Expertise and Specialization in Bank Trading*

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

November 4, 2026

Abstract

Using novel supervisory data, we document persistent and distinct trading specialties in large U.S. banks, despite their ostensible ability to trade the same instruments. We argue that this specialization is driven by traders being more productive where the bank already has installed organizational capital in their areas of expertise. Within a bank, specialist desks earn higher risk-adjusted returns and their risk models leave a smaller share of profits and losses unattributed. Job postings and employment data show that banks recruit more in their specialties when demand expands and that traders in peripheral activities leave more often. Trading-market segmentation reflects in large part how banks organize expertise.

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

The regularly repeated phrase "our assets go down in the elevators every night" was both an appreciation and a warning.

Beard (2007)

1 Introduction

In any industry, firms must choose which services to offer and which to leave to competitors. At first glance, banks' trading businesses appear to be one case where specialization matters least: prices are largely public, products are often standardized, electronic platforms minimize information asymmetries, and banks interact with one another with extreme frequency. In a frictionless benchmark with costless capital and freely available information, banks should share risks perfectly. In contrast to the predictions of this benchmark, banks have long had recognizable trading franchises in specific markets. The names of familiar market benchmarks are one visible example of these reputations: Goldman Sachs in commodities through the Goldman Sachs Commodity Index, J.P. Morgan in emerging-market sovereign debt through the Emerging Markets Bond Index family, and Morgan Stanley in equities through the Morgan Stanley Capital International indices.

In this paper, we provide systematic evidence of banks' specialization across activities, ask why they develop trading specialties, and examine how these choices give rise to market segmentation. To do so, we examine granular supervisory data on the largest U.S. banks' trading from the Federal Reserve's Form FR VV-1, collected from large banks to enforce the Volcker Rule's prohibition on proprietary trading. These data include each bank desk's profit and loss (P&L), its value-at-risk (VaR), the risk factors to which it attributes a share of that P&L, and the risk limits it operates under, including limits the bank sets based on its own forecast of reasonably expected near-term demand (RENTD). We link the bank data to job postings and to resume data on the traders they employ, so we can connect banks' trading specialties to the traders they employ and seek to hire. Our sample consists of 814 trading desks and 1.3 million desk-days, spanning roughly 93% of all U.S.-based bank trading activities.

Using these unique and detailed data, we document three main sets of facts. First, we show that banks specialize in intermediating different risks, and that they earn higher risk-adjusted returns and have more precise risk models in their areas of specialization. Second, as they approach their risk limits, banks' trading shifts towards their core specialization. Third, specialized labor tends to flow to banks where the associated activities are core from banks where those activities are more peripheral. These results suggest the importance of taking an organizational view of firm scope in understanding segmentation of banks into different activities. In particular, they suggest an important role for the complementarity between banks' organizational capital and specialized labor in driving market segmentation and deviations from benchmark theories of perfect risk-sharing.

To organize our empirical investigation, we present a simple model featuring a set of banks. Each bank is endowed with installed organizational capital, for example existing models, systems,

customer relationships, and experience, associated with trading specific asset classes or clusters of risk factors. Banks hire specialized traders as a complementary input to organizational capital, with the two jointly giving rise to cluster-specific expertise, which manifests as a bank's ability to manage risk more precisely and to face less risk in that cluster. Due to cluster-specific expertise, the model predicts that banks concentrate trading activity in clusters where their organizational capital is greater, and moreover, they earn higher risk-adjusted returns and have better risk assessment in their areas of concentration. We refer to these areas of higher activity as the bank's core and to areas with less activity as peripheral.

Additionally, the model makes a second set of predictions concerning banks' responses to changes in customer demand for intermediation and to binding risk constraints on trading activities, and motivates hypotheses about the movement of specialized traders across banks. Namely, it further predicts that (1) when expected customer demand for intermediation rises in an activity, banks with greater organizational capital in that activity respond with larger increases in hiring than banks with less, and banks respond more in their core activities than in their peripheral activities; and (2) when their constraints tighten, banks contract positions in their peripheral activities proportionally more than positions in their core activities. The complementarity between existing organizational capital and new hires also motivates the hypothesis that specialized traders within an asset class tend to flow from banks where that asset class is peripheral to banks where it is core, particularly when expected customer demand in that asset class rises. Unifying each of these predictions is the underlying idea that traders and trading activity have higher marginal value in banks that specialize in the corresponding areas. These predictions provide testable implications of the role of expertise in banks' trading activities.

With our model predictions in hand, we turn to analyzing the data. We begin by documenting that banks have distinct and persistent trading specialties. In most asset classes, such as equities and foreign exchange, activity is concentrated among a few banks. To examine how banks divide activity within asset classes, we group correlated risk factors into clusters that let us compare related trading activities across banks. Most clusters are dominated by only one or two banks. Each bank, in turn, concentrates its own activity in a small set of clusters: half of a typical bank's activity sits in its top 14 clusters. These allocations persist for years.

We exploit the granularity of the data to measure specialization at three levels of aggregation: the share of a bank's total absolute P&L accounted for by (1) a desk, (2) the desks in an asset class, or (3) the risk factors in a cluster of similar risk factors. Specialist desks earn higher risk-adjusted returns. When we compare desks within the same bank on the same day and adjust for common asset-class conditions, a one-standard-deviation increase in a desk's share of its bank's activity is associated with higher daily profits of about 6% of the desk's lagged VaR. At the sample's average VaR, the estimated difference is about $333,000 per day, or 45% of average daily profits.

In the model, specialists in a cluster are more profitable in that cluster because their expertise lowers the risk they bear. A unique aspect of the regulatory Form VV-1 is that on it, banks must

attribute a share of their profits to self-selected risk factors and separately report unattributed P&L, providing us with a non-return based way to test the prediction that banks manage risk more precisely in their areas of specialization. Specialist desks have more precise risk models, leaving a smaller share of P&L unattributed in their attribution across risk factors. A one-standard-deviation increase in specialization is associated with a reduction of 1 percentage point (pp) in that share, a large effect given that the median unexplained share is 0.2%.

Banks do not operate in a vacuum, and they endogenously allocate their trading capacity—both human capital and financial capital—to where they expect demand. Another unique feature of the Form VV-1 data is that, for each risk factor, banks report forecasts of RENTD associated with that factor.[1] At the cluster level, a one-standard-deviation increase in average RENTD limits is associated with about one additional desk at peak participation over the following month. Within a bank, a one-standard-deviation increase in risk-factor specialization is associated with about 66% larger RENTD limits, after accounting for conditions in the factor's cluster. Banks also operate closer to their limits where they specialize. Moreover, the bank raises those limits by more over the following months.

We next study the model's predictions regarding trading desk behavior in response to binding constraints. We find that banks reallocate toward the core when constraints on their trading capacity bind, consistent with the mechanism in the model that core activities where the bank has the most expertise are costlier to abandon when capacity tightens. In particular, when a bank runs closer to its own reported limits, its activity shifts toward core risk factors relative to peripheral ones.

Finally, we turn to the model's hiring predictions and examine trader mobility. In the model, specialized traders and organizational capital jointly give rise to expertise, so productivity depends on where traders work. At a common staffing level, a trader's knowledge is worth more at a bank with high organizational capital in that trader's area of expertise. Linking banks' trading records to job postings and traders' employment histories lets us examine this organizational account through both sides of the employment relationship: where banks seek additional expertise and where traders take their skills.

We document that in response to higher expected demand, banks are more responsive in making job postings for traders in asset classes in which the bank already specializes. Consider two asset classes at the same bank whose specialization differs by one standard deviation. A one-standard-deviation increase in expected demand in each asset class is associated with a larger increase in postings in the more specialized class. The corresponding difference is about 14% of average monthly postings. The comparison accounts for bank-wide hiring cycles and conditions common to each asset class. Given that a posting reflects a bank's demand for an additional trader before hiring, it gives us evidence on where banks seek to build their trading businesses. These monthly hiring decisions complement the daily evidence on how banks allocate trading activity and risk limits.

[1] These RENTD numbers are used for risk limits that desks adhere to as part of the Volcker Rule's prohibition on proprietary trading.

Additionally, a trader's expertise can be peripheral at one bank and central at another. We document that traders are more likely to leave when their asset class accounts for a small share of their employer's trading. Among traders who change banks but remain in the same asset class, the change in their employer's trading share in that asset class is 6.9 pp greater for those leaving peripheral activities than for those leaving core activities. Part of this difference reflects peripheral traders' lower starting shares. The difference is larger when expected customer demand in their asset class rises, although the demand comparison rests on a small sample of moves. Following traders who keep their asset class lets us examine how a given trading specialty fits into different banks' businesses. The sorting patterns are consistent with the idea that the value of expertise depends on the organization in which it is employed, with the same expertise being more valuable in banks that already have complementary organizational capital installed.

