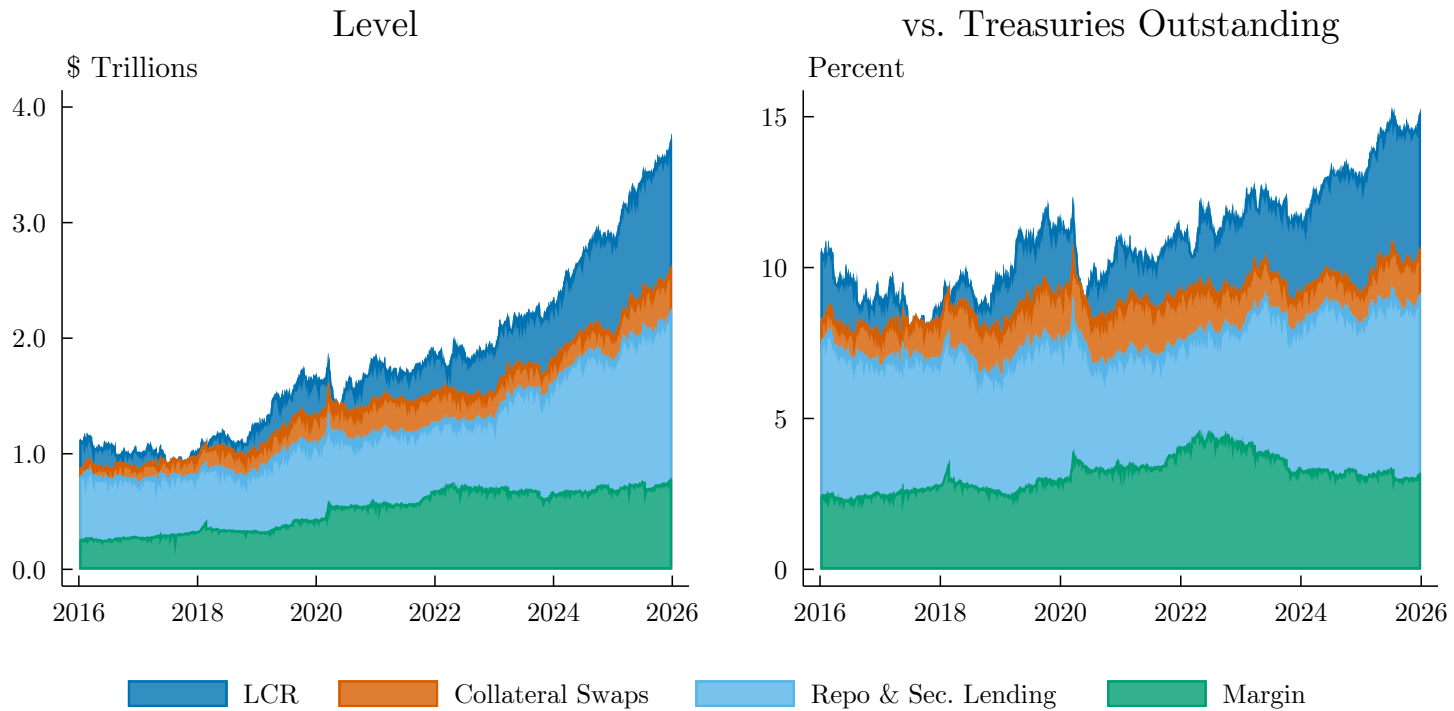


Figure 3: Collateral Sinks. Figure plots the total amount of sunk collateral using the methodology described in Section 4, aggregating across all collateral types and all banks in our sample. Left panel shows the level and right panel compares to the market value of publicly held Treasuries outstanding.




Figure 4: Marginal Explanatory Power of Collateral Sinks. Figure plots the percent increase in R^2 from adding each regressor to a specification of daily changes in convenience yields (PC1) that includes all other regressors. Panel A includes only proxies for sovereign credit risk (Δ US CDS) and Treasury market liquidity (Δ 5yr Bid-Ask Spread) alongside collateral sinks. Panel B adds the VIX, BAA-AAA spread, breakeven inflation, the effective federal funds rate, and the remaining implied volatility and equity return controls.

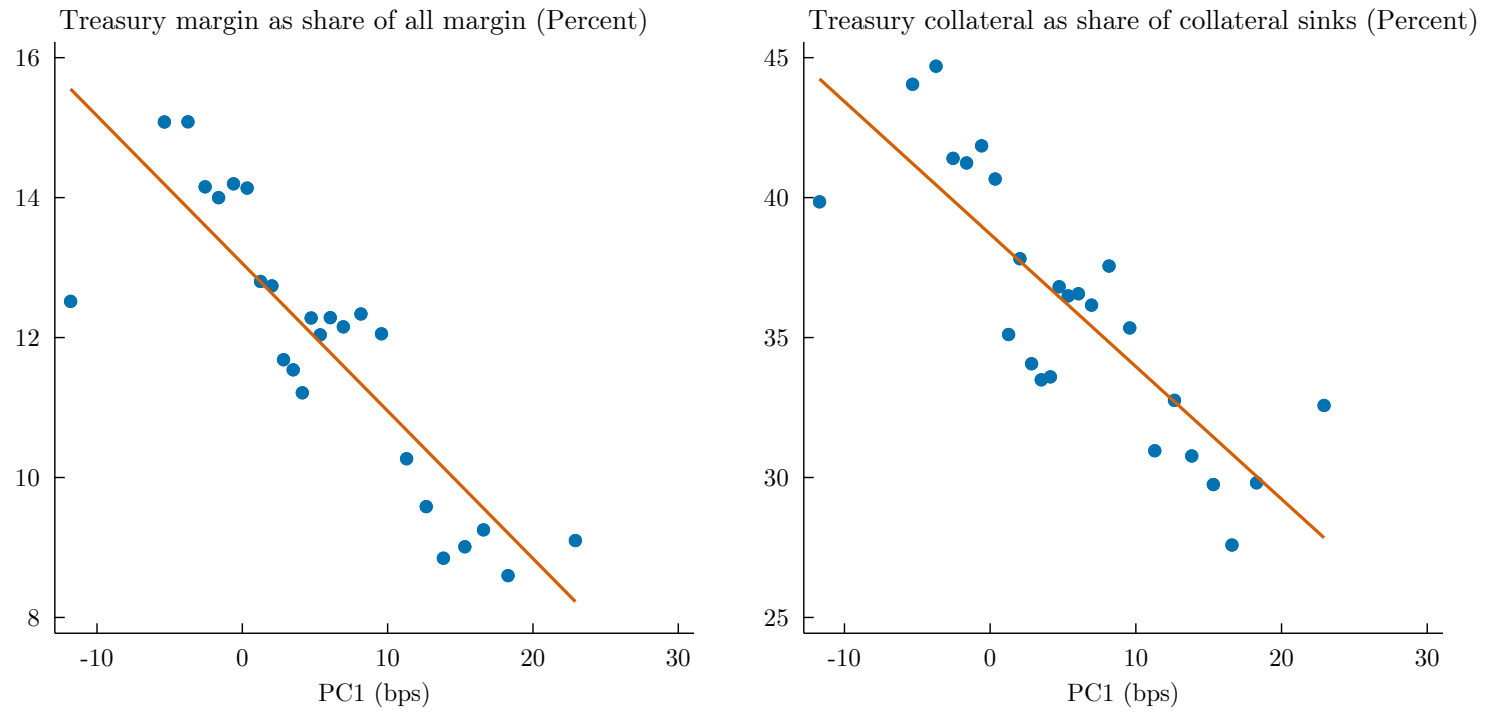


Figure 5: Treasury Share of Margin and Collateral Sinks vs. PC1. Left panel provides binscatter plot of Treasury margin as a share of all margin against PC1; right panel plots the Treasury collateral share of collateral sinks.

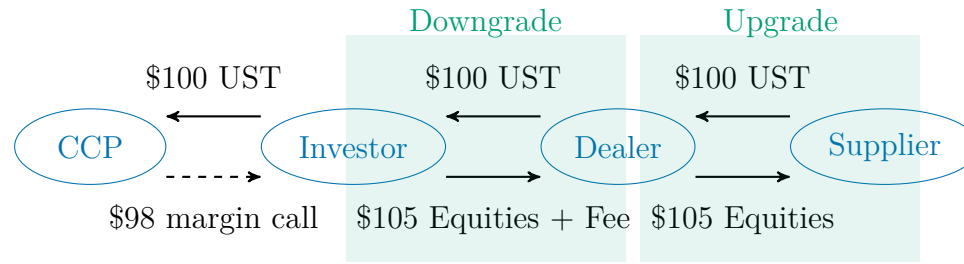


Figure 6: Collateral Swap Example. Figure provides an example collateral swap in which an investor transforms equities into Treasuries (UST).




