Expertise and Firm Boundaries: Evidence from Bank Trading*

Chase P. Ross† Sharon Y. Ross‡ Kaushik Vasudevan§ Martin M. Waibel¶

August 28, 2026

Abstract

Firms must choose which activities to pursue and which to leave to competitors, but empirical examination of the determinants of organizational scope remains limited by measurement challenges. Using novel supervisory data on individual trading desks at U.S. banks, we find that, distinct from prevailing theories of diversification driving scope, banks concentrate information-processing capacity in clusters of related activities. We document that banks specialize in activity clusters where they earn higher profits and have more precise assessment models. Examining reallocations, we find that specialized labor flows to expert banks, particularly in response to changes in risk limits tied to forecasted demand. Finally, banks contract non-core activities as they approach constraints. Our results highlight the importance of labor and organizational frictions for trading.

JEL Codes: G1, G12, G14, G2, G24, L2, L22

Keywords: firm organization, firm boundaries, specialization, segmentation, trading, human capital, risk factors

*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 of the Federal Reserve System.

†Warrington College of Business, University of Florida. Email: chase.ross@warrington.ufl.edu

‡Warrington College of Business, University of Florida. Email: sharon.ross@warrington.ufl.edu

§Carroll School of Management, Boston College. Email: kaushik.vasudevan@bc.edu

¶Board of Governors of the Federal Reserve System. Email: martin.m.waibel@frb.gov

1 Introduction

Firms choose which activities to perform and which to leave to competitors. What determines this scope of activities is an important question across industries, with implications for market structure, pricing, and resource allocation. However, analyzing the determinants of scope is empirically challenging because doing so requires fine-grained data on firms' activities, and measures of their marginal and total productivity in those activities.

We study scope and its determinants in the trading activities of the largest U.S. banks using unique supervisory data. Our main dataset, which comes from regulatory form FR VV-1, covers the trading activities of banks with at least $20 billion in average gross trading assets and liabilities, which accounted for roughly 93% of total U.S.-based trading activity as of 2019. These data provide detailed daily information on banks' allocations across different activities and asset classes, and also their profits, risk limits, measures of reasonably expected near-term customer demand, and profit attribution models. These data allow us to study, at a detailed level, how banks allocate capital across various risk factors and asset classes, and how these allocations respond to customer demand and bank constraints. We link these supervisory data to resume data on the traders employed by banks. Our combined dataset further allows us to explore the link between specialized labor and firm scope.

We document four main findings. First, banks specialize in their trading activities. Second, banks earn higher returns in their areas of specialization, and have more precise profit attribution models in those areas. Third, in response to binding constraints, banks contract peripheral activities and focus on core areas of expertise. Fourth, traders tend to reallocate toward banks where their expertise is more core, particularly following shocks to expected customer demand in those areas.

Our results are a departure from the predominant diversification-based perspective on trading activities, and indicate the importance of specialized expertise and labor in the scope of trading activities banks pursue. They highlight the role of increasing returns to scale in related activities that generate knowledge spillovers to one another (or decreasing returns to scope at the bank level). When faced with binding constraints on how large they can grow and with increasing returns to scale in related areas, banks specialize in different clusters of activities. Our evidence underscores the importance of taking an organizational- and labor-based perspective on trading, and particularly highlights the role of specialized labor and its higher marginal value in specialist banks.

We begin the paper by presenting a theoretical model to organize our empirical analysis and to provide guidance on the economic mechanisms at play. In the model, a continuum of ex-ante identical banks hire labor, where a unit of labor provides a bank with more precise information about a specific cluster of factors; for example, hiring a crude oil trader improves a bank's expertise in energy futures trading. The two key features in the model are that banks face a constraint on their overall headcount, and that their labor allocations have increasing returns to scale in producing information at the cluster-level. We solve the model in two stages: first, prices are determined as

a standard noisy rational expectations equilibrium taking banks' hiring as given; anticipating the corresponding pricing equilibrium, banks make their hiring decisions.

The first prediction of the model is that in equilibrium, banks specialize in different clusters of factors, due to a combination of increasing returns to scale and constraints on headcount. The second and third predictions of the model are direct consequences of cluster-level increasing returns to scale: that specializing banks should earn higher returns in their areas of specialization, and have more accurate models in those areas. The next set of model predictions concerns reallocations of activity. In response to binding constraints, specializing banks contract their peripheral activities, which is less costly to them than contracting core activities due to scale effects. And finally, banks scale up activity more in response to shocks to expected customer demand in their areas of specialization, because their already-developed expertise makes additional scale more valuable to them.

With our theoretical model in hand, we turn to the data. Before testing the model's predictions, we present summary statistics on banks' trading. The sample spans 814 trading desks across banks, with average desk profitability of $630,000 per day and a standard deviation of $5.5 million; summed across desks in our sample, this comes to $334 million daily. The typical desk has 9 different types of risk limits, including VaR limits, position limits, sensitivity limits, and RENTD limits (limits based on reasonably expected near-term demand).

Our first main result is that banks specialize in their trading activities. By regulation, banks are required to attribute their profit and loss to a set of risk factors; illustrative examples for risk factors for an equity desk include "U.S. large cap delta," "short-term interest rate delta," and "U.S. equity index basis." Most asset classes have hundreds of risk factors, which we sort into different "clusters" of about 130 related risk factors per asset class. Analyzing these clusters, we find that they tend to be highly concentrated, with an average cluster HHI of 8,800. This concentration provides evidence of banks choosing to focus on different trading strategies. Moreover, within banks, the top 14 clusters account for 50% of a bank's absolute P&L on average, indicating specialization by banks. Directly analyzing overlap, we find that it is low in top clusters across banks, providing further evidence that banks choose to focus on different trading activities. Finally, we document that core clusters are stable within a bank over time, with AR(1) coefficients of around 0.7 at horizons of 1-50 days and 0.5 at a 750-day horizon, providing evidence for the persistence of specialization.

Our second main result is that specialists outperform non-specialists. In our preferred specification, we find that a one standard deviation increase in specialization corresponds to a 6.1% higher ratio of daily returns per unit of VaR, an economically meaningful difference. Additionally, we analyze banks' profit attribution models, which attribute the share of desk-level P&L to several components, including movements in risk factors, components not mapped to a named risk factor, and mechanical components related to carry, cash flows, and changes in reserves. We compute the *unattributable* share of P&L as the ratio of unexplained P&L to total P&L for each desk at each bank. We show that this unattributable share is strictly decreasing in the share of activity that the desk contributes to the overall bank and the tenure of the desk at the bank. These results

underscore that banks have *better* information in their areas of specialization, which translates into stronger performance.

Our third result shows that banks retrench into their core areas of specialization when constraints bind. To test this prediction, we use our unique data on risk limits at the desk level. For each limit faced by a desk, we compute the desk's proximity to that limit, and compute the maximum proximity across different limits for each desk. We then compute a bank measure of constraints as the sum of desk-level proximities, weighted by the desks' VaR. We find that as the proximity to constraints tightens, banks immediately respond by retracting from non-core areas in the near term. This result emphasizes the relatively higher marginal value of additional resources in core activities versus peripheral activities for the bank.

Our final result focuses on trader reallocations. We obtain resume data from Revelio Labs and link it to our bank-level trading data. We show that, on average, traders tend to move to banks where their expertise is closer to the banks' specialization. Moreover, we show that following positive shocks to expected customer demand in a particular area (measured from RENTD limits), specialist banks are more responsive in hiring expert traders than non-specialist banks. This result indicates that the expertise and information-processing capacity of specializing banks are embodied in human capital, and is consistent with expert human capital having more marginal value in banks that operate heavily in matching areas.

Taken together, our results push against the traditional, diversification-based perspective on trading activities, and indicate the importance of specialized expertise in banks' trading activities. Moreover, they highlight the important role of human capital for expertise and information-processing capacities in trading. Finally, we show theoretically that our results are consistent with two important frictions: increasing returns to scale in clustered activities, and constraints at the bank level in information processing capacity. This link highlights the importance of labor- and organizational frictions in financial institutions' trading.

Related literature. Our paper relates most closely to work on specialization in financial markets. Prior work empirically documents industry specialization in banks' loan-making activities (Paravisini et al., 2023; Blickle et al., 2025). In lending, specialization is naturally associated with expertise because banks' core tasks are screening borrowers ex-ante and monitoring them ex-post, which are inherently considered to be information-intensive, and subject to potentially increasing returns to knowledge. Our paper documents similar specialization in banks' trading behavior, which cuts against the traditional intuition that investors typically choose portfolios to capture the gains from diversification. Against that backdrop, our results on specialization in trading provide direct empirical evidence for the importance of informational and organizational expertise for trading, consistent with a prior body of theoretical work (Van Nieuwerburgh and Veldkamp, 2010b; Eisfeldt et al., 2023).[1]

[1]In a related vein, Kacperczyk et al. (2005) document stronger performance in concentrated industries among mutual funds and Wittwer and Uthemann (2025) document specialization by Canadian dealers in stocks, bonds, derivatives

Our results also relate to prior empirical work showing that markets are segmented, with limited participation by intermediaries in arbitrage trading, options markets, and currency markets (Siriwardane et al., 2025; Bryzgalova et al., 2025; Moskowitz et al., 2026). Specialization of the sort that we document naturally gives rise to segmentation. We empirically tie such segmentation to expertise and specialization in information production. Moreover, our theoretical results highlight that specialization can arise when related activities exhibit locally increasing returns to scale, but banks face constraints on the total amount of activity they can pursue. In this regard, our paper also relates to a literature in organizational economics emphasizing that firms' scope reflects organizational frictions, including returns to specialization, coordination costs, and span-of-control constraints (Lucas, 1978; Becker and Murphy, 1992; Garicano, 2000; Stein, 1997, 2002; Garicano and Hubbard, 2007), and emphasizes that such frictions are important for understanding financial institutions' trading.

Finally, our paper relates to work on labor frictions in financial intermediation and asset pricing. Work on intermediary asset pricing and delegated asset management emphasizes the importance of managerial incentives (He and Krishnamurthy, 2013; Cen et al., 2023), and a broader literature studies managerial skill, matching, and mobility (Berk and van Binsbergen, 2015; Berk et al., 2017; Bonelli, 2025). Distinct from the literature's focus on managers, we focus on the importance of traders. Our results provide evidence for the importance of specialized traders within banks, where different traders have different strategy-specific productivities, and where putting related traders together generates organizational scale economies. Our analysis more broadly suggests the potential importance of specialized labor for understanding financial market and asset pricing outcomes.

The rest of our paper proceeds as follows. Section 2 presents our model. Section 3 describes the supervisory data we use. Section 4 presents evidence for specialization and the returns to specialization. Section 6 relates bank activity to customer demand. Section 7 tests the model's predictions on the reallocation of capital and labor in response to demand shocks and binding constraints. Section 8 concludes.

2 Model

In this section, we present a stylized model. The model highlights the roles of increasing returns to scale in related activities and bank constraints in giving rise to market specialization, and derives empirical predictions that we take to the data.

2.1 Model Setup

Environment. There is a continuum of ex-ante identical banks indexed by $b \in [0, 1]$. There are $C > 2$ clusters (asset classes), each consisting of L risk factors; a risk factor is a portfolio of traded

corresponding to stronger trading performance. In addition to documenting a similar empirical result for bank trading, we also provide direct evidence linking such performance improvements to information and its marginal value by analyzing banks' profit attribution models, risk limits, and hiring.

assets.

Each risk factor f trades for an endogenously determined price of p_f, and has a fundamental value of $\theta_f \sim N(0, \sigma_\theta^2)$, which is independent across factors. The terminal payoff from investing in factor f is equal to $\theta_f + \eta_f$, where $\eta_f \sim N(0, \sigma_\eta^2)$, reflects unlearnable noise (for example, execution risk).

Customers exogenously demand $D_f \sim N(0, \sigma_D^2)$ of risk factor f, which is independent across factors and from fundamentals and payoffs. Bank b takes a position of $z_{b,f}$ in f, and in aggregate, banks absorb customer demand, i.e., for every f, $\int_{b=0}^{1} z_{b,f} = D_f$.

Banks' Hiring and Portfolio Choice Problems. Bank b acquires information about risk factors by choosing labor supply in each cluster c, $n_{b,c}$, subject to a bank-level constraint $\sum_{c=1}^{C} n_{b,c} \leq N$. This labor supply choice gives rise to *expertise* in cluster c,

$$x_{b,c} = \alpha n_{b,c} + \frac{\beta n_{b,c}^2}{1 + \lambda n_{b,c}}, \quad \alpha, \beta, \lambda > 0.$$

Given their hiring choice, for each factor f in cluster c, bank b produces a private signal

$$q_{b,f} = \theta_f + u_{b,f}, \quad u_{b,f} \sim N(0, x_{b,c}^{-1}),$$

where $u_{b,f}$ is independent across factors and banks, and independent from customer demand. The form of expertise indicates increasing returns to scale at the cluster level. [2]

Each bank b has identical exponential utility with risk-aversion γ. After its hiring choices and after signals are observed, denoting the vector of bank b's expectations of payoffs as $m_b = \mathbb{E}_b(\theta_1 - p_1, \ldots, \theta_F - p_F)'$, and b's posterior variance-covariance matrix of risks as Σ_b, we can write b's portfolio choice problem as:

$$\max_{z_b} \; z_b' m_b - \frac{\gamma}{2} z_b' \Sigma_b z_b, \tag{1}$$

where $\mathbf{z}_b = (z_{b,1}, \ldots, z_{b,f})$ is a vector of its positions. Anticipating this portfolio choice problem, we can write bank b's hiring problem as

$$\max_{\{n_{b,c}\}} \mathbb{E}\left[\{\max_{z_b} \; z_b' m_b - \frac{\gamma}{2} z_b' \Sigma_b z_b\}\right] \quad s.t. \quad \sum_{c=1}^{C} n_{b,c} \leq N, \tag{2}$$

where expectations ($\mathbb{E}$) are formed over future realizations of signals and customer demand.

[2] We interpret $x_{b,c}$ as capturing the idea that putting more related workers together improves expertise, especially at first, though the gains taper as the desk grows. This is in the spirit of Becker and Murphy (1992), where specialization raises productivity but is limited by coordination costs.

Equilibrium. *An equilibrium is a collection*

$$\Big(\{n_{b,c}\}_{b\in[0,1],\, c=1,\ldots,C}, \{p_f\}_f, \{z_{b,f}\}_{b\in[0,1],\, f}\Big)$$

such that:

1. *given the price system* $\{p_f\}_f$*, each bank* b *chooses* $\{n_{b,c}\}_{c=1}^{C}$ *to solve* (2)*;*
2. *conditional on its hiring choices, realized signals, and realized prices, each bank forms posterior beliefs by Bayes' rule and chooses its portfolio to solve* (1)*;*
3. *for every factor* f*, factor markets clear:*

$$\int_0^1 z_{b,f}\, db = D_f;$$

4. *the price system* $\{p_f\}_f$ *is consistent with the beliefs and portfolios induced by the equilibrium hiring profile and clears all factor markets.*

2.2 Model Solution

We solve the model first by solving the asset market equilibrium conditional on banks' hiring decisions. Then we solve for equilibrium in hiring choices, and discuss conditions under which banks specialize in different clusters. Given that the primary goal of the model is to organize our empirical analysis, we provide a brief description of the model solution here and relegate the mathematical details to the appendix.

Proposition 1 (Linear Equilibrium Pricing Function)**.** *Given banks' hiring choices, factors in any cluster where a positive measure of banks hire admits a unique linear equilibrium pricing function of the form*

$$p_f = A_c\theta_f - B_c D_f, \qquad A_c, B_c > 0.$$

Proof. All proofs in Appendix **??**. □

A key aspect of the pricing equilibrium in Proposition 1, arising from banks' first order conditions, is that banks with higher expertise in a factor take larger positions in that factor, because they face less risk. Given the linear equilibrium pricing function, which treats banks' hiring choices as exogenous, we then compute conditions under which banks specialize in different clusters.

Proposition 2 (Specialization under intermediate returns to scale)**.** *There is a* $\bar{N} > 0$ *such that, for* $N < \bar{N}$*, there are finite and positive values* $\underline{\beta}(N) < \overline{\beta}(N)$ *such that for every* $\beta \in (\underline{\beta}(N), \overline{\beta}(N))$*,*

1. *there is no equilibrium in which all banks choose the same hiring strategy across asset classes;*
2. *there is an equilibrium in which banks specialize by hiring differently across clusters.*

The proof of Proposition 2 proceeds in two steps. First, we rule out equilibria in which all banks choose the same hiring strategy across clusters. Second, we construct an equilibrium in which each bank concentrates its hiring and banks are distributed across clusters in equilibrium. The relevant equilibrium object is a symmetric mixing distribution over ex-ante identical agents, which captures how different banks specialize in different clusters in the cross-section.

The logic of Proposition 2 is as follows. The β parameter controls the returns to scale in cluster expertise. The lower bound ensures that hiring additional workers in the same cluster raises expertise quickly enough that banks prefer to concentrate hiring; the upper bound eliminates an equilibrium where all banks crowd into the same cluster because the gains from concentration are sufficiently strong. The restriction that N be sufficiently small means that hiring capacity is scarce enough for this tradeoff to matter: banks must decide where limited expertise is most valuable.[3]

2.3 Model Predictions

With our model in hand, we next turn to its predictions. We present formal prediction statements and proofs in the appendix.

Our first set of predictions concerns banks' specialization and the average performance of the banks in the areas they specialize in.

Prediction 1 (Specialization)**.** *Banks specialize across different clusters.*

Prediction 2 (Specialist Outperformance)**.** *Banks earn higher risk-adjusted returns in the clusters they specialize in.*

Prediction 3 (Specialists and Model Precision)**.** *Banks have more precise models in the clusters they specialize in.*

In the model, banks have increasing returns to scale at the cluster level, resulting in higher risk-adjusted returns in the cluster they specialize in (Prediction 2). The specific mechanism in the model is that banks' expertise takes the form of more precise information (Prediction 3), which we directly observe a proxy for given our unique supervisory data.

Our next set of predictions has to do with how banks allocate resources and reallocate them in response to demand and binding constraints.

Prediction 4 (Expected Demand and Positions)**.** *Clusters with higher expected customer demand attract more resources.*

Prediction 5 (Tightening Risk Limits and Peripheral Contraction)**.** *When banks' risk limits tighten, they contract positions in areas outside of their core expertise.*

[3]The modeling and proof approach here relates to Van Nieuwerburgh and Veldkamp (2010b), who show that specialization arises in a model of rational inattention where investors face entropy-based costs to acquire information. There, the entropy-based cost function makes investors prefer to acquire information in a concentrated area. We model information as arising from hiring, with the benefits of concentrating expertise stemming from increasing-returns-to-scale within a cluster.

Prediction 6 (Expert Labor Flows)**.** *Expert traders move to firms that specialize in their areas of expertise, particularly following positive shocks to expected customer demand.*

Prediction 4 is straightforward: we should observe more activity by banks in areas where customer demand is forecast to be higher. This prediction is particularly useful to empirically test our measures of expected demand, which we use in Prediction 6.

Predictions 5 and 6 provide further evidence of one of the core mechanisms of our model: that on the margin, capital and expert workers carry higher value to firms within their areas of specialization than outside of them. Prediction 6 also highlights the important role of expert labor as the vehicle of banks' expertise.

3 Data

We use data covering the largest banks' trading desks collected by the Federal Reserve to implement the Volcker Rule's prohibition on proprietary trading, from form FR VV-1. The data is available daily from 2014 to 2025, with the most recent data focused on large U.S. and foreign banks that have at least $20 billion of average gross trading assets and liabilities, excluding U.S. government securities. Although the data covers trading activities at only the largest banks, large banks dominate trading so the data spans roughly 93 percent of total U.S.-based trading activity as of 2019. We focus on the set of banks that consistently provide data through the full sample. We describe how we clean the data in the Internet Appendix.[4]

We use the data to form three panels that underlie the bulk of our empirical work:

1. Date × desk: Banks have trading desks, and each desk is required to provide data on its profitability, value-at-risk (VaR), and trading volume, along with the desk's names and a description of its trading activities. For example, a bank may have a desk named "U.S. equity market making" with a description "The desk makes markets in U.S. equities, indices, and delta-one linked derivatives and conducting related hedging activities" (this is not an actual desk, but an illustrative example).

2. Date × desk × risk-factor: Banks are required to attribute their profit and loss to a set of risk factors; the bank provides both the name of the risk factor and a brief description. For example, the risk factors for the equities desk example above might be "U.S. large cap delta," "short-term interest rate delta," and "U.S. equity index basis" (again, these are illustrative examples).

[4]Other academic papers that have used these data include Falato et al. (2025), who show bank curtailed equity-market trading risk after the Volcker rule; Lu and Wallen (2025), who show that bank trading desks behave like toll-taking intermediaries and bear little market risk; Barbiero et al. (2024), who show that dealer risk limit shocks affect currency market prices; and Bräuning and Stein (2024) and Li et al. (2025), who study dealer risk limits and Treasury markets.

3. Date × desk × risk-limit: Banks are required to enumerate the risk-limits that constrain their risk, giving both the limit and the current utilization of that limit. Limits vary substantially, including VaR limits or limits related to underlying risks (which can be, but do not have to be, the same as the risk factors mentioned in the previous panel). In the second half of the sample, banks were required to report additional details on the limits, including the type of limit (VaR, position limit, sensitivity, inventory age, etc.) and the source of the limit (risk appetite, regulatory capital, reasonably expected near term demand (RENTD), risk reduction/mitigation, etc.). For example, the equities desk may report both a VaR limit in dollar terms of $100, a position limit in dollar terms based on its expectation of near-term demand of |$500|, and a sensitivity limit based on an equity index basis in basis points. To make risk limits more comparable across samples and limit types, we construct a risk limit *proximity* variable, which ranges from 0 (far from limit) to 100 (at limit), and values above 100 indicate a breach; we describe its construction in the Internet Appendix.

These three datasets create three levels of data: the desk, risk factor, and risk limit data. Desks are within banks, so all the data is at the "bank-desk" level. A desk can have multiple risk factors, and a risk factor can have multiple risk limits. We will use two other groupings of the data: asset class and cluster. Desks are assigned to asset classes, and risk factors are assigned to clusters within an asset class (this data are at the bank-cluster level so that we can assess bank specialization and concentration). We use the different datasets depending on the available variables (for example, VaR and returns are not directly available for risk-factors or clusters but are available at the desk-level; we can't calculate the number of desks in a risk factor, but we could for a cluster). Next, we discuss how we sort desks into asset classes, and section 3.2 discusses how we sort risk factors into clusters.

Sorting Desks into Asset Classes Banks do not provide uniform classifications of their desks and risk factors, so comparison across banks is non-trivial. We classify desks into asset classes, and we describe in Section 3.2 how we sort factors into clusters (within an asset class).

We use a standard TF-IDF approach to classify each desk into one of ten asset classes based on the desk's name and description: credit, equities, rates, foreign exchange, securitized and structured products, collateral and financing, commodities, overlay and counterparty risk, municipal, and legacy. Each bank may have multiple desks that are classified as equity desks. We describe the asset class classification procedure in the Internet Appendix.

Returns It is not possible to directly calculate a risk factor's return because no data are reported on the size of the desk's portfolio or the share of the portfolio committed to a specific risk factor. That is, if we know a risk factor earns $100 in profits, we do not know if that is on a portfolio worth $1,000 or $1 billion. For this reason, we define returns motivated in spirit by Falato et al. (2025):

$$Return_t = PnL/\sigma(PnL)$$

where $\sigma(PnL)$ is the volatility of the risk factor's PnL over the previous three months. Henceforth, any reference to the return of a risk factor or desk refers to this return proxy.

Sorting Desks into Asset Classes Banks do not provide uniform classifications of their desks and risk factors, so comparison across banks is non-trivial. We classify desks into asset classes, and we describe in Section 3.2 how we sort factors into clusters (within an asset class).

We use a standard TF-IDF approach to classify each desk into one of ten asset classes based on the desk's name and description: credit, equities, rates, foreign exchange, securitized and structured products, collateral and financing, commodities, overlay and counterparty risk, municipal, and legacy. Each bank may have multiple desks that are classified as equity desks. We describe the asset class classification procedure in the Internet Appendix.

3.1 Summary Statistics

Table 1 reports trading desk summary statistics. Panel A summarizes profitability and VaR by asset class, spanning 814 trading desks and 1.3 million desk–trading day observations. There are 176 credit trading desks over the full sample, followed by equities (151), rates (134), and securitized/structured products (128). The remaining desk types constitute smaller shares—ranging from 26 to 82 desks—but all constitute economically meaningful shares in trading activity. On any given day, there are 100 equities desks, 83 credit desks, and 74 rates desks, among others.

