Cross-Asset Limit Allocation of Trading Desks*

Falk Bräuning[1], Wenxin Du[2], and Hillary Stein[1]

[1]Federal Reserve Bank of Boston
[2]Harvard Business School

March 26, 2026

Abstract

We study the interrelation between trading desks' risk limits and asset market dynamics, across multiple asset markets. We first show that desk limits respond endogenously to risk-adjusted intermediation income. We then document significant spillovers between limits of different asset classes, using a causal identificaion scheme that hinges on exogenous desk closures and associated loss in intermediation capacity. We show that in response to idiosyncratic desk closures, remaining desks in a given asset market experience higher trading margins, and banks increase trading limits in these assets, but the increase comes at the expense of a limit reduction in other asset classes.

Keywords: banks, market making, financial intermediation, risk constraints

JEL Codes: G12, G14, G21

*Kathryn Peaquin and Lilia Qian provided excellent research assistance. The views expressed here are solely those of the authors and do not reflect the opinions of the Federal Reserve Bank of Boston, or the Federal Reserve System. Falk.Braeuning@bos.frb.org, Hillary.Stein@bos.frb.org

Introduction

The intermediation of financial markets is crucial for efficient market functioning and price discovery. Nevertheless, market-making is concentrated in a small subset of dealer-banks, and recent literature has highlighted how these dealers' constraints affect market liquidity (Barbiero et al., 2024; Bräuning and Stein, 2024). Given that the same dealers intermediate multiple markets, a dealer-level shock in one asset class may have implications for the liquidity of other asset classes. This can affect all key asset markets, including government bonds, corporate debt, asset-backed securities (ABS), equities, and foreign exchange. More generally, understanding how dealers allocate their balance sheet resources across trading desks and asset classes is essential for understanding overall asset market dynamics.

In this paper, we provide a comprehensive analysis of the interrelation between trading desks' risk limits and asset market dynamics across multiple asset classes. Our study focuses on aggregate dynamics but leverages granular supervisory desk-level data covering more than 90% of US-based trading activity, collected in accordance with the reporting requirements of the Volcker Rule.[1] This rich dataset includes key variables such as turnover, income from new (including intraday) and existing positions, overall position size, and dealer constraint usage, measured as Value-at-Risk (VaR) usage relative to its limit. Using this dataset we ask the following questions: What prompts limits to increase, and how are limits reallocated across desks? Given the extensive information about trading desk activity that the supervisory data provides, it allows for novel insights into these questions. Specifically, we use the information of both prices (income) and quantities (VaR, positions, and turnover) to assess how desks manage their risk in response to market shocks and adjust limits depending on profitability.

We first show that trading desk limits respond endogenously to intermediation profits. Using impulse responses, we find that trading desks experience increased turnover and higher income from bid-ask spreads in response to market volatility. This is consistent with our finding that desks derive the majority of their profit from bid-ask spread and commission income, rather than from market gains. In response to a surge in market volatility, trading desks also hold larger positions and utilize a higher share of their available risk limits. On net, however, intermediation income increases more than risk, such that desks face higher *risk-adjusted* intermediation income—measured by a modified sharp ratio as intermediation income over value at risk—which in turn leads banks to allocate more risk limits to their trading desks in subsequent periods. Importantly, we document a clear lagged relationship between limit increases and increases in risk-adjusted intermediation profits.

We next examine how risk limits are allocated across desks. A variance decomposition

[1] `https://www.federalreserve.gov/apps/reportingforms/Report/Index/FR_VV-1`.

exercise shows that about one-third of the variation in desk-level limit changes can be explained by bank-level changes over time, potentially due to changing levels of banks' equity capital. On the other hand, much less of the variation in limits can be explained by market-level shocks to an asset class. Using regressions with various fixed effects, we document significant "spillover" effects of limit changes within a bank, both across different asset classes and, within a given asset class, across desks. Specifically, we show how, within a given bank, an increase in limits of one asset class comes at the expense of a limit reduction in all other asset classes. We find a similar reallocation across desks within the same asset class. This provides suggestive evidence that banks first allocate limits to their trading business unit and then to asset classes within the trading business, implying that asset classes are interrelated through limit reallocation.

These findings carry two broader implications. First, any bank-wide capital shock, such as a large trading loss, a regulatory capital surcharge, or a ratings downgrade, can simultaneously affects liquidity across every asset class that bank intermediates. This constitutes a systemic amplification mechanism that is invisible when individual asset markets are studied in isolation. Second, our results provide direct micro-evidence on the internal structure of dealer-bank capital allocation: banks appear to operate a hierarchical internal capital market, first setting a total risk budget for the trading unit as a whole, then allocating to asset classes. This organizational structure distinguishes our paper from theories that model dealers as facing separate, independent constraints per asset class, and connects to the broader corporate finance literature on internal capital markets (Stein, 1997).

In order to more carefully understand how limits are reallocated within a bank, we turn to a causal identification strategy, which hinges on the exogenous management decisions of *other* banks. Specifically, we find that banks allocate a higher limit to an asset class when other banks shut down desks in that asset class. As we discuss, this shut-down decision is often a high-level management decision about bank specialization rather than linked to profitability of an asset class, yet it significantly impacts the profitability of financial intermediation for the remaining banks in that asset class. In fact, we show that these desks at other banks experience increased turnover, larger positions, and higher modified Sharpe ratios, indicating changes in risk-adjusted performance. Thus, they see their limits raised following the shock to accommodate increased intermediation. In turn, we trace out how this increase in limits in one asset class leads to a decrease in limits of other asset classes at that bank. These results underscore how limits are first fixed across the trading unit, leading to the inter-dependence of risk limits across asset markets.

Related Literature This paper most closely relates to the literature on constrained intermediaries. An established theoretical literature demonstrates the implications of intermediary constraints on asset pricing (Grossman and Miller, 1988; Bernanke and Gertler, 1989; He and Krishnamurthy, 2013; Brunnermeier and Sannikov, 2014), while an empirical literature provides support for these theories, generally using publicly available bank balance sheet data (Adrian, Moench and Shin, 2010; Adrian, Etula and Shin, 2015). An additional literature links intermediary constraints and market liquidity (Gârleanu and Pedersen, 2007; Breckenfelder and Ivashina, 2023; Duffie et al., 2023). In previous work, Barbiero et al. (2024) and Bräuning and Stein (2024) provide empirical support for both strands of this literature, linking dealer constraints to both asset price movements and liquidity in the exchange rate and Treasury markets, respectively. This paper asks a natural follow-up question: What in turn determines dealer constraints?

In considering cross-asset implications of the fact that the same small number of dealers intermediate multiple asset classes, our paper also relates to those that connect bank constraints with returns of multiple asset classes. He, Kelly and Manela (2017) show that bank capital ratios predict cross-sectional variation in expected asset returns, while Du, Hébert and Huber (2023) show a similar result for covered-interest parity deviations, which they argue measure the extent to which constraints bind. Our paper makes a similar point: Dealer constraints should affect multiple asset markets. However, our paper seeks a sharper identification in measuring spillovers between asset markets. It also provides direct micro-evidence behind the mechanism by which dealer constraints would affect multiple assets. In this sense, our paper also relates to the literature on cross-asset spillovers (Greenwood, Hanson and Liao, 2018), providing a novel mechanism that focuses on intermediaries rather than generalist investors.

This paper also relates to Anderson, McArthur and Wang (2023), Li, Petrasek and Tian (2024), and Lu and Wallen (2024), three papers that also use the VV-1 data and provide descriptions of limit usage at the trading-desk level. Both Anderson, McArthur and Wang (2023) and Li, Petrasek and Tian (2024) provide evidence in favor of binding limits in times of stress. Anderson, McArthur and Wang (2023) show a sharp increase in the VaR limits of corporate bond desks during the COVID period, while Li, Petrasek and Tian (2024) focus on Treasury desks to show that during the COVID crisis, dealers that were closer to their limits sold more Treasury securities and accepted lower prices. Unlike these two papers, we focus on other desk-level metrics in addition to limit usage, and we provide evidence for how VaR limits respond to these other metrics. We also take a more wholistic view of the trading business and consider market-making desks of all asset classes. Specifically, we consider both the separation of individual asset classes and the spillovers between them. Lu and Wallen (2024) also look across asset classes; however, the focus of their paper is on the

documentation of dealers as toll-taking intermediaries rather than risk-bearers, a finding that we also document in our Stylized Facts section using different metrics and a larger subset of banks. In addition to documenting facts about trading desks, our paper primarily focuses on the dynamics of limit allocation.