Together, our model and evidence show how expertise shapes banks' trading businesses. Organizational capital makes specialized traders more productive in related activities, giving banks an incentive to build on their existing expertise. Banks earn more in their specialties, shift activity toward them as they approach their risk limits, and increase recruiting more in those activities when expected customer demand rises, while traders in peripheral activities leave more often.

Related literature. Our paper relates most closely to work on specialization in financial markets. Earlier work documents industry specialization in banks' lending (Paravisini et al., 2023; Blickle et al., 2025). In lending, specialization is naturally tied to expertise: screening borrowers and monitoring them are information-intensive and subject to increasing returns to knowledge. We document similar specialization in banks' trading, which cuts against the intuition that investors choose portfolios to capture the gains from diversification. Our results give empirical evidence on the role of organizational expertise in trading, consistent with a recent body of theoretical work (Van Nieuwerburgh and Veldkamp, 2010; Eisfeldt et al., 2023). 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, and derivatives, corresponding to stronger trading performance. In addition to documenting a similar empirical result for bank trading, we tie performance to banks' risk assessments and hiring, both of which provide novel evidence on the nature of trading expertise.

Our results also relate to recent work showing that, in contrast to classical models, 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 we document gives rise to segmentation, which we tie to banks' expertise in managing risk. We further theoretically link banks' tendencies to build on existing expertise to traders being more productive in areas where the bank already has installed organizational capital. Our paper therefore also relates to organizational economics, where 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 indicates those frictions matter for financial institutions' trading.

Our paper also relates to work using supervisory data on banks' trading. Prior work shows that bank trading desks behave like toll-taking intermediaries and bear little market risk (Lu and Wallen, 2025), that banks curtailed their exposure to equity market risk after the Volcker Rule (Falato et al., 2025), and that tighter risk limits reduce dealers' intermediation and affect prices in Treasury and currency markets (Bräuning and Stein, 2024; Barbiero et al., 2024; Li et al., 2025). Bräuning et al. (2026) further show that banks shift risk limits toward asset classes where intermediation is more profitable. We use risk limits, including those banks set based on their own forecasts of RENTD, to show that a bank directs its trading capacity—both human capital and financial capital—toward activities where its expertise is most valuable. Comparing activities within the same bank, after accounting for conditions in each market, we find that banks assign larger limits to the risks they specialize in and expand them more over time, shift activity toward their core activities as they approach those limits, and increase recruiting more in their specialties when expected demand rises.

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 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). We focus on traders rather than managers, and we show evidence consistent with traders having different strategy-specific productivities and with a trader's knowledge being worth more at a bank with high organizational capital in that trader's area of expertise.

2 Model

We present a stylized model to organize and interpret our empirical investigation.

2.1 Setup

We consider $2 \leq B < \infty$ banks, indexed by b, and $2 \leq C < \infty$ groups of related trading activities (e.g., asset classes), indexed by c. We refer to these groups as clusters. Each cluster consists of a finite set of risk factors, where a risk factor is a portfolio of traded assets. Each factor f belongs to one cluster c. We model banks' decisions in two stages. In the first stage, banks choose hiring anticipating the customer demand for factor exposures that they will intermediate in the second stage.

Customer Demand for Intermediation. Customers demand exposure to factor f, supplied by banks. We denote bank b's position in factor f as $z_{b,f} \geq 0$, the price per unit of intermediation as

p_f, and customer demand as $d_f \geq 0$. The aggregate inverse demand curve for customers is

$$p_f = ad_f - \ell \sum_j z_{j,f}, \tag{1}$$

where $a > 0$ determines how much customer demand raises the price, and $\ell \geq 0$ is a price impact parameter that determines how quickly the price falls as banks intermediate more customer exposure. Customer demand follows a joint distribution with positive, finite second moments, known to all banks when they make their hiring decisions.

Banks' Expertise and Hiring. Each bank b is endowed with organizational capital $O_{b,c} > 0$ in cluster c, reflecting its existing models, systems, customer relationships, and experience. Supplementing its organizational capital, b hires specialists to work in cluster c at an exogenously given wage of $w > 0$. We denote hiring as $n_{b,c} \geq 0$ and define cluster-specific expertise from organizational capital and hiring as $E_{b,c} \equiv O_{b,c}\sqrt{1 + n_{b,c}}$.[2]

Greater expertise lowers bank b's residual payoff risk from managing factor exposures, as could arise from more effective hedging. Bank b's residual payoff per unit of exposure to factor f is

$$\varepsilon_{b,f} = \kappa R_f + u_{b,f},$$

where κR_f is an "explained payoff" component—captured by banks' risk models—and $u_{b,f}$ is an unexplained payoff component. The payoff shock R_f to factor f, common to banks and captured by their risk models, is standard normal. Its loading, $\kappa > 0$, determines the scale of the explained payoff component. The payoff shock $u_{b,f}$ that bank b's risk model leaves unexplained is normally distributed with mean zero. We define the variance of unexplained risk at bank b for each factor in cluster c as $\sigma^2_{u,b,c} \equiv (\tau_0 + E_{b,c})^{-1}$, where $\tau_0 > 0$ is the baseline precision of unexplained payoffs. For bank b, the explained and unexplained components are independent, so $\varepsilon_{b,f}$ is normally distributed with mean zero. Its total residual variance in cluster c, $v_{b,c}$, is

$$v_{b,c} = \kappa^2 + \frac{1}{\tau_0 + E_{b,c}} = \kappa^2 + \frac{1}{\tau_0 + O_{b,c}\sqrt{1 + n_{b,c}}}. \tag{2}$$

For each bank, residual payoffs are independent across factors and independent of customer demand.

In Equation (2), greater existing organizational capital and specialist hiring lower residual risk. Organizational capital increases the contribution of an additional worker to the bank's expertise $E_{b,c}$. Experts use their human capital to manage risk more precisely and face less risk than non-experts, as in Eisfeldt et al. (2023).

Banks' Problem. Each bank has mean-variance preferences with risk aversion $\gamma > 0$. Bank b's

[2] A substantial portion of organizational capital is, of course, embodied in labor. We primarily seek to distinguish between existing labor (organizational capital) and new hiring. Our assumed functional form makes higher organizational capital initially lead to banks hiring more specialists, while also including diminishing returns to additional hires that keep optimal hiring finite.

trading profit from intermediating $z_{b,f} \geq 0$ units of customer exposure is

$$\pi_{b,f} = z_{b,f}(p_f + \varepsilon_{b,f}).$$

We model each bank's behavior in two stages. First, it chooses hiring before observing customer demand. Then, it observes customer demand and chooses its positions in each cluster.

Bank b's portfolio choice problem in the second stage is

$$\max_{z_{b,f} \geq 0} \; z_{b,f} \left(a d_f - \ell \sum_j z_{j,f} \right) - \frac{\gamma}{2} v_{b,c} z_{b,f}^2. \tag{3}$$

Its hiring problem in the first stage, which anticipates this portfolio choice problem, is

$$\max_{\{n_{b,c} \geq 0\}} \; \mathbb{E}\left[\sum_c \sum_{f \in c} \left\{ z_{b,f} p_f - \frac{\gamma}{2} v_{b,c} z_{b,f}^2 \right\} \right] - w \sum_c n_{b,c}. \tag{4}$$

Equilibrium. *We define a subgame perfect equilibrium as a collection*

$$\left(\{n_{b,c}\}_{b,c}, \{p_f\}_f, \{z_{b,f}\}_{b,f} \right)$$

such that:

1. *given other banks' hiring choices, each bank b chooses $\{n_{b,c}\}_{c=1}^C$ to solve* (4)*, anticipating the subsequent equilibrium in prices and portfolios;*

2. *for every set of hiring choices and realization of customer demand, each bank chooses its portfolio to solve* (3)*, taking other banks' portfolios as given;*

3. *for every factor f, markets clear:*

$$p_f = a d_f - \ell \sum_{b=1}^B z_{b,f}.$$

We establish existence and uniqueness of equilibrium for every $\ell \geq 0$ in Appendix IA.B.

Prices, Banks' Positions, and Banks' Hiring. The banks' first order conditions, together with customer demand in Equation (1), imply that for factor f in cluster c

$$p_f = \frac{a d_f}{1 + \ell \sum_j (\ell + \gamma v_{j,c})^{-1}}, \qquad z_{b,f} = \frac{p_f}{\ell + \gamma v_{b,c}}. \tag{5}$$

Specialist Hiring. Banks' hiring does not have a general closed form solution. When $\ell = 0$ and hiring is positive, bank b chooses hiring to satisfy

$$\sqrt{1+n_{b,c}}\left[1+\kappa^2\left(\tau_0+O_{b,c}\sqrt{1+n_{b,c}}\right)\right]^2=\frac{a^2\sum_{f\in c}\mathbb{E}[d_f^2]\,O_{b,c}}{4\gamma w}, \tag{6}$$

where $\mathbb{E}[d_f^2]$ is the second moment of anticipated customer demand for factor f. The left hand side increases strictly with hiring, so hiring increases with anticipated customer demand and decreases with the wage. Hiring additional cluster c specialists reduces b's residual risk from trading factors in c. With positive hiring, bank b hires until its marginal expected utility from trading equals w. The bank hires no additional specialists when the marginal value at zero is at or below w.[3]

2.2 Model Predictions

The first three predictions arise from differences in banks' organizational capital and their hiring. Due to the complementarity between organizational capital and hiring, banks concentrate hiring in activities where their organizational capital is greater. We refer to these activities as the bank's core and to activities with less organizational capital as peripheral. A bank can have several core activities. Within a cluster, we refer to banks with greater organizational capital as specialist banks. For our predictions, we consider comparable demand across activities, and customers that are sufficiently price sensitive (or equivalently, sufficiently small $\ell > 0$).