Desk profitability varies substantially, averaging \$630,000 (equal-weighted across asset classes) with an average standard deviation of \$5.5 million. The average desk is profitable, with 90th-percentile profits of \$3.2 million compared to 10th-percentile profits of -\$1 million. Across all desks within an asset class, equity desks have the largest average daily profits at \$146 million (100 desks $\times$ \$1.46 million), followed by foreign exchange at \$50 million. All asset classes generate positive average profits. On average, trading desks in our sample generate about \$334 million in profits each day (summing the dot products of average desk per day and mean profitability across all asset classes), amounting to \$56 million per bank per day or \$84 billion on an annual basis across all banks in the sample. The large dollar values in the table show that the largest trading desks contribute a material share of banks' overall revenues.

The table also shows that more profitable desks tend to operate with higher VaR. Forex desks average \$7.6 million in VaR with a large standard deviation of \$15.8 million, followed by rates desks with an average \$7.4 million in VaR and standard deviation of \$15.5 million. The 10th percentile for most desks is 0, implying that desks do not always operate with material risk exposures, although most asset classes have 90th percentile VaR near \$13 million. Across all banks, total VaR is no more than \$2.6 billion (the dot product of desks per day and the mean VaR), but this overstates risk because VaR is not strictly additive as exposures across desks within a bank can offset and hedge one another.

Table 2 summarizes the risk limit panel, restricted to the post-2020 sample where banks are

required to report the type and source of each limit. The typical desk carries 9 limits on a given day, though the distribution is wide: desks at the 90th percentile face 16 simultaneous constraints. The volume of limits reflects the layered structure of risk governance at large dealers, where a single trading desk may be subject to a VaR limit, one or more sensitivity limits, a position limit, and a RENTD limit, each potentially originating from a different source within the bank's risk framework.

Panel B reports the distribution of our risk limit proximity measure, which captures how close each limit's current utilization is to the allowable bound (Appendix IA.A.2). The mean proximity across all limit-date observations is 20 percent, indicating that the typical limit has substantial headroom. The distribution varies meaningfully across limit types. Other than "other" limit types, position limits have the largest average usage (39%), followed by VaR (28%). Sensitivity and scenario limits tend to be less binding. The breakdown by limit source reveals that RENTD—which reflects the bank's ex-ante forecast of expected customer demand— is by far the most common form of analysis used to set risk limits. Figure IA.4 plots weekly averages of proximity and the share of limits near capacity over the sample period.

Panel C sharpens the picture by reporting the share of observations exceeding specific proximity thresholds. Across all limit types, only 14% of limit-days have proximity above 50 and 0.2% are above 100%. The rates differ substantially across types: position limits are above 50 on 37% of days, compared to 15% for sensitivity limits and 19% for VaR limits. By source, RENTD limits are above 50 on 13% of days, compared to 22% for risk appetite limits and 5% for risk reducing limits, although the latter two are much less common. The relatively low utilization of RENTD limits is consistent with banks calibrating expected demand for wide buffers: banks set these limits based on expected customer flow, and actual flow rarely approaches or exceeds the forecast.

Panel D shows that binding constraints are typically transient. The median near-capacity spell—defined as consecutive days with proximity above 80—lasts 2 days, with the 75th percentile at 5 days. Such short spells further indicate that limits are binding.

Panel E documents how frequently limits are adjusted, broken out by both source and type. On any given day, 0.5% of limits experience a change in size, with increases and decreases occurring at roughly similar rates. The median spell between consecutive changes is 183 days across all limit sources, and 162 days for RENTD limits, indicating that limit changes are relatively infrequent. By type, aging limits change most frequently at 1.4% of days, while VaR limits change on 0.5% of days. The variation in change frequency across sources is central to our later demand shock tests, which rely on changes in RENTD limit size as a signal that the bank's expectations about customer demand have shifted.

3.2 Sorting Factors into Clusters

We sort risk factors into clusters based on their comovement. Our analysis in the paper focuses on clusters of related activities, which allow us to understand which desks and banks specialize in which risks and to what extent those risks are correlated.

We cluster risk factors within each asset class and across banks. Suppose Bank A has 2 equity desks and 3 rates desks, and Bank B has 1 equity desk and 4 rates desks. We would form clusters using the 3 equity desks and, separately, the 5 rates desks. Conceptually, this would allow us to capture a "market" cluster in each asset class, and how much each bank and desk is exposed to that market risk. For example, clustering may yield distinct market clusters for equities and commodities. Bank C might have 4 equity desks. The first desk may have 1 risk factor, the second 2, the third 3, and the fourth 4. Clustering could result in the last two equity desks each having a risk factor assigned to the market cluster, whereas the other two equity desks' risk factors are assigned to different clusters.

We cluster each desk's reported risk factors within each asset class using the correlations of the risk factors' three-month return proxies. We adopt a two-stage taxonomy and assignment procedure to handle the disparate lifespans of trading strategies. In the first step, we identify the stable structure of the market. We restrict our sample to "stable" risk factors with at least a year of data and compute pairwise correlations only when two factors have at least a year of overlapping observations. We apply hierarchical agglomerative clustering using average linkage and a distance metric of one minus the absolute correlation motivated by Jensen et al. (2023). We define clusters by cutting the resulting dendrogram at a distance threshold equivalent to a minimum within-cluster absolute correlation of 0.10. This choice guarantees that stable risk factors that comove are grouped into the same cluster. We will shortly discuss the robustness of the 0.10 threshold choice.

In the second stage, we populate these clusters with short-lived risk factors (those with at least 10 daily observations). Because short-lived factors may not overlap temporally with a cluster's median or centroid, we employ a best overlap member approach: a short-lived factor is assigned to a cluster if it exhibits statistically significant comovement with any stable member of that cluster with which it overlaps. For each assignable factor, we attempt to match it to an existing cluster based on realized comovement: within each candidate cluster, we identify the cluster member with which the short-lived factor has the largest temporal overlap and compute the correlation over that overlap. We assess significance using a Fisher z-test against the null of zero correlation, while also requiring the correlation magnitude to meet the cluster's minimum threshold (0.10). Factors that fail to match any existing cluster are labeled as unclassified, representing idiosyncratic or transient risks that do not statistically significantly load on underlying factors. This two-stage approach has the advantage of allowing us to include a wide set of risk factors, but relying on the longer-lived factors to create the correlation-based clustering, where those correlations are better estimated relative to the correlations estimated for the shorter-lived factors.

For exposition, we then label each cluster using TF-IDF on the risk factors' names and descriptions, which provides interpretable labels such as "sovereign/credit/european" that capture the shared characteristics within each cluster. We find that the text-based cluster labels are relatively poor descriptors because many risk factor names and descriptions provided by the firms are too broad or general. Therefore, we find that the best information about the factor stems from its

realized profits. The process yields risk factor classifications that allow us to compare trading strategies across banks, which we plot in Figure 1 for the 10 largest risk factors in each asset class.

Table 3 reports factor summary statistics. Panel A shows the characteristics; most asset classes have hundreds of risk factors, and the average desk is exposed to 8 factors on a given day. Given that each desk is exposed to several risk factors, the average profits attributable to a factor are small, ranging from -$21,000 (securitized and structured products) to $34,000 (equities). The risk factors for equities and rates tend to have larger tails, with 90th-percentile profit of $400,000 in equities and $363,000 in rates.

Table 3 Panel B shows the summary statistics for the clusters to which we sort risk factors. The average asset class has about 130 unique clusters. The number of clusters aligns with our priors about the relative complexity of asset classes: municipal bond trading has the fewest clusters, about 30, whereas credit and equities both have the most, given the wide range of tradable products in those asset classes. Again, the stability in the number of clusters across the two subsamples shows that the estimation is consistent and stable. The last two columns show that even though there are many clusters, most have more than a single risk factor. On average, 6 percent of clusters are singletons.

The large number of distinct clusters suggests that many of the factors are esoteric. Figure 2 confirms this by calculating the average correlation within an asset class and the average correlation of risk factors across asset classes: virtually none of the buckets have a meaningful average correlation coefficient, even when looking at the average correlation of risk factors in the same asset class. This likely reflects the wide range of risk factors, many of which are both esoteric and small. The figure also shows the average correlations if we limit to the largest cluster (or clusters) in each asset class as measured by the sum of profit volatility across the factors in the cluster, and the correlation coefficients now become much larger, partly because we cluster using correlations. Within each asset class, a small set of clusters is widely traded, and these widely traded factors tend to be the largest in terms of profitability volatility. However, the long tail of smaller factors tends to be relatively more esoteric.

Table 4 studies how well common aggregate measures explain the trading desks' returns with Fama and French (1993) style regressions against aggregate indices. We calculate an aggregate return index from the risk factors using either equal-weighted returns or value-weighted returns, weighted by the standard deviation of the previous month's PnL. We calculate analogous indices for each asset class. The first two columns show the mean and standard deviation of these indices; recall, because we do not observe the portfolio sizes, we infer returns by comparing the daily PnL relative to the volatility of PnL over the previous three months, hence, while they are return proxies, they are unitless and not directly comparable to percentages. The aggregate value-weighted index has a daily return proxy of 0.005, meaning that the average PnL is 0.005 times its volatility over the past three months, with a standard deviation of 0.067.

The remaining columns in Table 4 show that aggregate measures do a poor job of explaining

variation across the risk factors. Columns (5)–(7) show the median beta, t-stat, and R^2 across risk factors in an asset class when regressing that risk factor's return proxy on the value-weighted index. Rates have the largest beta to the index, 0.38, but R^2 is only 0.25% (not 25%), low relative to Fama and French (1993)'s R^2s for equity portfolios in the range of 70 to 90%; and the t-stat shows that the average risk factor has no significant relationship with the index. The result does not change if we use value-weighted indices constructed using only the factors within an asset class. Even with these more targeted asset-class-specific indices, betas are largely less than 0.2, and R^2 is still well below 1%. The last three columns show that the standard Fama–French market index does no better at explaining risk factors' variation, with all betas near zero—even for risk factors at equity desks.

Sorting Robustness Two choices are particularly important in our risk factor sorting. First, we focus on the absolute value of correlations because we are interested in whether a desk trades a specific risk factor, not whether they are long or short the factor. Trading desks can typically switch between long and short expressions of a risk factor depending on their preference to hedge, make markets, manage their inventory, or satisfy changing client demand. A signed-correlation sorting procedure would erroneously separate desks active in the same risk but trading on opposite sides.

Second, we choose a correlation threshold of $\rho = 0.10$. At face value, this seems like a relatively low value. We choose the cutoff by balancing the statistical power of identifying true non-zero relationships against economic intuition about the number of clusters that are reasonable within an asset class. Figure 3 and Table 5 show the tradeoff: equities desks have about 1,500 risk factors, sorting into 273 clusters with the 0.10 threshold. But if we increase the threshold to 0.15, there are 420 clusters; to 0.25 there are 707, and if we impose strong comovement of $\rho = 0.75$ there are 988 clusters—a large share of which are singleton clusters containing a single risk factor. We choose 0.10 because it balances the number of clusters with correlations that are statistically different from zero, since a correlation of 0.10 is roughly equal to 1.6 standard deviations: under the null of zero correlation and $n = 252$, then $\mathbb{E}[\rho] = \sqrt{2/\pi}\sqrt{n-1} \approx 0.05$, so $|\rho| = 0.10$ is roughly twice the noise floor we would expect from random variation alone.[5]

The immediate concern is that correlations of 0.10, even if yielding a reasonable number of clusters, are insufficiently large to correctly cluster similar risks and might errantly cluster together risk factors that are spuriously correlated. Figure 4 plots the distribution of correlations across pairs of risk factors within an asset class, comparing risk factors within and between clusters. The mass of the distribution for the between clusters spikes just above 0, showing that the correlation of risk factors in the same asset class but different clusters is not different from zero (since we take absolute values, uncorrelated pairs will necessarily exhibit a small positive average correlation). The mass of within cluster correlations is shifted well to the right, as expected, but importantly shows that even though 0.10 is the threshold there is a nontrivial share of correlations that are well above

[5]Figures IA.2 and IA.3 plot additional summary statistics of the correlations of risk factors across banks, asset classes, and clusters when clustering based on either the absolute value or signed value of the 90-day return proxy.

it. The mean correlation within a cluster is 0.17, with a 75th percentile of 0.30. By contrast, the average correlation of risk factors in the same asset class but different clusters averages 0.04, with a 75th percentile of 0.09.

We more formally test whether the correlation-based clustering correctly captures persistent comovement relative to noise around the threshold in Table 6. Since the clusters were formed based on correlation, it is partly mechanical that the average correlation within cluster is higher than across clusters. Instead, we show that the separation in correlations within versus across clusters is quantitatively strong and replicates across different subsamples. We summarize this separation using the area-under-the-curve (AUC) defined as $Pr(|\rho|_{within} > |\rho|_{between})$. An AUC of 0.5 implies no separation, while higher values show that pairs with higher correlations are consistently concentrated within a cluster. In the full sample, the average AUC across asset classes is 0.88, confirming that any given within-cluster pair has higher comovement than a randomly chosen between-cluster pair 88 percent of the time.

To ensure this pattern reflects meaningful separation rather than noise around the threshold, we compute the AUC using only odd or only even trading days by recomputing the (absolute) correlation from the filtered subset (e.g., ρ calculated only using odd days). If clustering were primarily driven by sampling noise around the threshold, then we would expect the within-cluster advantage to attenuate when randomly splitting the sample. The data reject this: the AUC for both odd and even days remains well above 0.5 and similar to the full-sample values. Moreover, the table shows the hit rate—defined as the percent of pairs with correlations exceeding the 0.10 threshold—to check the stability of the clustering around the threshold. If many within-cluster pairs are only slightly above 0.10 due to sampling error, then we would expect them to frequently fail to reach the threshold in the subsamples. Within-cluster pairs clear the threshold in both subsamples 42 percent of the time versus only 3 percent for between-cluster pairs, implying that the within-cluster pairs are about 14 times more likely to exhibit stable, above-threshold relationships across the subsamples compared to between cluster pairs. Together, these results show that the clustering captures persistent return comovement rather than spurious correlations near the threshold.

4 Banks Specialize

With our clusters of risk factors in hand, we turn to our model's empirical predictions. Section 4.1 documents specialization at the cluster level. Section 5 then asks whether banks earn higher risk-adjusted returns and have better risk-model performance in asset classes where they are more specialized. Together, the results suggest that the clustering of intermediation reflects real informational advantages rather than institutional inertia.

4.1 Concentration: Bank Activity Clusters in a Core (Prediction 1)

Prediction 1 states that trading desks specialize in different risk factors.

Such specialization is well understood by market participants: several of the largest U.S. dealers have reputations as specialists in particular segments. For example, some banks' indices serve as blue-chip indices for certain markets. Goldman Sachs is viewed as a commodities specialist, an expertise dating to its acquisition of the large commodities trading brokerage J. Aron and reflected in the leading commodities index, the GSCI (originally named the Goldman Sachs Commodity Index). In emerging-market sovereign debt, the J.P. Morgan EMBI family serves as the canonical hard-currency benchmark. Morgan Stanley has a reputation for equities intermediation, spanning its equity and equity derivative intermediation, equity financing, and equity capital market underwriting businesses.[6]

To test Prediction 1, we define a bank as specialized if its risk-taking is concentrated in a bank-specific core set of clusters and the allocation is stable over time. We define specialization using three allocation-based components:

1. Concentration: bank b's activity is disproportionately concentrated in a small set of clusters of risk factors (its "core").

2. Persistence: core clusters are stable over time.

3. Heterogeneity: core clusters differ across banks.

Throughout the paper, we measure a bank's specialization as the share of its activity allocated to a given unit of risk. Depending on the outcome variable being studied, the unit is a risk factor, a trading desk, or an asset class; we use the finest unit at which the relevant outcome is observed. For bank b and unit ℓ on date t, define

$$\text{Share}_{b,\ell,t} \;=\; \frac{|\text{P\&L}|_{b,\ell,t}}{\sum_{\ell' \in b} |\text{P\&L}|_{b,\ell',t}}, \tag{3}$$

where $|\text{P\&L}|_{b,\ell,t}$ is the absolute value of profit and loss observed at unit ℓ, and the denominator sums across all of bank b's units of the same type. Our specialization measure is the trailing average of this share, standardized to have zero mean and unit variance:

$$\text{Specialist}_{b,\ell,t} \;=\; \text{Z}\left(\frac{1}{90} \sum_{s=t-90}^{t-1} \text{Share}_{b,\ell,s} \right). \tag{4}$$

The window excludes date t, so the measure is predetermined relative to outcomes on that date. When we work at monthly frequency, as in the labor market tests, we construct the analogous measure using the trailing three months. A one-unit increase in $\text{Specialist}_{b,\ell,t}$ therefore corresponds

[6]Morgan Stanley's CEO recently said: "We have talked for many years about what it takes to run a global fully laid out equities business. The old nine boxes where you have Prime Brokerage, derivatives, and importantly cash, and then the three regions [Americas, Asia, EMEA]. And the cost to run that business, just dollars in the ground to turn on the lights every year, those costs, as you know, have only gone up. And so not surprisingly, the top of the competitive heap is able to endure operating leverage and above-market returns, and the rest have to slug it out just to sort of make the nut."

to a one-standard-deviation increase in the bank's trailing allocation to that risk factor, desk, or asset class.[7]

First, we test whether risk factor returns are more correlated when they are traded by the same bank, the same asset class, or the same cluster. The model suggests that banks specialize in specific risks, and if banks specialize, the risk factors a given bank trades should be more correlated with one another than with the risk factors traded at other banks. We calculate the correlation of returns ρ_{ij} for all pairs of risk factors i and j and estimate the regression

$$\begin{aligned} |\rho_{ij}| \times 100 = \beta_0 &+ \beta_1 \mathbb{I}_{\text{SameBank}_{ij}} + \beta_2 \mathbb{I}_{\text{SameCluster}_{ij}} + \beta_3 \mathbb{I}_{\text{SameAC}_{ij}} \\ &+ \beta_4 (\mathbb{I}_{\text{SameBank}} \times \mathbb{I}_{\text{SameCluster}})_{ij} \\ &+ \beta_5 (\mathbb{I}_{\text{SameBank}} \times \mathbb{I}_{\text{SameAC}})_{ij} + \varepsilon_{ij} \end{aligned} \tag{5}$$

Table 7 column 1 shows that risk factor returns within a bank are more highly correlated. Same-bank pairs have 11 percent higher absolute correlation than cross-bank pairs (1.00/8.78). Column 6 estimates the full specification: risk factors in the same bank, cluster, and asset class have 3.6 percentage points higher absolute correlation than risk factors in the same cluster and asset class but at different banks.

The correlation patterns are consistent with specialization but could also reflect bank-specific shocks that affect all risk factors within a bank, such as common funding shocks or shared risk models. We distinguish these explanations by examining the cross-sectional heterogeneity and persistence of bank allocations below.[8]

4.2 Heterogeneity: Core Clusters Differ Across Banks

We would see similar correlation patterns if banks allocated similarly across clusters but faced common bank-level shocks or model adjustments. Instead, we find that most clusters are dominated by a small set of banks, typically only one or two banks. We estimate the concentration of clusters by computing the cluster Herfindahl-Hirschman Index (HHI). For each cluster and date, we compute a HHI of bank shares, where a bank's share is the bank's fraction of the cluster's total P&L volatility (standard deviation of P&L). [9]

Figure 5 shows the average Cluster HHI with clusters grouped by asset class and sorted by size

[7] In the labor market tests, where the unit is a trader's asset class, we classify an asset class as *peripheral* at a bank if $\text{Share}_{b,a,t}$ falls below the bank's within-month median across asset classes, and as *core* otherwise.

[8] Table IA.5 shows the analogous table, but when clustering on signed returns (rather than the absolute value of returns). The within-bank correlation in Table IA.5 is consistent with banks trading both sides of that risk rather than taking directional bets, so risk factors for the specializing bank tend to hedge one another.

[9] We sum the trailing 90-day standard deviation of P&L across risk factors within each bank-cluster-date, compute each bank's share of the cluster total, and sum squared shares. We use the standard deviation of P&L as a proxy for value-at-risk, which banks report at the desk level but not the risk-factor level. Table IA.1 shows that desk-level VaR and the trailing 90-day standard deviation of desk P&L are strongly related, with a coefficient close to one. Because standard deviations are not additive across risk factors, the summed measure is an upper bound on a bank's activity in the cluster; the results are similar if we instead measure activity with absolute P&L.

within each row. Cluster size is measured as total P&L volatility summed across banks, so the largest clusters in each asset class appear on the left. These large clusters also tend to include more banks, as indicated by their lighter shading.

Table 8 reports summary statistics for the cluster-date panel. The average cluster HHI is close to its theoretical maximum, at which a single bank accounts for all of a cluster's risk. Figure IA.5 shows that many clusters have a HHI close to the maximum, which would mean the cluster is dominated by one bank. The heterogeneity we observe reflects natural specialization boundaries. Concentration declines with cluster size. Table 8 shows a strong negative correlation between the cluster HHI and the size of the cluster, measured by the log of absolute P&L. Figure IA.6 plots the relationship across clusters, averaging each cluster's HHI over time: larger, more active clusters attract more participants. Banks specialize in distinct niches, with only the largest clusters supporting multiple specialist participants, and the typical bank may be served by a single specialist bank.

Thus far, we have shown that clusters are dominated by a small set of banks. But each banks' activity is concentrated in a small set of clusters. Figure 6 shows the cumulative share of a bank's activity accounted for by its top-ranked clusters, where clusters are ranked within each bank form largest to smallest by the same activity measure and shares are aggregated over the sample period. We consider three measures of activity: absolute P&L, profits (positive P&L), and losses (negative P&L). Across all three measures, concentration is steep: the top 14 clusters accounts for half of a bank's absolute P&L, and 67 clusters account for 90% (vs. 206 clusters total). The pattern is similar for profits and losses considered separately, with 8 and 35 clusters reaching the 50% and 90% thresholds for profits, and 8 and 34 for losses, so concentration is not a result of a few large gains or losses.

The identity of these core clusters, moreover, differs across banks. Figure 7 shows that the concentration of activity in a few top-ranked clusters is bank-specific rather than system-wide: banks are not all concentrated in the same clusters. For each pair of banks, we compute the Jaccard similarity of their top-ranked cluster sets, the ratio of the number of clusters in both banks' sets to the number in either set. We rank each bank's clusters by total absolute P&L over the sample period and define its top-ranked cluster set as its K largest clusters. The Jaccard similarity is the nubmer of clusters in boht banks' sets divided by the number in either set, so it equals one for identical sets and zero for disjoint sets. Figure 7 reports the pairwise similarities for $K = 3$ and $K = 10$. The median pairwise Jaccard similarity is low, indicating that banks largely specialize in distinct clusters.

4.3 Persistence: Core Allocations are Stable Over Time

Finally, the allocations are stable over time. For horizons τ ranging from [TK one month] to [TK several years], we regress each cluster's share of the bank's activity, $\text{Share}_{b,c,t}$ from equation (3), on

its lagged value:

$$\text{Share}_{b,c,t} = \alpha + \beta \, \text{Share}_{b,c,t-\tau} + \varepsilon_{b,c,t}. \tag{6}$$

Figure 8 plots the coefficients β_τ against the horizon τ. Clusters' shares of bank activity exhibit strong persistence: the coefficient is 0.7 at short horizons (1-50 days) and remains at 0.4 even at horizons of around 750 days. Banks maintain stable allocations to their core clusters rather than churning across risks.

These results establish the three components of specialization: bank activity is concentrated in a small set of clusters (Figures IA.5 and 6), the core clusters differ across banks (Figure 7), and allocations to them are persistent (Figure 8).

4.4 From Clusters to Asset Class Specialization

We define specialization as persistent, bank-specific concentration in a core set of clustered risk factors; and we have documented specialization using cluster-level data. Several of the outcome variables we study below—risk-adjusted returns, P&L attribution, and value-at-risk—are observed at the desk rather than the risk-factor level, so the corresponding tests use desks grouped into asset classes. We verify that the specialization facts hold at this coarser level.