More broadly, our paper also relates to the literature on internal capital markets in banking organizations. In a seminal paper, Stein (1997) shows theoretically that banking organization allocate scarce resources to competing projects across their organization in an internal capital market. Empirically, Cetorelli and Goldberg (2012*a*) and Cetorelli and Goldberg (2012*b*) document that global banks use internal capital markets to absorb shocks by reallocating liquidity across offices in different countries. Similarly, Bräuning and Ivashina (2020) show that banks' global lending decisions are interlinked through capital constraints leading to monetary spillovers. Our paper shows that such inter-linkages resulting from resource constraints on banks' trading business are also key for asset market intermediation.

Our paper proceeds as follows. Section 1 discusses the VV-1 data, the variables we use, and how we clean and label the data. Section 2 documents stylized facts about trading desks, which we learn from the VV-1 data. Section 3 provides evidence for hierarchical internal capital markets and "winner-picking." Section 4 explores how this structure connects otherwise disparate asset classes. Section 5 concludes.

1 Data Description

The primary dataset used in this paper is the Regulation VV Quantitative Measurements dataset (henceforth, VV-1), which is collected as part of the Volcker Rule of the Dodd-Frank Act. The data were intended to provide regulators with the ability to monitor the trading activities of market-making desks of global banks. The dataset covers bank holding companies (BHCs) and foreign banking organizations (FBOs) with average gross trading assets and liabilities over the prior four calendar quarters equal to $20 billion or more, excluding obligations issued or guaranteed by the United States or a U.S. agency. Information is collected on a daily basis. Reporting was streamlined when the original 2013 Volcker Rule was updated in 2019; the original Volcker dataset covers July 2019 through December 2020. The new dataset covers January 2021 through the present—our sample goes through December 2023. There are some differences in the available variables in these two datasets, which we detail below.

Variables. Banks self-report their desks according to their own internal organization. They give each desk an identifier (mostly consistent across time) as well as a textual name and

description. At a daily frequency, they report for each desk:

- Profit and loss of both new and existing positions. Existing positions are those that were held as of the end of the prior day. New positions include intraday trading activity and include both fee and commission income or expenses and market gains or losses.[2]
- Market value of long and short security positions and derivatives receivable and payable.
- Gross market value of transactions, split by securities and derivatives.
- 99% Value-at-Risk (VaR) of aggregated positions for a 1-day holding period.

Banks also report their internal limits for each desk. Internal limits often take the form of risk measures, such as VaR and risk factor sensitivities, but they may also take the form of other observable criteria, such as net open positions or inventory aging. Desks often have multiple limits. Banks provide textual names and descriptions of each limit. While limits have a consistent ID across time in the new data, there is no limit ID in the old data, making it nontrivial to track a given limit across time. At a daily frequency, banks report for each limit:

- Upper and lower bounds for the limit size.
- The limit usage.

In the vast majority of cases where both upper and lower bounds are reported, limits are symmetric in the sense that upper and lower bounds for the limit sizes are the same in absolute value. Moreover, in several cases, desks only report a lower limit, consistent with the aim of limiting losses rather than profits. In our analysis, we collapse the data by working with the maximum of the absolute value of the upper and lower bounds, which we refer to as "limit". Percent limit usages is than constructed as limit usage over limit (in percent).

Cleaning steps. We started by ensuring consistent desk IDs within a bank over time. This brought the number of desks in the full dataset from 2631 to 2488. We then categorized each market-making desk into one or more of the following asset classes: asset-backed securities (ABS), commodities, corporate bonds, equity, exchange rates (FX), government bonds, municipal bonds. We did so based on the desk's name, description, and identifier (which sometimes includes descriptive information); each desk's classification was checked by hand. Overall, we classified 1705 desks; 1556 were classified into a single category, and 149

[2]Note that these profits do not include internal capital charges and might thus be better viewed as income rather than profit.

were classified into multiple categories (for example, it is common for a desk to specialize in both FX and government bond trading).

Because limit types can vary widely, and because limit information is often not descriptive enough to know whether two limits may be compared across desks or banks, we identify the limit for each desk that corresponds to its reported 99% VaR. This allows us to identify the bounds of this 99% VaR. We find that 1619 desks across 18 banks have a 99% VaR as one of their internal limits. We include additional desks by identifying the 95% VaR; we do so by assuming normal returns and simply multiplying by the ratio of the z-scores. We then convert this 95% VaR back into a 99% VaR. This brings our total to 1870 desks across 18 banks. When we focus on market-making desks with an asset classification and for which we can identify the 99% VaR limit and use, we get 1406 desks across 17 banks.

Turning to turnover, there are some changes between variable definitions in the new and original Volcker datasets. Under the original Volcker rule, banks reported for each desk on each day the value of total transactions with customers and non-customers over the past 30 days. Under the updated rule, this was streamlined such that banks reported for each desk on each day the market value of transactions on that day. Additionally, the new data splits this transaction value by customers, non-customers, and internal/affiliated counter parties as well as by securities and derivatives. For both datasets, we first sum transaction value across customer segments for total transaction value. In order to use a comparable variable across time, we then construct a rolling 30-day sum of transaction value in the new data.[3] Thus, when we refer to "turnover," we are actually referring to a 30-day rolling sum of the value of total transactions.

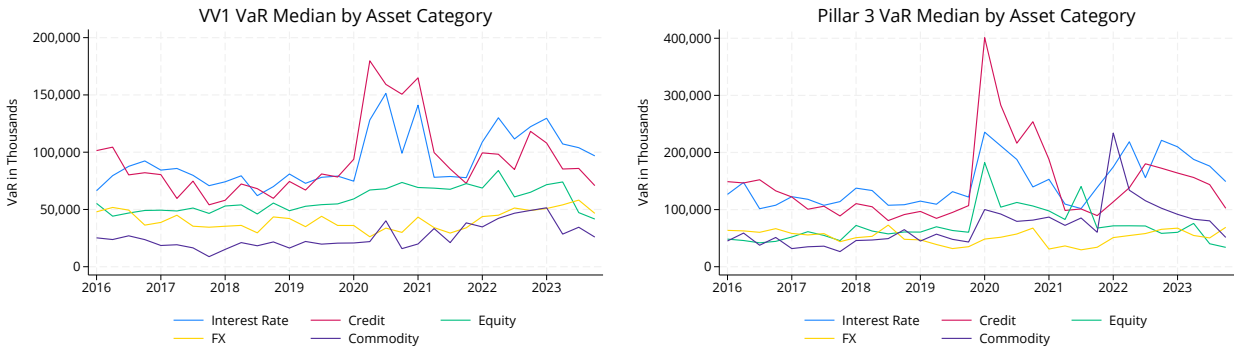
Additionally, we examine the profit and loss of both new and existing positions, as well as the market value of gross and net positions. We calculate gross position as the sum of the market values of all long security positions, short security positions, and derivatives receivable and payable. We calculate net position as the market values of long security positions and derivatives receivables minus the market values of short security positions and derivatives payable.[4] In total, 1114 classified market-making desks across 16 banks with an identified 99% VaR limit report positions, turnover, and profit and loss some time between 2016 and 2023.

[3]We construct this 30-day sum by assuming 22 working days per every 30 days.

[4]In the new data, desks report the market value of all long security positions, short security positions, and derivatives receivable and payable on a daily basis. In the original dataset, desks reported the value of security assets and liabilities, as well as the gross notional values of derivative assets and liabilities. In practice, we notice a break in the timeseries when the data shifts from the old to new reporting requirements, and so we splice position data by adding to the new data the difference between the means in the years 2020 and 2021.