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

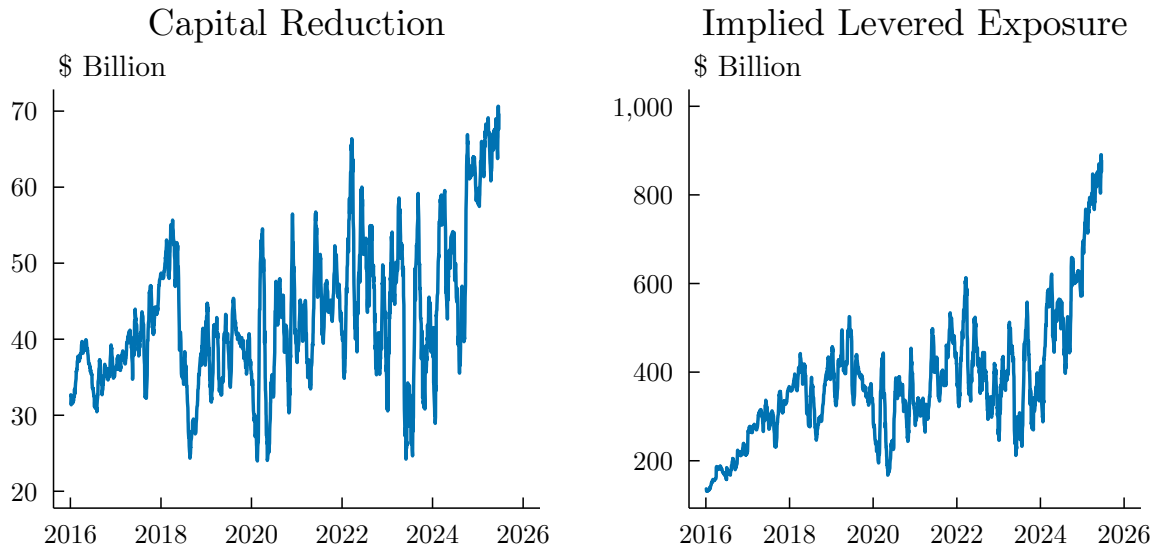


Figure 8: Estimated Margin Capital Reduction and Implied Notional Exposure from Collateral Swaps. The left panel calculates the capital reduction using CME haircuts, time-varying dealer reverse repo haircuts, and the volume of risky collateral posted to banks through collateral swaps per (19). The right panel multiplies the capital reduction by the equal-weighted average gross leverage ratio of the 11th through 50th largest qualifying hedge funds, as reported by the OFR Hedge Fund Monitor from SEC Form PF filings. See section IA.C for methodology details.

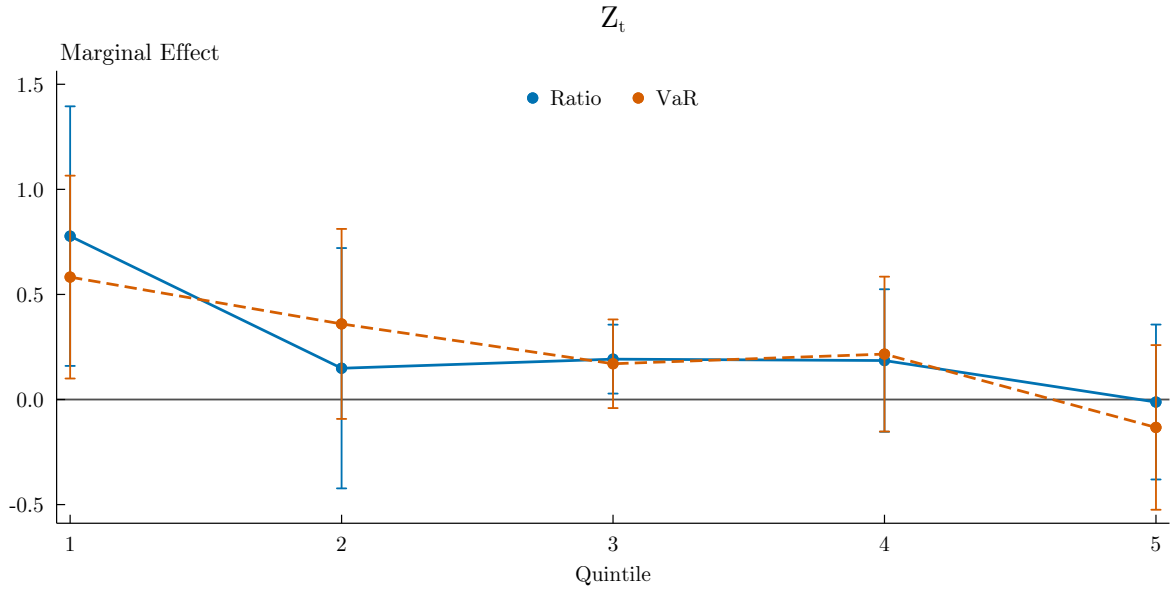


Figure 9: State-dependent Pass-through. Figure plots the total marginal effect of Z_t on $\Delta PC1_t$ by bank constraint state, which are calculated either using the ratio of total collateral pledged by banks in collateral swaps to total margin posted by or to banks or alternatively by the total value-at-risk across banks' trading desks.

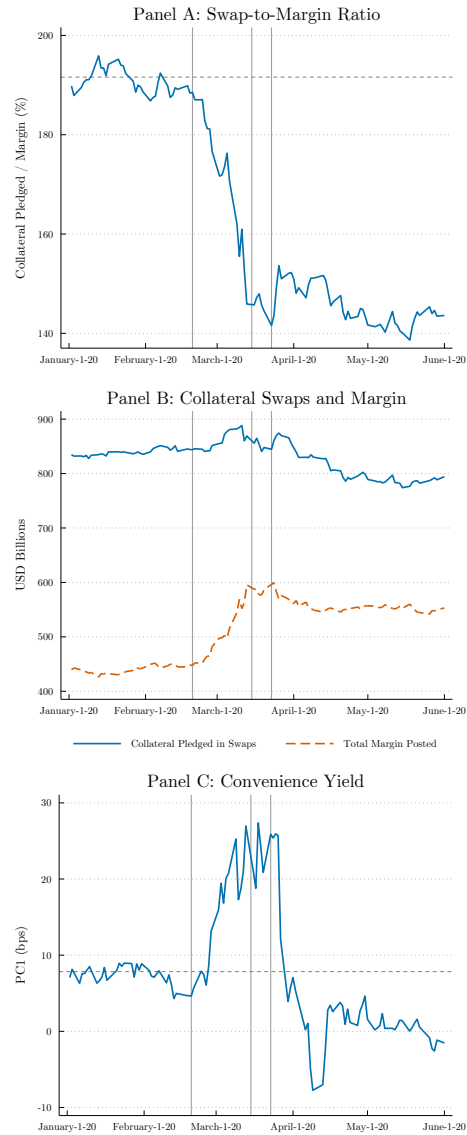


Figure 10: Collateral Swap Intermediation, Margin Demand, and Convenience Yields During COVID-19. Panel A plots the ratio of collateral pledged by banks in collateral swaps to total margin posted. Panel B plots the levels of collateral pledged in swaps and total margin posted, both in billions of dollars. Panel C plots the first principal component of convenience yield measures. The sample window is January–June 2020. Dashed horizontal lines indicate January 2020 averages. Vertical lines mark February 20 (first stock market crash), March 15 (emergency Fed rate cut), and March 23 (announcement of Fed emergency facilities and asset purchases).