The facts do hold, though naturally in attenuated form: asset classes are coarse aggregates of many clusters, so a bank can dominate a cluster while holding a modest share of the asset class. Table IA.2 shows that the average asset class HHI is around 3,000, which would be equivalent to roughly, equivalent to 3 banks equally participating in the average asset class.[10] The asset class HHI is negatively correlated with activity in the asset class, as at the cluster level (Figure IA.7). Table IA.3 shows the HHI by asset class. Most asset classes have HHI in the 2,000 to 2,500 range, indicating broad participation across the covered banks; repo is the exception, with an HHI indicating that only a couple of banks are large players in that market. Figure IA.8 and Figure IA.9 mirror the cluster-level results: a few asset classes account for most of the bank's P&L, and an asset class's share of bank activity is persistent.

Our specialization measure is defined *within* the bank: it asks how much of a bank's own activity sits in an asset class. An alternative notion, following Blickle et al. (2025), defines specialization *across* banks: a bank is specialized where its portfolio weight in an asset class exceeds the market-wide weight. [11]

Specialized banks have more activity in the asset class compared to the market-wide asset class share. Figure IA.10 shows that the two measures are closely related: a regression of one on the other yields a significantly positive coefficient. Because the two notions are conceptually distinct—a

[10] The asset class HHI is constructed analogously to the cluster HHI, with a bank's share measured by the sum of value-at-risk across its desks in the asset class.

[11] We measure the Blickle et al. (2025) analogue as $\frac{Activity_{b,a,t}}{\sum_a Activity_{b,a,t}} - \frac{Activity_{at}}{\sum_a Activity_{at}}$, where activity is desk-level value-at-risk to the bank–asset-class level.

small bank could allocate most of its activity to an asset class in which it holds a trivial market share—we use the across-bank measure as a control in performance regressions below, to verify that our results reflect within-bank allocation rather than market share.

5 Specialist Outperformance (Predictions 2 and 3)

Having documented that banks specialize, we ask whether specialization is associated with better performance. We test whether desks perform better in asset classes where they are more specialized in two ways. First, we ask whether specialist desks earn higher risk-adjusted returns. Second, we ask whether banks' risk models better explain P&L at specialist desks.

5.1 Specialists Earn Higher Risk-Adjusted Returns (Prediction 2)

We measure specialization at the desk level using $\text{Specialist}_{b,d,t}$ from equation (4), where the unit is the desk: desk d's trailing 90-day average share of bank b's total absolute P&L, standardized to a z-score. Table 9 shows that desks earn higher risk-adjusted returns in asset classes where their bank is more specialized. We estimate

$$r_{b,d,t} = \beta \, \text{Specialist}_{b,d,t} + \Gamma' \mathbf{X}_{b,d,t} + \alpha_{at} + \delta_{bt} + \varepsilon_{b,d,t},$$

where $r_{b,d,t}$ is the risk-adjusted return, estimated by desk d's P&L divided by its VaR on date $t-1$, winsorized at the 5th and 95th percentiles. α_{at} and δ_{bt} denote asset-class-by-date and bank-by-date fixed effects, and standard errors are two-way clustered by date and bank–asset class.

Our preferred specification is column 4, which includes asset-class-by-date fixed effects and bank-by-date fixed effects. The asset-class-by-date fixed effects ensure that we compare returns across desks within the same asset class on the same date, absorbing common returns to a particular asset class on a particular day. The asset-class-by-date fixed effects ensure that we compare returns across desks within the same asset class on the same date, absorbing common returns to a particular asset class on a particular day. The bank-by-date fixed effects absorb bank-specific daily profitability, ensuring that the coefficient does not reflect specialist banks happening to have good days for unrelated reasons. These fixed effects isolate a narrow and clean test: within a given bank on a given day, desks in which the bank specializes earn higher risk-adjusted returns than desks in which the bank does not specialize. Across all specifications, greater specialization is associated with higher risk-adjusted returns. In our preferred specification, a one-standard-deviation increase in specialization is associated with a 6% higher daily return. To put the magnitude in context, the daily coefficient of 0.06 corresponds to an annualized risk-adjusted return difference of approximately 95 percent ($0.06 \times \sqrt{252}$), suggesting specialists earn an economically meaningful premium.

Columns 5 and 6 provide within-desk robustness. Column 5 adds desk fixed effects to the asset-class-by-date specification, and this helps absorb time-invariant desk characteristics including

the desk's bank, asset class, and unobserved attributes like management quality. Column 6 further includes bank-by-date fixed effects, which helps to account for bank characteristics like bank size. The coefficient is identified from within-desk variation over time: the same desk earns higher risk-adjusted returns in periods when it accounts for a larger share of its bank's activity. Column 7 adds the across-bank specialization measure of Blickle et al. (2025) described in Section 4.1. The Specialist coefficient is essentially unchanged, indicating that the premium reflects the desk's position *within* its bank rather than the bank's market share in the asset class. Desks earn higher returns when the desk is more central to the bank's activity.

5.2 Specialists Have More Accurate Risk Models (Prediction 3)

We test the first attribution implication of (Prediction 3) and exploit the structure of the trading desk data to test whether banks have better P&L attribution at desks where they are more specialized. Attribution data are observed at the desk rather than risk-factor level, so specialization is measured at the desk level using equation (4).

The trading desk data requires banks to attribute their profitability to several components, including movements in named risk factors, and P&L not mapped to any named risk factor, along with other mechanical components related to carry, cash flows, and changes in reserves. For each desk-day, we construct an *unattributable share*, the ratio of unexplained P&L to total valuation-related P&L from existing positions:

$$\text{Unattributable Share}_{b,d,t} = \frac{|\text{Other Unattributed Profit \& Loss}_{b,d,t}|}{|\text{Risk Factor P\&L}_{b,d,t}| + |\text{Carry P\&L}_{b,d,t}| + |\text{Other Unattributable Profit \& Loss}_{b,d,t}|} \tag{7}$$

The ratio is bounded between 0 and 1, with smaller values indicating that the desk's risk factor model more fully describes its profitability. The median unattributable share is 0.2%, indicating that desks' risk factor models describe total P&L well.[12]

Figure 9 provides a first look. The left panel plots the unattributable share against $\text{Specialist}_{b,d,t}$: more specialized desks within a bank have lower unattributable shares. The right panel plots the desk's tenure, a complementary proxy for accumulated expertise. We note the natural selection bias: banks will close unprofitable desks. We measure tenure as the cumulative number of trading days on which the desk reports P&L, expressed in years, excluding desks present on the bank's first reporting date, as those desks existed before the data collection began (the history is left-censored). Longer-lasting desks have lower unattributable shares. The steady decline is consistent

[12] "Other Unattributed Profit and Loss" (Item 10) is "any other profit and loss on existing positions that is not reported [in the other items]." "Risk Factor P&L" (Item 4) is "... the profit and loss from existing positions due to changes in the specific risk factors monitored and managed as part of the trading desk's overall risk management policies and procedures ...". "Carry P&L" (Item 7) is "... the profit and loss due to changes in carry. Generally, this item includes funding costs ..."

with learning-by-doing: a desk's risk factor coverage improves as it accumulates experience.

Table 10 presents the corresponding regression estimates. The dependent variable is the desk's unattributable share; columns 1 to 3 use our main definition and columns 4 to 6 a wider definition that adds $|\text{Other Attributable P\&L}_{b,d,t}|$ to the denominator.[13] Column 1 includes bank and date fixed effects, and column 2 replaces the date fixed effects with asset-class-by-date fixed effects. Our preferred specification is column 3, which includes asset-class-by-date and bank-by-date fixed effects. The asset-class-by-date fixed effects remove the possibility that some asset classes experience hard-to-explain returns across all banks, and the bank-by-date fixed effects absorb variation in banks' overall model quality. The two sets of fixed effects isolate a within-bank-day comparison: on a given day, does a bank's risk model explain a greater share of P&L at its specialist desks than at its peripheral ones? A one-standard-deviation increase in specialization is associated with a reduction in the unattributable share by half, and the estimates are similar under the wider definition. Standard errors are two-way clustered by date and bank–asset class. Table IA.4 shows the results are similar if we include a control for for specialization across banks, as in Blickle et al. (2025).

The unattributable share measures how well a desk's risk factor model describes its profitability, but it doesn't describe the underlying informational advantage. In theories of information choice, investors with limited capacity decide which risks to learn about, and increasing returns lead them to concentrate learning on the assets they specialize in (Van Nieuwerburgh and Veldkamp, 2010a; Kacperczyk et al., 2016). In these models, specialization is not just a portfolio tilt but an allocation of scarce attention: specialists monitor more signals about their core risks and consequently face lower posterior variance. The granularity of a desk's risk model provides a direct observable for this allocation. Desks choose the set of risk factors they monitor and manage, and a desk that decomposes its exposures more finely is devoting more of the bank's information-processing capacity to a more detailed map of those risks. We measure model granularity as the number of distinct risk factors to which the desk attributes nonzero P&L on a given day.

Table 11 reports the results. The first three columns regress the log number of risk factors on specialization with increasingly demanding fixed effect specifications. The asset-class-by-date fixed effects are important because factor counts differ mechanically across asset classes—a rates desk monitors points along entire yield curves, while an equity desk may track a handful of names—so the comparison must be made within an asset class at a point in time. Our preferred specification in column 3 adds bank-by-date fixed effects, which absorb differences in how banks' risk systems name and subdivide factors, so the coefficient is identified from within-bank-day comparisons: a bank's specialist desks decompose their risks into more factors than its peripheral desks. Standard errors are two-way clustered by date and bank–asset class. The estimate implies that a one standard

[13] "Other Attributable Profit and Loss" (Item 5 in the FR VV-1 data) is "... any portion of the daily profit and loss produced by revaluing the positions held at the end of the previous day using the market data at the end of the current day that is not specifically attributed to changes in risk factors monitored and managed as part of the trading desk's overall risk management policies and procedures."

deviation increase in specialization is associated with a 12% increase in the number of monitored risk factors ($(e^{0.113} - 1) * 100$).

The final three columns connect model granularity to model accuracy. Column 4 repeats our earlier unattributable share regression on the matched sample, column 5 adds the log number of risk factors, and column 6 includes the factor count alone. Finer factor decomposition is associated with lower unattributable shares. Part of this relationship is close to mechanical—P&L attributed to a named factor is, by construction, not unattributed—but that is precisely what makes the variable a good measure: it isolates the component of attribution quality that comes from the granularity of the risk model itself. The economic content is in the specialization coefficient, which attenuates when the factor count enters. Specialists' superior attribution operates in part through maintaining more granular risk models. Viewed through the lens of information choice, the factor count captures the attention banks allocate to a risk, and the unattributable share captures the posterior uncertainty that remains. Specialists monitor more, and what they monitor, they explain.

6 Demand and Allocation of Resources (Prediction 4)

The previous sections characterize the cross-sectional organization of intermediation and show that it reflects genuine expertise. We now ask how this structure responds to changes in the economic environment. The model predicts that higher demand for a risk factor attracts more intermediation capacity and that, conditional on a demand shock, specialists respond more aggressively than non-specialists. This section tests both predictions. We first show that factors with greater demand are served by more banks and desks, consistent with the model's entry margin. We then examine how existing participants respond to demand shocks and find that specialists increase their positions by more, consistent with the intensive margin channel in which lower posterior variance makes each unit of spread more valuable to a better-informed bank.

6.1 High Demand Clusters Attract More Desks and Banks (Prediction 4)

If specialization reflects banks' equilibrium allocation of limited trading capacity, then high-demand risks should attract more activity. We test this prediction by studying whether higher demand corresponds to more desks and banks subsequently participating in the cluster.

We estimate demand using four proxies. The first two are based on realized trading activity: the trailing three-month average of daily total absolute risk factor P&L across all banks in the cluster ("activity level") and its standard deviation over the same window ("activity volatility").

The last two measures exploit the richness of the Volcker data's risk limit reporting; banks often report limits on specific risk factors. We isolate sensitivity-based limits denominated in USD that are based on the banks' internal forecast of RENTD, or reasonably expected near-term customer demand. A single risk factor can be subject to many limits: for example, a risk factor might have a limit on its delta exposure, its vega exposure, and a stress scenario loss, each denominated in

USD but measuring a distinct dimension of the bank's expected customer-facing behavior. We measure limit-derived demand capacity as the maximum absolute value of the upper and lower bounds. We then aggregate to the cluster-date level in two stages. In the first stage, within each bank-cluster-date, we compute both the average limit size across the bank's limits and the total limit size across those limits; averaging within the bank ensures the first measure is not mechanically inflated for banks with more granular limit structures. In the second stage, we average each measure across banks within the cluster-date, weighting banks equally, so that the cluster's demand reflects the typical bank's assessment rather than that of the largest banks. The resulting proxies—the average RENTD limit and the total RENTD limit—capture the scale of customer demand the typical bank expects in the cluster along each risk dimension and in aggregate, respectively.

For cluster c in asset class $a(c)$ on date t, we estimate

$$\textit{Number of Desks}_{c,(t,t+h]} = \beta\, \textit{Demand}_{c,t} + \alpha_c + \gamma_{a(c),t} + \varepsilon_{c,t}, \tag{8}$$

where $\textit{Number of Desks}_{c,(t,t+h]}$ is the maximum daily number of distinct desks (or banks) active in cluster c over the subsequent h days, $h \in \{30, 90\}$; $\textit{Demand}_{c,t}$ is one of the four demand proxies, in logs and standardized to zero mean and unit variance; and α_c and $\gamma_{a(c),t}$ are cluster and asset-class-by-date fixed effects. Standard errors are two-way clustered by cluster and date.

Table 12 shows that higher demand attracts more desks. The fixed effects are important for interpreting the dynamics between risk demand and banks' allocation. Without cluster fixed effects, the regression would ask if permanently larger risk clusters (perhaps major Treasury market-making desks) have both high demand and many active banks, reflecting a nearly mechanical feature of the cross-sectional size distribution rather than a dynamic response. The cluster fixed effects absorb these permanent differences and ensure that identification comes from within-cluster variation over time. This tests whether more desks participate when a cluster experiences higher-than-usual demand relative to its own history and reflects whether demand attracts entry, rather than just reflecting that some clusters are large. The asset-class-by-date fixed effects further absorb common time-varying conditions within each asset class.

A one-standard-deviation increase in demand increases the number of desks active in a cluster at both the one-month and three-month horizons. The RENTD results are particularly informative because the demand measure is derived from the bank's own risk management process and set before trading occurs, mitigating the concern that the relationship is driven by reverse causality from realized activity to measured demand. A one standard deviation increase in RENTD limits corresponds to roughly an additional desk added over the one-month horizon.

Coefficients are larger at the one-month horizon than at the three-month horizon. If entry were gradual, the ordering would be reversed. Instead, the pattern indicates that the desk response is concentrated within the first month and does not build thereafter: consistent with our proxies commingling transitory and persistent demand shocks, desks arrive quickly when demand rises. The

consistency of results across all four proxies, spanning both realized activity and ex-ante demand forecasts, provides robust evidence that high-demand clusters attract intermediation resources. Because the daily forward windows overlap, the final two columns of each table re-estimate the RENTD specifications on an end-of-month panel, in which each cluster contributes one observation per month and the outcome is measured over the following month; the results are similar, indicating that the daily estimates are not directly stemming from overlapping observations.

By contrast, if we assess the affect of demand on the number of distinct *banks* as the outcome, the number of banks responds less substantially less than the number of desks to the cluster-level demand. Participation changes are infrequent at both margins: the number of desks in a cluster changes on only 1.6% of cluster-days, and the number of banks changes even less often. A bank already active in an asset class can redirect an existing desk toward a high-demand cluster, whereas a bank absent from the cluster must build one. Setting up a desk requires investing in specialized labor and building cluster-level expertise, and the stickiness of desk counts is evidence that intermediation capacity is not fungible.

Combined, these results indicate that intermediation capacity flows toward demand: when a cluster's demand rises relative to its own history, more desks participate. The response is modest, and participation changes are infrequent. The coefficients imply an additional desk per standard deviation of demand.

This muted extensive margin is what the model predicts. Entry is costly. Expertise takes time to build, and a new desk competes against incumbents with better information. And entry is self-limiting: with a finite set of clusters and a headcount constraint, each additional participant produces additional signals, raising price informativeness and compressing the rents that expertise earns. If there is a sufficiently large shock to demand, a desk may enter a cluster, but as more desks participate, the cluster is less specialized. A demand shock therefore induces entry only when it is large enough to overcome both the fixed cost and the dilution of rents.

If entry absorbs only a small part of demand variation, the remainder must be absorbed by incumbents expanding capacity. Entry is costly because expertise takes time to build, and it is not simple to open a new desk and compete with incumbents. When demand rises, the response is that existing specialists expand rather than new participants enter. Section 7.2 examines this intensive margin directly and shows that it is incumbent specialists who expand most.

6.2 Specialists Respond to Demand Shocks

RENTD limits represent each desk's ex-ante forecast of expected customer demand for a specific risk factor. When actual demand approaches or exceeds this forecast, the bank can either raise the limit to accommodate additional flows, or leave it in place and turn customers away. We test whether specialists operate closer to risk limits and whether specialists change their risk limits when they near capacity.

To study this, we calculate each RENTD risk limit's proximity score, which captures how

close current usage is to the allowable bound (Appendix IA.A.2); we focus on RENTD limits that are measured in USD terms that map to specific risk factors. Each risk factor's proximity is the maximum proximity score across the risk limits for the risk factor. Proximity is scaled from 0 to 100, with 100 indicating usage exactly at the limit, and value above 100 indicating a breach. We also construct a binary indicators for limits near capacity (proximity above 80) and breached limits (proximity>100). Specialization is the risk factor's share of the bank's trailing 90-day average gross risk factor activity, standardized to a z-score as before. We measure each risk factor's RENTD demand as the total limit capacity across the factor's limits, that is, the sum of each limit's capacity as defined above.

Table 13 shows that specialists are more likely to have higher limits and operate near capacity, but not breach, their forecasted customer demand limit. This is consistent with specialist expertise in forecasting customer demand and operating near that limit.

Prediction 3 says that specialists have better models. We test the prediction by regressing the percentage change in the RENTD limit size from the current day to the average over the subsequent h days:

$$\Delta\text{Limit}_{it,h} = 100 \times \left[\ln\left(\overline{\text{Limit}}^{\text{fwd}}_{it,h} + 1\right) - \ln(\text{Limit}_{it} + 1)\right]$$

where i indexes the risk factor and t is the date. Our unit of observation is a specific risk factor (bank $\times$ desk $\times$ risk factor) on a given day. We examine horizons of 14, 30, and 90 days. We measure demand pressure using the risk factor's RENTD limit proximity score, as defined above. We focus on RENTD limits that are measured in USD terms that map to specific risk factors. Specialization is the bank's trailing 90-day average share of gross risk factor activity allocated to the limit's risk factor, standardized to a z-score as before.

We show that although specialists have larger changes in risk limits, specialists don't increase limits by more when the are close to their risk limits. Our regression takes the form:

$$\Delta\text{Limit}_{it,h} = \beta_1 \,\text{Proximity}_{it} + \beta_2 \,(\text{Proximity}_{it} \times \text{Specialist}_{it}) + \beta_3 \,\text{Specialist}_{it} + \mathbf{X}'_{it}\gamma + \alpha_{ct} + \delta_{bt} + \eta_i + \varepsilon_{it}$$

where α_{ct} are cluster-by-date fixed effects, δ_{bt} are bank-by-date fixed effects that absorb common time variation within banks, η_i are risk factor fixed effects that absorb permanent differences across limits (such as the size or active management), and $\mathbf{X}_{it}$ includes the log of current limit size and the trailing volatility of limit changes. The cluster-date fixed effects absorbs the common limit change for cluster on a given day and controls for a sector-wide demand shift that leads everyone active in that cluster to revise limits in the same direction. This avoids the concern that the result is driven by specialists happening to sit in clusters that all banks expanding that day. The bank-date fixed effect absorbs the bank's overall limit-setting that day, which would control for a firm-wide changes that shifts all the bank's limits at once. The log limit size controls for the mechanical relationship between limit size and percentage changes insofar as smaller limits can more easily have

larger percent changes; the trailing volatility controls for whether the limit is actively managed or effectively static.

Table 14 presents the results across horizons. Columns 1 to 3 show that specialists have greater changes in their risk limits over a horizon of 14, 30 or 90 days. A one standard deviation increase in specialization in a risk factor is associated with a 0.1 to 0.5 percentage point larger limit expansion over the 30 day or 90 day horizon. The effect is larger at 90 days.

The main coefficient of interest is β_2. A positive estimate implies that when demand pressure is high, specialists expand their limits more than non-specialists. Columns 4 through 6 show that the interaction term is not statistically significant. Specialists don't disproportionally increase their limits when close to their capacity. Instead, proximity to the limit is an important driver for limit changes for non-specialists.

Combined, these results suggest that specialists often increase their limits related to customer demand, however they don't increase their limits as they near the limit. This is consistent with specialists having better models to forecast customer demand so that they don't suddenly need to increase the limit as they near their limit.

7 Reallocations Within and Across Banks (Predictions 5 and 6)

The final set of predictions concerns what happens at the margins of the firm: how banks adjust their scope under stress, and where displaced expertise ends up. The model predicts that when balance sheet capacity tightens, banks shed peripheral activities before core ones, retreating toward their cluster of specialization. It further predicts that traders working in a bank's peripheral factors generate higher surplus at firms where those factors are core, creating a force for efficient sorting across firms. This section tests both predictions. We first show that banks facing binding constraints disproportionately reduce activity outside their primary clusters. We then use trader transition data to show that traders in peripheral roles are more likely to move to firms where their expertise is central, and that this pattern is stronger following positive demand shocks in the relevant factors.

7.1 Constrained Banks Retreat to Their Core (Prediction 5)

If specialization reflects a deliberate allocation of limited intermediation capacity, then banks facing binding constraints should contract activity in areas outside of their core expertise. The logic is straightforward: when a bank's risk budget tightens, it must reduce exposure somewhere. If the bank has invested in building expertise in certain clusters—and if that expertise is costly to rebuild once lost—it should preserve activity in those core clusters and cut from the periphery. We now show that when risk budgets tighten, banks preserve their core expertise while shedding peripheral activity, consistent with Prediction 5.

We measure bank constraints using the risk limit data. We want to calculate VaR-weighted utilization and VaR-weighted near-capacity share as the bank constraint measure, so we aggregate

the risk limits first to the desk-level, where we have VaR. We calculate each desk's maximum proximity and an indicator of whether the desk's maximum proximity is above 80 (percent). We use desk VaR to value-weight the desk-level constraints.

$$\text{Constraint}_{bt} = \sum_{d \in b} w_{dt} \text{Maximum Proximity}_{dt}, \text{ where } w_{dt} = \frac{VAR_{d,t-1}}{\sum_{d \in b} VAR_{d,t-1}}.$$

We define the forward change in a risk factor's activity as:

$$\Delta\text{Activity}_{bft,h} = 100 \times \Big[\ln\Big(\overline{\text{Activity}}^{\text{fwd}}_{bft,h} + 1\Big) - \ln\Big(\text{Activity}_{bft} + 1\Big)\Big]$$

where $\overline{\text{Activity}}^{\text{fwd}}_{bft,h}$ is the average |PnL| over h forward periods and Activity_{bft} is the trailing 90 day average of |PnL|. We regress the change in risk factor activity on the bank's specialization in a risk factor and the bank's constraints:

$$\Delta\text{Activity}_{bft} = \beta_1 \text{Constraint}_{bt} + \beta_2 \text{Specialist}_{bft} + \beta_3 \text{Constraint}_{bt} \times \text{Specialist}_{bft} + \alpha_b + \delta_{ct} + \varepsilon_{bft}.$$

We include cluster-by-date fixed effects to absorbs everything common to risk factors in a cluster on a given day (for example, cluster-wide demand or volatility shocks), so the comparison is across banks and factors in the same cluster-date. We also include bank fixed effects to absorb average differences across banks (for example, how aggressively a bank manages limits). This allows us to understand changes in activity allocation across-factors, within cluster, across banks. The results are similar if we included bank×date fixed effects but that would absorb our proximity measures.

Table 15 shows the results, where forward activity measured over 3 horizons—1 day, 3 days, and 5 days—relative to the activity level in the prior 90 days. We report specifications using VaR-weighted utilization and near-capacity share as the constraint measure. In all six specifications, the interaction coefficient is positive and statistically significant. The magnitude in column 1 implies that on days when the VaR-weighted share of the bank's desks near capacity increases one standard deviation (0.304), the gap activity growth widens by around 4.3 percentage points (14*0.304). The more specialized the bank is in a given risk factor, the more it preserves activity there when constrained.