Relationship between VV-1 data and Pillar 3 disclosures. As a validation exercise of the VV-1 data, we compare VaR limits with public Basel Pillar 3 disclosures by asset class. In order to do so, we first sum the 99% VaR limits across desks in each asset class within each bank in the VV-1 data. This gives us an estimate of the total VaR within market making desks of a given asset class by bank. In Figure 1, we plot the median VaR limit by asset class from the VV-1 data on the left and the median VaR by asset class from the Pillar 3 data on the right. We can see the broad patterns by asset class seem to match well. We can also see that the VaR in the Pillar 3 disclosures is larger than in the VV-1 data, reflecting the fact that market risk can sit in other areas of the bank in addition to market-making trading desks. Appendix Table A.1 compares asset class-bank-quarter observations, showing significant comovement between the datasets.

Figure 1: VaR in VV-1 Data and Pillar 3 Disclosures



Notes: This figure shows how VaR limits at the asset class level in the VV-1 data and in Pillar 3 disclosures vary over time. Observations are at the asset class-bank-quarter level. The sample covers 2016 to 2023 and includes banks for which we have a complete panel in both datasets. The VV-1 sample is comprised of market-making desks that are classified into at least one of our seven asset categories. *Sources*: FR VV-1, Pillar 3 disclosures, authors' calculations.

2 Stylized Facts about Trading Desks

In this section, we present a few stylized facts about trading desks in order to give some insight into their operations.

Fact 1. *Desks make the majority of their profit from intraday trading, implying profits primarily come from bid–ask spreads and commission income rather than market gains.*

Recall that banks separately report their profit from "new positions," which includes intraday trading activity, and "existing positions," which refer to positions held at the close

of the previous day. Crucially, intraday trading activity includes fee and commission income earned from buy and sell orders over the course of the day. For each profit category, we calculate the *modified Sharpe ratio*, which we define as a desk's profit divided by its 99% VaR. This gives us a measure of profit normalized by risk.

Summary statistics of these modified Sharpe ratios, as well as our other variables of interest, are presented in Table 1. In order to calculate summary statistics at the desk-level, we first aggregate each variable at the monthly level, and we then take the median by desk across time. We can see that the modified Sharpe ratio of new positions is higher than that of existing positions across the distribution, indicating that these desks are compensated for their market-making via higher commissions and bid–ask spreads, rather than via longer-term capital gains. Figure 2 reinforces this point, showing how the profit from new positions is consistently larger each year.

Table 1: Desk-Level Summary Statistics

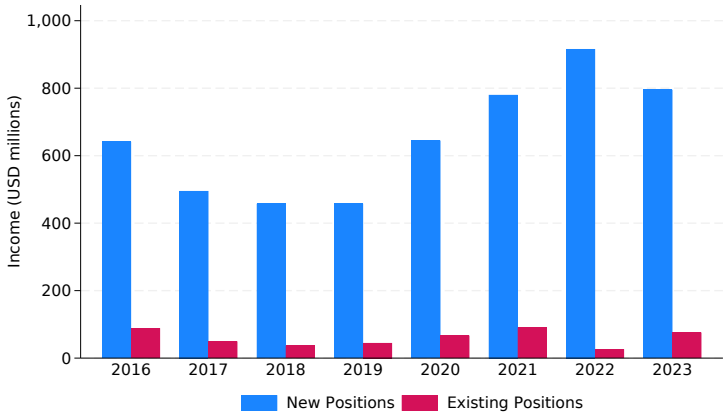
	mean	p10	p25	p50	p75	p90	count
Turnover	476.431	0.043	0.938	9.525	101.227	540.441	1123
Gross Position over Turnover	10.365	0.011	0.238	1.943	9.687	40.887	1098
Net Position over Gross Position	0.149	-0.036	0.000	0.024	0.180	0.627	1112
Sharpe (New)	0.535	0.000	0.021	0.093	0.262	1.267	1104
Sharpe (Existing)	0.024	-0.102	-0.031	0.002	0.046	0.154	1104
Limit	32.961	0.009	0.246	1.280	3.686	8.874	1123
VaR over Net Position	0.009	0.000	0.000	0.002	0.007	0.025	1109
Pct. Limit Use	0.252	0.023	0.111	0.245	0.372	0.483	1118
Income New Pos over Turnover	0.376	0.000	0.011	0.080	0.332	1.138	1098
Income Exist Pos over Net Position	0.131	-3.983	-0.428	0.001	0.630	2.966	1109

Notes: This table shows summary statistics for a variety of variables and ratios for market-making desks that are classified into at least one of our seven asset categories. The Turnover variable is expressed in USD Billions while Limit Use is in USD millions. Both Income New Position over Turnover and Income Existing Position over Net Position are displayed in bps. The sample covers 2016 through 2023, and all variables are winsorized at the 5th and 95th percentiles. In order to calculate summary statistics at the desk-level, we first aggregate each variable at the monthly level, and we then take the median by desk across time. *Sources*: FR VV-1, authors' calculations.

In fact, desks often make losses on existing positions. We calculate a desk's return on existing positions as the profit from existing positions divided by the absolute value of its net position. Table 1 shows how this value is often negative on a monthly basis. Additionally, this return can differ considerably across desks, with the 10th percentile desk earning a median monthly loss of 4.2 basis points and the 90th percentile desk earning a median monthly gain of 3.3 basis points.

That said, desks do seem to be compensated with higher returns on existing positions

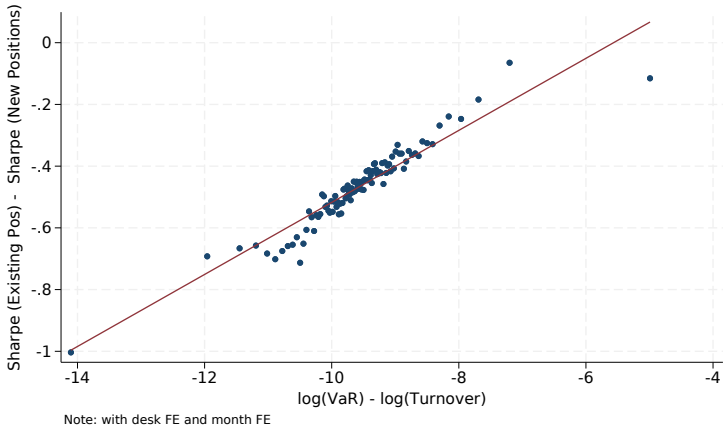
Figure 2: Mean Income from New and Existing Positions over Time



Notes: This figure shows the mean monthly income (PnL) each year from either new or existing positions across desks. It only includes market-making desks that are classified into at least one of our seven asset categories. *Sources*: FR VV-1, authors' calculations.

when those positions are riskier. Figure 3 provides evidence for this with a binned scatter plot of observations at the desk-month level, with desk and month fixed effects partialled out. Specifically, the figure shows how higher riskiness (measured as a higher ratio of VaR usage relative to turnover) is correlated with higher income on existing positions (measured as a higher Sharpe ratio on existing positions relative to new positions). On the other hand, when VaR is lower relative to turnover, banks make more income from new positions, that is through bid-ask spreads and commissions.

Figure 3: Desk Income and Riskiness



Notes: This binscatter uses desk-month-level data from 2016 to 2023 with constant sample of desks. It only includes market-making desks that are classified into at least one of our seven asset categories. Desk and month fixed effects are included in the binscatter. *Sources*: FR VV-1, authors' calculations.

We also calculate a desk's spread income on new positions as the profit from new positions divided by turnover.[5] On average, the average desk earns a median monthly spread of 0.38 basis points, and the median desk earns a median monthly spread of 0.08 basis points. This median monthly spread is positive for the vast majority of desks. In this context, it is important to note that we only observe income on new positions at the daily frequency. Hence, our proxy of spread—income on new positions relative to turnover—could be contaminated by intra-day price movements (capital gains) of new positions.

Consistent with the fact that desks' primary compensation comes in the form of commissions and bid–ask spreads, desks hold relatively small positions relative to their turnover. This may not come as a surprise due to our focus on market making desks, which by their line of business focus on facilitating trade between buyers and sellers rather than having a large proprietary trading book. In fact, Volcker regulation generally prohibits proprietary trading for market makers.