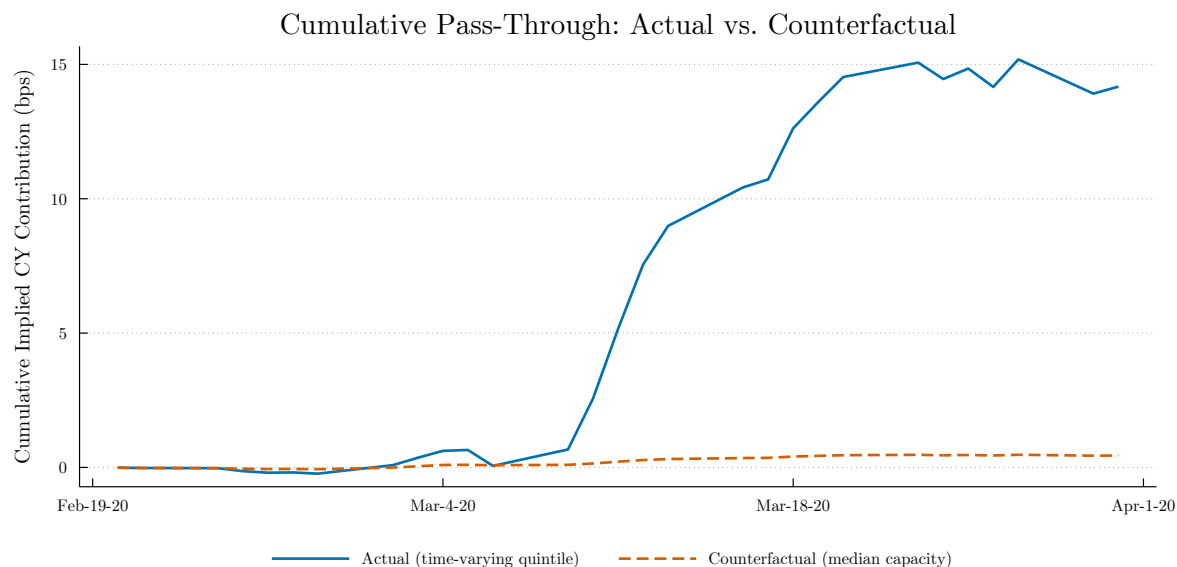


Figure 11: Cumulative Pass-Through of Margin Shocks to Convenience Yields: Actual vs. Counterfactual. Figure plots the cumulative implied contribution of daily margin shocks Z_t to convenience yields (PC1) over March 2020 under two scenarios. The solid line applies each day's quintile-specific slope estimated analogously to the state-dependent regression in Table IA.19, where quintiles are assigned using expanding-window sorts on the collateral swap constraint ratio using only data available as of each date. The dashed line applies the median-quintile (Q3) slope to every day, representing the counterfactual in which banks maintain median collateral swap intermediation capacity throughout the episode. The gap between the two lines measures the amplification of margin shocks attributable to dealer withdrawal from the collateral swap market.

9 Tables

<i>Panel A: Margin Posted (\$ billions)</i>		Obs.	Mean	SD	Min	Max
Cash	2,463	2,463	376	71	247	571
USTs	2,463	2,463	70	31	21	146
HQLA	2,463	2,463	470	163	231	726
Non-HQLA	2,463	2,463	102	44	26	178
Total	2,463	2,463	572	131	359	822

<i>Panel B: ΔMargin Posted (\$ billions)</i>		Obs.	Mean	SD	Min	Max
Cash	2,430	2,430	0.1	5.1	-43.2	50.8
USTs	2,430	2,430	0.0	1.5	-15.8	9.1
HQLA	2,430	2,430	0.2	5.5	-48.5	62.7
Non-HQLA	2,430	2,430	0.0	3.1	-109.3	21.9
Total	2,430	2,430	0.2	5.9	-46.6	59.5

<i>Panel C: Margin Composition (Percent)</i>		Obs.	Mean	SD	Min	Max
Cash	2,463	2,463	67	8	49	80
USTs	2,463	2,463	12	3	6	20
HQLA	2,463	2,463	80	11	64	96
Non-HQLA	2,463	2,463	20	11	4	36

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 2025. 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. All values are aggregated across all banks in the sample before calculating moments.

<i>Repo and Securities Lending (\$ billions)</i>	Obs	Mean	SD	Min	Max
Treasuries	2,474	232	136	77	605
Other HQLA	2,474	288	92	191	555
Other Non-HQLA	2,474	213	51	136	375
<i>Total</i>	2,474	732	268	468	1,482

<i>Collateral Swaps (\$ billions)</i>	Obs	Mean	SD	Min	Max
Treasuries	2,474	29	17	5	106
Other HQLA	2,474	124	36	44	192
Other Non-HQLA	2,474	62	22	20	145
<i>Total</i>	2,474	215	65	72	375

<i>Margin (\$ billions)</i>	Obs	Mean	SD	Min	Max
Cash	2,474	221	59	150	407
Treasuries	2,474	104	44	34	193
Other HQLA	2,474	120	70	20	248
Other Non-HQLA	2,474	81	38	37	190
<i>Total</i>	2,474	526	175	251	788

<i>LCR (\$ billions)</i>	Obs	Mean	SD	Min	Max
Treasuries	2,474	351	280	0	1,141

<i>Aggregated (\$ billions)</i>	Obs	Mean	SD	Min	Max
Cash	2,474	221	59	150	407
Treasuries	2,474	715	457	159	1,993
Other HQLA	2,474	532	167	289	930
Other Non-HQLA	2,474	356	105	233	700
<i>Total</i>	2,474	1,824	722	919	3,773

<i>vs. Treasuries Outstanding (percent)</i>	Obs	Mean	SD	Min	Max
Repo & Sec. Lending	2,460	5.0	0.7	3.4	6.1
Collateral Swaps	2,460	1.3	0.3	0.7	1.9
Margin	2,460	3.2	0.6	2.3	4.6
LCR	2,460	2.0	1.1	0.0	4.6
<i>Total (S_t)</i>	2,460	11.5	1.7	8.0	15.3

Table 2: Collateral Sink Summary Statistics. Table provides summary statistics for the collateral sink data described in Section 4. Aggregated reflects the sum of sunk collateral across repo/securities lending, collateral swaps, margin, and LCR-related holdings. 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			ΔPC1 _t		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Collateral Sink Measures</i>						
Collateral Sinks/Treasuries Outstanding _t	4.099*** (3.04)	4.209*** (3.39)	4.491*** (2.74)			
ΔCollateral Sinks/Treasuries Outstanding _t				2.646** (2.28)	2.799** (2.46)	2.746** (2.34)
Equity Implied Vol. _t		0.016 (0.23)	0.049 (0.48)			
Treasury Implied Vol. _t		-0.634 (-1.07)	-0.743 (-1.52)			
FX Implied Vol. _t		0.589 (0.94)	0.383 (0.77)			
Oil ETF Implied Vol. _t			0.009 (1.09)			
Effective Fed Funds Rate _t			24.785*** (4.69)			
BAA-AAA Spread _t			-0.701 (-0.21)			
US CDS _t			18.473 (1.56)			
5yr Treasury Bid-Ask Spread _t			1016.887 (1.59)			
5y5y Breakeven Inflation _t			-2.812 (-0.69)			
ΔEquity Implied Vol. _t				0.008 (0.23)	-0.105 (-1.26)	
ΔTreasury Implied Vol. _t				0.242 (0.67)	-0.100 (-0.23)	
ΔFX Implied Vol. _t					-0.157 (-0.49)	
ΔOil ETF Implied Vol. _t					0.004 (0.57)	
S&P 500 Return _t					-22.952 (-1.63)	
ΔEffective Fed Funds Rate _t					19.423*** (4.57)	
ΔBAA-AAA Spread _t					-3.712 (-0.97)	
ΔUS CDS _t					22.695** (2.28)	
Δ5yr Treasury Bid-Ask Spread _t					1121.093** (2.47)	
Δ5y5y Breakeven Inflation _t					-12.671*** (-3.96)	
<i>N</i>	2,441	2,387	2,381	2,407	2,349	2,334
<i>R</i> ²	0.01	0.02	0.21	0.00	0.00	0.09
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 heteroskedastic and autocorrelation consistent standard errors using the Newey and West (1994) automatic lag selection procedure where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

<i>Panel A: Treasury Swap Spreads</i>							
	2Y	5Y	10Y	30Y	PC1	PC1	PC1
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Δ Collateral Sinks/Treasuries Outstanding _{<i>t</i>}	1.007*** (3.34)	1.005*** (3.11)	0.915*** (3.63)	0.783*** (2.65)	0.076*** (4.31)	0.071*** (4.13)	0.060*** (3.89)
Δ PC1 FX Basis _{<i>t</i>}						0.030** (1.99)	0.028** (2.41)
<i>N</i>	2,424	2,424	2,424	2,424	2,424	2,368	2,299
<i>R</i> ²	0.00	0.01	0.01	0.00	0.01	0.02	0.06
Controls	No	No	No	No	No	No	Yes
<i>Panel B: Non-Treasury Arbitrage Spreads</i>							
	FX Basis			FX Basis PC1		CDS-Bond Basis	
	EUR	JPY	GBP	PC1	PC1	IG	HY
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Δ Collateral Sinks/Treasuries Outstanding _{<i>t</i>}	0.806 (0.60)	3.213 (0.92)	0.904 (0.58)	0.117 (0.73)	0.094 (0.80)	-1.309 (-0.88)	-3.618 (-0.90)
<i>N</i>	2,408	2,418	2,392	2,368	2,299	2,424	2,422
<i>R</i> ²	0.00	0.00	0.00	0.00	0.14	0.00	0.00
Controls	No	No	No	No	Yes	No	No