The coefficient for for the proximity measures alone β_1 are negative and larger than the interaction term β_3, indicating that a banks pull back in activity when they are constrained, but they pull back less in risk factors that they specialize in.

The horizon comparison shows the dynamics of the reallocation. The interaction coefficient falls from 14pp at the 1-day horizon to 7pp at the 5-day horizon, suggesting that short-term changes happen outside the bank's core clusters, but longer-horizon changes are more equally distributed. The results are similar whether we measure constraints using the continuous utilization score or the near-capacity share, confirming that the finding is not sensitive to how we define "constrained." Table IA.7 shows that the results are robust to including bank-by-date FE, in which case the bank constraint measures drop out of the regression.

7.2 Specialist Competitors Absorb Released Demand

Prediction 5 concerns the constrained bank's own reallocation. Retrenchment by one bank is also an opportunity for others: when a constrained intermediary sheds activity in a cluster, the customer demand it served does not disappear. The model implies that this released demand should flow disproportionately to the remaining banks with expertise in the cluster, for whom additional scale is most valuable. We test this prediction by asking whether banks expand in clusters where their competitors become constrained, and whether specialists expand by more.

For each bank b, cluster c, and date t, we measure competitor constraint tightness as the weighted average limit proximity of the other banks active in the cluster, with weights given by each competitor's lagged 90-day share of cluster activity, so that the measure captures the fraction of the cluster's intermediation capacity that is near its limits. We require at least one other bank with limit data in the cluster-date. We estimate:

$$\Delta\text{Activity}_{bct,h} = \beta_1 \text{Proximity}_{ct}^{-b} + \beta_2 \text{Specialist}_{bct} + \beta_3 \text{Proximity}_{ct}^{-b} \times \text{Specialist}_{bct} + \alpha_{bc} + \gamma_{bt} + \delta_{ct} + \varepsilon_{bct}.$$

$\text{Proximity}_{ct}^{-b}$ measures the proximity of other banks' in the cluster. Because competitor constraints vary at the bank-cluster-date level, we include bank-by-date fixed effects, which absorbs the bank's own condition on a given day (for example, its capital, its overall risk budget, its own constraint tightness). Bank-by-cluster fixed effects absorb the persistent level of specialization. Cluster-by-date fixed effects absorb common cluster conditions, including the average level of competitor tightness, so that β_3 is from a within-cluster-day comparison.

Table 16 shows that when a cluster's capacity tightens, the activity expands more for the more specialized of the remaining banks. A one-standard-deviation tightening of competitor constraints widens the growth gap between a bank one standard deviation above the mean in specialization and the mean bank by 5 percentage points at the 1-day horizon.

Together with the previous subsection, these results trace a complete reallocation: constrained banks retreat to their core, and the demand they release is absorbed by the remaining specialists. Bräuning et al. (2026) show that constraints propagate within banks across asset classes through the internal allocation of risk limits, and our results show that expertise governs how intermediation capacity redistributes across banks within an activity.

7.3 Expert Traders Sort Across Banks (Prediction 6)

If specialization reflects genuine expertise that is costly to build, then the labor market should respect these investments (Prediction 6). In the model, expertise is embodied in traders, and the marginal expert trader is worth more at a bank that specializes in the matching area—both because the bank extracts more value from the trader's information and because demand shocks in that area raise the value of that information further. A trader working in a peripheral asset class at one bank holds expertise that is possibly underinvested in at the current employer but potentially valuable

elsewhere.

We test this prediction in two steps that examine the two sides of the labor market. First, we observe banks' labor *demand* directly using job postings: when expected customer demand expands in an asset class, do specialist banks disproportionately recruit? Because a posting precedes any match, this evidence isolates the bank's own assessment of the marginal value of expertise, independent of trader supply. Second, we examine *realized* trader moves using resume data: do traders in peripheral roles flow toward banks where their expertise is core, and does this flow intensify following demand shocks? The two pieces of evidence are complementary. Postings show that banks act on the higher marginal value of expertise; moves show that the labor market completes the match. Together they trace the mechanism from demand shock to recruiting to reallocation of expert labor across firms.

7.3.1 Labor Demand: Job Postings

The trader-level moves we study below are equilibrium outcomes: a move requires both a bank seeking expertise and a trader willing to supply it. To isolate the demand side—the bank's assessment of the marginal value of an expert trader before any match occurs—we use data on banks' job postings. A posted vacancy is a direct, forward-looking statement that the bank values an additional trader in that area at the prevailing wage. The model's prediction is that because expert labor has higher marginal value at banks that specialize in the matching area, positive shocks to expected customer demand should raise labor demand disproportionately at specialist banks.

We construct a bank–asset-class–month panel of trading-related job postings, which we describe in Section 3. We exclude internship and campus-program postings, which follow a fixed annual recruiting calendar and recruit workers who have not yet developed specialized expertise. We estimate:

$$\text{Postings}_{b,a,t} = \beta_1 \,\text{Specialization}_{b,a,t-1} \times \text{Demand Shock}_{a,t} \tag{9}$$

$$+ \beta_2 \,\text{Specialization}_{b,a,t-1} + \beta_3 \,\text{Demand Shock}_{a,t} + \alpha_{b,a} + \delta_{b,t} + \varepsilon_{b,a,t} \tag{10}$$

where b indexes banks, a asset classes, and t months. The dependent variable is either the number of postings by bank b in asset class a during month t or an indicator equal to one if the bank posts any opening in that asset class that month. $\text{Demand Shock}_{a,t}$ is the market-level expansion in RENTD limit capacity for risk factors in asset class a relative to its trailing average—the same measure we use in the trader-move regressions below—so that increases mechanically reflect banks' increased expectations of near-term customer demand.

We measure specialization two ways: *Asset Class Share* is the bank's share of gross risk-factor activity in the asset class in the prior month, and *Specialist* is its trailing three-month average, a smoothed measure less sensitive to transitory swings in activity. All independent variables are standardized to have zero mean and unit variance.

The fixed effects determine what variation identifies β_1. All specifications include bank-by-asset-class fixed effects, which absorb the mechanical relationship that banks post more openings in areas where they are permanently large—without them, the interaction would partly reflect that specialists are big rather than that they respond differently—and bank-by-month fixed effects, which absorb bank-wide hiring cycles: expansions, freezes, and layoffs common to all of a bank's desks. Identification therefore comes from *within-bank* variation across asset classes in the same month: when demand expands in a given asset class, does the bank tilt its recruiting toward that class more strongly where it is a specialist? The even columns additionally include asset-class-by-month fixed effects, which absorb the demand shock's main effect along with any other common conditions in the asset class that month—volatility, market-wide hiring booms, or regulatory changes. In these specifications, β_1 is identified purely from the *differential* response of more- versus less-specialized banks to the same shock, so the estimate cannot reflect asset-class-level conditions that happen to coincide with RENTD expansions.

Table 17 shows that when demand expands in a given asset class, bank do tilt its recruiting toward that class more strongly where it is a specialist. The first four columns use the number of postings as the dependent variable, and the last four the indicator for any posting; within each block, the first pair of columns uses Asset Class Share and the second pair the smoothed Specialist measure. Column 1 shows that demand expansions raise recruiting across all banks: a one-standard-deviation demand shock raises a bank's postings in the asset class by 0.02 per month, compared to a sample median of 0 postings and sample mean of 0.88 postings. The interaction is the model's prediction: the same shock raises postings by an additional 0.06 at a bank whose activity share in the class is one standard deviation above its mean. Column 2 adds asset-class-by-month fixed effects, absorbing the shock's main effect; the interaction coefficient is similar, suggesting that the differential response of specialists is not driven by asset-class-level conditions coinciding with demand expansions. Columns 3 and 4 show similar estimates using the smoothed *Specialist* measure. The last four columns repeat the analysis for the extensive margin: a one-standard-deviation demand shock raises a specialist bank's probability of posting any opening in the asset class by 0.02, and the estimate is stable to the inclusion of asset-class-by-month fixed effects.

The postings evidence complements the realized moves we examine next. Postings measure banks' recruiting intent before a match forms, so the result cannot reflect the supply side of the labor market—traders' willingness to move—or the availability of displaced traders from contracting banks. Together with the trader-level results below, the evidence traces the full chain: demand expansions raise the marginal value of expertise at specialist banks, those banks respond by recruiting, and expert traders flow to them.

7.3.2 Realized Trader Flows

If specialization reflects genuine expertise that is costly to build, then the labor market should respect these investments (prediction 6). A trader working in a peripheral asset class at one bank—e.g., a

commodities trader at a bank that primarily intermediates rates—holds expertise that is possibly underinvested in at the current employer but potentially valuable elsewhere. The prediction is that such traders are more likely to move, and conditional on moving, more likely to land at a bank where their asset class is core. Two channels could amplify this sorting: demand shocks that raise the value of the trader's expertise at receiving banks (a pull channel), and binding constraints at the origin bank that force it to shed peripheral activity first (a push channel).

We test this using a trader-level panel constructed from Revelio Labs data, matched to the banks in our sample. The panel tracks individual traders at monthly frequency, recording the trader's employer, position and job description, and seniority level. We use these values to map traders to the most likely asset class, which we describe in the Internet Appendix. In total, the data provides information on 25,000 individuals with trading roles. We identify a move as a change in employer between consecutive months and observe two types of departures: there are 5,448 hires by banks in our sample, and 3,374 hires where the employee is moving between two covered banks. About 400 of the moves overlap with periods in which we observe trading data for the origin and destination bank.

We measure a bank's asset class share in month t as the monthly average of the daily share of the bank's total trading activity attributable to that asset class, where activity is the absolute value of desk-level comprehensive P&L aggregated to the asset class. We classify a trader's asset class as *peripheral* at the origin bank if the bank's asset-class share in that month falls below the bank's median share across its asset classes that month, and *core* otherwise. We define core at the destination symmetrically.

Table 18 presents the transition matrix for all moves between covered banks. Of the 97 traders who move from a peripheral position, 70% land at a bank where their asset class is core. By contrast, among traders moving from a core position, 86% end up at a bank where they are also core. Traders disproportionately sort toward banks in which their asset class specialization is core at the destination bank.

To test this more formally, we regress an indicator for whether the trader moved on the trader's asset class share at its origin bank:

$$\mathbf{I}(\text{Moved})_{i,t} = \alpha + \beta\, \text{Origin Asset Class Share}_{i,t} + \mathbf{X}'_{i,t}\gamma + \varepsilon_{i,t}$$

where i indexes traders, t indexes months, and Origin Activity Share is the amount of gross risk factor activity in the trader's asset class relative to the total activity at the trader's bank. Higher values indicate a more core position within the bank, and Origin Activity Share is z-scored within a month. $\mathbf{X}_{it}$ includes seniority controls and fixed effects.

Table 19 Columns 1 to 4 show that periphery traders are more likely to leave their bank. A trader in an asset class that is one-standard-deviation below the bank's average asset class share is 0.04pp more likely to move. This is a small but statistically significant effect, in part because

there's a limited number of moves in the sample. The estimate is stable across specifications that add month × asset class fixed effects (columns 2 and 3) to absorb common shocks (e.g., regulatory changes, profitability shocks) that affect all traders in a given asset class-month. Column 3 shows that senior traders are less likely to move overall. Column 4 restricts to senior traders only and finds the same.

Peripheral traders could leave because of "push" or "pull" effects. A constrained origin bank shedding peripheral activity could push them out, or rising demand for their expertise elsewhere could pull them away. Column 5 tests the push mechanism by interacting the origin asset class share with the origin bank's constraint, measured as the bank's VaR-weighted share of desks with limit proximity above 80. The interaction is insignificant. The uninteracted constraint coefficient is positive, meaning traders leave constrained banks at a higher rate, consistent with constrained banks experiencing high turnover. But peripheral traders are no more likely than core traders to be the ones who leave. We find no evidence that the push channel drives the sorting.

Next, we test the "pull" interpretation in two ways. First, Column 7 shows that traders move to meet demand at the destination by regressing the indicator for moving on the asset class's RENTD demand. Higher demand at the destination makes traders more likely to move. A one standard deviation in the level of market demand is associated with a substantial increase in the probability of moving. Column 8 shows that the demand shock does not have a significant effect on moving, in part because of the limited number of moves, that moves may take more than one period, and the interaction with the effect of demand with seniority.

Second, we study the movers within an asset class and test whether peripheral traders increase their asset class share when they move. We define ΔAsset Class Share$_{i,t}$ as the difference between the trader's destination asset class share and origin asset class share, where the asset class share is the trader's asset class share at destination or origin bank. We regress the ΔAsset Class Share$_{i,t}$ on an indicator for being a peripheral trader:

$$\Delta\text{Asset Class Share}_{i,t} = \alpha + \beta\, \mathbb{I}(\text{Peripheral at Origin Bank}_{i,t}) + \mathbf{X}'_{i,t}\gamma + \varepsilon_{i,t}.$$

A positive β indicates that peripheral traders disproportionately move to banks where their asset class is more central. $\mathbf{X}_{it}$ includes seniority controls and asset class and month fixed effects.

Table 20 Columns 1 and 2 show that traders moving from peripheral positions experience a 7pp increase in asset class share at the destination relative to traders moving from core positions. The effect is large considering that the asset class share of the average periphery trader at their bank is 6%, hence moving periphery traders more than double their asset class share. The estimate is similar when controlling for seniority.

Columns 3 and 4 test whether sorting intensifies during demand shocks. We study movers within an asset class to see whether a demand shock drives more peripheral traders to move and increase their asset class share. We interact the peripheral indicator with a standardized measure

of market-level capacity expansion: the smoothed change in aggregate RENTD limit size for risk factors in the trader's asset class relative to its trailing average.

$$\begin{aligned}\Delta\text{Asset Class Share}_{i,t} = \alpha + \beta_1 \mathbb{I}(\text{Peripheral}_{i,t}) + \beta_2 \,\text{Demand Shock}_{a,t} \\ + \beta_3 \mathbb{I}(\text{Peripheral}_{i,t}) \times (\text{Demand Shock}_{a,t}) + \mathbf{X}'_{i,t}\gamma + \varepsilon_{i,t}.\end{aligned}$$

Increased RENTD limits mechanically reflect increased demand expectations, raising demand for specialized traders. A demand shock at the destination leads movers to sort to banks where with larger asset class shares in their specialty. In columns 3 and 4, the interaction coefficient β_3 is positive and significant, suggesting that peripheral traders are more affected the demand shock than specialist traders. Accounting for seniority, the main demand shock effect is concentrated in peripheral traders. Peripheral disproportionately move to better-matched banks when market-level capacity in their asset class expands.

Combined, these results paint a picture of pull-driven reallocation. Peripheral traders leave not because they are forced out by constrained employers, but because the labor market offers them better matches elsewhere. The sorting is efficient: traders flow from banks that undervalue their asset class to banks that value it more, and this flow intensifies when demand for their expertise is high. The implication is that specialization is not merely an organizational feature of banks but is partly embodied in the human capital of individual traders, and the labor market acts as a mechanism that reinforces and sharpens the pattern of specialization over time.

7.3.3 Internal Movers

In addition to asset-class specialists switching banks and staying in their asset class specialty, there are internal movers, or traders that remain at the same bank but change their asset class. Above, we show that movers that are peripheral at their bank are pulled by banks facing demand shocks and desiring their expertise. Below we show that internal movers do not face the same pull due to a demand shock.

Traders sorting across firms within an asset class differ substantially from traders moving within a firm. Our data cannot identify traders moving across desks in a bank, and we consider internal movers to be traders that have changed asset classes within a firm, condition on having at least a year of experience to account for rotational programs. Table 20 Columns 5 to 8 show that peripheral traders are more likely to move internally, similar to traders switching banks; however, internal movers are less likely to move amid a demand shock, and peripheral traders aren't pulled by a demand shock. This result is unsurprising and is consistent with the development of expertise. Peripheral traders within the bank don't have expertise in an asset class, so when there is demand shock in the destination asset class, they are unable to fulfill it. By contrast, movers across banks have developed that expertise.

8 Conclusion

We study the scope of bank intermediation, which risk factors a desk chooses to trade, and why. We show that banks organize their trading around clusters of related risks, that this clustering reflects genuine knowledge complementarities, and that complementarities shape how the financial system prices risk and reallocates resources. Specialists earn higher returns, produce more accurate forecasts, and respond more aggressively to demand. When capacity tightens, banks shed peripheral activities first. Displaced traders sort toward firms where their expertise is central. Some of these patterns are individually consistent with alternatives: clustering with risk diversification, specialist returns with selection, contraction ordering with capital allocation. But two findings are not. Denser coverage of related factors improves forecasts for existing activities, which diversification and incentive theories do not predict. And expertise is non-portable across firms, which selection on individual ability does not predict. A simple model with clustered signal precision and a balance sheet constraint generates each of these patterns without entry costs, learning lags, or search frictions.

The results suggest that the knowledge structure of intermediaries is a first-order determinant of how risk is priced and allocated. If the organization of expertise determines who intermediates what, then policies that reshape bank scope—capital requirements or activity restrictions—reshape the knowledge structure of the market itself. Understanding those consequences requires taking seriously that intermediation capacity is not fungible. Our evidence shows it is not.

References

Omar Barbiero, Falk Bräuning, Gustavo Joaquim, and Hillary Stein. Dealer risk limits and currency returns. Technical Report 24-11, Federal Reserve Bank of Boston, 2024.

Gary S. Becker and Kevin M. Murphy. The division of labor, coordination costs, and knowledge. *The Quarterly Journal of Economics*, 107(4):1137–1160, 1992.

Jonathan B. Berk and Jules H. van Binsbergen. Measuring skill in the mutual fund industry. *Journal of Financial Economics*, 118(1):1–20, 2015.

Jonathan B. Berk, Jules H. van Binsbergen, and Binying Liu. Matching capital and labor. *The Journal of Finance*, 72(6):2467–2504, 2017.

Kristian Blickle, Cecilia Parlatore, and Anthony Saunders. Specialization in banking. Technical report, 2025.

Maxime Bonelli. Up the capital ladder: The role of labor mobility in the mutual fund industry, 2025.

Falk Bräuning and Hillary Stein. The effect of primary dealer constraints on intermediation in the treasury market. Technical Report 24-7, Federal Reserve Bank of Boston, 2024.

Falk Bräuning, Wenxin Du, and Hillary Stein. Cross-asset limit allocation of trading desks. 2026. Working paper, Federal Reserve Bank of Boston and Harvard University.

Svetlana Bryzgalova, Anna Pavlova, and Taisiya Sikorskaya. Strategic arbitrage in segmented markets. *Journal of Financial Economics*, 166:104008, 2025.

Xiao Cen, Winston Wei Dou, Leonid Kogan, and Wei Wu. Fund flows and income risk of fund managers, 2023. NBER Working Paper 31986.

Andrea L Eisfeldt, Hanno Lustig, and Lei Zhang. Complex asset markets. *Journal of Finance*, 78 (5):2519–2562, 2023.

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.

Eugene F Fama and Kenneth R French. Common risk factors in the returns on stocks and bonds. *Journal of financial economics*, 33(1):3–56, 1993.

Luis Garicano. Hierarchies and the organization of knowledge in production. *Journal of Political Economy*, 108(5):874–904, 2000.

Luis Garicano and Thomas N. Hubbard. Managerial leverage is limited by the extent of the market: Hierarchies, specialization, and the utilization of lawyers' human capital. *Journal of Law and Economics*, 50(1):1–43, 2007.

Zhiguo He and Arvind Krishnamurthy. Intermediary asset pricing. *American Economic Review*, 103 (2):732–770, 2013.

Theis Ingerslev Jensen, Bryan Kelly, and Lasse Heje Pedersen. Is there a replication crisis in finance? *The Journal of Finance*, 78(5):2465–2518, 2023.

Marcin Kacperczyk, Clemens Sialm, and Lu Zheng. On the industry concentration of actively managed equity mutual funds. *Journal of Finance*, 60(4):1983–2011, 2005.

Marcin Kacperczyk, Stijn Van Nieuwerburgh, and Laura Veldkamp. A rational theory of mutual funds' attention allocation. *Econometrica*, 84(2):571–626, 2016.

Dan Li, Lubomir Petrasek, and Mary Tian. Risk-averse dealers in a risk-free market: The role of trading desk risk limits. Technical Report 2025-034, Board of Governors of the Federal Reserve System, 2025.

Lina Lu and Jonathan Wallen. What do bank trading desks do? Technical report, Harvard Working Paper, 2025.

Robert E. Jr. Lucas. On the size distribution of business firms. *The Bell Journal of Economics*, 9 (2):508–523, 1978.

Tobias J Moskowitz, Chase P Ross, Sharon Y Ross, and Kaushik Vasudevan. Risk, specialization, and covered-interest parity. *Journal of Finance, Forthcoming*, 2026.

Daniel Paravisini, Veronica Rappoport, and Philipp Schnabl. Specialization in bank lending: Evidence from exporting firms. *Journal of Finance*, 78(4):2049–2085, 2023.

Emil N Siriwardane, Adi Sunderam, and Jonathan Wallen. Segmented arbitrage. *The Journal of Finance*, 80(5):2543–2590, 2025.

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

Jeremy C. Stein. Information production and capital allocation: Decentralized versus hierarchical firms. *The Journal of Finance*, 57(5):1891–1921, 2002.

Stijn Van Nieuwerburgh and Laura Veldkamp. Information acquisition and under-diversification. *Review of Economic Studies*, 77(2):779–805, April 2010a. doi: 10.1111/j.1467-937X.2009.00583.x.

Stijn Van Nieuwerburgh and Laura Veldkamp. Information acquisition and under-diversification. *Review of Economic Studies*, 77(2):779–805, 2010b.

Milena Wittwer and Andreas Uthemann. Specialization in financial markets. 2025.

9 Figures

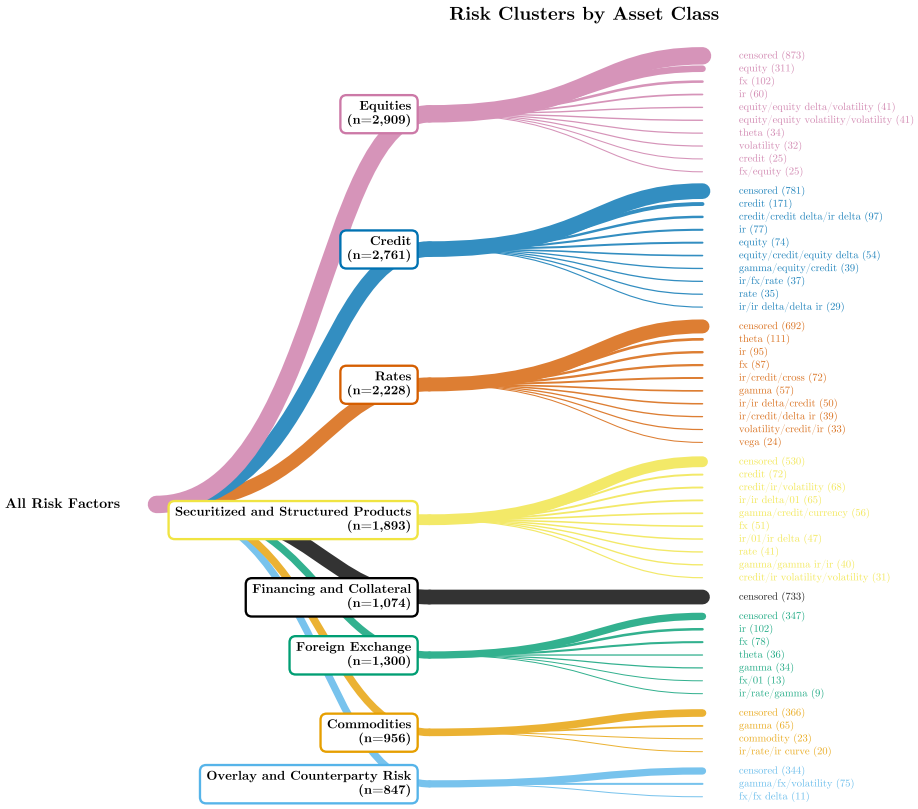


Figure 1: Desk and Risk Classification. Figure shows the classification of desks into asset classes and the 10 largest risk factor clusters within each asset class that include all banks in our sample. Figure provides text labels for risk factors only when the text label is derived from risk factors provided by the full set of banks; if a risk factor has too few banks or the text label is used by too few banks, we label it "censored" given anonymity requirements for confidential data.