Fact 2. *On average, desks operate with a healthy buffer in terms of their VaR limit.*

Table 1 shows that in general, desks' 99% VaR limit usage remains far from the limit itself, as shown by the summary statistics on the median monthly VaR use over VaR limit. Figure 4 shows the distribution of this fraction at the desk-month level and demonstrates considerable more heterogeneity. It seems that, although the median desk in the median month only uses about a quarter of its limit, some desks in some months use a considerable higher fraction of their allocated limit. Additionally, as discussed in Barbiero et al. (2024) and Bräuning and Stein (2024), desks can be constrained by their limit even when they have not reached their limit precisely. This is because banks or desks may choose a target VaR which is lower than their allocated maximum limit.

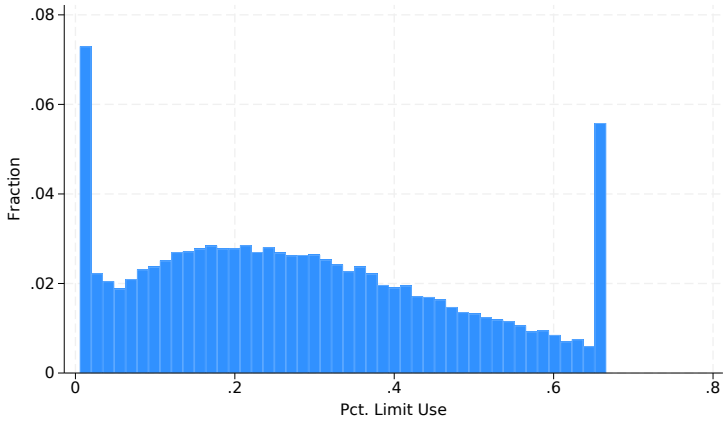
Fact 3. *Desks differ by asset classes in terms of the risk they take and their turnover.*

Table 2 shows summary statistics split by asset categories.[6] There is considerable heterogeneity between asset categories. For example, we can see how FX desks have the highest modified Sharpe ratios for new positions, indicating that these desks make the most income off of bid–ask spreads relative to the risk they take. Nevertheless, they do not charge relatively high spreads. It seems FX desks make their income due to their high transaction volumes and their relatively low risk. Somewhat similarly, equity desks also take relatively low risk and have the highest modified Sharpe ratios for existing positions. On the other end

[5]This measure resembles an effective spread, as income on new positions may also encompass commissions, or other fixed incomes per transaction, in addition to spread income.

[6]This table allows for overlap between the groups, as a trading desk may focus on multiple asset classes. Appendix Table **??** shows the same summary statistics by asset category for non-overlapping groups.

Figure 4: Fraction Use of 99% VaR Limit



Notes: This table shows a histogram of Percent Limit Use as calculated by the ratio of the VaR use over VaR limit for market-making desks that are classified into at least one of our seven asset categories. The sample covers 2016 through 2023, and each observation is a month-desk. The Percent Limit Use presented has been winsorized at the 5th and 95th percentiles. *Sources*: FR VV-1, authors' calculations.

of the spectrum, ABS desks have much higher VaR usage. This asset category is much less liquid, with relatively low turnover and a more difficult price discovery process.

Despite the heterogeneity across a range of variables, indicating different business models actross different asset categories, it is notable that VaR use over VaR limit is remarkably consistent across asset categories. It seems that management allocates limits, in general, in such a way that businesses can be profitable while maintaining a healthy VaR limit buffer.

Table 2: Desk-Level Summary Statistics by Asset Class

	ABS		Comm		Corp		Equi		FX		Gov		Muni	
	mean	p50	mean	p50	mean	p50	mean	p50	mean	p50	mean	p50	mean	p50
Turnover	293.037	1.989	221.582	8.803	170.418	5.398	139.797	16.828	931.003	75.894	1039.380	31.975	33.533	0.938
Gross Position over Turnover	8.027	1.197	11.221	3.654	12.339	3.682	4.455	0.683	5.935	1.434	16.380	5.745	11.990	0.490
Net Position over Gross Position	0.247	0.128	0.097	0.010	0.170	0.033	0.134	0.031	0.081	0.005	0.088	0.017	0.454	0.367
Sharpe (New)	0.216	0.034	0.374	0.081	0.436	0.079	0.542	0.112	1.258	0.347	0.314	0.086	0.652	0.056
Sharpe (Existing)	0.034	0.010	0.049	0.016	0.012	-0.001	0.067	0.010	0.005	0.000	0.002	-0.002	-0.008	0.004
Limit	149.397	1.645	3.052	1.289	27.311	1.644	14.483	1.030	2.981	0.779	76.319	1.644	7.375	0.648
VaR over Net Position	0.007	0.003	0.006	0.001	0.007	0.003	0.017	0.003	0.011	0.001	0.002	0.000	0.006	0.004
Pct. Limit Use	0.283	0.283	0.252	0.266	0.264	0.265	0.247	0.258	0.198	0.184	0.268	0.258	0.249	0.216
Income New Pos over Turnover	0.472	0.204	0.291	0.084	0.509	0.217	0.282	0.072	0.361	0.031	0.249	0.027	0.754	0.312
Income Exist Pos over Net Position	1.484	0.300	2.542	0.018	0.147	-0.005	-0.017	0.060	-0.752	0.000	-0.220	-0.000	-0.550	0.075

Notes: This table shows summary statistics for a variety of variables and ratios for market-making desks that are classified into at least one of our seven asset categories. The Turnover variable is expressed in USD Billions while Limit Use is in USD millions. Both Income New Position over Turnover and Income Existing Position over Net Position are displayed in bps. The sample covers 2016 through 2023, and all variables are winsorized at the 5th and 95th percentiles. In order to calculate summary statistics at the desk-level, we first aggregate each variable at the monthly level, and we then take the median by desk across time. Desk groupings here are not mutually exclusive; a desk can be in several groups. *Sources*: FR VV-1, authors' calculations.

3 The Hierarchy of Limit Allocation

In this section, we provide evidence that banks seem to operate with hierarchical internal capital markets. We first examine how banks allocate risk to the market-making unit as a whole. We then show how risks are allocated at the desk level in response to ex-post desk profitability.

Overall market-making risk allocation. Banks are comprised of multiple organizational layers. Banks as a whole are subject to regulatory capital requirements, which may be risk-sensitive or risk-insensitive. These constraints are then allocated across different divisions according to both managerial preferences and regulatory requirements. Table 3 shows how the sum of 99% VaR limits across market-making desks (of all asset classes) are related to the Tier 1 capital ratio, even accounting for bank and quarter fixed effects. Thus, when banks are better capitalized, they allocate more risk to their market-making business, providing evidence for such a heirarchical allocation structure.

Table 3: Relationship between VV-1 and Pillar 3 VaR

	Log(Limit)	
	(1)	(2)
SLR	0.038	
	(0.287)	
Tier 1 Capital Ratio		0.216***
		(0.081)
Time FE	Yes	Yes
Bank FE	Yes	Yes
R^2	0.739	0.751
R^2 within	0.000	0.045
N	187	187

Notes: This table shows how limit use and percent limit use at the asset class level varies with reported value at risk in Pillar 3 disclosures. Observations are at the asset class-bank-quarter level. The sample covers 2016 to 2023 and is comprised of market-making desks that are classified into at least one of these asset categories. Robust standard errors are clustered at the desk level. *** p<0.01, ** p<0.05, * p<0.1. *Sources*: FR VV-1, Pillar 3 disclosures, authors' calculations.

Desk-level risk allocation. Theories of internal capital markets (e.g., Stein, 1997) emphasize how managers allocate internal capital, given their limit, according to profitability. A natural question in this context is, thus: does desk profitability translate into a higher limit

allocation? Because endogeneity concerns complicate the analysis, we seek to mitigate some of these concerns by using an instrument for profitability. In the next section, we utilize an even stronger causal identification scheme in delving into reallocation across asset classes.