Table 4: Collateral Sinks and Arbitrage Spreads. Panel A regresses daily changes in LIBOR swap spreads at the 2-, 5-, 10-, and 30-year tenors, as well as their first principal component (PC1), on changes in collateral sinks relative to Treasuries outstanding. Columns (6) and (7) add the first principal component of the FX basis spreads from Panel B as a control; column (7) further includes the full set of controls. Panel B regresses daily changes in non-Treasury arbitrage spreads on the same measure: 3-month covered interest parity deviations for EUR, JPY, and GBP (SOFR-based), their first principal component, and investment-grade and high-yield CDS-bond bases. Column (5) adds the full set of controls to the FX basis PC1 specification. *t*-statistics shown using heteroskedastic 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)
$\Delta \widehat{\text{Sunk Margin}}/\text{UST Out.}_t$	3.57** (2.44)	3.34** (2.53)				
$\Delta \widehat{\text{Sinks (Excl. LCR)}}/\text{UST Out.}_t$			2.20*** (2.60)	1.92*** (2.79)		
$\Delta \widehat{S}_t$					2.52** (2.41)	2.43** (2.43)
N	2,094	2,094	2,094	2,094	2,094	2,094
Controls	No	Yes	No	Yes	No	Yes

Panel B: First Stage						
	$\Delta \text{Margin}/\text{UST Out.}$		$\Delta \text{Sinks (Excl. LCR)}/\text{UST Out.}$		$\Delta \text{Sinks}/\text{UST Out.}$	
	(1)	(2)	(3)	(4)	(5)	(6)
Z_t	0.09*** (4.56)	0.08*** (4.35)	0.14*** (6.17)	0.15*** (4.93)	0.12*** (5.57)	0.12*** (4.00)
N	2,094	2,094	2,094	2,094	2,094	2,094
$F - stat$	20.79	18.92	38.06	24.32	31.05	15.97
Controls	No	Yes	No	Yes	No	Yes

Panel C: OLS						
	Dependent Var.: $\Delta PC1_t$					
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \text{Sunk Margin}/\text{UST Out.}_t$	0.23*** (2.78)	0.22*** (2.75)				
$\Delta \text{Sinks (Excl. LCR)}/\text{UST Out.}_t$			0.30** (2.36)	0.30** (2.33)		
ΔS_t					0.25** (2.40)	0.25** (2.42)
N	2,166	2,094	2,166	2,094	2,166	2,094
R^2	0.00	0.09	0.01	0.09	0.00	0.09
Controls	No	Yes	No	Yes	No	Yes

Table 5: CME Residual Regression. Table presents the first-stage, second-stage, and OLS estimates described in Section 5.1. The first two columns instrument for sunk margin relative to Treasuries outstanding, the middle two for collateral sinks excluding LCR-related components, and the last two for the full collateral sink measure. Controls include 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. Margin data runs from 2016 to 2024. Constant omitted. t -statistics shown using heteroskedastic 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)
$\Delta \widehat{\text{Sunk Margin}} / \text{UST Out.}_t$	3.47*** (2.95)	2.68*** (2.74)				
$\Delta \widehat{\text{Sinks (Excl. LCR)}} / \text{UST Out.}_t$			2.28*** (2.83)	1.82*** (2.63)		
$\Delta \widehat{S}_t$					3.56** (2.24)	2.72** (2.26)
<i>N</i>	1,210	1,155	1,210	1,155	1,210	1,155
Controls	No	Yes	No	Yes	No	Yes

Panel B: First Stage						
	$\Delta \text{Margin} / \text{UST Out.}$		$\Delta \text{Sinks (Excl. LCR)} / \text{UST Out.}$		$\Delta \text{Sinks} / \text{UST Out.}$	
	(1)	(2)	(3)	(4)	(5)	(6)
GIV_t	0.08*** (4.10)	0.10*** (4.95)	0.13*** (3.78)	0.14*** (4.40)	0.08*** (2.61)	0.10*** (3.11)
<i>N</i>	1,210	1,155	1,210	1,155	1,210	1,155
<i>F - stat</i>	16.84	24.53	14.25	19.34	6.79	9.65
Controls	No	Yes	No	Yes	No	Yes

Panel C: OLS						
	Dependent Var.: $\Delta PC1_t$					
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \widehat{\text{Sunk Margin}} / \text{UST Out.}_t$	0.02 (0.31)	0.01 (0.09)				
$\Delta \widehat{\text{Sinks (Excl. LCR)}} / \text{UST Out.}_t$			0.22 (1.62)	0.20 (1.35)		
ΔS_t					0.17 (1.61)	0.17 (1.46)
<i>N</i>	1,210	1,155	1,210	1,155	1,210	1,155
R^2	0.00	0.05	0.01	0.05	0.00	0.05
Controls	No	Yes	No	Yes	No	Yes

Table 6: Granular IV. Table presents the first-stage, second-stage, and OLS estimates described in Section 5.2. The first two columns instrument for sunk margin relative to Treasuries outstanding, the middle two for collateral sinks excluding LCR-related components, and the last two for the full collateral sink measure. Controls include 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 heteroskedastic and autocorrelation consistent standard errors using the Newey and West (1994) automatic lag selection procedure where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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. Sections IA.B provides additional institutional details and results related to margin. Section IA.C provides additional institutional details and results related to collateral swaps, including additional summary statistics and no arbitrage collateral swap pricing derivation. Section IA.D presents additional tables and figures.

IA.A Data

IA.A.1 FR2052a 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 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 quarterly Y-9C reports.

FR VV-1 data provides notional exposures for trading desks until 2020. Following a change in reporting requirements, banks switched to reporting 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 years). The variables correspond to BHCK582 to BHCK623.

The two datasets provide related but different measures of notional exposures. The VV-1 data spans derivatives booked on covered trading desks, while the Y-9C 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 Y-9C, 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 Y-9C 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 Internet Appendix Figure IA.10.

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 22 bps of their notional exposure on average, with cash accounting for 15 bps and Treasuries 3 bps. 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 perfectly. We decompose the log variance of the ratio, such that $\text{Var}(\ln(\text{Margin Ratio})) = \text{Var}(\ln(\text{Margin Posted})) + \text{Var}(\ln(\text{Notional exposures})) - 2 \text{Cov}(\ln(\text{Margin Posted}), \ln(\text{Notional exposures}))$. $\text{Var}(\ln(\text{Margin Posted}))$ explains 172% and $\text{Var}(\ln(\text{Notional exposures}))$ 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 use data from TreasuryDirect to net buybacks. 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 to 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 Treasury bid-ask spread. Based on data provided by the interdealer broker community.
- 5-year/5-year breakeven inflation rate. Calculated using hot-run yields from Morgan Markets.

IA.B Margin Details

An example of the lifecycle of a derivative trade makes the margin mechanics clear. Suppose an investor goes long \$1 million of notional exposure to a 30-year U.S. Treasury future settling in September 2025 (10 UBU5 contracts each with \$100,000 face value). The contract trades at 117'25 (the digits after the tick are units of 32nds), or 117.78125 per contract with face value of \$100,000. 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 by 12 ticks to 117'13, 117.40625. The trader loses 12/32 of a point for each contract for a total loss of \$3,750. The equity in their account is now $51,150 - 3,750 = 47,400$. Assuming that the initial margin requirement is equal to the maintenance margin requirement, the CCP issues a margin call and the trader needs to post an additional \$3,750 of collateral.

The six largest U.S. banks alone posted nearly \$800 billion in recent years. 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 centrally cleared margin 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.14). 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 to prefund. While this also requires large balances—CCPs required about \$170 billion in 2024 globally—the number tends to be slower moving.

Table IA.11 shows 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). Figure IA.11 plots the time series of notional exposures.

IA.C Collateral Swap Market Details

Additional Institutional Details 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; Figure 6 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 in the 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.

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.

Collateral Swap Summary Statistics Table IA.12 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.

The top half of Figure IA.12 shows that banks largely pledge collateral through the swap to banks (\$351 billion) and investment companies (\$113 billion). The bottom half of Figure IA.12 shows that the swaps transfer high quality collateral from insurance companies and public sector firms to intermediaries and non-regulated funds.

Haircut Arbitrage Derivation 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_{ccp}^ℓ , 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_d^ℓ in exchange for Treasuries; post the Treasuries to the CCP at haircut h_{ccp}^t , 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 received²¹

For simplicity, we assume the dealer sizes the Treasuries delivered in the collateral swap to be equal across the two collateral types, meaning there is no cash exchanged for haircut

²¹We 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 10 bps on the post-haircut collateral value.

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:

$$\underbrace{\left(M - (1 - h_{ccp}^{\ell})V_{\ell}\right) (r_c - r_m)}_{\text{option (1)}} = \underbrace{\left(M - (1 - h_{ccp}^t)(1 - h_d^{\ell})V_{\ell}\right) (r_c - r_m)}_{\text{option (2)}} + f(1 - h_d^{\ell})V_{\ell}.$$

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

$$\frac{\text{Savings}}{V_{\ell}} = (r_c - r_m) \left[(1 - h_{ccp}^t)(1 - h_d^{\ell}) - (1 - h_{ccp}^{\ell}) \right] - f(1 - h_d^{\ell}). \quad (18)$$

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 makes collateral swap more profitable; we plot the tradeoff between these two forces in Figure IA.13.