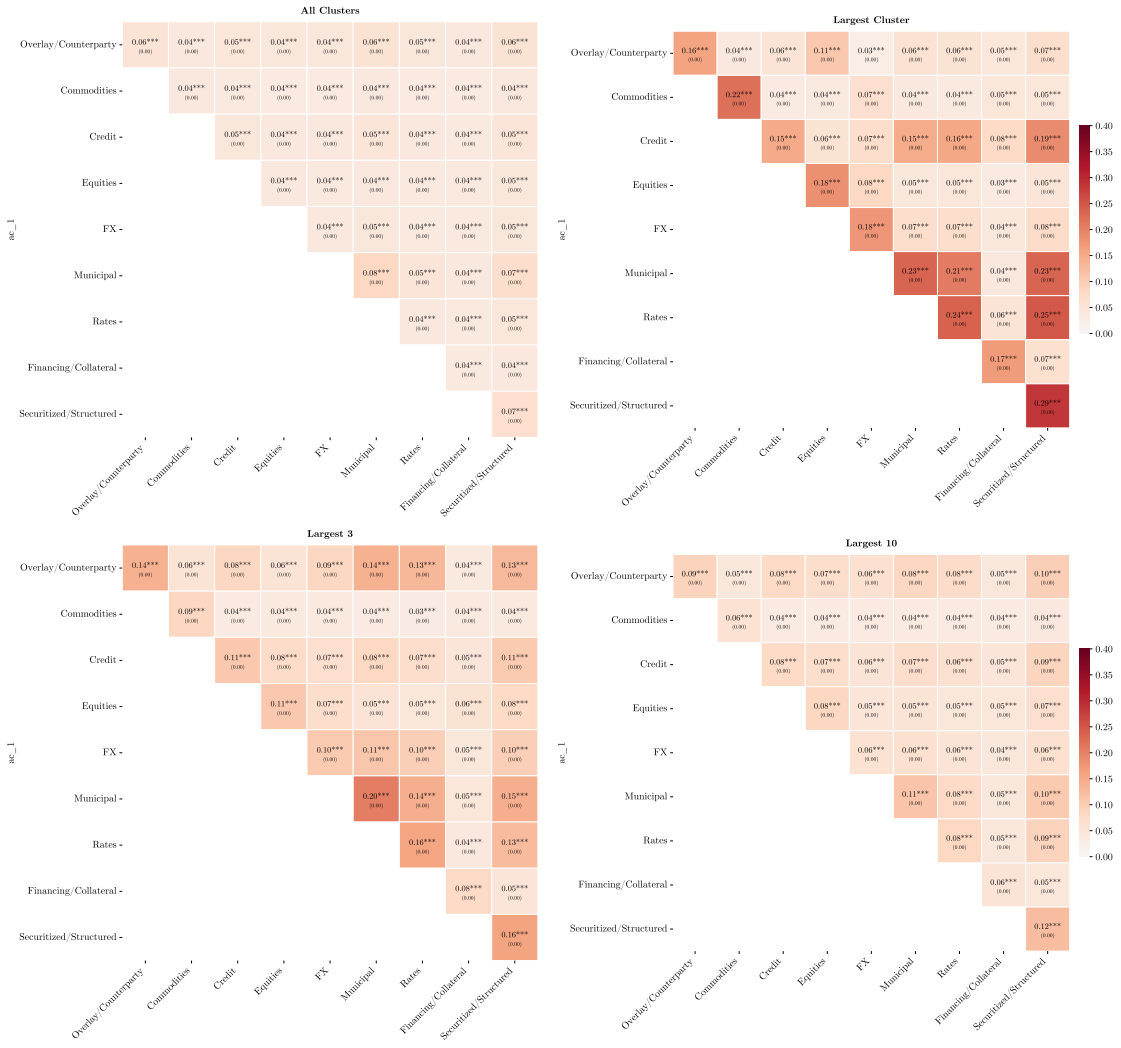


Figure 2: Correlation of Risk Factors Across Asset Classes. Each cell reports the average absolute correlation of returns (defined above) between risk factor pairs across asset classes. Correlations are aggregated using Fisher's z-transformation, weighting by both the inverse variance of each pairwise estimate and by risk factor volatility. Cluster size is defined as the sum of constituent risk factors' P&L standard deviations. p-values in parentheses; *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively. Sample restricted to post-2021 observations.



Figure 3: Clusters and Singletons across Correlation Thresholds. Figure plots the number of resulting clusters, along with number of singleton clusters that contain only a single risk factor, when categorizing risk factors based on the absolute value correlation threshold of our 90 day return proxy. Figure based on post-2020 sample clustering.

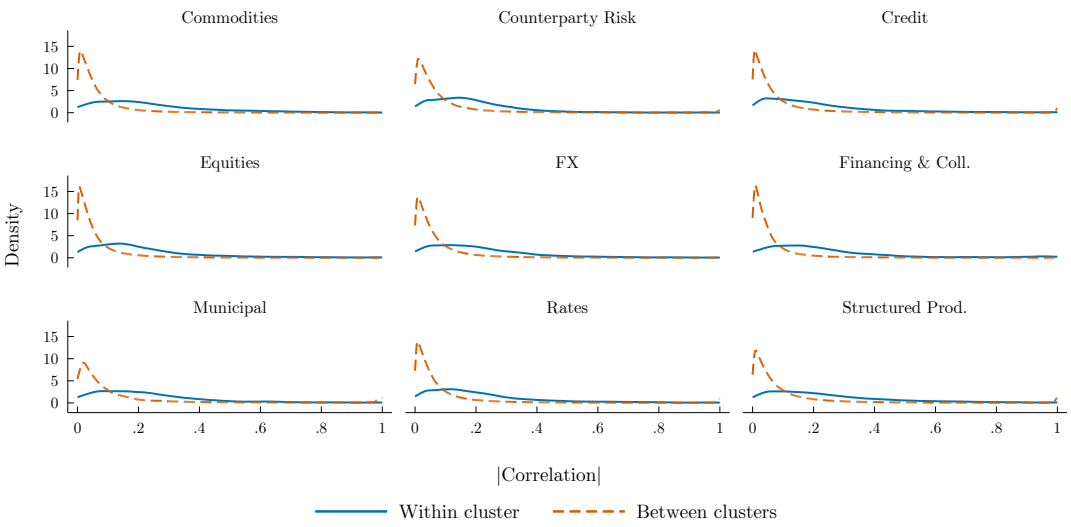


Figure 4: Distribution of $|\rho|$ across risk factor pairs within an asset class. Figure shows the distribution of $|\rho|$ for pairs of risk factors in the same asset class that are either in the same cluster or in separate clusters. We require the risk factor pair have at least one year of overlapping data.

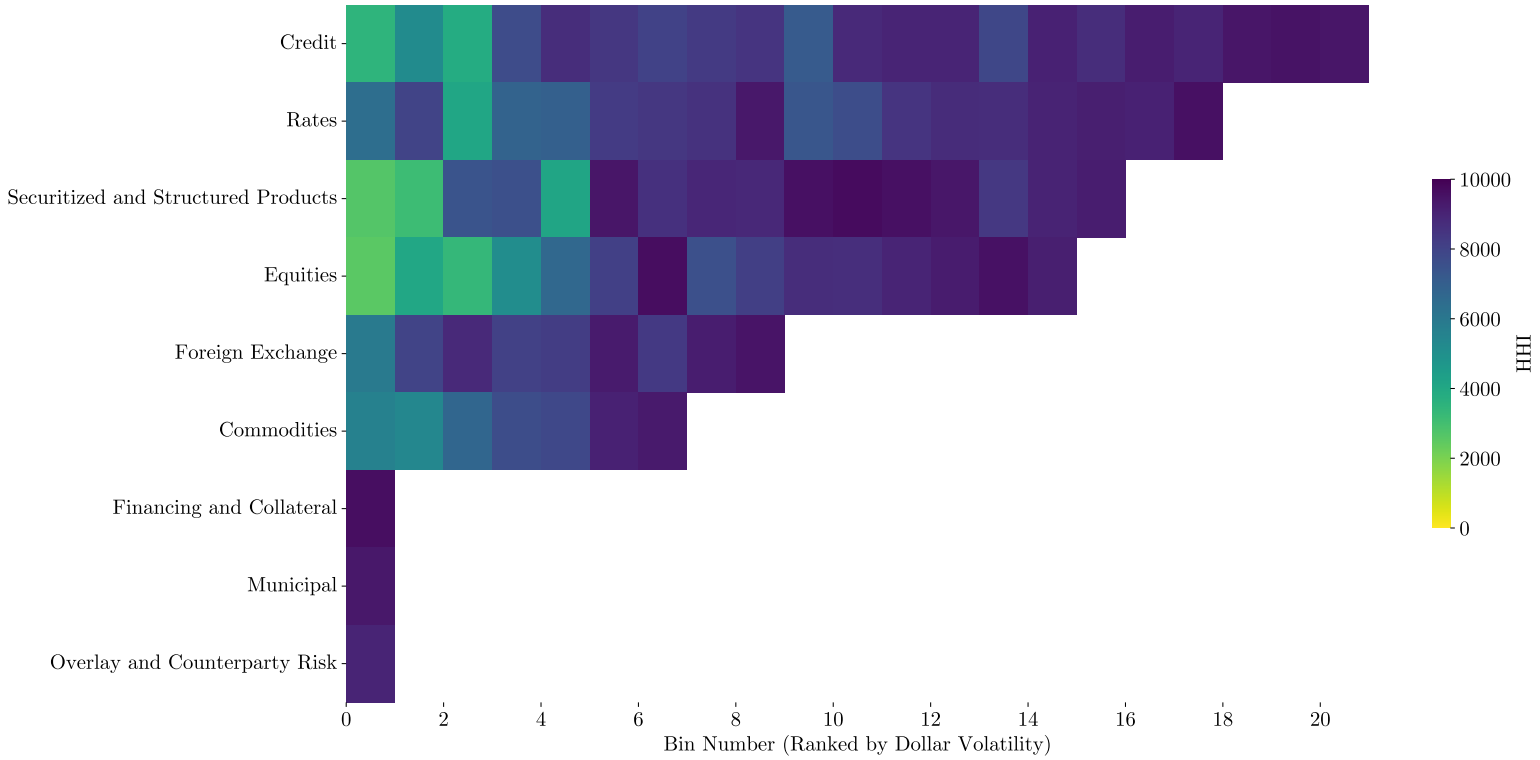


Figure 5: Bank Concentration within Clusters and Cluster Size. This figure displays the degree of bank specialization across risk factor clusters, organized by asset class. Each row represents an asset class. Within each asset class, clusters are ranked from left to right by average monthly dollar PnL volatility (highest to lowest). Adjacent clusters are grouped into bins such that each bin contains observations from all banks in our sample. The color intensity represents the Herfindahl-Hirschman Index (HHI) across banks within each bin, averaged over the constituent clusters. Darker colors indicate higher concentration (specialized risk-taking by fewer banks), while lighter colors indicate lower concentration (more diversified risk-taking across banks). Market shares are calculated as each bank's standard deviation of PnL within a cluster over the past 3 months divided by the sum of all banks' standard deviations in that cluster over the past 3 months. The plotted HHI is the time-series average of daily values. Based on risk clusters estimated in the post-2020 sample.

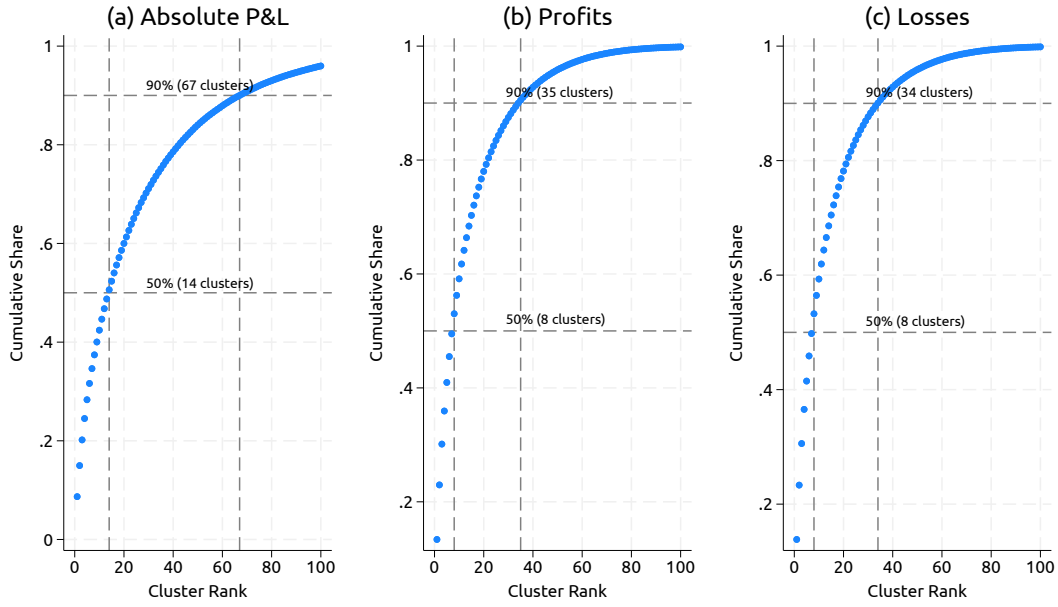


Figure 6: Concentration of Activity Across Risk Factor Clusters. Figure shows the cumulative share of bank activity accounted for by the top 100 risk factor clusters. For each bank-day, we rank clusters by their share of total activity and compute the cumulative share at each rank. We then take the median across days within each bank, followed by the median across banks. Panel (a) uses the absolute value of P&L as the activity measure, Panel (b) uses profits (positive P&L only), and Panel (c) uses losses (absolute value of negative P&L). Dashed lines indicate the number of clusters required to account for 50% and 90% of activity.

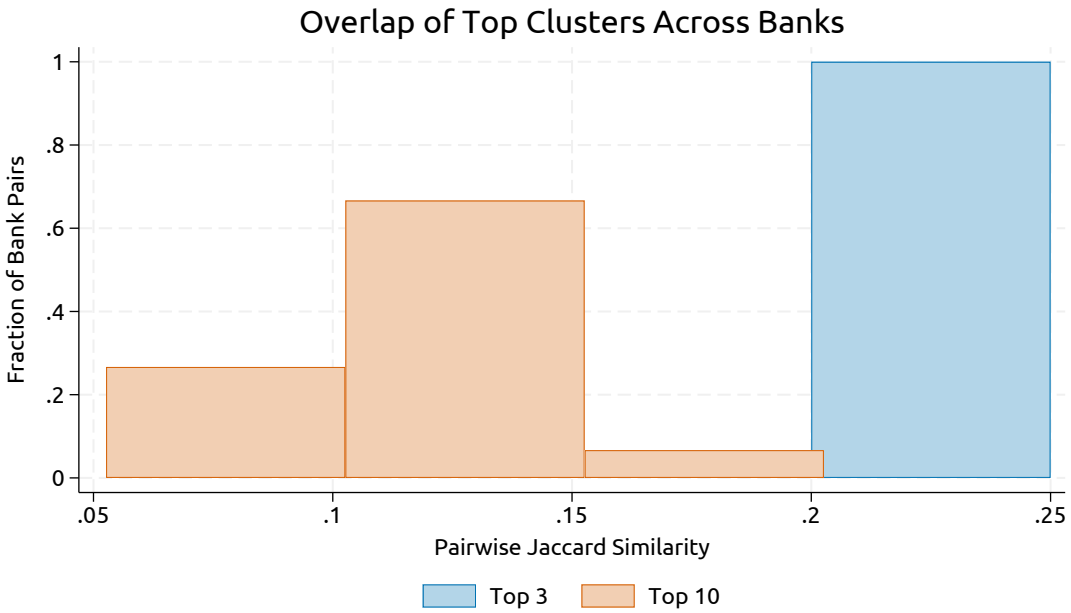


Figure 7: Overlap of Top Clusters Across Banks. Figure shows the distribution of pairwise Jaccard similarity scores across all bank pairs for the top 3 and top 10 risk factor clusters by activity share. For each bank-month, we rank clusters by their share of total absolute P&L and select the top K. For each bank pair within a month, we compute the Jaccard similarity (size of intersection divided by size of union) of their top-K sets, then take the median across months for each pair. The histogram displays the resulting cross-pair distribution. Values near zero indicate that banks concentrate their activity in largely distinct sets of clusters; a value of one would indicate identical top-K sets.

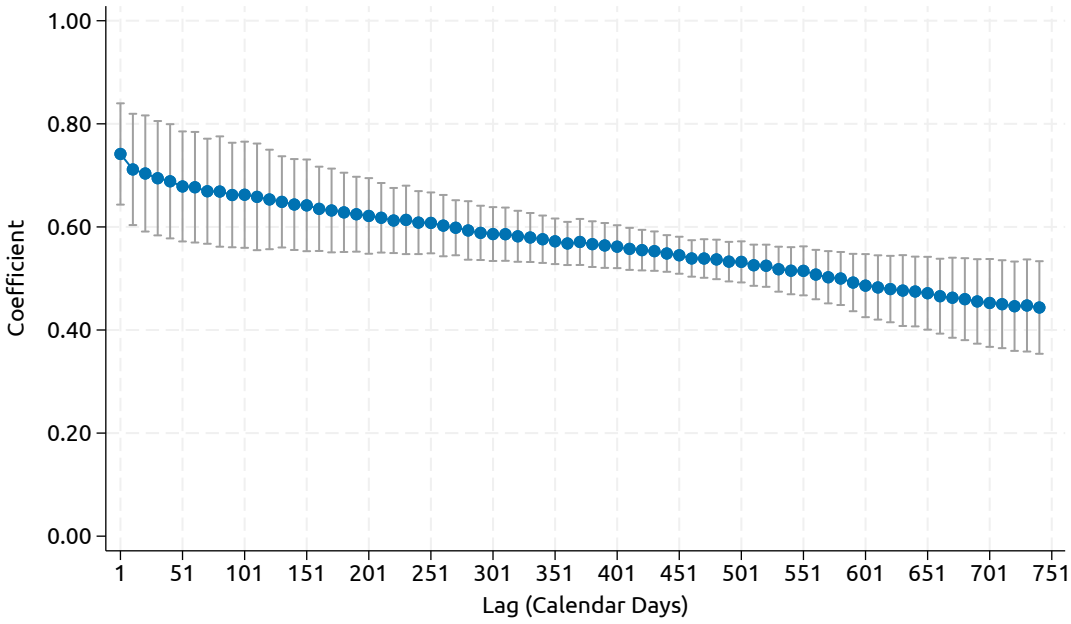


Figure 8: Cluster Share Persistence. Figure plots the coefficient from the regression $\text{Share}_{b,c,t} = \alpha + \beta\,\text{Share}_{b,c,t-\tau} + \varepsilon_{b,c,t}$ where $\text{Share}_{b,c,t}$ is bank b's percentage of total risk (standard deviation of PnL over previous three months) in cluster c on date t. Based on risk clusters estimated in the post-2020 sample.

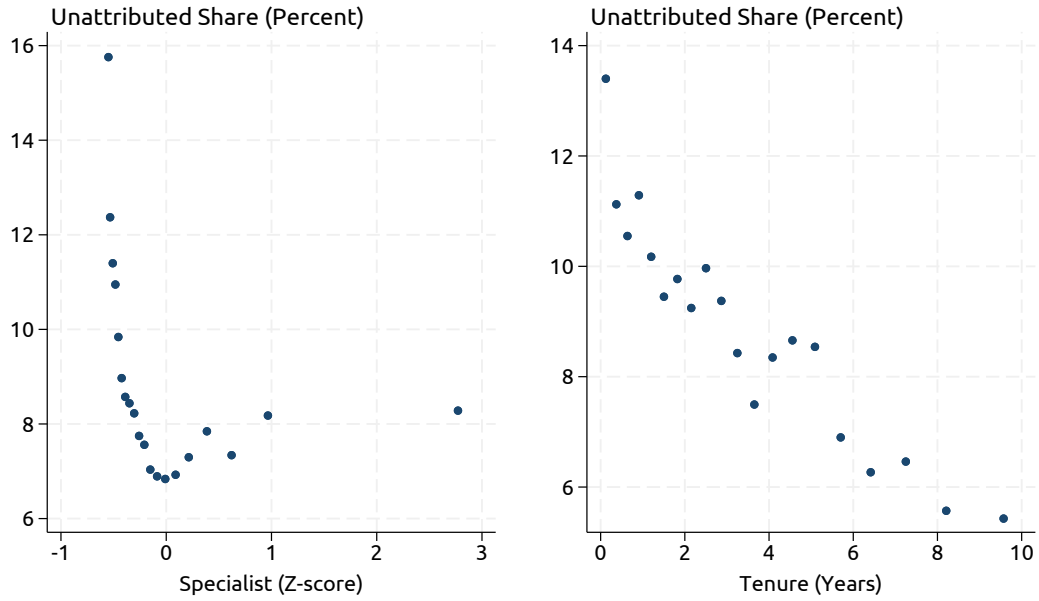


Figure 9: Specialization and Model Accuracy. Figure plots the relationship between the unattributable share of P&L and two measures of bank expertise. The unattributable share is the share of unexplained P&L (Item 10 under the FR VV-1 reporting requirements) to the P&L from existing positions' risk factors and carry, with all components entering in absolute value. Panel (a) is a binscatter of the unattributable share against the standardized specialization measure, defined as the trailing 90-day average share of a bank's activity in each desk. Panel (b) is a binscatter of the unattributable share against the desk's cumulative number of active trading years. The sample in Panel (b) excludes desks that were already active on the bank's first reporting date to avoid left-censored tenure.

10 Tables

				Profitability ($ mn)					Value-at-Risk ($ mn)				
Asset Class	Count	Desk–Day N	Avg. Desks/Day	Mean	StDev	10th	90th	Skew	Mean	StDev	10th	90th	Skew
Commodities	45	79,780	28	0.7	4.4	-1.5	3.5	-1.6	6.8	8.5	0.4	15.2	4.2
Credit	155	202,006	72	0.7	5.0	-0.5	2.6	-12.1	4.4	11.2	0.0	9.9	11.6
Equities	152	278,357	99	1.5	4.4	-0.6	5.0	-1.1	5.2	7.3	0.0	12.8	3.9
Financing and Collateral	59	96,213	34	0.6	3.6	-0.8	2.7	-3.9	4.2	5.4	0.3	8.9	4.4
Foreign Exchange	82	106,366	38	1.3	4.5	-1.0	5.1	-3.1	7.5	16.0	0.0	15.6	5.8
Municipal	24	42,245	15	0.3	1.6	-0.2	1.2	-2.7	3.0	5.3	0.0	9.3	2.9
Overlay and Counterparty Risk	32	56,252	20	0.0	6.2	-2.5	2.7	-18.6	4.9	8.8	0.0	14.9	3.7
Rates	137	214,126	76	0.2	16.1	-1.2	3.6	-17.7	7.2	15.3	0.0	15.6	7.2
Securitized and Structured Products	128	190,392	68	0.4	4.4	-0.6	2.1	-10.2	6.7	13.1	0.0	17.2	5.8

Table 1: Trading Desk Summary Statistics. Table reports summary statistics at the trading-desk level. We report the number of desks by asset class and bank, along with the distribution of desk-level P&L. Statistics are computed over the full sample period; see data section for details.