We measure bank b's trading activity in factor f as $|z_{b,f}\kappa R_f|$ and cluster activity by summing across factors.

Prediction 1 (Specialization)**.** *Banks allocate more specialist resources to activities where their organizational capital is greater. These activities account for a larger share of the bank's expected trading activity.*

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

Prediction 3 concerns the accuracy of banks' risk models. Bank b's risk model attributes $z_{b,f}\kappa R_f$ to known exposures and leaves the payoff component $z_{b,f}u_{b,f}$ unexplained. Bank b observes this residual after returns arrive, but its risk model does not explain it through known exposures. For a positive position, we define bank b's unattributed share as

$$\frac{|z_{b,f}u_{b,f}|}{|z_{b,f}\kappa R_f|+|z_{b,f}u_{b,f}|}=\frac{|u_{b,f}|}{|\kappa R_f|+|u_{b,f}|}. \tag{7}$$

[3]For our predictions, we consider expertise levels satisfying $\kappa^2 E_{b,c} < (1+\kappa^2\tau_0)/\sqrt{3}$, which ensures that additional expertise remains sufficiently valuable for hiring to increase with organizational capital and for banks with greater organizational capital to increase hiring more when anticipated demand rises. We derive the hiring choice and its properties for sufficiently small $\ell > 0$ in Appendix IA.B.

Prediction 3 (Specialists and Model Precision)**.** *Banks manage risk more precisely in the clusters they specialize in, leaving a smaller expected share of P&L from factor payoffs unattributed.*

Predictions 1, 2, and 3 arise from differences in banks' expertise in managing risk. Banks hire more in activities where their organizational capital is greater, raising expertise and lowering risk in Equation (2). Lower risk allows the bank to take larger positions and earn higher risk-adjusted returns. Greater expertise also lowers the expected unattributed share, as in Equation (7).

Prediction 4 concerns hiring when anticipated customer demand changes. We consider a local proportional increase in d_f for all factors in a cluster, anticipated by banks before they hire, and compare equilibrium hiring with and without this increase. This change in customer demand applies to all banks trading factors in the cluster.

Prediction 4 (Expected Demand and Specialist Hiring)**.** *An increase in anticipated customer demand in a cluster raises hiring at banks that already hire in that cluster, with a larger increase at banks with greater organizational capital in the affected cluster. Within a bank, the increase is larger in its core clusters than in its peripheral clusters.*

Prediction 4 can be understood by the fact that an increase in anticipated customer demand in a cluster raises the marginal value of additional hiring. Banks that already hire specialists in that cluster therefore increase hiring further. Banks with greater organizational capital hire more to begin with, and since expertise grows with $\sqrt{1+n_{b,c}}$, raising their expertise by a given percentage requires more additional specialists. When expertise is low enough that unexplained risk remains large relative to explained risk, this effect outweighs the diminishing benefit of reducing residual risk, so these banks increase hiring by more.

The next prediction concerns banks' positions after making hiring decisions. We consider how a bank adjusts its positions when its own risk budget tightens.

To study bank constraints, we consider an internal risk budget $\overline{R}_b$ on bank b,

$$\sum_c \sqrt{v_{b,c}} \sum_{f \in c} z_{b,f} \leq \overline{R}_b. \tag{8}$$

Bank b's factor f position has a risk charge equal to its P&L standard deviation, $\sqrt{v_{b,c}} z_{b,f}$. The budget sums these charges without a diversification adjustment. For Prediction 5, we consider a tightening specific to bank b and unanticipated when banks hire. Only bank b has a binding budget, while all other banks' budgets remain slack.

Prediction 5 (Tightening Bank Constraints and Peripheral Contraction)**.** *When a bank's risk budget tightens, it contracts its peripheral positions by a larger percentage than its core positions while both remain positive.*

A bank's greater expertise in core activities reduces its risk budget usage per unit of core exposure. The bank can therefore shed peripheral activities at a lower cost.

Our final prediction concerns the value of specialized traders in banks where their specialization is core versus peripheral.

Prediction 6 (Specialist Sorting)**.** *Under the model's parameter restrictions, an additional specialist has greater marginal trading value at the bank with more organizational capital in the specialist's cluster, holding all banks' hiring fixed and comparing banks with the same staffing. A proportional increase in anticipated customer demand in the cluster raises this difference.*

Organizational capital increases the contribution of an additional specialist to the bank's expertise, so, at a common staffing level, the same specialist adds more expertise at a bank with greater organizational capital in the specialist's cluster. Higher anticipated demand raises the value of this expertise and so widens the difference between banks. Together, these effects motivate examining whether specialized traders flow from banks where their specialization is peripheral towards banks where their expertise is core.

With our model's predictions in hand, we next turn to describing the data used to test them.

3 Data

The Federal Reserve collects daily data on the largest banks' trading desks on form FR VV-1 to enforce the Volcker Rule's prohibition on proprietary trading. Our data runs from 2014 to 2025. The most recent data covers large U.S. and foreign banks with at least $20 billion of average gross trading assets and liabilities, excluding U.S. government securities. Large banks dominate trading, so the data spans roughly 93% of total U.S.-based trading activity as of 2019. We focus on banks that provide data consistently through the full sample. We describe how we clean the data in Appendix IA.A.1.[4]

The data gives us three related but distinct panels:

1. Date × desk: Each desk reports its profitability, VaR, and trading volume, along with its name and a description of its trading activities. An illustrative 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 conducts related hedging activities."

2. Date × desk × risk-factor: Banks attribute a portion of the P&L from *existing positions*—positions held at the end of the previous day—to named risk factors, each with a brief description. They report P&L from *new positions* arising from the current day's trading separately, in aggregate, without a required risk-factor breakdown. New-position P&L includes commissions, fee income and expenses, and market gains or losses on that day's transactions.

[4]Other papers using this data include Falato et al. (2025), who show banks 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.

The illustrative equities desk's risk factors might be "U.S. large cap delta," "short-term interest rate delta," and "U.S. equity index basis." VaR is not available at the risk-factor level.

3. Date × desk × risk-limit: Each desk reports its risk limits and the current utilization of each limit. Limits take many forms, including VaR limits and limits on underlying risks (which can be, but do not have to be, the same as the risk factors in the previous panel). In the second half of the sample, banks also report each limit's type (VaR, position limit, sensitivity, inventory age, etc.) and source (risk appetite, regulatory capital, RENTD, risk reduction/mitigation, etc.). The illustrative equities desk may report 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 that runs from 0 (far from limit) to 100 (at limit); values above 100 indicate a breach. We describe its construction in Appendix IA.A.4.

Each bank reports several desks, a desk reports several risk factors, and a risk factor can carry several risk limits. Because of the different levels of granularity of the three panels, the available variables determine the unit of each test. VaR is reported for desks but not for risk factors or clusters, and the number of desks active in a unit is defined for a cluster but not for a risk factor, which belongs to one desk. The VV-1 data is also split into two samples given a change in the reporting requirements in 2020, though we focus on the banks that consistently report across both periods.

Reported risk-factor P&L accounts for about 11% of comprehensive P&L. This is a material component of trading profits—roughly $8 billion per year—but most P&L comes from new positions, for which banks do not report a risk-factor breakdown. New trades and overnight inventory are managed by the same desks and draw on the same trading mandates and expertise. We therefore use the reported factors to identify desks' broader risk-taking specialties, on the premise that the main risks a desk intermediates during the day are reflected in its overnight inventory.

Returns Banks report neither the size of a desk's portfolio nor the share committed to each risk factor, so we cannot calculate conventional percentage returns. We instead scale P&L by the measure of risk exposure available at each level. At the desk level, banks report VaR, so we divide daily P&L by the previous day's VaR in the spirit of Falato et al. (2025):

$$DeskReturn_t = \frac{\text{P\&L}_t}{\text{VaR}_{t-1}}$$

and we winsorize the variable at the 5th and 95th percentiles to reduce the influence of outliers.

At the risk-factor level, banks do not report attributable VaR. We instead proxy a factor's economic risk exposure with the standard deviation of its P&L over the previous three months.

Dividing the factor's daily P&L by that standard deviation gives

$$FactorReturn_t = \frac{\text{P\&L}_t}{\sigma(\text{P\&L})}.$$

Both measures express profits per unit of risk.