A simple numerical example makes clear. Suppose the lower quality collateral are 10-year corporate bonds, which are swapped for 10-year Treasuries, and $h_{ccp}^t = 4.5\%$, $h_{ccp}^{\ell} = 25\%$, $h_d^t = 1\%$, $h_d^{\ell} = 10\%$, 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.²² 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 IA.13, 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 capital reduction. We define capital reduction for the investor as the difference in the amount of cash the investor has to get externally as

$$\left[(1 - h_{ccp}^t)(1 - h_d^{\ell}) - f(1 - h_d^{\ell}) - (1 - h_{ccp}^{\ell}) \right] V_{\ell}. \quad (19)$$

The savings are harder to estimate since it requires assumptions about investors' cost of

²²See the OFR “Hedge Fund Monitor” at <https://www.financialresearch.gov/hedge-fund-monitor/categories/risk-management/chart-69/>.

capital, but using the assumptions in the above example yields an average savings rate (as measured in (18)) of 1 percent, amounting to about \$5.5 billion in 2024. The bulk of the savings stems from the difference in banks' equity repo haircuts of 10 percent compared to the CME's 30 percent. The Internet Appendix provides additional details.

To gauge aggregate implications, note that the capital reduction represents equity that investors would otherwise need to source externally. For levered investors, freed equity supports a multiple of its face value in notional exposures. The right panel of Figure 8 reports gross leverage ratios from the OFR Hedge Fund Monitor, computed as the ratio of gross assets to net assets from SEC Form PF filings. Large qualifying hedge funds (the 11th through 50th largest by gross assets) have operated at average gross leverage ratios of roughly 10:1 in recent years. Applying this ratio implies that the \$70 billion capital reduction could support approximately \$700 billion in additional notional exposure capacity. Even under the conservative assumption that only half of collateral swaps serve margin-related motives, the implied capacity remains large at \$350 billion. Moreover, the effect is large given the hedge fund industry likely has on the order of \$3 trillion in total assets.²³

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 $1 - \text{SFT loan value} / \text{SFT collateral value}$. We exclude reverse repos and securities borrowing transactions that are likely driven by short-sale motives by requiring the SFT have 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.

²³See the Federal Reserve's Flow of Funds series "Hedge Funds: Total Financial Assets."

IA.D Appendix Figures and Tables

IA.9




Figure IA.1: Secured Financing Transactions. Figure provides the pre-haircut value of all collateral pledged by banks in our sample in repos, securities lending transactions, and collateral swaps. Plot uses week-end values.

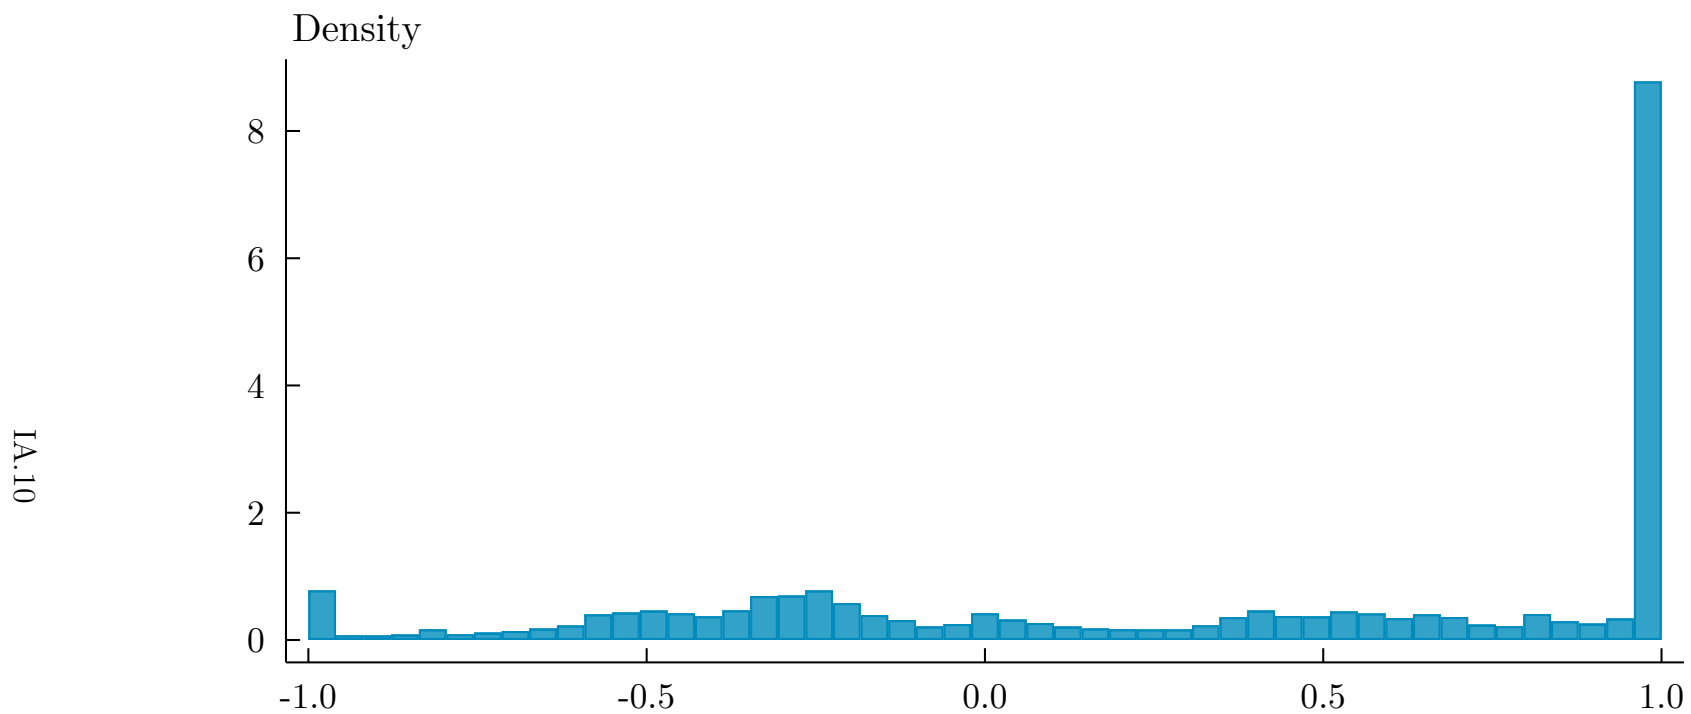


Figure IA.2: Distribution of Collateral Flow Imbalances. Figure is a histogram of the collateral flow imbalance variables described in section 4. Data aggregates across all banks in our sample, and the imbalance measure is at the date by counterparty by settlement type level.




Figure IA.3: Median Collateral Flow Imbalances by Counterparty and Settlement. Figure provides the median collateral imbalance by counterparty and settlement type; imbalance variable described in section 4. 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.4: Alternative Collateral Sink Measures. Figure plots four specifications of the aggregate collateral sink. “Main” is the benchmark measure, which includes repo and securities lending collateral pledged to sink counterparties, collateral swap outflows, non-rehypothecatable and centrally cleared margin, and the LCR Treasury residual lower bound. “Excl. LCR” removes all LCR-related components. “Full LCR” replaces the Treasury residual lower bound with total estimated HQLA required (the sum of estimated net cash outflows). “Mgmt Buffer” adjusts net cash outflows by the bank-specific 5th percentile of observed LCR ratios before computing the Treasury residual. Left panel shows levels in trillions of dollars. Right panel shows each measure as a percent of the market value of Treasuries outstanding, where both numerator and denominator are five-day moving averages.




Figure IA.5: 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.6: Derivative Beta Estimates by Asset Class. Figure plots $\beta_{der,A}^a$ and $\beta_{der,L}^a$ in (14) by asset class.




Figure IA.7: 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.8: Magnitude Comparison. Figure plots the average size of several markets in 2024.