Section 2 showed how trading desks make the majority of their income from commissions and bid–ask spreads, rather than capital gains. Thus, anything that increases the volume of transactions should increase their profits from their market-making business, keeping spreads constant. We therefore instrument VaR limits with market volatility, proxied with the VIX. First, we directly examine the correlation between changes in the VIX and desk-level outcome by running regressions of the following form:

$$y_{d,b,m} = \alpha_{d,b} + \beta \Delta VIX_m + \sum_{j=1}^{12} \mu_j \Delta VIX_{m-j} + \sum_{j=1}^{12} \zeta_j y_{d,b,m-j} + \varepsilon_{d,b,m}, \tag{1}$$

for various outcome variables y, where d denotes desk, b denotes bank, and m denotes month. Notably, the regression includes desk fixed effects. Results are shown in Table 4. Column (1) shows that when volatility increases, turnover increases as well. Desks are doing more business and increase both their gross position (column 2) and the size of their net position (column 3). In keeping with this, their limit usage increases as a percent of their allocated limit size (column 4), even as they are allocated a larger limit (column 5). In line with increased turnover, their modified Sharpe ratios of new positions increases (column 7), such that their income increases even in relation to an increased risk size. However, their modified Sharpe ratios of existing positions decreases (column 6), such that capital gains fall relative to the desk's increased risk.

Table 4: Correlation between VIX Changes and Trading Outcomes

	(1) Turnover	(2) Gross Position	(3) Net Position	(4) Pct. Limit Use	(5) Limit	(6) Sharpe (Existing)	(7) Sharpe (New)
Δ VIX	0.039***	0.010***	0.024***	0.009***	0.004**	-0.016***	0.028***
	(0.005)	(0.003)	(0.008)	(0.002)	(0.002)	(0.003)	(0.007)
LHS lags	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RHS lags	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Desk FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.966	0.994	0.939	0.827	0.992	0.397	0.891
N	8,951	8,807	8,907	8,964	8,964	8,964	8,964

Notes: This table shows the correlation of changes in the VIX index with various desk-level variables. For this analysis, the logarithms of the Turnover, Gross Position, Net Position, and Limit variables are used. Furthermore, the absolute value of the net position variable is utilized. Observations are at the desk-month level. The sample covers 2016 to 2023 and is comprised of market-making desks that are classified into at least one of our seven asset categories. The sample is constant across time. Robust standard errors are clustered at the desk level. *** p<0.01, ** p<0.05, * p<0.1. *Sources*: FR VV-1, Haver, authors' calculations.

In order to better understand the dynamics of these effects, we run local projections (Jordà, 2005) in which we consider how market volatility affects desk level outcomes over time. Although the majority of our analysis in this paper is at the monthly frequency, here we examine data at a weekly frequency, in order to zoom in more on the dynamic relationship between key trading variables.[7] Specifically, for weekly horizons $h \in [0, 6]$ we run:

$$y_{d,b,m+h} = \alpha_{d,b} + \beta_h VIX_m + \sum_{j=1}^{12} \mu_j VIX_{m-j} + \sum_{j=1}^{12} \zeta_j y_{d,b,m-j} + \varepsilon_{d,b,m}. \tag{2}$$

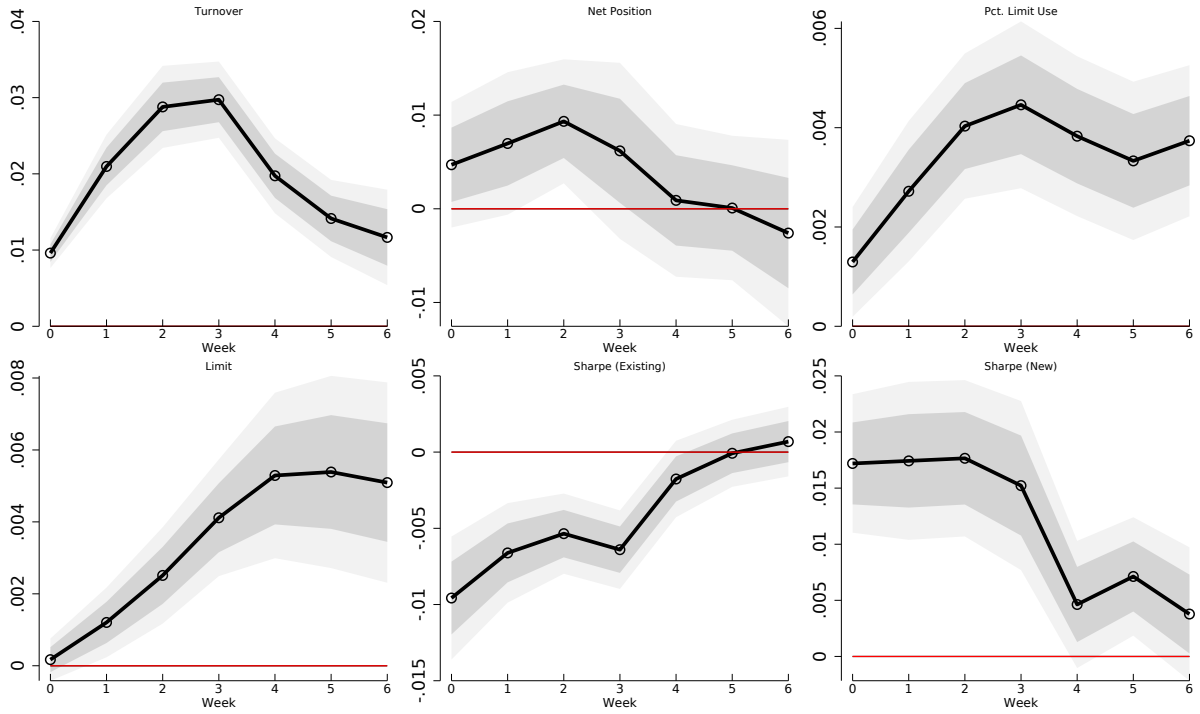
Figure 5 plots the results of these local projects. The initial impact matches the results in Table 4. However, while turnover returns to baseline within six weeks, the increase in the limit size of the desk is persistent.

VIX changes thus appear to be a valid instrument for desk profitability; when market volatility increases, desks experience higher turnover and higher modified Sharpe ratios of new positions. It seems that management increases a desk's limit in response; however, we can more directly test this by instrumenting the modified Sharpe ratio of new positions with the VIX. The results are shown in Figure 6. We can see that an increase in the modified Sharpe ratio of new positions indeed seems to increase the limit size of the desk in a persistent manner. Overall, this provides evidence that management is engaging in "winner-picking" in their internal capital market allocation.

It is important to remember in these analyses that the VIX is not, in fact, exogenous to trading desks, as these desks are instrumental in determining price movements of the assets they trade. For example, Barbiero et al. (2024) show that limit shocks to FX trading desks amplify movements in exchange rates, and Bräuning and Stein (2024) show the same for Treasury trading desks. The evidence presented here is suggestive that market movements, in turn, prompt limit changes; the final section will discuss a more robust way of identifying this mechanism.

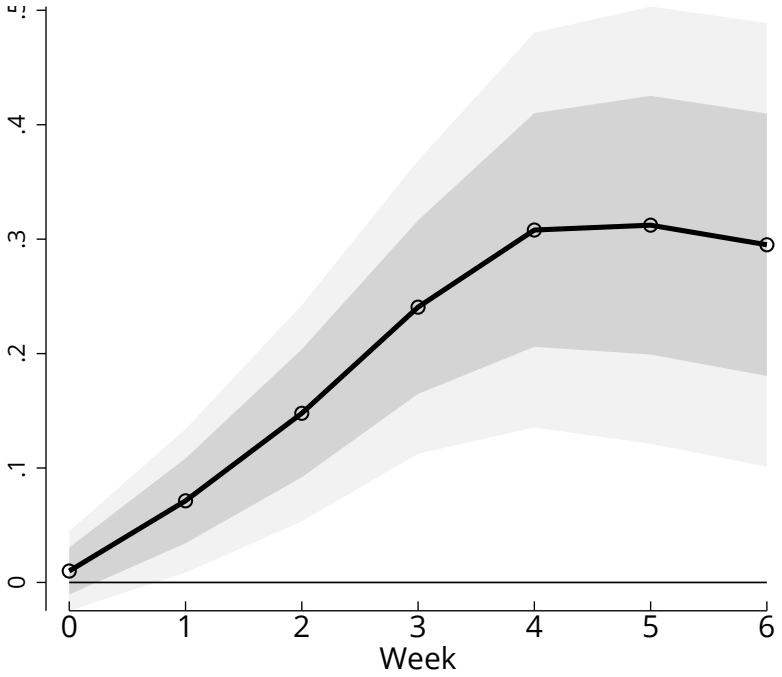
[7] Recall, turnover is constructed as a rolling 30-day sum. This is true regardless of whether the unit of analysis is the week or the month. Thus, turnover will be appear to be slower moving than other variables for analyses at the weekly level.