Sorting Desks into Asset Classes The data does not, however, provide asset-class classifications or any grouping of risk factors across banks. We describe how we address these groupings next. We assign each desk to one of ten asset classes from its name and description: credit, equities, rates, foreign exchange, securitized and structured products, collateral and financing, commodities, overlay and counterparty risk (XVA), municipal, and legacy. We combine keyword matching with supervised classification. We build a keyword list for each asset class by hand and assign a desk directly when its text matches a keyword that unambiguously signals one class. When the text matches keywords from more than one class, we break the tie with term frequency–inverse document frequency (TF-IDF) similarity. Appendix IA.A.2 describes the procedure in detail. A bank can have many desks in the same asset class.

Sorting Risk Factors into Clusters We group risk factors by their return comovement to compare trading activities across banks. Banks choose their own factor names and descriptions, so the same risk can carry different labels at different banks. We pool factors from all desks and banks within each asset class and use correlations of signed daily P&L scaled by its volatility over the previous three months. Absolute correlations let factors with opposite exposures to the same risk belong to the same cluster: our interest is in which risks a desk trades, regardless of whether that exposure is long or short. We now provide a short sketch of the sorting procedure, and we provide the full details of the sorting procedure in Appendix IA.A.3 along with robustness checks.

The clustering proceeds in two stages. We first use risk factors observed for at least a year and require at least a year of overlapping observations to estimate the correlation of all resulting pairs. We use a hierarchical algorithm to group these factors by successively merging the groups with the strongest average absolute return correlation, stopping at a chosen correlation cutoff. We set the cutoff at 0.10, so groups merge only if their average absolute return correlation reaches that threshold. We then assign shorter-lived factors to the resulting clusters. For each candidate cluster, we compare the factor with the member whose history overlaps most with its own and require the correlation to clear both a statistical significance test and the correlation cutoff. Unmatched factors remain unclassified.

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. Appendix Figure IA.1 and Appendix Table IA.1 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 against sampling noise, since a correlation of 0.10 is roughly equal to 1.6 standard deviations: under the null of zero correlation with $n = 252$ independent observations, $\mathbb{E}[|\rho|] \approx \sqrt{2/\pi}/\sqrt{n-1} \approx 0.05$, so $|\rho| = 0.10$ is roughly twice the noise floor we would expect from random variation alone.[5]

3.1 Summary Statistics

The desk panel spans 814 trading desks and 1.3 million desk–trading day observations (Table 1). Credit has the most desks over the full sample at 155, followed by equities (152), rates (137), and securitized and structured products (128). The remaining asset classes have between 24 and 82 desks. On a given day, 99 equities desks, 76 rates desks, and 72 credit desks report.

Desk profitability varies widely, averaging \$630,000 (equal-weighted across asset classes) with an average standard deviation of \$5.6 million. The average desk is profitable, with 90th-percentile profits of \$3.2 million against 10th-percentile profits of −\$1 million. Equities desks are the most numerous on a given day and the most profitable, with 99 desks earning \$1.5 million a day on average, followed by foreign exchange at 38 desks and \$1.3 million. Average profits are positive in every asset class except XVA, which rounds to zero. Summing the product of desks per day and mean profitability across asset classes, trading desks in our sample generate about \$335 million in profits each day, or \$84 billion a year.

More profitable desks tend to operate with higher VaR. Foreign exchange desks average \$7.5 million in VaR with a standard deviation of \$16.0 million, and rates desks average \$7.2 million with a standard deviation of \$15.3 million. The 10th percentile is zero in seven of the nine asset classes, so desks do not always carry material risk exposures; the 90th percentile ranges from \$8.9 million to \$17.2 million. Total VaR across all banks is no more than \$2.6 billion (the dot product of desks per day and mean VaR), but that overstates risk because VaR is not strictly additive: exposures across a bank's desks offset and hedge one another.

Desks operate under several limits at once. Table 2 summarizes the risk limit panel in the post-2020 sample, when banks must report each limit's type and source. The median desk carries 6 limits on a given day and the mean is 9; desks at the 90th percentile face 16 at once. A single desk can be subject to a VaR limit, several sensitivity limits, a position limit, and a RENTD limit, each set by a different part of the bank's risk framework.

The typical limit has substantial headroom. Panel B reports mean proximity of 20 across all limit-date observations on the 0–100 proximity scale (Appendix IA.A.4). Proximity differs across limit types: other than the small residual "other" category, position limits run closest to their

[5] Appendix 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 absolute or signed correlations of the same signed daily risk-factor return proxies.

bounds, at an average proximity of 39, followed by VaR limits at 28, while sensitivity and scenario limits sit lower. RENTD is the most common source of a limit, covering 3.0 million of the 4.5 million limit-days in the panel. Appendix Figure IA.4 plots weekly averages of proximity and the share of limits near capacity over the period shown.

Limit breaches are rare. Across all limit types, 14% of limit-days have proximity above 50 and 0.2% above 100 (Panel C). The rates differ sharply across types: position limits are above 50 on 37% of days, compared with 15% for sensitivity limits and 19% for VaR limits. By source, RENTD limits are above 50 on 13% of days, compared with 22% for risk appetite limits and 5% for risk reducing limits, although both of those sources are much less common. Banks set RENTD limits from expected customer flow, and realized flow rarely approaches the forecast.

Panel D shows that near-capacity spells on a single limit are short. The median spell of consecutive days with proximity above 80 lasts two days, with a 75th percentile of five days and a mean of seven. These spells are measured for a single limit, not for the VaR-weighted bank-level constraint used in Section 7.

Limit sizes also change infrequently. On any given day, 0.5% of limits change in size, and increases and decreases occur at roughly similar rates (Panel E). The median spell between consecutive changes is 183 days across all limit sources and 162 days for RENTD limits. By type, aging limits change most frequently, on 1.4% of days, against 0.5% for VaR limits. We use RENTD limit levels to measure expected demand in Section 6 and deviations of log RENTD limits from their trailing mean to measure demand shifts in the labor tests of Sections 6 and 7.

3.2 Risk-Factor Clusters

Figure 1 shows the asset-class groupings and up to ten clusters within each displayed asset class, ranked by the number of member factors and subject to disclosure filtering. Table 3 Panel A summarizes the risk factors we group into clusters. Most asset classes have hundreds of risk factors, and the average desk is exposed to eight factors on a given day. Because each desk carries several factors, the average profit attributable to one factor is small, ranging from −$21,000 in securitized and structured products to $34,000 in equities. Equities and rates factors have larger tails, with 90th-percentile profits of $400,000 and $363,000.

Panel B of Table 3 reports the cluster counts. The average asset class has about 130 unique clusters; municipal trading has the fewest, about 30, and equities the most in both eras, matching the range of tradable products in each asset class. About 6% of clusters contain a single risk factor.

Risk factors within an asset class need not move together. Figure 2 reports average absolute return correlations within and across asset classes. Average correlations are low when all clusters are included and rise when the sample is restricted to the largest clusters, measured by the sum of member factors' P&L volatilities. Part of that increase follows from forming clusters on correlations.

3.3 Reporting Incentives

The VV-1 data is connected to two regulations that constrain bank trading: restrictions on proprietary trading and capital requirements for market risk. Regulators collect the data to monitor compliance with the Volcker Rule. Supervisors examine desks' risk exposures, P&L, and customer-facing activity to assess whether trading is consistent with permitted activities, such as market-making and underwriting. Unusual patterns or discrepancies can lead to further scrutiny.

The data is also related to market-risk capital requirements. Under the market-risk rule, banks must hold capital against their trading exposures through charges based on VaR, stressed VaR, and other risks. The VaR-based charges depend on recent average risk estimates multiplied by a factor, starting at three. That factor rises when realized trading losses exceed the model's one-day VaR too often. These *backtesting exceptions* can push the multiplier as high as four, unless the supervisor directs a different adjustment. Repeatedly understating risk is therefore costly: the higher multiplier can raise the capital required to support the bank's trading activities. Although the desk VaRs in our data do not themselves determine the bank's capital charge, banks must calculate them consistently with the models and methods used for market-risk capital.

These regulatory uses make banks' reporting choices consequential. Banks choose their risk factors and have discretion over their risk models, which naturally leads to concerns that the dynamics we study reflect artificial modeling or classification choices. But much of banks' reporting discretion is subject to supervisory review. A bank must receive prior written approval before using an internal model to calculate market-risk capital and must notify its supervisor of specified changes, including nontrivial changes to modeling assumptions or changes that meaningfully affect risk-weighted assets. Hence, the capital requirements subject the models underlying reported risk to oversight beyond the Volcker reporting requirement.

The P&L reporting requirements provide another check on the validity of banks' reported data. Banks report each desk's total daily P&L and separate it into new and existing positions. P&L from existing positions must reconcile across the prescribed attribution categories, including an explicit unattributed residual. Changing a risk-factor description cannot, by itself, remove the associated P&L from the desk's total. Amounts omitted from one attribution category must appear elsewhere for the report to reconcile. These requirements give supervisors a way to understand discrepancies between reported trading outcomes and the sources to which banks attribute them.