Panel A: Structure	*N*	Mean	SD	p10	p25	p50	p75	p90
Limits per desk-date	519,109	8.6	8.4	2	3	6	12	16
Panel B: Risk Limit Usage (%)	*N*	Mean	SD	p10	p25	p50	p75	p90
All Limits Types	4,451,870	19.6	25.0	0.0	8.2	32.9	58.0	58.0
Aging	423,099	24.3	33.8	0.0	0.2	48.1	80.3	80.3
Other	1,528	45.5	27.1	2.3	29.1	49.7	69.6	76.2
Position	151,924	38.8	29.4	12.0	36.8	62.5	80.3	80.3
Scenario	1,528,811	8.0	18.1	0.0	0.0	3.8	33.5	33.5
Sensitivity	1,566,027	23.6	23.4	4.3	16.5	36.9	59.1	59.1
VaR	779,949	27.8	24.4	6.6	23.7	43.1	61.9	61.9
All Limit Sources	4,451,870	19.6	25.0	0.0	8.2	32.9	58.0	58.0
RENTD	3,011,972	17.6	23.8	0.0	5.1	29.6	55.1	63.7
Risk Appetite	903	28.5	26.5	6.5	20.6	46.0	72.5	55.1
Risk Reducing	410	20.5	17.0	6.3	23.0	30.8	35.8	72.5
Other	1,438,585	23.8	26.9	0.0	15.2	38.4	63.7	35.8
Panel C: Proximity Thresholds (%)	Share >50	>80	>90	>95	>100			
All Limits Types	13.9	3.0	1.4	0.9	0.2			
Aging	24.1	10.1	6.4	4.8	0.6			
Other	49.8	5.8	1.3	0.2	0.0			
Position	37.0	10.2	4.2	2.2	0.2			
Scenario	5.1	0.9	0.4	0.2	0.1			
Sensitivity	15.1	2.6	0.9	0.4	0.2			
VaR	18.5	2.8	1.1	0.6	0.3			
All Limit Sources								
RENTD	12.5	2.3	0.8	0.5	0.2			
Risk Appetite	22.1	3.9	2.7	2.2	0.1			
Risk Reducing	5.4	0.5	0.2	0.2	0.2			
Other	16.9	4.6	2.5	1.7	0.3			
Panel D: Persistence	*N*	Mean	SD	p10	p25	p50	p75	p90
Near-cap spell (>80%, days)	19,046	7	24	1	1	2	5	14
Panel E: Limit Changes								
By Limit Source	All Limit Sources	RENTD	Risk Appetite	Risk Reducing	Other			
% days with any change	0.54	0.45	0.22	0.24	0.73			
% days with increase	0.33	0.30	0.00	0.24	0.39			
% days with decrease	0.33	0.30	0.22	0.00	0.39			
Med. $\|\Delta\|$ (%, cond.)	20	25	19.5	100	12.5			
Med. days between changes	183	162	19	61	216			
By Limit Type	All Limits Types	Aging	Other	Position	Scenario	Sensitivity	VaR	
% days with any change	0.53	1.44	0.53	0.43	0.41	0.43	0.53	
% days with increase	0.33	0.77	0.49	0.30	0.20	0.37	0.27	
% days with decrease	0.33	0.69	0.44	0.27	0.21	0.36	0.30	
Med. $\|\Delta\|$ (%, cond.)	20	5	27	20	30	22	20	
Med. days between changes	183	230	66	145	159	187	159	

Table 2: Risk Limit Summary Statistics. Each observation is a specific risk limit on a specific trading desk on a given day from the FR VV-1 reporting form. Panel A reports the distribution of limits per desk-date. Panel B reports the distribution of proximity scores overall and by limit type (VaR, sensitivity, position, scenario, aging) and by limit source (risk appetite, RENTD, risk reducing). Proximity measures how close current usage is to the allowable bound, on a scale where 0 indicates maximum slack, 100 indicates usage at the limit, and values above 100 indicate a breach (see Appendix IA.A.2). Panel C reports the share of limit-date observations exceeding various proximity thresholds, by limit type and source. Panel D reports the duration of near-capacity spells, defined as consecutive days with proximity above 80. Panel E reports the frequency and magnitude of limit changes, broken out by limit source and by limit type; the conditional median absolute change and median days between changes are computed on the subset of limit-days with a nonzero change. Limited to the Post-2020 sample.

Panel A: Characteristics	Avg. RF/Desk/Day	Factors (Count)		Profitability ($ mn)			
		Pre-2021	Post-2020	Mean	StDev	10th	90th
Commodities	9.1	592	364	0.02	1.42	-0.25	0.31
Credit	9.0	1,668	1,093	-0.01	1.02	-0.15	0.12
Equities	8.6	1,607	1,302	0.03	1.66	-0.30	0.40
Financing and Collateral	8.8	598	476	0.01	0.99	-0.10	0.11
Foreign Exchange	8.2	864	436	0.00	1.78	-0.34	0.33
Municipal	5.2	209	145	-0.01	0.86	-0.13	0.13
Overlay and Counterparty Risk	11.8	450	397	-0.01	1.79	-0.31	0.30
Rates	7.5	1,233	995	-0.01	2.14	-0.39	0.36
Securitized and Structured Products	6.7	1,120	773	-0.02	1.69	-0.31	0.26

Panel B: Clusters	Classified Factors (Count)		Clusters		Singleton Clusters (%)	
	Pre-2021	Post-2020	Pre-2021	Post-2020	Pre-2021	Post-2020
Commodities	580	353	103	93	9	12
Credit	1610	1032	207	180	1	2
Equities	1567	1238	274	254	3	6
Financing and Collateral	563	454	128	126	5	10
Foreign Exchange	814	415	110	102	1	3
Municipal	174	124	30	28	10	14
Overlay and Counterparty Risk	441	373	70	80	1	11
Rates	1190	940	180	194	3	4
Securitized and Structured Products	1071	719	133	97	5	1

Table 3: Risk-Factor Cluster Summary Statistics. Table summarizes the risk factors and their clusters; data calculated from daily factor returns. Count is the number of the set of unique risk factors-desk observations reported across all desks and asset classes before factors that last for less than a year. Profitability moments are calculated from the date × desk × risk factor panel. Classified factors are the count of factors that are classified; the difference between total factors and classified factors are those that do not have data for at least a year. Clusters is the number of risk factor clusters we estimated following the methodology described in section 3.2. Singleton clusters columns give the number of clusters that contain a single factor only.

	Value-Weighted Index		Equal-Weighted Index		vs. Aggregate VW Index			vs. AC VW Index			vs. Fama-French Mkt-RF		
	Mean	Stdev	Mean	Stdev	b	t-stat	R^2 (%)	b	t-stat	R^2 (%)	b	t-stat	R^2 (%)
Commodities	0.039	0.203	0.023	0.113	0.35	0.32	0.22	0.14	0.41	0.22	0.00	-0.03	0.00
Credit	-0.036	0.139	-0.012	0.081	0.31	0.25	0.26	0.11	0.19	0.31	-0.01	-0.13	0.00
Equities	0.057	0.107	0.021	0.077	0.30	0.31	0.19	0.25	0.39	0.20	0.00	0.01	0.00
Financing and Collateral	-0.042	0.353	-0.003	0.145	0.22	0.24	0.17	0.11	0.37	0.21	0.00	-0.02	0.00
Foreign Exchange	0.006	0.163	0.004	0.113	0.38	0.35	0.25	0.15	0.35	0.29	-0.01	-0.14	0.00
Municipal	-0.040	0.300	-0.033	0.211	0.07	0.07	0.49	0.02	0.05	0.62	0.01	0.16	0.01
Overlay and Counterparty Risk	-0.018	0.237	-0.023	0.120	0.31	0.33	0.25	0.10	0.25	0.61	-0.01	-0.12	0.00
Rates	-0.008	0.141	-0.007	0.085	0.38	0.36	0.25	0.15	0.29	0.24	0.00	-0.09	0.00
Securitized and Structured Products	-0.033	0.146	-0.048	0.116	0.37	0.30	0.31	0.18	0.26	0.49	-0.01	-0.15	0.01
Aggregate	0.005	0.067	-0.004	0.049									

Table 4: Risk-Factor Variation Regressions. Table summarizes risk factor indices and their ability to explain variation across risk factors. The first four columns present the moments from the return indices calculated either equal weighted (across all risk factors, or across all risk factors within an asset class) and an analogous value-weighted index which uses the standard deviation of a risk factor's profits in the previous month as weights, for risk factors with at least 15 trading days of data in the previous month. Remaining columns show the median beta and R^2 across factors within an asset class from regression of Risk Factor$_t^i = \alpha + \beta \text{Index}_t + \varepsilon_t$, where Index$_t$ is either the value-weighted aggregate risk factor index, the value-weighted risk factor index calculated across risk factors within the same asset class, or the Fama–French equity market excess return. R^2 columns are in percent (e.g., 0.20 = 0.20%).

Asset Class	Period	Correlation Threshold						
		0.05	0.1	0.15	0.2	0.25	0.5	0.75
Commodities	Pre-2021	37	103	155	186	224	300	316
	Post-2020	30	93	139	190	218	291	301
Credit	Pre-2021	89	230	350	470	591	920	984
	Post-2020	68	180	277	374	460	676	736
Equities	Pre-2021	103	273	420	576	707	960	988
	Post-2020	102	254	407	544	636	849	875
Foreign Exchange	Pre-2021	46	117	177	231	281	387	401
	Post-2020	43	102	167	212	259	318	327
Financing and Collateral	Pre-2021	39	92	135	162	183	238	241
	Post-2020	53	126	183	228	273	332	340
Municipal	Pre-2021	14	34	45	54	64	89	95
	Post-2020	10	28	40	51	58	82	90
Overlay and Counterparty Risk	Pre-2021	27	76	121	166	206	321	344
	Post-2020	28	80	124	168	195	263	275
Rates	Pre-2021	70	175	281	375	464	672	709
	Post-2020	64	194	298	397	480	646	674
Securitized and Structured Products	Pre-2021	56	133	192	240	301	469	533
	Post-2020	37	97	152	200	252	388	444

Table 5: Cluster Count by Correlation Cutoff Threshold. Table shows the number of clusters resulting from the clustering process described above; all columns are based on cluster using the absolute value of correlations of the trailing 3-month return proxy within an asset class. The columns show the resulting clusters when adjusting the required (absolute value) correlation threshold when sorting risk factors into clusters.

	Pairs		AUC			Within Hit Rate			Between Hit Rate		
Asset Class	Within	Between	Full	Odd	Even	Odd	Even	Both	Odd	Even	Both
Commodities	2,549	148,937	0.84	0.75	0.75	0.70	0.68	0.48	0.29	0.29	0.14
Credit	13,555	1,084,396	0.78	0.71	0.70	0.66	0.63	0.45	0.32	0.32	0.17
Equities	9,628	1,454,589	0.83	0.73	0.74	0.64	0.68	0.44	0.27	0.27	0.13
Financing and Collateral	2,035	203,223	0.86	0.78	0.73	0.72	0.64	0.46	0.25	0.25	0.12
Foreign Exchange	3,959	253,993	0.80	0.72	0.69	0.68	0.61	0.43	0.32	0.31	0.15
Municipal	973	17,407	0.76	0.69	0.68	0.69	0.72	0.55	0.39	0.41	0.23
Overlay and Counterparty Risk	4,720	126,026	0.79	0.74	0.68	0.74	0.65	0.51	0.34	0.34	0.18
Rates	9,466	828,626	0.81	0.74	0.70	0.70	0.63	0.46	0.31	0.32	0.16
Securitized and Structured Products	11,981	515,050	0.79	0.72	0.72	0.71	0.72	0.57	0.36	0.37	0.20

Table 6: Within-Cluster Comovement. Table reports diagnostics that validate the correlation-based clustering of risk factors within each asset class. For each asset class, the first two columns report the number of factor pairs classified as within the same cluster ("Within") versus in different clusters ("Between"), restricting to pairs within the same asset class. Columns *AUC* report the area-under-the-curve statistic $\Pr(|\rho|_{\text{within}} > |\rho|_{\text{between}})$ computed using absolute correlations estimated on the full sample, odd trading days only, and even trading days only. Columns *Within Hit Rate* report the share of within-cluster pairs with $|\rho| \geq 0.10$ when correlations are estimated on odd days, on even days, and on both subsamples simultaneously; *Between Hit Rate* reports the same quantities for between-cluster pairs. AUC values substantially above 0.5 and large gaps in hit rates between within- and between-cluster pairs indicate that clusters capture persistent return comovement rather than correlations arising from sampling noise around the $|\rho| = 0.10$ cutoff.

	$\lvert\rho_{ij}\rvert$					
	(1)	(2)	(3)	(4)	(5)	(6)
Same Bank	1.00***			1.05 ***	0.97***	1.05 ***
	(11.75)			(11.70)	(11.41)	(11.70)
Same Asset Class		0.08 *		0.11 **		−0.03
		(1.81)		(2.13)		(−0.69)
Same Cluster			14.08***		13.22***	13.25 ***
			(54.53)		(45.35)	(45.03)
Same Bank × Same AC				−0.33 ***		−0.46 ***
				(−3.58)		(−4.98)
Same Bank × Same Cluster					2.16***	2.54 ***
					(5.73)	(6.75)
Constant	8.78***	8.96***	8.95***	8.76 ***	8.76***	8.76 ***
	(109.46)	(109.32)	(109.84)	(109.10)	(109.20)	(109.10)
N	34,582,464	34,582,464	34,582,464	34,582,464	34,582,464	34,582,464
R^2	0.001	0.000	0.002	0.001	0.002	0.002

Table 7: Bank Specialization in Risk Factors. Table reports regressions of risk factor pair (absolute value of) correlations on dummies for whether the risk factors are in the same bank, same asset class, or same cluster. Sample includes all pairs that have at least one year of overlapping daily observations. Standard errors are two way clustered on the two factors in the pair. Coefficients are multiplied by 100 for readability.

	N	Mean	StdDev	Min	Max
Cluster HHI	2930779	8870.57	1909.71	1758.26	10000.00
ln(1+|PnL|)	2930779	10.53	4.02	0.00	19.89
Correlation			-0.295***		

Table 8: Bank Specialization in Clusters. Table reports summary statistics for two cluster-level measures computed at the cluster-by-day level. Cluster HHI is the Herfindahl-Hirschman Index of bank market shares within each cluster on each date, where market shares are defined as a bank's contribution to the cluster's value-at-risk proxy on a given day. The cluster's value-at-risk is proxied by the sum of member risk factors' 90-day P&L volatilities. $\ln(1 + |\text{PnL}|)$ is the log of one plus the |PnL| in a cluster on a given day, capturing cluster size in dollar risk terms. The bottom panel reports the pairwise correlation between the two measures where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Based on risk clusters estimated in the post-2020 sample.

	Return						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Specialist	0.0803**	0.0751**	0.0546*	0.0577*	0.0723***	0.0895	0.0923***
	(2.21)	(2.18)	(1.79)	(1.81)	(5.38)	(.)	(7.41)
Excess Specialization							−0.4827***
							(−2.93)
Date FE	Yes						
Bank FE			Yes				
Desk FE					Yes	Yes	Yes
Asset Class × Date FE		Yes	Yes	Yes	Yes	Yes	Yes
Bank × Date FE				Yes		Yes	Yes
N	1,115,060	1,115,060	1,115,060	1,115,060	1,115,058	1,115,058	1,114,983
R^2	0.004	0.004	0.002	0.002	0.001	0.002	0.003

Table 9: Specialist Returns. Table reports estimates from regressions of risk-adjusted returns on a measure of desk specialization. The dependent variable is daily P&L divided by lagged Value-at-Risk at the desk level. *Specialist* is the desk's average share of the bank's absolute P&L over the previous 3 months (excluding date t), standardized to have zero mean and unit variance. Returns are winsorized at the 5th and 95th percentiles. Within R^2 reported. t-statistics shown using robust standard errors clustered by date and desk where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Unattributable Share					
	Main			Wider Definition		
	(1)	(2)	(3)	(4)	(5)	(6)
Specialist	−0.011**	−0.011**	−0.010**	−0.011**	−0.011**	−0.010**
	(−2.63)	(−2.36)	(−2.33)	(−2.57)	(−2.36)	(−2.32)
Bank FE	Yes	Yes		Yes	Yes	
Date FE	Yes			Yes		
Asset Class × Date FE		Yes	Yes		Yes	Yes
Bank × Date FE			Yes			Yes
N	1,163,375	1,163,375	1,163,375	1,168,093	1,168,093	1,168,093
R^2	0.003	0.003	0.003	0.004	0.004	0.003

Table 10: Specialization and Model Accuracy. Table reports regressions of the unattributable share on measures of bank expertise. The main unattributable share is the unexplained P&L (Item 10 under the FR VV-1 reporting requirements) as a share of P&L calculated from existing positions' risk factors, carry, and the unattributable component, with all components entering in absolute value. The wider definition also includes "other attributable" P&L. Specialist is the standardized trailing 90-day average share of a bank's activity allocated to each desk. Specialist is standardized to have zero mean and unit variance within each panel. Within R^2 reported. t-statistics shown using robust standard errors clustered by date and desk where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Number of Risk Factors			Unattributable Share		
	(1)	(2)	(3)	(4)	(5)	(6)
Specialist	0.114**	0.114**	0.113**	−0.010**	−0.006*	
	(2.35)	(2.49)	(2.47)	(−2.26)	(−1.72)	
Num. of Risk Factors					−0.015*	−0.017**
					(−2.00)	(−2.26)
Bank FE	Yes	Yes				
Date FE	Yes					
Asset Class × Date FE		Yes	Yes	Yes	Yes	Yes
Bank × Date FE			Yes	Yes	Yes	Yes
N	1,106,419	1,106,419	1,106,419	1,148,267	1,091,913	1,092,706
R^2	0.029	0.030	0.030	0.003	0.006	0.005

Table 11: Specialization and Risk Model Granularity. Table reports regressions of desk-level risk model granularity and attribution quality on specialization. In columns (1)–(3) the dependent variable is the log number of distinct risk factors to which the desk attributes P&L over the trailing 90 days. In columns (4)–(5) the dependent variable is the unattributable share, defined as in Table 10. Specialist is the desk's trailing 90-day average share of the bank's gross risk factor activity, standardized to have zero mean and unit variance within the panel. All columns control for the log of the desk's VaR. t-statistics shown using robust standard errors clustered by date and desk where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	# Distinct Desks over Horizon									
	Level		Volatility		Average RENTD		Total RENTD		Monthly	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Activity Level	0.511***	0.471***								
	(20.95)	(19.76)								
Activity Volatility			0.429***	0.395***						
			(21.40)	(20.05)						
Average RENTD Limit					1.042***	0.890***			1.120***	
					(2.92)	(2.63)			(2.90)	
Total RENTD Limit							1.013***	0.821***		1.107***
							(3.20)	(2.86)		(3.20)
Horizon	1m	3m	1m	3m	1m	3m	1m	3m	1m	1m
Cluster FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Asset Class × Date FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
N	2,947,104	2,901,627	2,944,676	2,899,199	262,276	252,986	262,276	252,986	12,540	12,540
R^2	0.072	0.065	0.057	0.052	0.037	0.033	0.041	0.033	0.039	0.045

Table 12: Demand and Resource Allocation. Table reports regressions of future cluster participation on measures of demand. The dependent variable is the maximum number of distinct desks active in the cluster over the subsequent 30 or 365 days. Three demand proxies are considered: *Activity Level* is the log trailing 90-day average of total gross risk factor P&L across all banks in the cluster. *Activity Volatility* is the log trailing 90-day standard deviation of the same measure. *Average RENTD Limit* is derived from banks' internal position limits justified by Reasonably Expected Near-Term Demand analysis. For each USD-denominated RENTD limit, we take the maximum absolute value of the upper and lower bounds as a measure of demand capacity. We aggregate to the cluster-date level by averaging across limits within each bank-cluster and then across banks. *Total RENTD Limit* aggregates to the cluster-date level by averaging across limits within each bank-cluster and then adding across banks. All independent variables are log-transformed and standardized to have zero mean and unit variance. Independent variables are standardized to have zero mean and unit variance within each panel. Within R^2 reported. t-statistics shown using robust standard errors clustered by date and cluster where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	RENTD Limit		Proximity		𝕀(Near Capacity)		𝕀(Near Cap, not Breach)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Specialist	0.34***	0.51***	0.11***	0.18***	0.01***	0.02***	0.01**	0.02 ***
	(3.06)	(4.12)	(2.78)	(5.11)	(2.63)	(3.11)	(2.38)	(2.94)
Cluster × Date FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Bank × Date FE		Yes		Yes		Yes		Yes
N	198,217	198,217	198,217	198,217	198,217	198,217	198,217	198,217
R^2	0.032	0.074	0.017	0.045	0.005	0.011	0.004	0.009

Table 13: Specialists and Risk Limits. Table reports regressions of the risk factor RENTD limit size and proximity on the bank specialization in the risk factor. Each observation is a risk factor (bank × desk × risk factor) on a given day. The dependent variables are the RENTD limit, proximity to the limit, and indicator variables related to proxmity. Proximity measures how close current usage is to the allowable bound (0 = maximum slack, 100 = at limit). 𝕀(Near Capacity) is an indicator for proximity above 80. 𝕀(Near Cap, not Breach) is an indicator for proxmitity between 80 and 100. Specialist is the risk factor's trailing 90-day share of the bank's gross risk factor activity, standardized to mean zero and unit variance. Controls include the log of current limit size and the trailing 90-day standard deviation of log limit levels. Limited to USD measured RENTD limits which map to risk factors. All specifications include cluster-by-date fixed effects and some include bank-by-date fixed effects. Within R^2 reported. t-statistics in parentheses using standard errors clustered by date and risk limit where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	ΔLimit					
	Overall			Near Limit		
	14-Day	30-Day	90-Day	14-Day	30-Day	90-Day
	(1)	(2)	(3)	(4)	(5)	(6)
Specialist	0.05 **	0.15 **	0.50 **	0.03	0.10 *	0.40 *
	(2.07)	(2.13)	(2.08)	(1.46)	(1.71)	(1.81)
Proximity × Specialist				−0.02	−0.04	−0.06
				(−1.63)	(−1.52)	(−0.77)
Proximity				0.25 ***	0.55 ***	1.11 ***
				(9.18)	(7.87)	(6.59)
Controls						
ln(Limit Size)	−0.25 ***	−0.76 ***	−2.79 ***	−0.31 ***	−0.90 ***	−3.09 ***
	(−4.35)	(−4.70)	(−5.11)	(−5.36)	(−5.44)	(−5.59)
Vol(ln(Limit))	−0.05	−0.07	0.17	−0.21 **	−0.42	−0.55
	(−0.45)	(−0.21)	(0.20)	(−2.02)	(−1.37)	(−0.67)
Limit FE	Yes	Yes	Yes	Yes	Yes	Yes
Cluster × Date FE	Yes	Yes	Yes	Yes	Yes	Yes
Bank × Date FE	Yes	Yes	Yes	Yes	Yes	Yes
N	197,456	197,459	197,469	197,456	197,459	197,469
R^2	0.005	0.011	0.034	0.017	0.025	0.047

Table 14: Specialists Near Risk Limits. Table reports regressions of the percentage change in the risk factor RENTD limit size on the interaction between proximity to the limit and bank specialization in the risk factor. Each observation is a risk factor (bank × desk × risk factor) on a given day. The dependent variable is $100 \times [\ln(\overline{\text{Limit}}_h^{\text{fwd}} + 1) - \ln(\text{Limit} + 1)]$, where the forward limit is averaged over the subsequent 14, 30, or 90 days. Proximity measures how close current usage is to the allowable bound (0 = maximum slack, 100 = at limit). Specialist is the risk factor's trailing 90-day share of the bank's gross risk factor activity, standardized to mean zero and unit variance. Controls include the log of current limit size and the trailing 90-day standard deviation of log limit levels. Limited to USD measured RENTD limits which map to risk factors. All specifications include cluster-by-date and bank=by-date fixed effects. Within R^2 reported. t-statistics in parentheses using standard errors clustered by date and risk limit where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	ΔActivity$_{bft}$					
	1-Day		3-Day		5-Day	
	(1)	(2)	(3)	(4)	(5)	(6)
Near Capacity$_{bt}$ × Specialist$_{bft}$	14.232 ***		10.938 ***		6.589 ***	
	(3.96)		(3.68)		(2.63)	
Utilization$_{bt}$ × Specialist$_{bft}$		9.040 **		6.081 *		1.030
		(2.12)		(1.80)		(0.39)
Near Capacity$_{bt}$	−34.117 ***		−31.720 ***		−28.390 ***	
	(−5.16)		(−5.45)		(−5.77)	
Utilization$_{bt}$		−38.324 ***		−35.729 ***		−31.163 ***
		(−4.60)		(−4.86)		(−5.02)
Specialist$_{bft}$	3.062 ***	0.931	2.717 ***	1.670	3.454 ***	4.643 **
	(3.48)	(0.35)	(3.29)	(0.77)	(3.90)	(2.45)
N	6,573,316	6,573,316	8,206,432	8,206,432	8,516,633	8,516,633
R^2	0.001	0.001	0.001	0.001	0.001	0.001