Figure 5: Dynamic Response to Changes in VIX, weekly



Notes: This figure shows the dynamic response of various desk-level variables to a change in VIX. For this analysis, the logarithms of the Turnover, Net Position, and Limit variables are used. Furthermore, the absolute value of the net position variable is utilized. Observations are at the desk-week level. The sample covers 2016 to 2023 and is comprised of market-making desks that are classified into at least one of our seven asset categories. The sample is constant across time. Robust standard errors are clustered at the desk level. The light shaded area shows a 95% confidence band. The dark shaded area shows 1 standard error band around the estimates. The red line excludes shocks in the year 2020. *Sources*: FR VV-1, Haver, authors' calculations.

Figure 6: Log Limit Response to Modified Sharpe Ratio of New Positions, Instrumented with Changes to VIX



Notes: This figure shows the dynamic response of the log limit to a change in the modified Sharpe ratio of new positions, where this ratio is instrumented with the change to the VIX. Observations are at the desk-month level. The sample covers 2016 to 2023 and is comprised of market-making desks that are classified into at least one of our seven asset categories. The sample is constant across time. Robust standard errors are clustered at the desk level. The light shaded area shows a 95% confidence band. The dark shaded area shows 1 standard error band around the estimates. The red line excludes shocks in the year 2020. *Sources*: FR VV-1, Haver, authors' calculations.

4 Limit Reallocation between Asset Classes

In the previous section, we provided evidence for a heirarichal internal capital market structure: the market-making division receives a higher limit allocation when the bank is better capitalized, and the division in turn increases desk-level limits ex-post according to intermediation profitability. Here, we delve deeper into how this structure impacts reallocation between asset classes, which carries implications for the interconnectedness of otherwise disparate asset classes. Although these asset classes may have different fundamental characteristics, it is important to remember that all of these trading desks sit within the same trading business unit at a relatively small number of large banks. Once the trading business unit is allocated a certain amount of risk in aggregate by senior management, managers must then allocate across asset classes and desks.

Variation in limit changes across asset classes Before delving into how limits are re-allocated across asset classes, we start by examining the dimentions along which limit changes tend to vary. Table 5 shows a variance decomposition of limit changes at the desk-month-asset class level. Including bank, asset, and time fixed effects separately only explains a small amount of variation in limit changes (column 1). However, when we include bank-time fixed effects, in addition to asset fixed effects, the portion explained jumps from 2.2 percent to 34.8 percent (column 2). On the other hand, when we include asset-time fixed effects, in addition to bank fixed effects, the portion explained only rises to 5.1 percent (column 3). Including asset-time fixed effects with bank-time fixed effects only increases the percent explained from 37.9 percent to 36.1 percent (column 4). However, including bank-asset-time fixed effects does increase the percent explained to 45.3 percent (column 5). Thus, it seems as about a third of the variation in desk-month limit changes can be explained with bank-level changes over time, potentially due to changing levels of banks' equity capital, while much less can be explained by market-level shocks to an asset class. Still, a large portion of the variation in limit changes is unexplained by shocks to the bank or asset class.

Cross-desk allocation of limits We next turn to understanding spillovers between desks within a bank. For each bank, let $\mathcal{D}_b$ be the set of all trading desks, and let $\mathcal{G}(d)_b$ be the set of desks with the same asset classification (FX, equity, etc.) as desk d. We define two quantities: the total limit size of other asset classes in the same bank as desk d, and the total

Table 5: Variation in Δ Log Limit

	(1) Δ Log(Limit)	(2) Δ Log(Limit)	(3) Δ Log(Limit)	(4) Δ Log(Limit)	(5) Δ Log(Limit)
Time FE	Yes				
Bank FE	Yes		Yes		
Asset FE	Yes	Yes			
Bank*Time FE		Yes		Yes	
Asset*Time FE			Yes	Yes	
Bank*Asset*Time FE					Yes
R^2	0.022	0.348	0.051	0.361	0.453
R^2 within	39,657	39,646	39,657	39,646	38,499

Notes: This table shows a variance decomposition of limit changes. Observations are at the desk-month-asset class level. The sample covers 2016 to 2023 and is comprised of market-making desks that are classified into at least one of our seven asset categories. *Sources*: FR VV-1, authors' calculations.

limit size of other desks in the same asset class in the same bank as desk d:

$$\text{Limit of other assets in bank}_{d,b,m} \equiv \sum_{j \in \mathcal{D}_b} \text{limit}_{j,b,m} - \sum_{j \in \mathcal{G}(d)_b} \text{limit}_{j,b,m}, \tag{3}$$

$$\text{Limit of other desks in asset}_{d,b,m} \equiv \sum_{j \in \mathcal{G}(d)_b} \text{limit}_{j,b,m} - \text{limit}_{d,b,m}. \tag{4}$$

Table 6 examines how a desk's limit changes correlate with changes in these two quantities with different sets of fixed effects. We can see how when we include bank-time fixed effects, a desk's limit decreases when desks in other asset classes at the bank receive a larger limit. When we further include bank-asset-time fixed effects, a desk's limit decreases when when desks within its own asset class at the bank receive a larger limit. This provides evidence for limit constraints both at the bank level and then at the asset class level, and suggest that managers first allocate risk limits to the trading business unit as a whole, and then to asset classes within the trading unit, and then finally to individual desks. We next seek to strengthen our confidence in this mechanism with a causal identification scheme.

Causal identification of limit reallocation Consider the following example: Bank A decides to move out of the municipal bond business and drops a large number of municipal bond desks. Bank B also has a municipal bond business, as well as desks in other asset classes. In allocating more to risk to its municipal bond desks, Bank B might take risk away from its corporate bond desks. This constitutes an exogenous shock to Bank B's corporate bond business. In order to analyze this situation, we construct a leave-one-out metric that measures the net risk limit lost due to desk dropouts for all *other* banks in a given asset class.

Table 6: Limit Reallocation Within Banks

	(1) Δ Log(Limit)	(2) Δ Log(Limit)	(3) Δ Log(Limit)	(4) Δ Log(Limit)
Δ Log(Limit of other assets in bank)	0.253*	-0.236***	-0.231***	
	(0.138)	(0.055)	(0.053)	
Δ Log(Limit of other desks in asset)	0.033	-0.004	-0.004	-0.457**
	(0.022)	(0.007)	(0.006)	(0.192)
Asset*Time FE	Yes	No	Yes	-
Bank*Time FE	No	Yes	Yes	-
Bank*Asset*Time FE	No	No	No	Yes
R^2	0.116	0.343	0.357	0.491
R^2 within	0.071	0.013	0.012	0.069
N	38,462	38,460	38,460	38,499

Notes: This table shows how a desk's limit changes vary with the total limits of desks in other asset classes within the same bank and with the total limits of other desks in the same asset class within the same bank. Observations are at the desk-month level. The sample covers 2016 to 2023 and is comprised of market-making desks that are classified into at least one of our seven asset categories. Robust standard errors are clustered at the desk level. *** p<0.01, ** p<0.05, * p<0.1.*Sources*: FR VV-1, authors' calculations.

Specifically, we compute for bank b, asset class a, at month m:

$$\text{VaR of closed desks}_{b,a,t} = \sum_{\iota != b} \Delta\text{Limit}_{\iota,a,t} \tag{5}$$

We normalize this by the total risk in an asset class across banks last period. To put some numbers on the example above, imagine that Bank A is the only bank that drops municipal bond desks this period and that these desks' risk had made up 10 percent of the municipal bond market in terms of allocated VaR risk. Then our leave-one-out metric would be 10 percent for all municipal bond desks in Bank B.