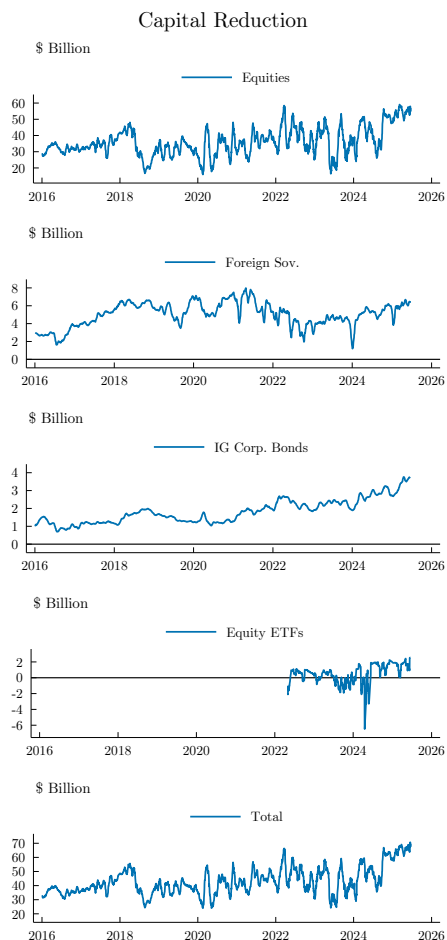


Figure IA.9: Estimated Margin Capital Reduction provided by Collateral Swaps. Figure uses CME haircuts, time-varying dealer reverse repo haircuts, and the volume of risky collateral posted to banks through collateral swaps and calculates the capital reduction using (19). See section IA.C for methodology details.




Figure IA.10: Estimated Notional Post 2020. Plots the actual VV-1 notional, estimated VV-1 notional, and actual Y-9C notional. See Internet Appendix for discussion of the estimation.




Figure IA.11: Notional and the *Effective Margin Ratio*. Top panel plots total notional exposure, along with exposures 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 IA.12: 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.

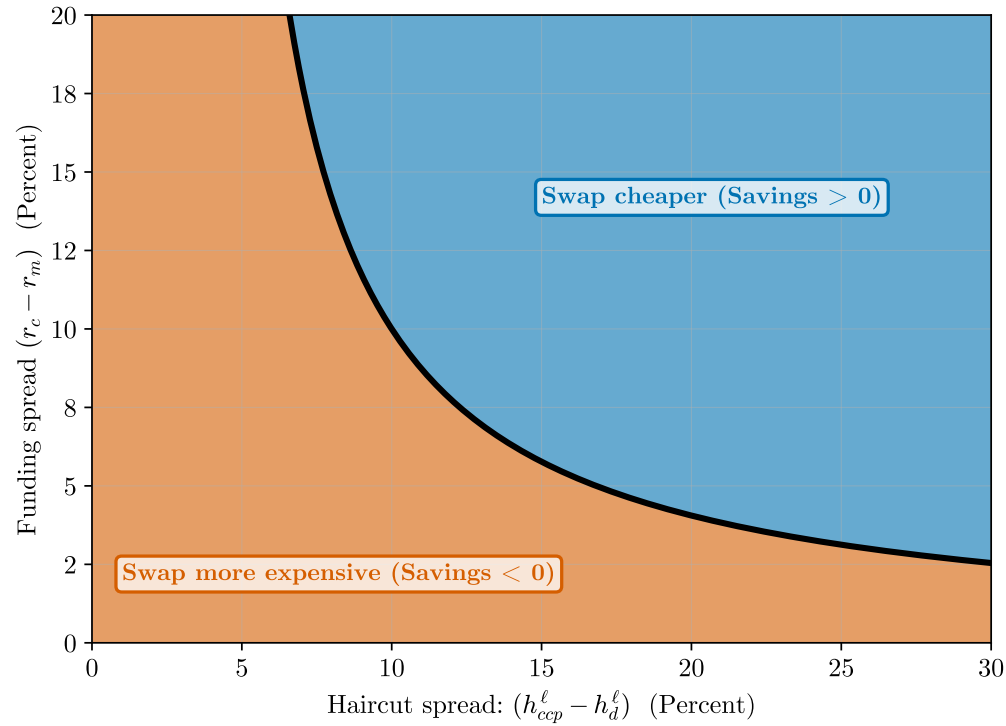


Figure IA.13: Regions of Collateral Swap Viability. Figure plots the regions where collateral swaps are economically profitable for an investor needing to post margin, as a function of the spread between the investor's cost of capital and the rate paid by the CCP on cash ($r_c - r_m$) against the difference in low-quality collateral haircuts. Figure assumes $h_{ccp}^t = 4.5\%$, $h_d^t = 1\%$, $h_d^l = 10\%$, and $f = 75$ bps.




Figure IA.14: 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.15: 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.16: Net collateral pledged through collateral swaps. Figure plots the net market value of collateral pledged by banks using collateral swaps.




Figure IA.17: Net Treasuries pledged through collateral swaps. Figure plots the net market value of Treasuries pledged by banks using collateral swaps.




Figure IA.18: Upgrades and Downgrades through Collateral Swaps. Figure plots the total notional of collateral swaps that are upgrades, downgrades, or flat for the bank.

Savings




Figure IA.19: Estimated Savings provided by Collateral Swaps. Figure plots the estimated savings using (18) 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.C for methodology details.

Savings As Percent of Risky Collateral




Figure IA.20: Estimated Savings provided by Collateral Swaps. Figure plots the estimated savings using (18). 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.C for methodology details.




Figure IA.21: Ratio of Collateral Swaps to Margin. Figure plots the ratio of collateral pledged by banks to margin posted to or by banks.




Figure IA.22: VaR. Figure plots the total VaR across the banks in our sample.

	Excl. LCR	Full LCR	Mgmt Buffer
	(1)	(2)	(3)
$\Delta\text{Main Measure}_t$	0.720*** (35.76)	0.888*** (46.53)	1.159*** (21.27)
N	2,424	2,424	2,424
R^2	0.73	0.73	0.20

Table IA.1: Comparison of Alternative Collateral Sink Measures. Table regresses daily changes in three alternative collateral sink specifications on daily changes in the main measure, where all variables are scaled by the market value of Treasuries outstanding. “Excl. LCR” removes all LCR-related components from the sink. “Full LCR” replaces the Treasury residual lower bound with total estimated HQLA required (the sum of estimated net cash outflows). “Mgmt Buffer” adjusts net cash outflows by the bank-specific 5th percentile of observed LCR ratios to proxy the effective regulatory floor before computing the Treasury residual; the average management buffer is 14%, corresponding to an effective LCR of 114%. t -statistics shown using heteroskedastic and autocorrelation consistent standard errors using the Newey and West (1994) automatic lag selection procedure where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Parsimonious		Kitchen Sink	
	(1)	(2)	(3)	(4)
<i>Collateral Sink Measure</i>				
Δ Collateral Sinks/Treasuries Outstanding _{<i>t</i>}		2.528** (2.13)		2.746** (2.34)
<i>Controls</i>				
Δ US CDS _{<i>t</i>}	21.034* (1.90)	21.709** (2.22)	22.029** (2.28)	22.695** (2.28)
Δ 5yr Treasury Bid-Ask Spread _{<i>t</i>}	874.137* (1.69)	825.185* (1.92)	1177.313*** (2.60)	1121.093** (2.47)
Δ BAA-AAA Spread _{<i>t</i>}			-3.380 (-0.88)	-3.712 (-0.97)
Δ 5y5y Breakeven Inflation _{<i>t</i>}			-12.916*** (-3.95)	-12.671*** (-3.96)
Δ Effective Fed Funds Rate _{<i>t</i>}			19.485*** (4.52)	19.423*** (4.57)
Δ Equity Implied Vol. _{<i>t</i>}			-0.105 (-1.27)	-0.105 (-1.26)
Δ Treasury Implied Vol. _{<i>t</i>}			-0.108 (-0.25)	-0.100 (-0.23)
Δ FX Implied Vol. _{<i>t</i>}			-0.149 (-0.47)	-0.157 (-0.49)
Δ Oil ETF Implied Vol. _{<i>t</i>}			0.002 (0.36)	0.004 (0.57)
S&P 500 Return _{<i>t</i>}			-21.247 (-1.47)	-22.952 (-1.63)
<i>N</i>	2,401	2,401	2,334	2,334
<i>R</i> ²	0.01	0.01	0.09	0.09