4 Banks Specialize (Prediction 1)

Prediction 1 says that banks specialize in different clusters of risk factors. We document concentration from both sides: a few banks account for most of the activity in a unit of risk (either risk-factor clusters or asset classes), and a few units account for most of a bank's activity. Each bank's allocations persist, and its leading activities differ from its competitors'. These persistent, distinct specialties form the bank's *core*. We begin by documenting concentration and persistence at the

asset-class level, the coarsest grouping and the one least exposed to how banks name and bucket their risks. Clusters then show how banks divide activity within asset classes and whether their core activities differ.

4.1 Concentration and Persistence Across Asset Classes

Asset classes are concentrated across banks. We compute a Herfindahl–Hirschman index (HHI) from each bank's share of the asset class's total desk VaR; the maximum of 10,000 means one bank holds all of the asset class's risk. The mean HHI is 3,063, roughly what three equally sized banks would produce (Table 4). Appendix Table IA.2 reports the HHI by asset class. Repo is the most concentrated asset class, securitized products and rates the least, and five of the nine sit between 2,000 and 2,600, with foreign exchange and XVA falling in between. Concentration falls with size: the correlation between the HHI and the log of asset-class activity is -0.35.[6]

Each bank concentrates its activity in a few asset classes. We rank asset classes by their shares of a bank's activity and cumulate: a handful account for most of the bank's activity (Figure 3). The figure defines activity as the absolute value of net risk-factor P&L within each asset class.

Banks' asset-class shares also persist over time. We measure persistence using each activity's share of its bank's total absolute risk-factor P&L. For bank b, activity ℓ, and date t, define

$$Share_{b,\ell,t} = \frac{Activity_{b,\ell,t}}{\sum_{\ell' \in b} Activity_{b,\ell',t}}, \tag{9}$$

where ℓ indexes asset classes here and clusters in the next subsection. *Activity* is the sum of absolute risk-factor P&L within the asset class or cluster, and the denominator sums across the bank's activities at the same level.[7] We regress each asset class's share on its own lag, estimating a separate coefficient for each lag τ, measured in trading days:

$$Share_{b,a,t} = \alpha + \beta_\tau \, Share_{b,a,t-\tau} + \varepsilon_{b,a,t}. \tag{10}$$

Asset classes that account for a large share of a bank's activity today tend to remain important at long horizons (Figure 4, which plots β_τ against τ).

4.2 Specialization Within Asset Classes

A bank can dominate a specific risk and still hold a modest share of its asset class. Clusters—groups of correlated risk factors within an asset class—reveal these finer specialties.

Most clusters are dominated by one or two banks. We calculate HHIs analogous to the asset-class HHIs above, but since we do not have VaR at the risk-factor level, we build cluster shares from P&L volatility instead: we sum each risk factor's trailing 90-day standard deviation of P&L within a

[6] Appendix Figure IA.5 provides a scatter plot of asset classes by their HHI.
[7] For cluster shares, the denominator excludes factors not assigned to a cluster.

bank–cluster–date, compute each bank's share of the cluster total, and sum squared shares.[8] Table 4 reports a mean cluster HHI of 8,871, close to the maximum at which a single bank accounts for all of a cluster's risk.[9]

Concentration again falls with size: the correlation between the HHI and the log of cluster activity is -0.295, and Appendix Figure IA.7 plots each cluster's time-averaged HHI against its size. Only the largest clusters support several participants, and the typical cluster is served by one bank. Figure 5 shows the pattern within each asset class. Clusters are sorted from left to right by total P&L volatility, and the large clusters on the left are lighter, meaning more banks.

Each bank concentrates its own activity in a few clusters. To see this, we follow the construction used in Figure 3. For each bank-day we sum signed risk-factor P&L within each cluster and then take the absolute value, so offsetting gains and losses within a cluster cancel, and we rank clusters by their activity shares and cumulate before taking the median across days within a bank and then across banks. The top 14 clusters account for half of a bank's activity and 67 clusters account for 90%, out of 206 clusters (Figure 6). Profits and losses look alike taken separately, at 8 and 35 clusters for profits and 8 and 34 for losses.

Allocations to individual clusters also persist. As in the asset-class persistence test, we take each risk factor's absolute P&L before summing within clusters, so offsetting gains and losses do not cancel, matching the activity measure used in equation (9). We repeat the regression in equation (10) using each cluster's share of its bank's total absolute risk-factor P&L in place of the asset-class share. The coefficient in Figure 7 is about 0.7 at lags of 1 to 50 trading days and is still about 0.4 at lags near 750.

Banks also differ in which clusters make up their core. Figure 8 reports overlap between banks' top three and top ten clusters. We rank clusters by their share of a bank's total absolute risk-factor P&L within each bank-month. The reported index divides the number of shared clusters by $2K$ minus that count and takes the median across months with at least one shared cluster for each bank pair. When both sets contain K clusters, this index is their Jaccard similarity. The figure shows low overlap even among months in which the banks share a top-ranked cluster.

It might seem surprising that most clusters are dominated by a single bank. A natural concern is that this concentration reflects differences in reporting: one bank might combine two risk factors into a single reported factor while another reports them separately. But banks must report P&L from existing positions attributable to the specific risk factors "that are monitored and managed as part of the trading desk's overall risk management policies and procedures." Moreover, the regulation requires banks to report enough factors to explain the preponderance of P&L changes attributable to risk factors. These requirements weigh against combining risk factors that a desk monitors and manages separately. We cannot rule out differences in aggregation across banks, but

[8] Appendix Table IA.3 shows that desk VaR and the trailing 90-day standard deviation of desk P&L move together with a coefficient close to one.

[9] See Appendix Figure IA.6. Standard deviations do not add across risk factors, so the summed measure is an upper bound on a bank's activity in a cluster. The results are similar using absolute P&L.

the reporting requirements tie those differences to how banks monitor and manage their risks.

Two economic mechanisms may also help explain why so many clusters are dominated by a single bank, although we cannot test them directly with the available data. First, trading desks largely intermediate and make markets. Their inventories can reflect risks they are unable or unwilling to lay off, rather than broad positions that are easy to replicate and hedge. These residual exposures can differ across banks, so the associated risk-factor P&L need not move together. Second, risk factors can reflect exposures to bespoke, nonlinear, or illiquid instruments. Differences in these instruments across banks may also produce weak correlations in their P&L. Since residual and bespoke exposures are likely to be small, these mechanisms help explain why small clusters are dominated by one bank while the largest clusters include several banks.

Common bank shocks present a separate concern: factors might cluster together because they share funding conditions or valuation methods. Funding costs generally enter carry, which banks report separately from named risk-factor P&L. The instructions also call for consistent attribution methods for common factors across a bank's desks. These requirements limit a simple explanation based on common funding shocks, but they do not remove common shocks to positions or correlated marks generated by a bank's risk models.

Specialization in a cluster does not imply an undiversified portfolio. We form clusters using the absolute correlation of signed return proxies, $|\rho_{ij}|$, so factors with opposite exposures to the same risk can belong to the same cluster. A bank can therefore concentrate its trading in a few risks while holding offsetting positions within them. Cluster concentration alone does not establish how much net risk the bank bears.

Appendix Tables IA.4 and IA.5 compare pairwise comovement under two cluster constructions. The first uses clusters formed from absolute correlations of signed return proxies and an absolute-correlation outcome; the second uses clusters formed from signed correlations of the same proxies and a signed-correlation outcome. Each table regresses the pairwise correlation on indicators for whether the factors share a bank, an asset class, and a cluster, together with same-bank interactions. The same-cluster coefficients in column (6) are positive in both tables, 13.25 and 22.22, as are the same-bank-by-same-cluster interactions, 2.54 and 2.72. The cluster assignments differ across the tables, so these coefficients describe different groups of factor pairs. The within-cluster comovement is partly mechanical because the clustering itself uses comovement.

5 Specialist Outperformance (Predictions 2 and 3)

So far we have shown that banks specialize insofar as they put a disproportionate share of their trading books toward specific asset classes and risk factors, and those allocations are persistent. But that does not necessarily mean that banks are unusually talented at managing risk in those positions. The obvious corollary to specialization is that it should be profitable if it stems from expertise. Predictions 2 and 3 say that a bank should earn higher risk-adjusted returns and hold

more precise risk models where it specializes. We test both predictions at the desk level, where returns, P&L attribution, and VaR are reported.

We measure specialization by the share of a bank's own activity devoted to a desk, risk factor, risk-factor cluster, or asset class, using the share construction in equation (9). Our regression measure, *Specialist*, is the standardized trailing mean of this share:

$$\textit{Specialist}_{b,\ell,t} = \overline{\textit{Share}}_{b,\ell,[t-90,\,t-1]} \tag{11}$$

where b denotes the bank, ℓ the level of aggregation (asset class, desk, risk factor, or cluster), and t the date. The measure tells us if a larger share of the firm's absolute value of P&L is in a given exposure. The bar denotes the mean over available observations in the 90 calendar days before t, and we standardize the variable to have mean zero and unit standard deviation.