Table 15: Binding Constraints and Reallocation Toward Core Expertise. Table reports regressions of changes in bank-risk-factor activity on the interaction between bank-level constraint tightness and risk-factor-level specialization. The dependent variable is $100 \times [\ln(\overline{\text{Activity}}^{\text{fwd}} + 1) - \ln(\overline{\text{Activity}}^{\text{trail}} + 1)]$, where forward activity is averaged over the subsequent 1, 3, or 5 days and trailing activity over the previous 90 days. Specialist is the standardized trailing 90-day average share of a bank's gross risk factor activity allocated to each risk factor. Utilization is the VaR-weighted average of each desk's maximum limit proximity score across the bank, where proximity measures how close usage is to the allowable bound (0 = maximum slack, 1 = at limit). Near Capacity is the VaR-weighted share of desks with proximity above 80. All specifications include bank-by-date and cluster-by-date fixed effects. Within R^2 reported. t-statistics shown using robust standard errors clustered by date and bank-cluster where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	ΔActivity$_{bct}$					
	1-Day		3-Day		5-Day	
	(1)	(2)	(3)	(4)	(5)	(6)
Other Banks Proximity$_{bct}$ × Specialist$_{bct}$	5.134**		5.255**		5.618**	
	(2.06)		(2.29)		(2.20)	
Other Banks Near Capacity$_{bct}$ × Specialist$_{bct}$		3.395*		2.560*		2.769**
		(1.70)		(1.66)		(2.02)
N	180,573	298,752	224,065	371,159	231,117	383,560
R^2	0.001	0.000	0.002	0.000	0.003	0.000

Table 16: Competitor Constraints and Specialist Expansion. Table reports regressions of changes in bank-cluster activity on the interaction between the bank's cluster-level specialization and the constraint tightness of other banks in the cluster. The dependent variable is $100 \times [\ln(\overline{\text{Activity}}^{\text{fwd}} + 1) - \ln(\overline{\text{Activity}}^{\text{trail}} + 1)]$, where forward activity is averaged over the subsequent 1, 3, or 5 days and trailing activity over the previous 90 days. Specialist is the standardized trailing 90-day average share of a bank's gross risk factor activity allocated to each risk factor. Proximity measures how close usage is to the allowable bound (0 = maximum slack, 1 = at limit). Near Capacity is the VaR-weighted share of desks with proximity above 80. All specifications include bank-by-cluster, bank-by-date, and cluster-by-date fixed effects. Within R^2 reported. t-statistics shown using robust standard errors clustered by date and cluster where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Number of Postings				I(Any Posting)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AC Share × Demand Shock	0.064**	0.103*			0.014***	0.008		
	(2.15)	(1.92)			(3.47)	(1.16)		
Specialist × Demand Shock			0.085***	0.120**			0.018***	0.015**
			(3.47)	(2.18)			(4.54)	(2.24)
Demand Shock	0.022 *		0.017		−0.003		−0.006	
	(1.97)		(1.59)		(−0.31)		(−0.61)	
AC Share	0.175	0.120			0.022	0.007		
	(1.03)	(0.81)			(1.05)	(0.21)		
Specialist			0.308	0.102			0.013	−0.028
			(1.35)	(0.50)			(0.47)	(−0.63)
Bank × Asset Class FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Bank × Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Asset Class × Month FE		Yes		Yes		Yes		Yes
N	2,830	2,830	3,042	3,042	2,830	2,830	3,042	3,042
R^2	0.002	0.002	0.003	0.002	0.001	0.000	0.001	0.001

Table 17: Demand Shocks and Specialist Banks' Labor Demand. Table reports regressions of banks' job postings on specialization interacted with demand shocks, at the bank–asset-class–month level. The dependent variable is the number of the bank's trading-related job postings in the asset class during the month or an indicator equal to one if the bank posts any opening in the asset class that month. All independent variables are standardized to have zero mean and unit variance. Within R^2 reported. t-statistics in parentheses using standard errors clustered by month and bank-asset-class where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Destination			
Origin	Peripheral	Core	Total	% Core
Core	41	256	297	86%
Peripheral	29	68	97	70%
Total	71	324	394	

Table 18: Transition Matrix: Origin Peripherality and Destination Specialization. Table reports the number of trader moves between covered banks, classified by whether the trader's asset class was peripheral at the origin bank and core at the destination bank. A trader's asset class is classified as peripheral (core) if the origin bank's trailing share of gross risk factor activity in that asset class falls below (above) the bank's within-month median. Destination core is defined symmetrically. Sample includes all moves between covered banks where both origin and destination specialization data are available. The final column reports the share of movers landing at a core destination.

	I(Moved)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Origin AC Share	−0.024**	−0.040***	−0.040***	−0.048**	−0.039***			
	(−2.65)	(−3.94)	(−3.82)	(−2.20)	(−3.97)			
Origin AC Share × Bank Constraint					0.004			
					(0.59)			
Bank Constraint					0.031 *	0.032*		
					(1.69)	(1.81)		
Seniority			−0.042***		−0.042***		−0.063***	−0.063***
			(−6.57)		(−6.56)		(−6.66)	(−6.69)
Demand							2.218	
							(1.40)	
Demand Shock								1.191
								(1.44)
Asset Class FE	Yes							
Date FE	Yes							
Bank FE		Yes	Yes	Yes	Yes	Yes	Yes	Yes
Asset Class × Date FE		Yes	Yes	Yes	Yes	Yes	Yes	Yes
Sample	All	All	All	Senior	All	All	All	All
N	432,838	432,838	432,838	53,425	432,838	436,821	196,433	196,433
R^2	0.000	0.000	0.000	0.000	0.000	0.000	0.002	0.000

Table 19: Trader Sorting: Peripheral Traders Leave and Meet Demand. Table reports regressions of an indicator for departure to any firm on the origin bank's share of activity in the trader's asset class. The sample is all trader-months at covered banks, excluding final observations; column (4) restricts to senior traders (seniority ≥ 6). Dependent variables multiplied by 100 for readability. t-statistics in parentheses using standard errors clustered by month and bank–asset class where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Δ(AC Share)							
	Movers				Internal Movers			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Peripheral	6.937***	7.235***	4.377	4.909 *	18.268***	18.422***	15.011***	15.045***
	(2.76)	(2.93)	(1.70)	(1.85)	(14.78)	(14.90)	(10.21)	(10.60)
Demand Shock			−0.170	−1.075			−1.539***	−1.494***
			(−0.31)	(−0.95)			(−5.83)	(−5.47)
Peripheral × Demand Shock			2.828***	3.576***			−0.305	−0.288
			(3.93)	(3.35)			(−0.91)	(−0.91)
Seniority		−1.613 **		−1.694		−0.669 *		−1.078 *
		(−2.02)		(−1.16)		(−1.68)		(−1.93)
Asset Class FE	Yes	Yes						
Bank FE					Yes	Yes	Yes	Yes
Date FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	205	205	77	77	832	832	355	355
R^2	0.030	0.067	0.063	0.106	0.206	0.209	0.185	0.192

Table 20: Trader Sorting: Peripheral Traders Meet Demand and Improve Match Quality. Table reports regressions of an indicator for departure to any firm and the change in asset class share on demand. In columns (5) to (8), the sample is restricted to movers between covered banks that do not change their asset class. Dependent variables multiplied by 100 for readability. t-statistics in parentheses using standard errors clustered by month where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Internet Appendix to

Expertise and Specialization in Bank Trading

Chase P. Ross, Sharon Y. Ross, Kaushik Vasudevan

The Internet Appendix consists of three sections. Section IA.A provides additional details about the data. Section IA.B presents additional tables and figures. Section **??** presents proofs for the model.

IA.A Data Cleaning and Construction

We clean the data in several ways. We construct our analysis sample from regulatory filings submitted by banks under the Federal Reserve's Volcker Rule trading desk reporting requirements. The data come in two distinct reporting formats: the pre-2020 format and the post-2020 format, which we merge to create a consistent time series. The reporting instructions were updated in December 2020 to provide additional granularity on certain fields, particularly limit characteristics and desk classifications. We clean the data so observations before and after 2020 are directly comparable.

We focus on the banks that consistently report throughout the sample. All monetary values are converted to U.S. dollars using the bank-reported currency conversion rates. We exclude observations missing any of the key identifiers (bank, desk, or risk factor). We drop trading desks that never report meaningful activity throughout their listing period. Specifically, we exclude desks where both comprehensive P&L and VaR are either missing or zero for all dates. Similarly, we drop individual risk factors within a desk if their risk-move P&L is missing or zero for all dates.

We winsorize extreme observations to limit the influence of outliers. For risk factor P&L and desk-level variables (comprehensive P&L, VaR, and P&L components), we winsorize at the 0.1st and 99.9th percentiles. We balance the panel by identifying each desk's first and last active dates—defined as dates when the desk reports non-zero and non-missing P&L or VaR.

IA.A.1 Desk Classifications

We sort desks into ten groups: credit, equities, rates, foreign exchange, securitized and structured products, collateral and financing, commodities, overlay and counterparty risk, municipal, and legacy.

1. Credit: desks that trade corporate debt, credit derivatives, or support the firm's debt capital markets and leveraged loan businesses.

2. Equities: desks that trade equities, equity derivatives, or convertible products.

3. Rates: desks that trade government bonds, agency debt, interest rate swaps, swaptions, or macro products.

4. Foreign exchange: desks that trade currencies or currency derivatives.

5. Commodities: desks that trade commodities, including energy, metals, and agriculture.

6. Securitized and structured products: desks that trade asset-backed securities, securitized products (e.g., CLOs, CMOs, CDOs), non-agency mortgage products, or other structured products.

7. Municipal: desks that trade municipal, tax-exempt, or local government obligations.

8. Collateral and financing: desks that trade repos, collateral swaps, provide financing for internal counterparties or prime brokerage clients, or engage in securities financing.

9. Overlay and counterparty risk: desks that conduct valuation adjustment (VA) trading; including valuation adjustments to hedge counterparty risk (CVA), hedges to funding conditions (FVA), or desks that focus on centralized hedging of other desks' net positions or focus on regulatory capital or risk weighted asset optimization.

10. Legacy: desks that manage portfolios or lines of business that the firm is exiting or otherwise unwinding.

The classification uses a hybrid approach combining keyword matching with supervised machine learning, trained on a subset of desks where we have high confidence in the asset class assignment. For each asset class, we manually construct keyword lists based on domain knowledge of trading desk naming conventions. These lists include both standard terminology (e.g., "equity," "stock," "ipo" for the equity class) and specialized jargon (e.g., "sonia," "sofr," "eurodollar" for the rates class). We designate certain keywords as "dominant" when they unambiguously signal a particular asset class (e.g., "xva," "cva," "counterparty credit risk" for XVA desks). When a desk name or description matches a dominant keyword, we immediately assign that asset class without further disambiguation. For the remaining desks, we apply keyword matching to the combined desk name and description text. If text matches keywords from only one asset class, we assign that class. If text matches keywords from multiple asset classes we disambiguate using TF-IDF (term frequency-inverse document frequency) similarity. We exclude desks in the rare case they have no desk name or desk description from the classification, rather than setting them to unclassified.

IA.A.2 Measuring Risk Limit Utilization

We construct a unified measure of how close each trading desk is to its regulatory risk limits using data from the FR VV-1 reporting form. The measure captures the proximity of a desk's current risk exposure to its allowable bounds, expressed on a scale where 0 indicates maximum slack, 100

indicates usage exactly at the limit, and values above 100 indicate a breach. Our construction addresses reporting format changes across our sample period.

In the second subsample, banks report explicit upper and lower bounds along with the usage value for each limit, so the allowable range is directly observed. In the first subsample, banks report only a single limit size, a usage value, and a limit percent used field. To harmonize the two formats, we must determine whether each first-subsample limit represents a symmetric constraint (where exposure must stay within $\pm|\text{limit size}|$) or a directional constraint (where the limit size is a one-sided bound). For example, if a delta limit of +\$100M is symmetric, then usage of −\$80M represents 80% utilization; if it is directional, the desk is far from the constraint.

We exploit the bank's own reported limit percent used to classify each limit. When usage and the limit size have opposite signs, the signed ratio usage/limit size is negative, but a bank using an absolute value convention will report a positive limit percent used equal to |usage|/|limit size|. We flag a limit as symmetric whenever we observe a day on which the signed ratio is negative yet the reported limit percent used is approximately equal to the absolute value of that ratio (within a tolerance of one percentage point). If any such day exists for a given limit, we classify the entire limit as symmetric. Symmetric limits receive both bounds: an upper bound of $+|\text{limit size}|$ and a lower bound of $-|\text{limit size}|$. Remaining limits are treated as directional, with positive limit sizes assigned as upper bounds and negative limit sizes as lower bounds.

Given the harmonized bounds, we compute a proximity score for each limit-date observation based on how close the current usage is to the nearest binding constraint. For each limit, we compute the headroom as the distance from current usage to the nearest bound: for two-sided limits with upper bound u and lower bound l, this is $\min(u - x,\, x - l)$; for one-sided limits, it is the distance to the single reported bound. We normalize headroom by the half-width of the allowable range—$(u - l)/2$ for two-sided limits, or the absolute value of the bound for one-sided limits—so that the resulting proximity score equals 0 at the midpoint of the allowable range, 100 at either boundary, and exceeds 100 in the event of a breach. For one-sided limits, usage on the opposite side of the constraint can produce negative proximity values, indicating that the desk is moving away from the relevant bound. We clamp values at zero from below and winsorize the right tail of proximity at the 99.9th percentile, so that all observations far from their relevant constraint register as unconstrained.

We aggregate per-limit proximity scores to the desk-date level by taking the maximum across all limits on the desk, capturing the most binding constraint. As a robustness check, we also compute the mean of the top three proximity scores, which is less sensitive to outliers from data errors or discrete limit resets. To construct a bank-date constraint measure, we aggregate across desks using Value-at-Risk weights, so that a small desk near capacity contributes less than a large desk under pressure. We also construct binary indicators for whether any desk on the bank experienced a breach (proximity above 100) or was near capacity (proximity above 80) on a given day.

IA.A.3 Trader Employment Data

We use monthly trader employment data from Revelio Labs. The panel tracks individual traders at monthly frequency, recording the trader's employer, position and job description, and seniority level. We use these values to map traders to the most likely asset class. [Add sorting process]

IA.A.4 Job Postings Data

Our job posting data come from Dewey LinkUp, which compiles bank job postings scraped at daily frequency. For each posting we observe the hiring bank, job title, description, and the dates on which the job posting is first observed and subsequently removed.

We classify each posting into an asset class using a large language model applied to the posting text description. For 26 percent of postings that are unclassified after the LLM-classification, we use a keyword search over job titles. We drop postings that cannot be assigned to any asset class after this process (approximately 22 percent of postings).

We restrict our sample to trading roles, and we exclude internships, summer positions, and industrial-placement programs, identified from the job title. We also clean the data for re-uploaded chains of postings. Because postings are scraped daily, a single vacancy that is re-listed can appear as several consecutive records: a new record whose creation date coincides with the removal date of the prior record for the same bank, asset class, and title. To avoid double-counting, we chain such records together. Within each bank $\times$ asset class $\times$ title group, we treat a record as a continuation of the previous vacancy whenever its creation date equals the previous record's removal date, and as a new vacancy otherwise. We then collapse each chain to a single observation spanning the first creation date to the final removal date. This yields one observation per distinct vacancy.

We construct indicators for ten asset classes—XVA, commodities, credit, equity, foreign exchange, legacy, municipals, rates, repo, and secured products—and reshape the data to the posting $\times$ asset-class level. Each posting is attributed fractionally: a posting matching k asset classes contributes a weight of $1/k$ to each, so that the weights sum to one within a posting. We aggregate to the bank $\times$ asset-class $\times$ month level, summing the fractional weights to obtain the flow of new postings.

IA.B Figures and Tables

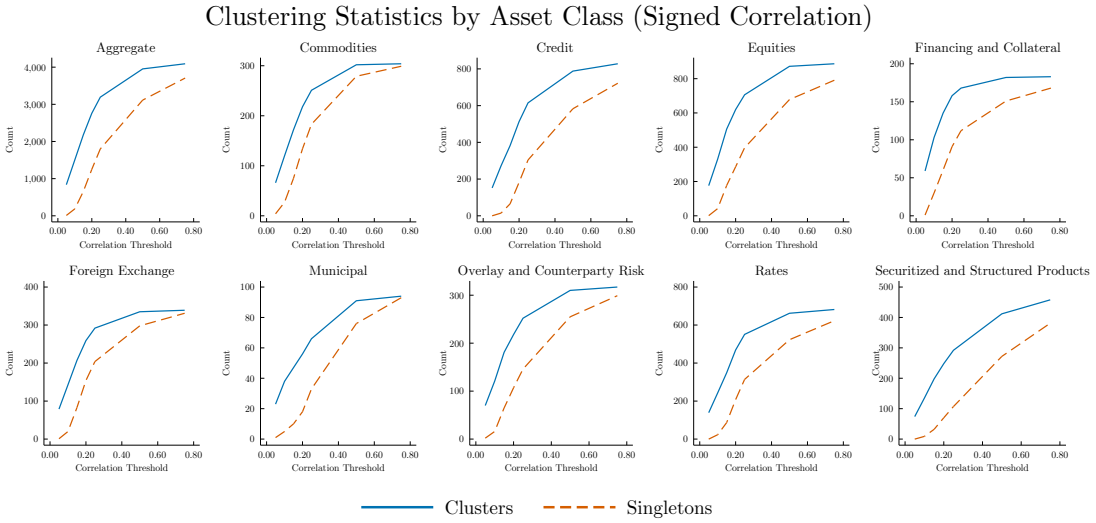


Figure IA.1: Clusters and Singletons across Correlation Thresholds. Figure plots the number of resulting clusters, along with number of singleton clusters that contain only a single risk factor, when categorizing risk factors based on the signed correlation threshold of our 90 day return proxy (as opposed to the benchmark clustering based on the absolute value of correlations). Figure based on post-2020 sample clustering.

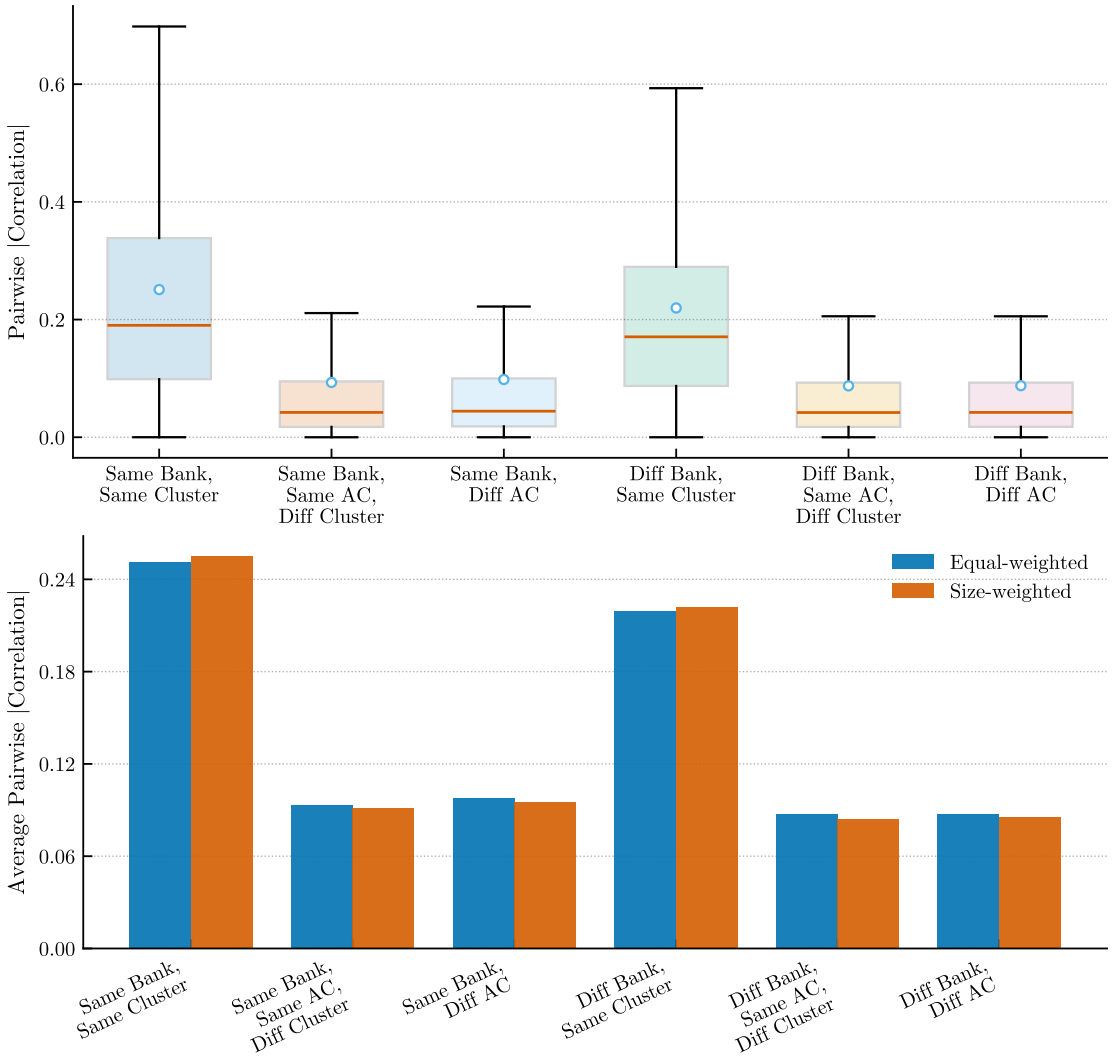


Figure IA.2: Correlations of Risk Factors. Figure plots the (absolute value) of correlations of risk factors across clusters, banks, and asset classes, where the clusters are estimated using the benchmark approach using the absolute value of our 90 day return proxy.

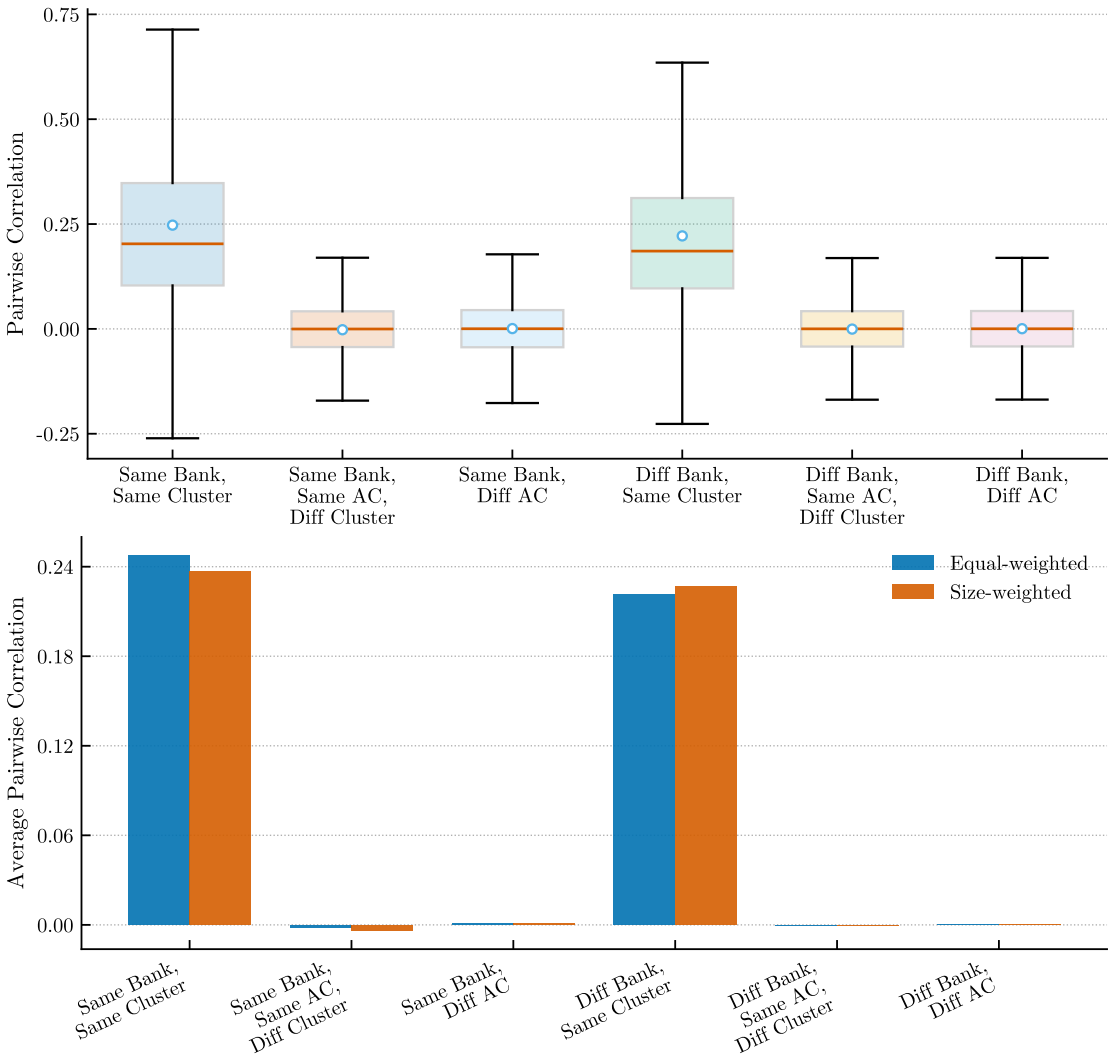


Figure IA.3: Correlations of Risk Factors when clustering with signed return proxy. Figure plots the (absolute value) of correlations of risk factors across clusters, banks, and asset classes, where the clusters are estimated using the benchmark approach using the correlations of signed 90 day return proxy.