Table IA.2: Convenience Yield Regressions: Incremental R^2 from Collateral Sinks. Table regresses daily changes in convenience yields (PC1) on changes in collateral sinks relative to Treasuries outstanding and controls. Columns within each group show the specification without and with the collateral sink measure. The parsimonious specification controls for sovereign credit risk and Treasury market liquidity; the kitchen sink specification further includes Treasury, FX, and oil implied volatility and equity returns. *t*-statistics shown using heteroskedastic 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)	(7)	(8)
<i>Collateral Sink Components</i>								
Δ Sunk Margin/Treasuries Outstanding _{<i>t</i>}	5.955*** (2.74)	5.849*** (2.70)						
Δ Sunk Repo Collateral/Treasuries Outstanding _{<i>t</i>}			2.514 (1.14)	2.644 (1.15)				
Δ Sunk Collateral Swap Collateral/Treasuries Outstanding _{<i>t</i>}					8.522 (1.60)	7.792* (1.73)		
Δ Sunk LCR Collateral/Treasuries Outstanding _{<i>t</i>}							0.393 (0.27)	0.529 (0.36)
<i>Controls</i>								
Δ Equity Implied Vol. _{<i>t</i>}		-0.104 (-1.25)		-0.103 (-1.23)		-0.110 (-1.33)		-0.105 (-1.27)
Δ Treasury Implied Vol. _{<i>t</i>}		-0.147 (-0.34)		-0.097 (-0.22)		-0.120 (-0.27)		-0.104 (-0.24)
Δ FX Implied Vol. _{<i>t</i>}		-0.167 (-0.53)		-0.171 (-0.53)		-0.116 (-0.37)		-0.147 (-0.47)
Δ Oil ETF Implied Vol. _{<i>t</i>}		0.002 (0.29)		0.003 (0.51)		0.002 (0.30)		0.002 (0.38)
S&P 500 Return _{<i>t</i>}		-23.051 (-1.60)		-21.267 (-1.48)		-22.036 (-1.54)		-21.355 (-1.48)
Δ Effective Fed Funds Rate _{<i>t</i>}		19.406*** (4.55)		19.486*** (4.55)		19.430*** (4.59)		19.484*** (4.52)
Δ BAA-AAA Spread _{<i>t</i>}		-3.757 (-0.96)		-3.431 (-0.89)		-3.580 (-0.93)		-3.386 (-0.88)
Δ US CDS _{<i>t</i>}		21.904** (2.21)		22.464** (2.29)		21.731** (2.22)		22.102** (2.31)
Δ 5yr Treasury Bid-Ask Spread _{<i>t</i>}		1137.579** (2.54)		1162.252** (2.55)		1180.333*** (2.60)		1172.882*** (2.61)
Δ 5y5y Breakeven Inflation _{<i>t</i>}		-13.000*** (-3.98)		-12.753*** (-3.89)		-12.720*** (-3.94)		-12.907*** (-3.95)
<i>N</i>	2,407	2,334	2,407	2,334	2,407	2,334	2,407	2,334
<i>R</i> ²	0.00	0.09	0.00	0.09	0.00	0.09	0.00	0.09
Controls	No	All	No	All	No	All	No	All
Time FE	No	No	No	No	No	No	No	No

Table IA.3: 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. *t*-statistics shown using heteroskedastic and autocorrelation consistent standard errors using the Newey and West (1994) automatic lag selection procedure where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Correlation of PC1 with:		
Treasury margin as share of all margin	-0.51***	(p=0.00; N = 2, 445)
Treasury collateral as share of collateral sinks	-0.40***	(p=0.00; N = 2, 442)

Table IA.4: 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	$\Delta\text{Margin Requirement}_t$			
	(1)	(2)	(3)	(4)
ΔVIX_t	3.288*** (12.87)			
$\Delta\text{VIX}_t \times \mathbf{1}_{\Delta\text{VIX}_t < 0}$		2.743*** (5.48)		
$\Delta\text{VIX}_t \times \mathbf{1}_{\Delta\text{VIX}_t > 0}$		3.662*** (15.29)		
Constant	0.019 (0.15)	-0.490** (-2.04)	5.526*** (14.91)	-4.590*** (-19.48)
N	5,178	5,178	2,358	2,818
R^2	0.39	0.40	0.00	0.00
Sample	All	All	$\Delta\text{Margin} > 0$	$\Delta\text{Margin} < 0$

Table IA.5: 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 heteroskedastic and autocorrelation consistent standard errors where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Z_t			
	(1)	(2)	(3)	(4)
	$\Delta \ln(\text{Central. Margin})$	$\Delta\text{Central. Ratio}$	$\Delta \ln(\text{Margin})$	$\Delta\text{Margin Ratio}$
Z_t	0.002** (2.45)	0.003** (2.19)	0.001*** (4.28)	0.003*** (3.94)
N	2,115	2,115	2,109	2,109
R^2	0.01	0.01	0.01	0.01

Table IA.6: Margin Requirements and CME Shock Measures. Table presents the relationship between different measures of margin and the Z_t measures described in section 5.1. $\Delta \ln(\text{Central. Margin})$ is the change in the log of sunk margin that is posted to a CCP or exchange-traded. $\Delta\text{Central. Ratio}$ is change in 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 heteroskedastic 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)
$\Delta \widehat{\text{Sunk Margin}}/\text{UST Out}_{t+1m}$	0.72 (0.42)	1.45 (0.94)				
$\Delta \widehat{\text{Sinks (Excl. LCR)}}/\text{UST Out}_{t+1m}$			4.92 (0.24)	8.08 (0.37)		
$\Delta \widehat{S}_{t+1m}$					-54.37 (-0.03)	12.65 (0.25)
<i>N</i>	2,049	1,977	2,049	1,977	2,049	1,977
Controls	No	Yes	No	Yes	No	Yes

Panel B: First Stage						
	$\Delta \text{Margin}/\text{UST Out.}$		$\Delta \text{Sinks (Excl. LCR)}/\text{UST Out.}$		$\Delta \text{Sinks}/\text{UST Out.}$	
	(1)	(2)	(3)	(4)	(5)	(6)
Z_{t+1m}	0.05 (1.41)	0.05 (1.50)	0.01 (0.27)	0.01 (0.38)	-0.00 (-0.03)	0.01 (0.26)
<i>N</i>	2,049	1,977	2,049	1,977	2,049	1,977
<i>F - stat</i>	1.98	2.24	0.08	0.14	0.00	0.07
Controls	No	Yes	No	Yes	No	Yes

Panel C: OLS						
	Dependent Var.: $\Delta PC1_t$					
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \text{Sunk Margin}/\text{UST Out}_{t+1m}$	0.02 (0.21)	0.02 (0.19)				
$\Delta \text{Sinks (Excl. LCR)}/\text{UST Out}_{t+1m}$			0.07 (0.56)	0.07 (0.57)		
ΔS_{t+1m}					0.06 (0.56)	0.04 (0.38)
<i>N</i>	2,120	2,047	2,120	2,047	2,120	2,047
<i>R</i> ²	0.00	0.09	0.00	0.09	0.00	0.09
Controls	No	Yes	No	Yes	No	Yes

Table IA.7: CME Residual Regression Placebo. Table presents the first-stage, second-stage, and OLS estimates described in Section 5.1 using the placebo shock, which shifts the Z_t shocks forward to one month to Z_{t+1m} . The first two columns instrument for sunk margin relative to Treasuries outstanding, the middle two for collateral sinks excluding LCR-related components, and the last two for the full collateral sink measure. Controls include 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. Margin data runs from 2016 to 2024. *t*-statistics shown using heteroskedastic 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)
$\Delta \widehat{\text{Sunk Margin}} / \text{UST Out.}_t$	-0.13 (-0.02)	0.22 (0.04)				
$\Delta \widehat{\text{Sinks (Excl. LCR)}} / \text{UST Out.}_t$			-0.05 (-0.02)	0.09 (0.04)		
$\widehat{\Delta S}_t$					-0.05 (-0.02)	0.08 (0.04)
<i>N</i>	1,210	1,155	1,210	1,155	1,210	1,155
Controls	No	Yes	No	Yes	No	Yes

Panel B: First Stage						
	$\Delta \widehat{\text{Margin}} / \text{UST Out.}$		$\Delta \widehat{\text{Sinks (Excl. LCR)}} / \text{UST Out.}$		$\Delta \widehat{\text{Sinks}} / \text{UST Out.}$	
	(1)	(2)	(3)	(4)	(5)	(6)
GIV Placebo _{<i>t</i>}	0.01 (0.37)	0.02 (0.48)	0.03 (1.03)	0.04 (1.20)	0.03 (1.04)	0.04 (1.53)
<i>N</i>	1,210	1,155	1,210	1,155	1,210	1,155
<i>F - stat</i>	0.14	0.23	1.06	1.44	1.09	2.33
Controls	No	Yes	No	Yes	No	Yes

Table IA.8: Granular IV Placebo. Table presents the first-stage and second-stage estimates described in Section 5.2. 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 5.2. The first two columns instrument for sunk margin relative to Treasuries outstanding, the middle two for collateral sinks excluding LCR-related components, and the last two for the full collateral sink measure. Controls include 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 heteroskedastic 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 $_t$	-0.0129 (-0.33)
N	1,255
R^2	0.00

Table IA.9: GIV vs. Placebo GIV. Table presents the regression of the actual GIV on the placebo GIV as described in section 5.2. t -statistics shown using heteroskedastic 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.: $\ln(\text{Margin Posted by Banks}_t)$		
	(1)	(2)	(3)
$\ln(s_{i,t})$	0.377*** (67.69)	0.203*** (34.58)	0.193*** (11.88)
$\ln(\text{VaR}_{i,t})$		0.466*** (43.88)	
N	7,100	7,100	7,100
R^2	0.55	0.61	0.83
FE	No	No	Yes

Table IA.10: Risk-Adjusted Notional vs. Margin Posted. Table regresses posted margin as reported in FR2052a on the risk-adjusted derivative notional $s_{i,t}$ described in section 5.2 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 heteroskedastic 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.