Specialization can describe a bank's importance in a market or an activity's importance within a bank. A bank may account for a large share of all banks' activity in an asset class, or it may devote a large share of its own book to that asset class or risk. We use the second definition. Large banks can hold large market shares in several asset classes simply because of their scale. Banks operate at different sizes for reasons beyond their allocation of trading expertise. We take bank size as given and ask how a bank allocates its trading activity across its own book.

An alternative measure of specialization, following Blickle et al. (2025), compares a bank's portfolio weights with those of the market. A bank specializes in an asset class when that class accounts for a larger share of its own portfolio than of the market portfolio. This measure also accounts for bank size, but asks whether the bank overweights an activity relative to the market. Appendix Figure IA.8 shows that it is positively related to our measure at the asset-class level.[10]

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

We now test if specialists have higher risk-adjusted profits by regressing our return proxy, $r_{b,d,t}$, on the specialist measure. We run this regression at the desk level, which is the most granular level for which we have VaR, the denominator in the desk-level return proxy. We estimate

$$r_{b,d,t} = \beta\, \textit{Specialist}_{b,d,t} + \alpha_{at} + \delta_{bt} + \varepsilon_{b,d,t},$$

where α_{at} are asset-class-by-date and δ_{bt} bank-by-date fixed effects.

Specialist desks earn higher risk-adjusted returns than other desks at the same bank. Column (4) of Table 5 reports our preferred specification. Bank-by-date effects remove shocks common to a bank's desks on a given day and asset-class-by-date effects remove common return movements within each asset class. The coefficient therefore compares a bank's desks on the same day, adjusted for

[10] We construct the Blickle et al. (2025) analogue as $\frac{\mathrm{VaR}_{b,a,t}}{\sum_{a'} \mathrm{VaR}_{b,a',t}} - \frac{\sum_{b'} \mathrm{VaR}_{b',a,t}}{\sum_{b',a'} \mathrm{VaR}_{b',a',t}}$, where $\mathrm{VaR}_{b,a,t}$ is desk-level VaR summed within bank b and asset class a on date t.

conditions in each desk's asset class. A one-standard-deviation difference in *Specialist* is associated with daily P&L higher by 0.06 units of lagged VaR, or 5.77% of lagged VaR.

To gauge the magnitude of the coefficient of about 0.06, we use the summary statistics in Table 1 to translate the estimate into dollars. Weighting the asset-class means by their desk-day observation counts gives average desk VaR of about $5.77 million and average daily P&L of about $745,000. At this level of risk, a one-standard-deviation increase in *Specialist* is associated with about $333,000 more daily P&L, or about 45% of average daily profits.

The association also holds as a given desk's specialization changes over time. Column (7) retains bank-by-date and asset-class-by-date effects and adds desk fixed effects and the Blickle et al. (2025) market-deviation control. Desk effects remove persistent differences across desks, so the coefficient relates changes in a desk's specialization to changes in its returns, net of common bank-day and asset-class-day movements. A one-standard-deviation increase in *Specialist* is associated with an increase of 0.09 in that desk's daily P&L per unit of lagged VaR, holding fixed the bank's portfolio weight in the asset class relative to the market's.

5.2 Specialists Have More Precise Risk Models (Prediction 3)

Specialization can raise profits through several mechanisms. Strong customer relationships or market power could make a bank's core activities more profitable without any advantage in technical expertise. Our model links outperformance to knowledge capital: specialized traders contribute to a bank's knowledge capital (expertise) in an activity, and that knowledge is reflected in more precise risk models (Prediction 3). We examine this mechanism using desks' P&L attribution, asking whether specialists have more precise risk models.

For existing positions, banks report P&L attributed to movements in named risk factors, along with carry, other attributable components, and a residual item. For each desk-day we build an *unattributable share*, the unattributed item divided by the sum of the included absolute P&L components:

$$\textit{Unattributable Share}_{b,d,t} = \frac{|\textit{Other Unattributed Profit \& Loss}_{b,d,t}|}{|\textit{Risk Factor P\&L}_{b,d,t}| + |\textit{Carry P\&L}_{b,d,t}| + |\textit{Other Unattributed Profit \& Loss}_{b,d,t}|} \tag{12}$$

The ratio runs between 0 and 1. A smaller value means that the unattributed component accounts for less of the absolute P&L components.[11]

Figure 9 plots a binscatter of the unattributable share against *Specialist*. The left panel shows

[11] "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 ..."

that specialist desks better account for their sources of P&L. The right panel plots the unattributable share against desk tenure, the cumulative number of trading days on which the desk reports P&L and a rough proxy for specialization, expressed in years. Longer-lived desks have lower unattributable shares, consistent with learning by doing.

The raw relationship could simply reflect that some banks generally have better models. We test the relationship between specialization and model accuracy within a bank in Table 6. Column (3) compares within a bank and asks whether the unattributed component accounts for less of the absolute P&L components at a bank's specialist desks than at its peripheral ones. A one-standard-deviation increase in *Specialist* lowers the unattributable share by about 1 pp.

The result changes little across outcome definitions and comparisons. Columns (1) to (3) use the main definition of the unattributable share and columns (4) to (6) an expanded denominator that also includes $|Other\ Attributable\ P\&L_{b,d,t}|$, leaving the numerator unchanged.[12] The coefficient ranges from -0.011 to -0.010 across the six columns. Columns (1) and (4) hold bank and date effects, comparing all desks reporting on the same date net of each bank's average level. Columns (2) and (5) hold bank and asset-class-by-date effects, comparing desks in the same asset class on the same date net of bank effects.[13]

Another way to see specialization and model accuracy is the fineness of the risk model itself. We measure reporting granularity by the number of distinct risk factors to which a desk attributes nonzero P&L on a given day. Information-choice models predict the same pattern: investors with limited capacity choose which risks to learn about, and increasing returns push them to concentrate learning on the assets they specialize in (Van Nieuwerburgh and Veldkamp, 2010; Kacperczyk et al., 2016).

Specialist desks report their risks in greater detail. Column (3) of Appendix Table IA.7 regresses the log number of risk factors a desk reports on its specialization measure, with bank-by-date and asset-class-by-date effects. A one-standard-deviation increase in *Specialist* is associated with a 12% increase in the number of monitored risk factors. Moreover, column (5) shows that the relationship between the unattributable share and the specialist measure continues to be negative—though smaller—when we control for the log number of risk factors the desk tracks.

6 Customer Demand and Banks' Resource Allocations

The results so far describe where banks concentrate their trading and how well those activities perform. We next ask how banks allocate their trading capacity across activities with different

[12] "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."

[13] Appendix Table IA.6 includes the market-deviation control. Its first three columns retain the main outcome definition; its last three also add absolute other attributable P&L to the numerator, so those columns change both the outcome and the controls.

levels of expected customer demand. The model predicts that capacity follows expected customer demand. Banks can allocate capacity by choosing which markets to participate in and recruiting traders. We examine both margins: how many desks and banks trade in high demand clusters, and where banks recruit traders.

6.1 Clusters with Higher Demand Have More Desks Active (Prediction 4)

High demand risks should draw more of banks' trading capacity, as higher anticipated demand raises specialist hiring in the model (Prediction 4). We test this prediction at the cluster level, treating the desks active in a cluster as the specialist capacity banks devote to it. Banks choose how to organize their desks, and desk classifications are comparable across banks only at the broad asset-class level. Our clustering method groups risk factors within each asset class by their comovement across banks, giving us common categories in which to measure demand and count participating desks and banks. We ask whether more desks and banks become active in a cluster over the following month or quarter when demand for that cluster is higher than usual.

Customer demand is difficult to observe directly, so we use four proxies based on realized trading activity and banks' own demand forecasts. Two use realized trading: the cluster's trailing 90-day average of daily total absolute risk-factor P&L across all banks (*Activity Level*) and its standard deviation over the same window (*Activity Volatility*). The other two use the risk limits banks set from their RENTD forecasts.

RENTD is a term of art covering the demands of clients, customers, or counterparties. Under the Volcker Rule's market-making exemption, a trading desk's activities must be designed not to exceed this demand, taking account of the liquidity, maturity, and depth of the relevant market. Internal limits calibrated to this demand are subject to supervisory review. Given their role in implementing the Volcker Rule, we view RENTD limits as reliable and accurate measures of banks' own expectations of demand.

Risk limits take several forms, including VaR, position, and sensitivity limits. A sensitivity limit constrains a desk's exposure to changes in a particular risk factor, such as an interest rate or exchange rate. The reporting form links sensitivity limits to the risk factors used to attribute P&L, allowing us to assign these limits to our clusters. We keep sensitivity limits that are denominated in U.S. dollars, linked to a reported risk factor, and justified by RENTD. We measure each limit's capacity as the larger of the absolute values of its upper and lower bounds. RENTD limits record expected demand before the trading they are meant to accommodate. That timing likely makes them more useful than proxies built from realized activity.