Table 7: Variation in Desk Dropouts: Monthly Level

	Desk Dropouts						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Bank FE	Yes	-	-	-	-	-	Yes
Asset FE	-	Yes	-	-	-	-	Yes
Time FE	-	-	Yes	-	-	-	Yes
Bank x Asset FE	-	-	-	Yes	-	-	Yes
Bank x Time FE	-	-	-	-	Yes	-	Yes
Asset x Time FE	-	-	-	-	-	Yes	Yes
R-squared	0.000	0.000	0.285	0.001	0.584	0.347	0.652
N	9,223	9,223	9,223	9,223	9,166	9,223	9,166

What are the drivers of desk closures and associated loss in risk limits available to intermediate the market? We start with a pure statistical decomposition of lost limits using standard linear fixed effects models. Table 7 reports the R^2 associated with different fixed effects specification for our normalized lost VaR limits variable on the LHS. The table reveals that bank-time fixed effects explain a large share of variation in lost VaR with an R^2 of 58.4% (column 5). Asset-time fixed effects explain about 34.7% of the variation (column 6). All fixed effect together explain nearly two-thirds of the variation at a monthly frequency.

While the confidential nature of the data prevents us from discussing a systematic narrative analysis, we have found broad patterns based on a review of articles in business newspapers. Anecdotally, desk closures are often higher level management decisions that are typically part of a bank-specific strategy overhaul. For example, desk closures could be part of a broader bank re-organization where certain business lines are consolidated or dropped in some geographic areas, which could happen for a variety of reasons. Multiple banks over the course of our sample sought to reorient themselves towards their more stable retail banking division and desks viewed as more central to firm priorities, and this reorientation often was accompanied by job losses and trading desk closures. Additionally, firm-specific litigation and fines at times led to workforce reduction in divisions impacted by settlements. Finally, firm acquisitions can result in the closure of existing desks.

We study the effects of desk closures on key outcome variables using the following regression model:

$$\Delta y_{b,t+h} = \alpha + \beta \text{ VaR of closed desks}_{b,a,t} + \alpha_{a,t} + \epsilon_{b,a,t} \tag{6}$$

Table 8 shows the effect of desk closures and associated limit losses on the trading book at *other* banks. In response to such an exogenous loss in limits available to intermediate an asset market, turnover at other banks increases consistent with remaining desks taking part of the business. Moreover, remaining banks take larger net positions (in absolute value) and those larger positions also come with a greater VaR usage. These effects are statistically significant and also economically meaningful. In response to a one-standard-deviation change in dropped VaR, turnover at remaining desks increases by 8.8% (column 2), net positions increase by 11.4% (column 4), and VaR increases by 5.8% (column 6). Note that these effects are identified using both bank-time fixed effects and asset-time fixed effects. Thus, the identified relationship is not driven by common asset-specific market dynamics, such as trading in one asset class being more profitable, nor by common bank-time factors, such as a bank generally increasing its trading activity.

Table 9 shows the response of VaR limits in the next period to such a trading shock. The first two columns show that desks increase their VaR limits when desk closures at other banks lead to a drop in total VaR limit availability in a given asset class. Column (2) points to a

Table 8: Desk Dropouts and Turnover, Positions, and Profits in Other Banks

	Turnover		Position		Sharpe (New)	
	(1)	(2)	(3)	(4)	(5)	(6)
VaR of Closed Desks	0.080**	0.075*	0.130***	0.150**	0.008**	0.009*
	(0.032)	(0.042)	(0.044)	(0.063)	(0.003)	(0.005)
Asset*Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Bank*Asset FE	Yes	Yes	Yes	Yes	Yes	Yes
Bank*Time FE		Yes		Yes		Yes
R^2	0.13	0.30	0.10	0.27	0.65	0.72
N	7,486	7,437	7,471	7,424	7,482	7,437

Notes: Turnover and Position are in log differences. Modified Sharpe ratio is computed as income from new positions relative to value at risk. "VaR of closed desks" is standardized to have zero mean and unit variance. Standard errors are clustered at bank-asset level. *** p<0.01, ** p<0.05, * p<0.1. *Sources*: FR VV-1, authors' calculations.

sizable 15.3% increase in limits. Columns (3) and (4) show results similar to Figure 6: when modified Sharpe ratios of new positions increase (now instrumented with desk closures), risk limits increase.

In the last two columns we zoom in on the reallocation of limits within bank in response to our asset-bank-time specific trading shock. Specifically, our outcome variable in these regressions is the limit change at desk in other asset classes and banks. The significantly negative estimates indicate that a bank decreases limits in those other asset classes to redeploy capital for trading in the affected asset classes. Results are particularly pronounced when including bank-time fixed effects in column (6) and point to a 13.8% limit reduction in those unaffected desks in response to a one-standard-deviation desk closure shock.

Table 9: Risk Limit Adjustments to Desk Closures

	Risk Limit				Limit in Other Assets	
			IV	IV		
	(1)	(2)	(3)	(4)	(5)	(6)
VaR of Closed Desks	0.157***	0.153***			-0.081***	-0.138***
	(0.036)	(0.034)			(0.021)	(0.028)
Sharpe (New)			23.625*	17.511		
			(13.450)	(10.891)		
Asset*Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Bank*Asset FE	Yes	Yes	Yes	Yes	Yes	Yes
Bank*Time FE		Yes		Yes		Yes
R^2	0.13	0.43	-70.19	-47.27	0.10	0.28
F-Stat	19.35	20.21	3.09	2.59	15.21	24.66
Lower CI						
Upper CI						
N	7,338	7,295	7,317	7,276	7,338	7,295

Notes: All LHS variables are in log differences and refer to the subsequent $(t+1)$ period. "VaR of closed desks" is standardized to have zero mean and unit variance. Columns (3) and (4) are IV, using VaR of closed desks as instrument for Modified Sharpe Ratio of New positions. Standard errors are clustered at the bank-asset level. *** p<0.01, ** p<0.05, * p<0.1. *Sources*: FR VV-1, authors' calculations.

5 Conclusion

All key asset markets are intermediated by a small set of market-making dealer banks subject to endogenous risk constraints. Using supervisory trading-desk data that cover the vast majority of U.S. trading activity, this paper studies how dealer banks allocate risk limits across different trading desks and how this internal risk tolerance allocation links the liquidity of different asset markets.

A key insight from our paper is that the allocation of limits to trading desks reflects bank-wide constraints and exhibits strong spillovers across desks. Limit increases in one asset class are associated with reductions in limits in other asset classes within the same bank, suggesting that dealer banks manage a fixed risk budget across their trading businesses. Exploiting variation generated by desk shutdowns at other banks, we show that improved intermediation opportunities lead to higher limits, turnover, and positions in the affected asset class, accompanied by reductions in limits allocated to the intermediation of other asset classes within the same bank.

Our results highlight a novel channel linking asset markets through the internal capital markets of dealer banks. Because a small set of intermediaries operate across multiple asset classes, shocks affecting dealer capital or trading opportunities in one market can propagate to others through internal risk reallocation. Understanding these internal allocation mechanisms is therefore essential for assessing market liquidity, cross-asset spillovers, and the broader role of intermediaries in financial markets.

References

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

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

Anderson, Christopher S., David C. McArthur, and Ke Wang. 2023. "Internal risk limits of dealers and corporate bond market making." *Journal of Banking & Finance*, 147: 106653.

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

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

Bräuning, Falk, and Hillary Stein. 2024. "The Effect of Primary Dealer Constraints on Intermediation in the Treasury Market." *Federal Reserve Bank of Boston Research Department Working Papers*, 24-7.

Bräuning, Falk, and Victoria Ivashina. 2020. "Monetary Policy and Global Banking." *Journal of Finance*, 75(6): 3055–3095.

Breckenfelder, Johannes, and Victoria Ivashina. 2023. "Bank Balance Sheet Constraints and Mutual Fund Fragility."

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

Cetorelli, Nicola, and Linda S. Goldberg. 2012*a*. "Banking Globalization and Monetary Transmission." *The Journal of Finance*, 67(5): 1811–1843.

Cetorelli, Nicola, and Linda S. Goldberg. 2012*b*. "Liquidity management of U.S. global banks: Internal capital markets in the great recession." *Journal of International Economics*, 88(2): 299–311.

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

Du, Wenxin, Benjamin Hébert, and Amy Wang Huber. 2023. "Are Intermediary Constraints Priced?" *The Review of Financial Studies*, 36(4): 1464–1507.