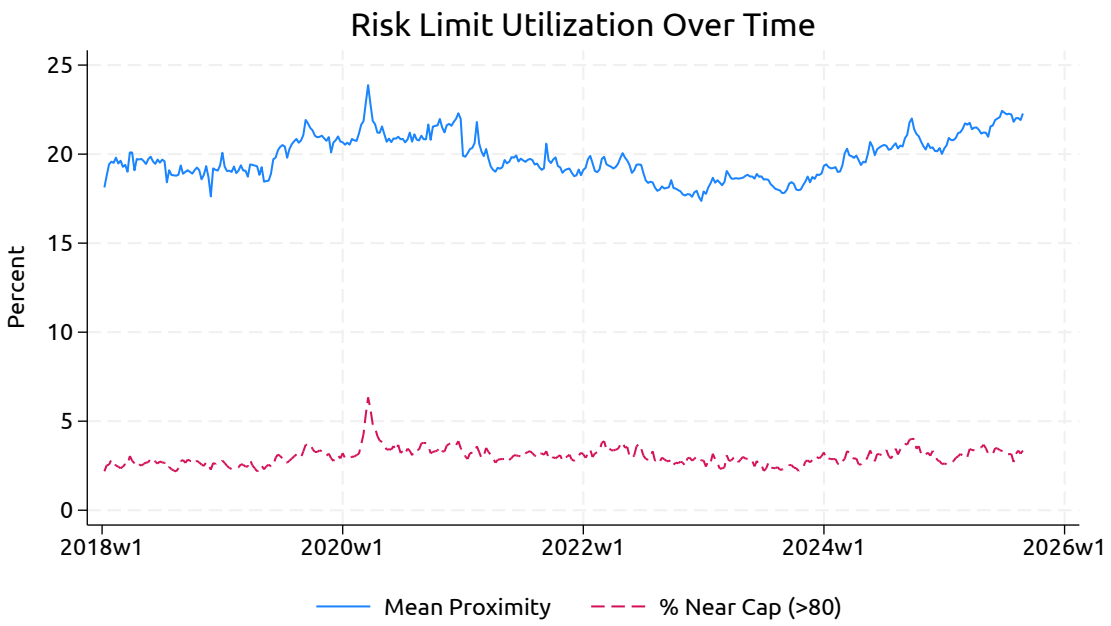


Figure IA.4: Risk Limit Utilization Over Time. Weekly averages of risk limit proximity (solid) and the share of limit-date observations near capacity (proximity above 80, dashed). Proximity measures how close current usage is to the allowable bound, on a scale where 0 indicates maximum slack and 100 indicates usage at the limit. The sample includes all limit types and sources from the FR VV-1 reporting form.

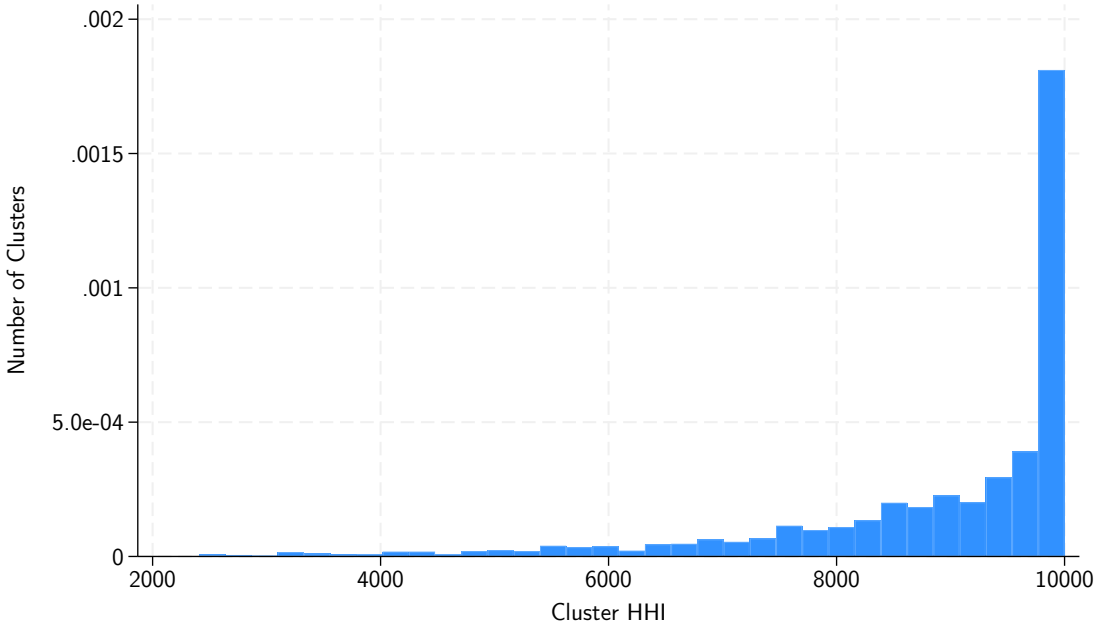


Figure IA.5: Number of Clusters vs. Cluster HHI. Figure shows a histogram of the number of clusters in each Cluster HHI bucket. Cluster HHI is the Herfindahl-Hirschman Index of bank market shares within each cluster on each date, where market shares are defined as a bank's contribution to the cluster's value-at-risk proxy on a given day. The cluster's value-at-risk is proxied by the sum of member risk factors' 90-day P&L volatilities.

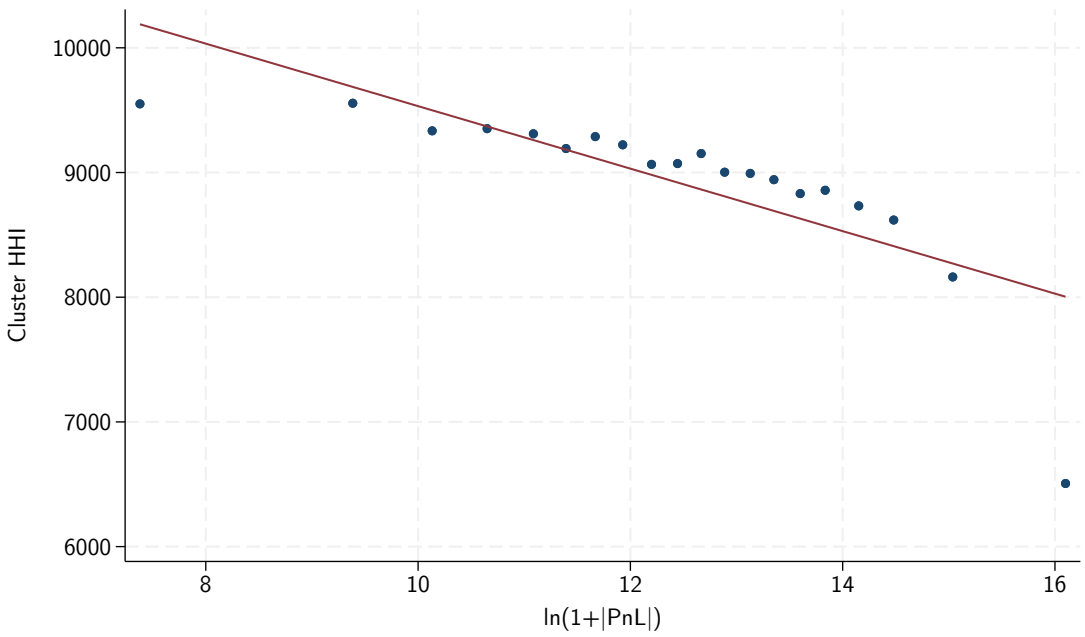


Figure IA.6: Cluster HHI vs Cluster Activity. Figure shows Cluster HHI bucket plotted against the cluster's (log of one plus) absolute PnL. Cluster HHI is the Herfindahl-Hirschman Index of bank market shares within each cluster on each date, where market shares are defined as a bank's contribution to the cluster's value-at-risk proxy on a given day. The cluster's value-at-risk is proxied by the sum of member risk factors' 90-day P&L volatilities.

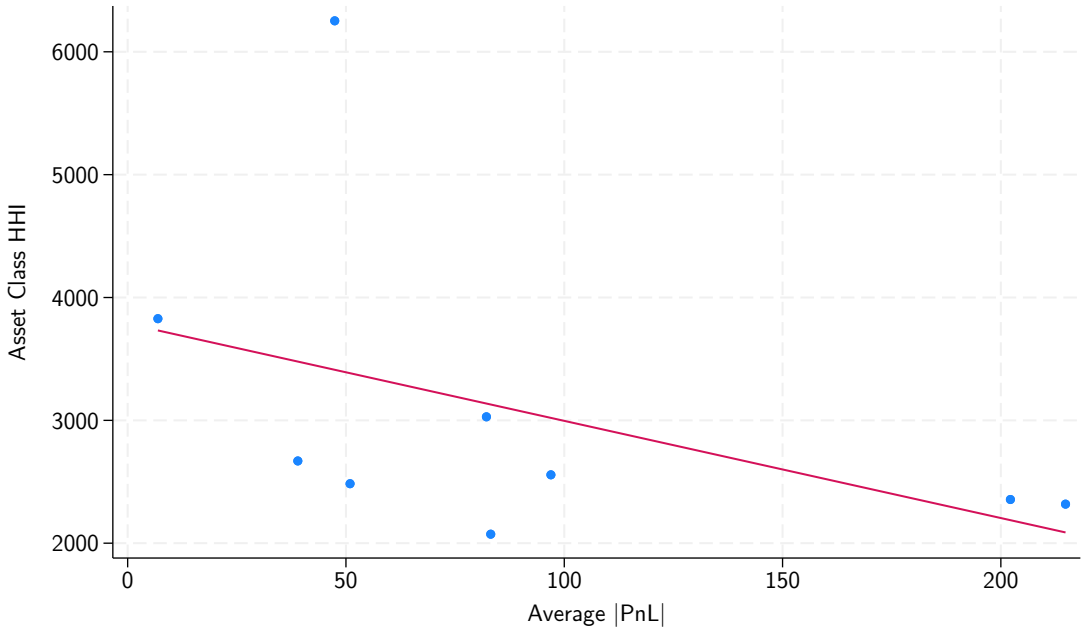


Figure IA.7: Asset Class HHI vs Asset Class Activity.

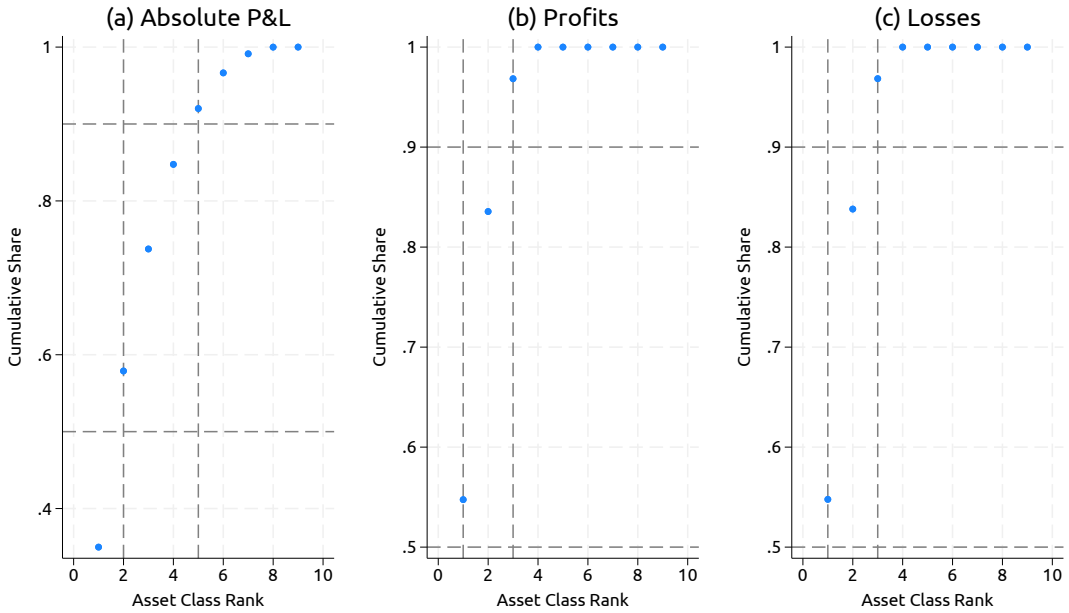


Figure IA.8: Asset Class Top K. .

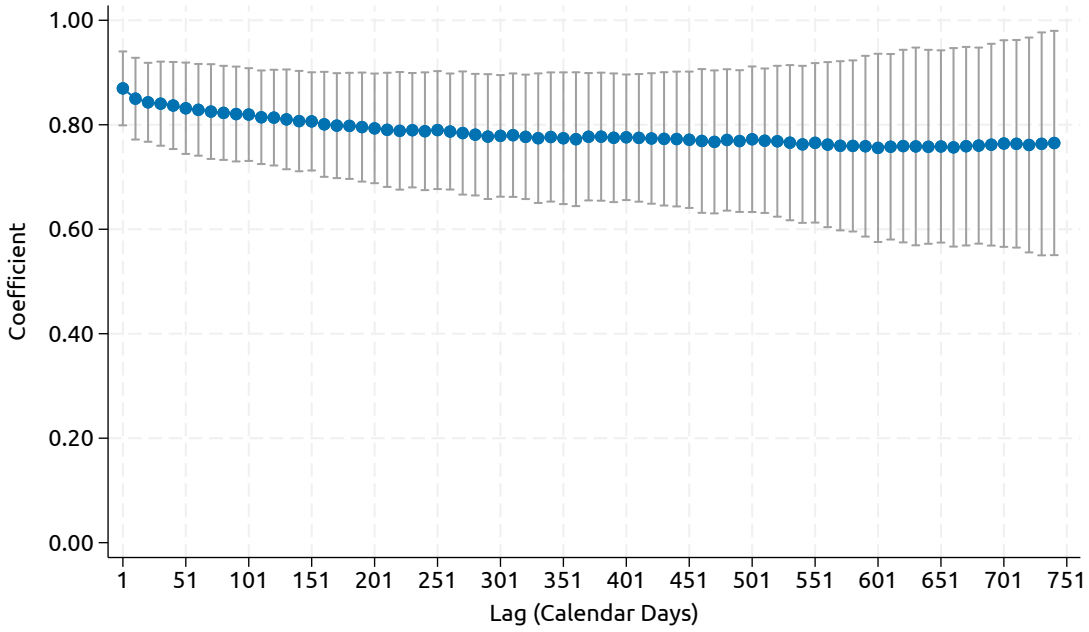


Figure IA.9: Asset Class Share Persistence. TK.

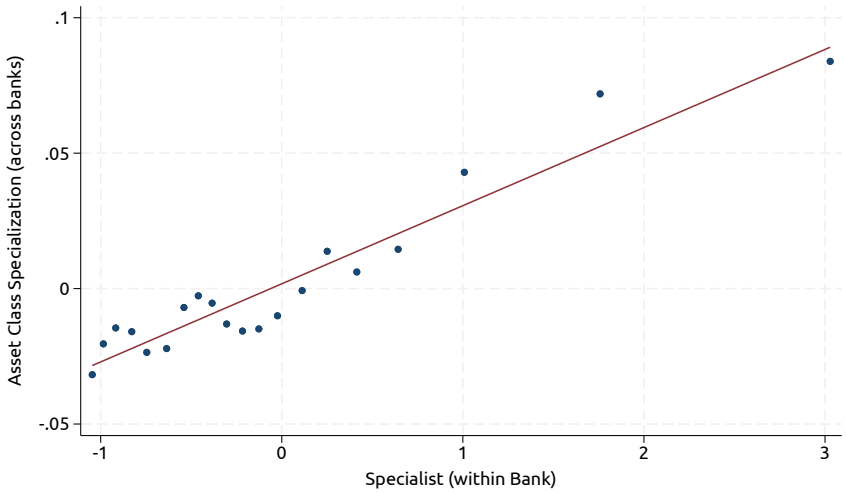


Figure IA.10: Asset Class Specialization. Y-axis is asset class specialization across banks, constructed analogously to Blickle et al. (2025). X-axis is our measure of *Specialist*, which is defined within a bank.

	Desk VaR			
	(1)	(2)	(3)	(4)
SD(Desk PnL)	0.888***	0.887***	1.141***	1.103***
	(3.19)	(3.19)	(4.53)	(4.15)
Desk FE	No	No	Yes	Yes
Date FE	No	Yes	No	Yes
N	1,190,984	1,190,984	1,190,984	1,190,984
R^2	0.278	0.279	0.217	0.203

Table IA.1: Correlation between Desk VaR and Standard Deviation of PnL. Table reports regressions of desk Value-at-Risk on the trailing 90-day standard deviation of desk PnL. Within R^2 reported. t-statistics shown using robust standard errors clustered by date and desk where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	N	Mean	StdDev	Min	Max
Asset Class HHI	25281	3063.00	1317.07	1687.43	8779.85
ln(1+|PnL|)	25281	17.82	1.12	13.08	20.92
Correlation			-0.346***		

Table IA.2: Asset Class HHI.

Asset Class	HHI	|PnL|
Repo	6252.30	47,425,071
Muni	3827.77	6,922,816
FX	3028.98	82,175,893
XVA	2669.52	38,971,319
Credit	2556.71	96,943,752
Commodities	2484.38	50,941,398
Equity	2355.88	202,188,204
Rates	2318.08	214,818,749
Securitized Products	2073.40	83,160,702

Table IA.3: Asset Class HHI.

	Unattributable Share					
	Main			Wider Definition		
	(1)	(2)	(3)	(4)	(5)	(6)
Specialist	-0.011^{**}	-0.010^{**}	-0.009^{**}	-0.010^{**}	-0.008	-0.009^{*}
	(-2.58)	(-2.23)	(-2.22)	(-2.26)	(-1.67)	(-1.86)
Excess Specialization	-0.034	-0.097	-0.080	-0.051	-0.141^{**}	-0.103
	(-0.57)	(-1.55)	(-1.32)	(-0.73)	(-2.09)	(-1.58)
Bank FE	Yes	Yes		Yes	Yes	
Date FE	Yes			Yes		
Asset Class × Date FE		Yes	Yes		Yes	Yes
Bank × Date FE			Yes			Yes
N	1,163,300	1,163,300	1,163,300	1,168,018	1,168,018	1,168,018
R^2	0.004	0.005	0.004	0.003	0.005	0.004

Table IA.4: Specialist Unattributable Share Robustness. Table reports regressions of the unattributable share on measures of bank expertise. The main unattributable share is the unexplained P&L (Item 10 under the FR VV-1 reporting requirements) as a share of P&L calculated from existing positions' risk factors, carry, and the unattributable component, with all components entering in absolute value. The wider definition also includes "other attributable" P&L. Specialist is the standardized trailing 90-day average share of a bank's activity allocated to each desk. Specialist is standardized to have zero mean and unit variance within each panel. Excess Specialization is the bank's specialization compared to the market. Within R^2 reported. t-statistics shown using robust standard errors clustered by date and desk where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	$\lvert\rho_{ij}\rvert$					
	(1)	(2)	(3)	(4)	(5)	(6)
Same Bank	0.02 *			0.03 ***	−0.01	0.03 ***
	(1.74)			(2.83)	(−0.89)	(2.83)
Same Asset Class		0.03***		0.05 ***		−0.10 ***
		(2.94)		(4.17)		(−8.90)
Same Cluster			22.97***		22.14***	22.22 ***
			(91.57)		(77.16)	(77.32)
Same Bank × Same AC				−0.07 ***		−0.18 ***
				(−3.57)		(−8.35)
Same Bank × Same Cluster					2.57***	2.72 ***
					(6.35)	(6.68)
Constant	0.05***	0.05***	0.03***	0.05 ***	0.03***	0.05 ***
	(9.71)	(9.95)	(6.44)	(8.43)	(6.29)	(8.43)
N	34,582,464	34,582,464	34,582,464	34,582,464	34,582,464	34,582,464
R^2	0.000	0.000	0.002	0.000	0.002	0.002

Table IA.5: Bank Specialization in Risk Factors when clustering on signed returns. Table reports regressions of risk factor pair correlations on dummies for whether the risk factors are in the same bank, same asset class, or same cluster. Identical to Table 7 except the clustering here is calculated by using signed returns rather than the absolute value of returns. Sample includes all pairs that have at least one year of overlapping daily observations. Standard errors are two way clustered on the two factors in the pair. Coefficients are multiplied by 100 for readability.

	# Distinct Banks over Horizon									
	Level		Volatility		Average RENTD		Total RENTD		Monthly	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Activity Level	0.205***	0.183***								
	(24.37)	(21.93)								
Activity Volatility			0.179***	0.160***						
			(24.13)	(21.67)						
Average RENTD Limit					0.139*	0.079			0.159*	
					(1.79)	(1.04)			(1.94)	
Total RENTD Limit							0.092	0.031		0.115
							(1.41)	(0.49)		(1.63)
Horizon	1m	3m	1m	3m	1m	3m	1m	3m	1m	1m
Cluster FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Asset Class × Date FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
N	2,947,104	2,901,627	2,944,676	2,899,199	262,276	252,986	262,276	252,986	12,540	12,540
R^2	0.061	0.049	0.052	0.042	0.005	0.002	0.002	0.000	0.006	0.004

Table IA.6: Demand and Resource Allocation Across Banks. Table reports regressions of future cluster participation on measures of demand. The dependent variable is the maximum number of distinct desks active in the cluster over the subsequent 30 or 90 days. Three demand proxies are considered: *Activity Level* is the log trailing 90-day average of total gross risk factor P&L across all banks in the cluster. *Activity Volatility* is the log trailing 90-day standard deviation of the same measure. *Average RENTD Limit* is derived from banks' internal position limits justified by Reasonably Expected Near-Term Demand analysis. For each USD-denominated RENTD limit, we take the maximum absolute value of the upper and lower bounds as a measure of demand capacity. We aggregate to the cluster-date level by averaging across limits within each bank-cluster and then across banks. *Total RENTD Limit* aggregates to the cluster-date level by averaging across limits within each bank-cluster and then adding across banks. All independent variables are log-transformed and standardized to have zero mean and unit variance. Independent variables are standardized to have zero mean and unit variance within each panel. Within R^2 reported. t-statistics shown using robust standard errors clustered by month and cluster where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	ΔActivity$_{bft}$					
	1-Day		3-Day		5-Day	
	(1)	(2)	(3)	(4)	(5)	(6)
Near Capacity$_{bt}$ × Specialist$_{bft}$	14.303***		10.860***		6.195 **	
	(3.88)		(3.57)		(2.44)	
Utilization$_{bt}$ × Specialist$_{bft}$		9.046**		6.045*		0.899
		(2.08)		(1.77)		(0.35)
Specialist$_{bft}$	3.193***	1.111	2.834***	1.817	3.644***	4.834***
	(3.57)	(0.41)	(3.41)	(0.83)	(4.14)	(2.60)
N	6,573,316	6,573,316	8,206,432	8,206,432	8,516,633	8,516,633
R^2	0.001	0.001	0.001	0.001	0.001	0.001

Table IA.7: Binding Constraints and Reallocation Toward Core Expertise, Robustness. Table reports regressions of changes in bank-risk-factor activity on the interaction between bank-level constraint tightness and risk-factor-level specialization. The dependent variable is $100 \times [\ln(\overline{\text{Activity}}^{\text{fwd}} + 1) - \ln(\overline{\text{Activity}}^{\text{trail}} + 1)]$, where forward activity is averaged over the subsequent 1, 3, or 5 days and trailing activity over the previous 90 days. Specialist is the standardized trailing 90-day average share of a bank's gross risk factor activity allocated to each risk factor. Utilization is the VaR-weighted average of each desk's maximum limit proximity score across the bank, where proximity measures how close usage is to the allowable bound (0 = maximum slack, 1 = at limit). Near Capacity is the VaR-weighted share of desks with proximity above 80. All specifications include bank-by-date and cluster-by-date fixed effects. Within R^2 reported. t-statistics shown using robust standard errors clustered by date and cluster where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.