A bank can have several limits within a cluster because the cluster can contain several of its risk factors, including factors reported by different desks. A single factor can also carry more than one limit. For each bank, cluster, and day, we calculate both the mean and the sum of limit capacities. We then average the bank-level means across reporting banks (*Average RENTD Limit*) and separately average the bank-level sums (*Total RENTD Limit*), giving each bank equal weight.

Suppose two desks at one bank report risk factors assigned to the same cluster, with limit capacities of $100,000 and $300,000, while another bank reports a single limit of $1 million in that cluster. The *Average RENTD Limit* is then $600,000, the average of the banks' means of $200,000 and $1 million, and the *Total RENTD Limit* is $700,000, the average of their totals of $400,000 and $1 million.

Using banks' forecasts reduces the concern that our demand measure simply reflects the realized activity we seek to explain. But these limits do not identify exogenous demand shocks: clusters with larger limits may differ from other clusters in ways we do not observe.

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

$$\textit{Number of Desks}_{c,(t,t+h]} = \beta\, \textit{Demand}_{c,t} + \alpha_c + \gamma_{a(c),t} + \varepsilon_{c,t}, \tag{13}$$

where $\textit{Number of Desks}_{c,(t,t+h]}$ is the maximum daily count of distinct desks active in cluster c over the following h calendar days, $h \in \{30, 90\}$, and $\textit{Demand}_{c,t}$ is one of the four proxies, logged and standardized to zero mean and unit variance. A desk counts as active in a cluster on a day when it reports at least one risk factor assigned to that cluster. A positive β therefore says that more desks participate in a risk cluster when demand is high.

After stripping out persistent differences across clusters (α_c) and daily conditions shared by clusters in the same asset class ($\gamma_{a(c),t}$), we compare each cluster with itself over time. We ask whether more desks participate over the following month when demand in that cluster is higher than usual, beyond what common movements in its asset class would predict.

Higher demand is associated with more desks participating in a cluster (Table 7). In column (5), a one-standard-deviation increase in log *Average RENTD Limit* is associated with 1.04 additional desks at peak participation over the next 30 days. That is about one additional desk, or roughly a one-third increase relative to a benchmark of three desks per cluster. The *Total RENTD Limit* measure gives a similar estimate of 1.01 additional desks. The activity-based estimates are smaller: 0.511 desks for a one-standard-deviation increase in log *Activity Level* and 0.429 for log *Activity Volatility*. The RENTD estimates use a smaller sample because limits are reported only in the post-2020 sample.

The desk results also hold over three months and in a monthly panel. The three-month estimates are somewhat smaller than the one-month estimates: 0.47 compared with 0.51 for *Activity Level*, and 0.89 compared with 1.04 for *Average RENTD Limit*. To reduce overlap between forward windows, columns (9) and (10) re-estimate the RENTD specifications using one observation per cluster each month, with participation measured over the next 30 calendar days. The coefficients are 1.120 for *Average RENTD Limit* and 1.107 for *Total RENTD Limit*, close to the daily estimates and both significant at the 1% level.

A desk is a standing organization with traders, systems, and reporting lines. Adding desk participation therefore differs from changing the capacity of desks already active in a cluster.

Expertise takes time to build, and a new participant competes against incumbents with already installed organizational capital.

The estimated differences in bank participation are smaller. Appendix Table IA.8 repeats the regressions using the maximum daily number of distinct *banks* (in lieu of desks). A one-standard-deviation increase in log *Activity Level* is associated with 0.2 additional banks over the following month. The one-month estimates are 0.14 banks for the *Average RENTD Limit* and 0.09 for the *Total RENTD Limit*. The RENTD coefficients are positive but mostly statistically insignificant. The sample contains relatively few banks because we require consistent reporting from 2014 to 2025, so the limited precision of the bank-count estimates is not surprising.

6.2 Specialists' Hiring is More Responsive to Customer Demand (Prediction 4)

If expected customer demand rises, banks should seek additional traders in the markets where their expertise is most valuable. We test this prediction using a bank–asset-class–month panel of trading-related job postings. The outcome is either the number of openings or an indicator for whether the bank posts any opening in that asset class during the month. We exclude internships and campus programs, which recruit workers before they have developed specialized trading expertise. Appendix IA.A.6 describes how we classify postings and link repeated advertisements for the same vacancy.

We use the same own-activity-share approach to specialization as in the earlier tests, here measured at the asset-class level using absolute desk-level comprehensive P&L. The table reports two versions: *AC Share* uses the previous month's average share, while *Specialist* averages the share over the previous three months. Both are expressed in standard-deviation units.

As in the earlier demand tests, RENTD capacity measures banks' expectations of near-term customer business. For each asset class and month, we construct a market-wide mean of RENTD capacity.[14] *Demand Shock* is the log of this mean minus its average log value over the current and previous two months, which we also standardize. A higher value indicates that banks expect more customer business than in recent months.

Note that we measure specialization before the posting month to capture where the bank already concentrates its trading business. We measure expected demand during the posting month to ask whether banks respond to growing customer business by recruiting more in their existing specialties.

We test whether banks recruit more in their specialties when demand expands:

$$\begin{aligned} \mathit{Postings}_{b,a,t} = {} & \beta_1 \, \mathit{Specialist}_{b,a,t-1} \times \mathit{Demand\ Shock}_{a,t} \\ & + \beta_2 \, \mathit{Specialist}_{b,a,t-1} + \alpha_{b,a} + \delta_{b,t} + \gamma_{a,t} + \varepsilon_{b,a,t}. \end{aligned} \tag{14}$$

Bank-by-asset-class effects $\alpha_{b,a}$ account for where each bank usually recruits. Bank-by-month

[14] We first average capacity across reported risk-factor limits within each bank–asset-class–day, then average these bank-day means within the asset class and month.

effects $\delta_{b,t}$ account for changes in its overall recruiting, and asset-class-by-month effects $\gamma_{a,t}$ account for hiring conditions shared by banks in the same market. We therefore compare how a bank directs recruiting across its asset classes, after accounting for conditions in each market. A positive interaction means that an increase in expected demand is associated with more additional recruiting where the bank specializes more.

Table 8 shows that banks increase recruiting more in their specialties when expected demand rises. We focus on the even-numbered columns, which include asset-class-by-month effects to account for hiring changes shared by banks in the same market and month. These effects absorb, for example, a general increase in commodities recruiting in June, allowing us to ask whether a bank's recruiting response favors its specialties beyond those common changes. They also absorb the demand measure's main effect. The odd-numbered columns omit these effects and separately estimate the response to demand at average specialization and the additional response where banks specialize more.

Consider two asset classes at the same bank whose *Specialist* scores differ by one standard deviation. In column (4), a one-standard-deviation increase in expected demand in each asset class is associated with an additional 0.12 postings per month, on average, in the more specialized class relative to the less specialized class. This difference is about 14% of average monthly postings, using the mean of 0.88 postings per bank and asset class in the broader postings sample. Making the same comparison using the previous month's *AC Share* gives a similar difference of 0.10 postings (column (2)).

The last four columns ask whether this recruiting response also makes banks more likely to advertise any opening. Using *Specialist*, the same increase in expected demand is associated with a 1.5 pp larger increase in the probability of posting an opening in the more specialized asset class (column (8)). Using *AC Share* gives a smaller difference of 0.8 pp that is statistically insignificant (column (6)).

A bank can also change the capacity behind a risk factor without changing which desks participate or whom it hires. We next examine this choice through the limits banks set for individual risk factors.

7 Risk Limits and Reallocation Within and Across Banks (Predictions 5 and 6)

Thus far, we have shown that banks earn more where they specialize and recruit more in those activities when expected demand rises. We now examine how banks allocate risk limits and reallocate activity and expertise. When its risk budget tightens, a bank should preserve its specialist activities relative to the rest of its book (Prediction 5). Expertise should also move toward banks where it is more productive: traders should move toward employers that specialize in their skills, particularly when expected demand rises (motivated by the complementarity in Prediction 6). We examine these

choices using daily risk limits and trading activity and monthly data on trader moves.

7.1 Specialization and Risk-Limit Allocation

Before examining how banks adjust trading when constraints tighten, we document how they allocate risk limits across their activities. The higher returns banks earn where they specialize give them a reason to put more capacity behind those risks. Their limit choices let us examine whether they do so. We ask whether banks set larger limits for their specialist activities, use more of that capacity, and expand it more over time. When usage approaches a limit, the bank also faces a choice about where to accommodate additional business. We ask whether banks are more willing to raise limits in the risks they specialize in.