Gârleanu, Nicolae, and Lasse Heje Pedersen. 2007. "Liquidity and Risk Management." *American Economic Review*, 97(2): 193–197.

Greenwood, Robin, Samuel G Hanson, and Gordon Y Liao. 2018. "Asset Price Dynamics in Partially Segmented Markets." *The Review of Financial Studies*, 31(9): 3307–3343.

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

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

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

Jordà, Òscar. 2005. "Estimation and inference of impulse responses by local projections." *American Economic Review*, 95(1): 161–182.

Li, Dan, Lubomir Petrasek, and Mary Tian. 2024. "Risk-averse Dealers in a Risk-free Market – The Role of Trading Desk Risk Limits."

Lu, Lina, and Jonathan Wallen. 2024. "What Do Bank Trading Desks Do?"

Stein, Jeremy C. 1997. "Internal capital markets and the competition for corporate resources." *The Journal of Finance*, 52(1): 111–133.

A Additional Tables

Table A.1: Relationship between VV-1 and Pillar 3 VaR

	Log(Limit Use) calculated from VV-1					Log(Percent Limit Use) calculated from VV-1				
	Credit	Interest Rate	Equity	FX	Commodity	Credit	Interest Rate	Equity	FX	Commodity
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Log(VaR) calculated from Pillar 3	0.738	0.380	0.498***	0.460***	0.527*	0.681**	0.355**	0.422***	0.225***	0.236*
	(0.564)	(0.255)	(0.112)	(0.147)	(0.316)	(0.266)	(0.148)	(0.085)	(0.061)	(0.126)
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Bank FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.139	0.317	0.941	0.382	0.325	0.289	0.236	0.877	0.242	0.217
R^2 within	0.005	0.004	0.119	0.024	0.015	0.020	0.011	0.143	0.033	0.019
N	360	571	184	441	221	360	571	184	441	221

Notes: This table shows how limit use and percent limit use at the asset class level varies with reported value at risk in Pillar 3 disclosures. Observations are at the asset class-bank-quarter level. The sample covers 2016 to 2023 and is comprised of market-making desks that are classified into at least one of these asset categories. Robust standard errors are clustered at the desk level. *** p<0.01, ** p<0.05, * p<0.1. *Sources*: FR VV-1, Pillar 3 disclosures, authors' calculations.

Table A.2: Desk Dropouts and their Impact in Other Banks in Period t

	Turnover		Gross Position		Net Position		VaR Use		Pct. Limit Use		Limit		Sharpe (Existing)		Sharpe (New)		Limit in Other Assets	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)
VaR of Closed Desks	0.080**	0.075*	0.054	0.081	0.130***	0.150**	0.024	0.024	0.002	0.004	0.022	0.020	-0.001	0.004	0.008**	0.009*	-0.001	-0.016
	(0.032)	(0.042)	(0.037)	(0.051)	(0.044)	(0.063)	(0.022)	(0.030)	(0.004)	(0.004)	(0.018)	(0.022)	(0.005)	(0.006)	(0.003)	(0.005)	(0.012)	(0.018)
Asset*Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Bank*Asset FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Bank*Time FE		Yes		Yes		Yes		Yes		Yes		Yes		Yes		Yes		Yes
R^2	0.13	0.30	0.09	0.29	0.10	0.27	0.14	0.37	0.60	0.70	0.12	0.41	0.36	0.47	0.65	0.72	0.10	0.26
N	7,486	7,437	7,436	7,368	7,471	7,424	7,471	7,426	7,504	7,456	7,504	7,456	7,482	7,437	7,482	7,437	7,504	7,456

Notes: Turnover, Position, Gross Position, VaR Usage, Limit, and Limit in Other Assets are in log differences. The Modified Share ratio (new) is computed as income from new positions relative to value at risk while the Modified Share Ratio (existing) is computed as income from existing positions relative to value at risk. Additionally, the absolute value of net position is taken. This table presents the variables utilizing their values from the present period t. "VaR of closed desks" is standardized to have zero mean and unit variance. Standard errors are clustered at bank-asset level. *** p<0.01, ** p<0.05, * p<0.1. *Sources*: FR VV-1, Haver, authors' calculations.

Table A.3: Desk Dropouts and their Impact in Other Banks in Period $t+1$

	Turnover		Gross Position		Net Position		VaR Use		Pct. Limit Use		Limit		Sharpe (Existing)		Sharpe (New)		Limit in Other Assets	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)
VaR of Closed Desks	0.084	0.112*	0.137**	0.111**	0.019	0.023	0.134**	0.141***	-0.006*	-0.001	0.157***	0.153***	0.002	0.002	-0.002	-0.007	-0.081***	-0.138***
	(0.063)	(0.065)	(0.056)	(0.050)	(0.073)	(0.086)	(0.052)	(0.048)	(0.004)	(0.004)	(0.036)	(0.034)	(0.004)	(0.006)	(0.006)	(0.006)	(0.021)	(0.028)
Asset*Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Bank*Asset FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Bank*Time FE		Yes		Yes		Yes		Yes		Yes		Yes		Yes		Yes		Yes
R^2	0.13	0.30	0.09	0.29	0.10	0.26	0.14	0.38	0.60	0.70	0.13	0.43	0.36	0.48	0.65	0.72	0.10	0.28
N	7,320	7,277	7,272	7,206	7,308	7,265	7,308	7,267	7,338	7,295	7,338	7,295	7,319	7,278	7,319	7,278	7,338	7,295

Notes: Turnover, Position, Gross Position, VaR Usage, Limit, and Limit in Other Assets are in log differences. The Modified Share ratio (new) is computed as income from new positions relative to value at risk while the Modified Share Ratio (existing) is computed as income from existing positions relative to value at risk. Additionally, the absolute value of net position is taken. This table presents the variables utilizing their values from the forward period $t+1$. "VaR of closed desks" is standardized to have zero mean and unit variance. Standard errors are clustered at bank-asset level. *** p<0.01, ** p<0.05, * p<0.1. *Sources*: FR VV-1, Haver, authors' calculations.

B Variable Summary

Table B.1: Definition of Variables

Variable	Definition
Gross Position	Gross Position is the sum of all securities and derivatives. For a given week or month the mean gross position is used. The unit of the variable is USD millions.
Income Existing Pos	The Income from Existing Positions is defined as all income derived from existing positions as noted in the profit and loss statement. Existing positions are defined as positions held by a trading desk as of the end of the prior day. The unit of the variable is USD millions.
Income New Pos	The Income from New Positions is defined as all income derived from new positions as noted in the profit and loss statement. New positions are defined as those positions newly recognized on a trading desk's profit and loss statement as a result of that day's trading activities. The unit of the variable is USD millions.
Limit	The Limit is the Value-at-Risk limit for the trading desk as expressed in USD millions.
Limit in Other Assets	The Limit in Other Assets is the ratio of the bank's total VaR limit over the bank's VaR limit for that trading desk's specific desk type. The name and description of the trading desk were utilized to determine the desk's categorization. The seven asset categories are asset-backed securities, commodities, corporate bonds, equities, FX, government bonds, and municipal bonds. The mixed desk type is characterized by a trading desk that is determined to be in more than one desk asset category.
Net Position	Net position is the absolute value of the sum of the net securities and net derivatives where net securities and derivatives are calculated as the long term positions minus the short term positions. The unit of the variable is USD millions.

Variable	Definition
Pct. Limit Use	The Percent Limit Use is defined as the ratio of the VaR Use over the VaR Limit.
Sharpe (Existing)	Sharpe (Existing) is the ratio of all existing positions over the current limit use. Existing positions are defined as positions held by a trading desk as of the end of the prior day.
Sharpe (New)	Sharpe (New) is the ratio of all new positions over the current limit use. New positions are defined as those positions newly recognized on a trading desk's profit and loss statement as a result of that day's trading activities.
Turnover	Turnover is the gross volume as measured in USD billions. For the purposes of this analysis gross volume is defined as the rolling sum of transactions from the last 30 days. For a given week or month the gross volume is amount as of the final day of the period.
VaR Use	VaR Use is defined as the total positions in USD millions that are under the provision of the Value-at-Risk limit of the trading desk.