<i>Panel A: Notional Exposure (\$ billions)</i>	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

<i>Panel B: Margin Ratio (Basis Points)</i>	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 IA.11: Notional and Margin Ratio 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 A 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. Panel B provides the margin ratio, which 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.

<i>\$ billions</i>	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 IA.12: Collateral Swap Summary Statistics. Table provides summary statistics on collateral swaps on average in 2024.

<i>Panel A: Initial Margin (\$ billions)</i>		Obs.	Mean	SD	Min	Max
	Cash	2,463	59	22	20	121
	USTs	2,463	52	23	9	100
	HQLA	2,463	169	77	37	300
	Non-HQLA	2,463	33	11	10	64
	Total	2,463	202	86	51	362

<i>Panel B: Variation Margin (\$ billions)</i>		Obs.	Mean	SD	Min	Max
	Cash	2,463	317	54	202	470
	USTs	2,463	19	9	8	51
	HQLA	2,463	300	91	147	497
	Non-HQLA	2,463	69	49	1	148
	Total	2,463	369	57	239	536

Table IA.13: 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 2025. All values are aggregated across all banks in the sample before calculating moments.

	S_t				ΔS_t			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Last 5 Days of Quarter _{<i>t</i>}	0.083 (0.34)				-0.016** (-2.29)			
Last Day of Quarter _{<i>t</i>}		0.054 (0.19)				0.001 (0.04)		
Last 5 Days of Year _{<i>t</i>}			0.252 (0.46)				-0.015 (-1.50)	
Last Day of Year _{<i>t</i>}				0.234 (0.40)				0.036 (1.05)
<i>N</i>	2,460	2,460	2,460	2,460	2,424	2,424	2,424	2,424
<i>R</i> ²	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Table IA.14: Quarter-End and Year-End Effects on Collateral Sinks. Table regresses collateral sinks relative to Treasuries outstanding on quarter-end and year-end indicators. The first four columns use the level S_t as the dependent variable; the second four columns use the daily change ΔS_t . Within each group, columns use indicators for the last 5 business days of the quarter, the last business day of the quarter, the last 5 business days of the year, and the last business day of the year, respectively. *t*-statistics shown using heteroskedastic and autocorrelation consistent standard errors using the Newey and West (1994) automatic lag selection procedure where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1)	(2)	(3)	(4)	(5)	(6)
	GCF-Tbill	FF-Tbill	OIS-Tbill	Neg. Z-Spread	TIPS-Tsy	OIS Swap-Tsy
Collateral Sinks/Treasuries Outstanding _{<i>t</i>}	0.064** (2.34)	0.049** (2.21)	0.034*** (2.63)	0.029 (1.49)	0.008 (0.87)	0.001 (0.14)
<i>N</i>	2,381	2,381	2,381	2,381	2,381	2,381
<i>R</i> ²	0.04	0.83	0.08	0.01	0.15	0.12
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes

Table IA.15: Collateral Sinks and Individual Convenience Yield Measures: Levels. Table regresses individual convenience yield measures on collateral sinks relative to Treasuries outstanding with day-of-week fixed effects. GCF-Tbill is the three-month general collateral financing repo rate minus the three-month T-bill yield. FF-Tbill is the effective federal funds rate minus the three-month T-bill yield, maturity-matched using the OIS curve. OIS-Tbill is the three-month overnight-indexed swap rate minus the three-month T-bill yield. Neg. Z-Spread is the average T-bill yield (4 to 26 weeks to maturity) minus the fitted Treasury curve, multiplied by negative one so that larger values indicate higher convenience, following Greenwood et al. (2015). TIPS-Tsy is the 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, following Fleckenstein et al. (2014). OIS Swap-Tsy is the 30-year overnight-indexed swap rate minus the maturity-matched nominal Treasury yield, following Feldhütter and Lando (2008) and Du et al. (2023). All specifications include the full set of controls. *t*-statistics shown using heteroskedastic and autocorrelation consistent standard errors using the Newey and West (1994) automatic lag selection procedure where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1)	(2)	(3)	(4)	(5)	(6)
	$\Delta\text{GCF-Tbill}$	$\Delta\text{FF-Tbill}$	$\Delta\text{OIS-Tbill}$	$\Delta\text{Neg. Z-Spr.}$	$\Delta\text{TIPS-Tsy}$	$\Delta\text{OIS Sw-Tsy}$
$\Delta\text{Collateral Sinks/Treasuries Outstanding}_t$	0.046** (2.02)	0.024** (2.45)	0.023*** (2.63)	0.001 (0.03)	0.007 (0.83)	0.009 (1.53)
N	2,334	2,334	2,334	2,334	2,334	2,334
R^2	0.01	0.73	0.04	0.01	0.21	0.08
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	No	No	No	No	No	No

Table IA.16: Collateral Sinks and Individual Convenience Yield Measures: Changes. Table regresses daily changes in individual convenience yield measures on daily changes in collateral sinks relative to Treasuries outstanding. See Table IA.15 for variable definitions. All specifications include the full set of controls. t -statistics shown using heteroskedastic and autocorrelation consistent standard errors using the Newey and West (1994) automatic lag selection procedure where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	ΔPC1_t		$\Delta\text{PC1}_t^{\text{Long}}$		$\Delta\text{PC1}_t^{\text{Short}}$	
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta\text{Collateral Sinks/Treasuries Outstanding}_t$	2.646** (2.28)	2.517** (2.15)	0.440 (1.03)	0.583 (1.29)	2.895** (2.17)	2.758** (2.05)
N	2,407	2,334	2,407	2,334	2,407	2,334
R^2	0.00	0.09	0.00	0.18	0.00	0.09
Controls	No	Yes	No	Yes	No	Yes

Table IA.17: Collateral Sinks and Convenience Yields by Maturity. Table regresses daily changes in convenience yield measures on daily changes in collateral sinks relative to Treasuries outstanding, separately for the all-maturity first principal component (PC1), the long-maturity component (PC1^{Long} , constructed from the 10-year TIPS-Treasury spread and the 30-year OIS swap-Treasury spread), and the short-maturity component ($\text{PC1}^{\text{Short}}$, constructed from the GCF repo-Tbill, OIS-Tbill, Fed funds-Tbill, and negative Z-spreads). Odd columns include no controls; even columns include the full set of controls and $\Delta \ln(\text{Treasuries Outstanding})$. t -statistics shown using heteroskedastic and autocorrelation consistent standard errors using the Newey and West (1994) automatic lag selection procedure where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	ΔPC1_t		$\Delta\text{PC1}_t^{\text{Long}}$		$\Delta\text{PC1}_t^{\text{Short}}$	
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \ln(\text{Collateral Sinks})_t$	28.376** (2.34)	27.748** (2.29)	6.647 (1.59)	5.839 (1.30)	30.680** (2.19)	30.621** (2.19)
N	2,407	2,334	2,407	2,334	2,407	2,334
R^2	0.00	0.09	0.00	0.18	0.00	0.09
Controls	No	Yes	No	Yes	No	Yes

Table IA.18: Collateral Sinks and Convenience Yields: Log Specification. Table regresses daily changes in convenience yield measures on changes in the log of collateral sinks, replacing the ratio specification with a log-level measure that does not scale by Treasuries outstanding. $\Delta \ln(\text{Treasuries Outstanding})$ is included as a separate control in even columns. See Table IA.17 for definitions of PC1, PC1^{Long} , and $\text{PC1}^{\text{Short}}$. t -statistics shown using heteroskedastic and autocorrelation consistent standard errors using the Newey and West (1994) automatic lag selection procedure where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Z_t	
	(1) Ratio	(2) VaR
Z_t	0.777** (2.47)	0.583** (2.37)
Quintile 2	0.133 (0.80)	0.017 (0.11)
Quintile 3	0.146 (1.03)	0.057 (0.40)
Quintile 4	0.071 (0.50)	0.117 (0.73)
Quintile 5	0.054 (0.39)	-0.060 (-0.43)
Quintile 2 $\times Z_t$	-0.629 (-1.46)	-0.223 (-0.69)
Quintile 3 $\times Z_t$	-0.585* (-1.79)	-0.412 (-1.44)
Quintile 4 $\times Z_t$	-0.592* (-1.71)	-0.366 (-1.18)
Quintile 5 $\times Z_t$	-0.789** (-2.12)	-0.715** (-2.12)
N	2,105	2,143
R^2	0.10	0.09

Table IA.19: Passthrough Regression. Table presents the regression described in section 6. 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 heteroskedastic and autocorrelation consistent standard errors using the Newey and West (1994) automatic lag selection procedure where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.