A large limit need not leave much unused capacity. We measure how closely a factor operates to its limits using *Proximity*, which equals 0 at maximum slack, 100 at the bound, and more than 100 when usage breaches the limit. When several limits apply to a factor, we use the highest value across them, so the measure records its tightest constraint. *Proximity* can rise because usage of the limit increases or because the bank lowers the limit. It therefore measures how close usage is to the reported bound. We classify a factor as near capacity when proximity exceeds 80. We describe the construction of *Proximity* in Appendix IA.A.4.

Table 9 relates limit size and proximity to specialization. Every column includes cluster-by-date fixed effects, which account for daily conditions shared by factors in the same cluster. Our preferred specifications, the even columns, also include bank-by-date fixed effects. These remove daily conditions shared by a bank's activities, such as a bank-wide change in risk appetite. We therefore compare the bank's own risk factors on the same day, after accounting for conditions in each factor's cluster.

Banks have larger limits where they specialize and operate closer to those limits. A one-standard-deviation increase in *Specialist* is associated with a log RENTD limit higher by 0.51 in column (2). This corresponds to about 66% larger limits, a substantial difference in the capacity banks assign to their activities.[15] The same increase in *Specialist* is associated with *Proximity* higher by 0.18 standard deviations in column (4). The probability of operating near capacity is higher by 2 pp (columns (6) and (8)), whether breaches are included in the near-capacity measure or excluded from it.

And specialization also predicts subsequent increases in capacity. For risk factor i on date t, we measure the change in its total limit capacity as

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

where $\overline{Limit}_{it,h}^{\text{fwd}}$ is the average limit over available observations in the following h calendar days, with $h \in \{14, 30, 90\}$. The outcome compares the average capacity available over the forward window

[15] The percentage difference is $100(e^{0.51} - 1)$.

with the capacity available today. Multiplying the log difference by 100 expresses the change in log pp. We winsorize the outcome at the 1st and 99th percentiles.

We estimate

$$\Delta Limit_{it,h} = \beta_1 \, Proximity_{it} + \beta_2 \, (Proximity_{it} \times Specialist_{it}) + \beta_3 \, Specialist_{it} + \mathbf{X}'_{it}\gamma + \eta_i + \alpha_{ct} + \delta_{bt} + \varepsilon_{it},$$

where i identifies a risk factor at a particular desk and bank, c its cluster, and b its bank. The interaction asks whether banks increase limits more in response to *Proximity* when the factor accounts for a larger share of their trading activity.

The fixed effects let us compare changes in capacity across a bank's own activities. Bank-by-date effects δ_{bt} remove common changes across the bank, and cluster-by-date effects α_{ct} remove common conditions in each cluster. Limit fixed effects η_i account for persistent differences across factors in subsequent limit growth. After removing these differences, we ask whether a factor has greater limit growth when it accounts for more of its bank's activity.

The controls $\mathbf{X}_{it}$ account for current limit size and past adjustments. We include the log of the current limit because the same dollar increase produces a larger proportional change in a small limit. We also include the trailing 90-day standard deviation of the log limit to account for differences in how much limits have varied recently.

Limit growth is higher in the risk factors where banks specialize. Columns (1)–(3) of Table 10 include *Specialist* along with the controls and fixed effects. In column (2), a one-standard-deviation increase in *Specialist* is associated with 0.15 log pp more growth in average limit capacity over the following 30 days. The estimates are 0.05 over 14 days and 0.50 over 90 days. Moreover, the size of the coefficients matches intuition, since there is likely some delay in the bureaucratic process required to increase a limit, so specialization better predicts increased limits at longer horizons.

Factors operating closer to their limits also have greater subsequent limit growth. Columns (4)–(6) add *Proximity* and its interaction with *Specialist*. At the average specialization in the construction panel, a one-standard-deviation increase in *Proximity* is associated with 0.55 log pp more limit growth over the following 30 days. The estimates are 0.25 over 14 days and 1.11 over 90 days. The insignificant interaction indicates that we do not detect a stronger response to *Proximity* in specialist activities.

Combined, the results show that banks assign larger limits to the risks they specialize in and operate closer to those limits. These limits also grow more over time, consistent with banks directing additional capacity toward activities where their expertise is most valuable. Limits grow more when usage approaches the bound, providing room for additional trading.

7.2 Constrained Banks Retreat to Their Core (Prediction 5)

If specialization is a deliberate allocation of limited intermediation capacity, a bank with less capacity to deploy should cut where the expertise is cheapest to give up. We ask whether it preserves activity in the risks it specializes in relative to the rest of its book. Each observation is a risk factor reported by a desk at a bank on a day. *Specialist* is the risk-factor measure in equation (11): the factor's daily share of the bank's total absolute risk-factor P&L, averaged over the previous 90 calendar days and standardized over the full risk-factor panel.

We measure pressure on the bank's trading capacity using all its reported risk limits. A bank-level measure lets us ask how the bank allocates activity across its own risks when more of its trading business approaches its limits. Let P_{dt} be the highest proximity across desk d's limits on date t, and let

$$w_{dt} = \frac{\mathrm{VaR}_{d,t-1}}{\sum_{d' \in b} \mathrm{VaR}_{d',t-1}}.$$

Our first measure is the share of the bank's lagged desk VaR accounted for by desks whose maximum proximity exceeds 80:

$$\textit{Near Capacity}_{bt} = \sum_{d \in b} w_{dt}\, \mathbf{1}\{P_{dt} > 80\}.$$

This share runs from zero to one. Our continuous measure, *Utilization*, is constructed from the VaR-weighted average of the desk proximity scores, $\sum_{d \in b} w_{dt} P_{dt}$. Weighting by VaR gives more weight to desks that account for more of the bank's measured trading risk. Desk proximity equals 100 at the limit and can exceed 100 when usage breaches it.

We measure trading activity by the absolute P&L attributed to a risk factor. For bank b's factor f on date t, we calculate

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

where the forward term is average absolute P&L over the following h calendar days and the trailing term is its average over the previous 90 calendar days. We use $h \in \{1, 3, 5\}$ and express the change in log pp.

These short windows let us examine daily adjustments in activity as banks approach their limits. Table 2, Panel D, reports a median near-capacity spell of two days for an individual limit. The spell length describes individual limits rather than the persistence of the bank-level constraint measure.

We estimate

$$\Delta \textit{Activity}_{bft,h} = \beta_1 \Big(\textit{Constraint}_{bt} \times \textit{Specialist}_{bft} \Big) + \beta_2\, \textit{Specialist}_{bft} + \delta_{bt} + \alpha_{ct} + \varepsilon_{bft},$$

where *Constraint* is either *Near Capacity* or *Utilization* and c is the factor's cluster. A positive

interaction coefficient means that activity shifts toward more specialized factors as the bank approaches its limits.

The fixed effects separate this allocation decision from changes shared by the bank or the market. Bank-by-date effects δ_{bt} remove bank-wide changes in activity, including the main effect of the constraint measure. Cluster-by-date effects α_{ct} remove conditions shared by risk factors in the same cluster. We therefore compare a bank's own factors on the same day, after accounting for conditions in their clusters. The interaction asks whether the difference in activity growth between specialist and peripheral factors becomes larger when the bank is closer to its limits.

Table 11 shows that banks expand activity more—or cut it back less—in their specialties than in their peripheral risks when they approach their limits. In column (1), the interaction between *Near Capacity* and *Specialist* is 14.30. Consider two risk factors at the same bank whose *Specialist* scores differ by one standard deviation. A ten-percentage-point increase in the bank's near-capacity share is associated with a 1.43 log percentage point increase in the difference in next-day activity growth between the more specialized and less specialized factors. The corresponding interaction coefficients over three and five days are 10.86 and 6.20.

The even-numbered columns use the continuous *Utilization* measure. Its interaction with *Specialist* is positive and statistically significant over one and three days. The five-day estimate is positive but statistically insignificant.

Appendix Table IA.9 replaces the main table's bank-by-date fixed effects with bank fixed effects while retaining cluster-by-date effects. This lets us compare activity in the same risk cluster on the same day across more- and less-constrained banks, after accounting for persistent differences across banks. At average specialization, activity is lower at banks closer to their limits, indicating that banks pull back while cutting proportionally less in their specialties.

When competitors approach their limits, they may have less room to accommodate customer trades, creating opportunities for other banks. Appendix Table IA.10 examines whether the relationship between specialization and activity growth varies with competitors' constraints. We compare changes across a bank's clusters, accounting for changes affecting the whole bank, conditions shared by all banks trading the same cluster, and persistent differences in the bank's activity growth across clusters.

The within-bank results show that activity shifts toward a bank's core as it approaches its limits. Bräuning et al. (2026) show that constraints propagate within banks across asset classes through the internal allocation of risk limits. Our results relate the allocation across a bank's risk factors to its specialization in those activities.

7.3 Expert Traders Sort Across Banks

In the model, specialized traders and organizational capital jointly give rise to expertise, and organizational capital raises the marginal value of a specialist at a common staffing level (Prediction 6). Having shown in Section 6 that banks recruit more in their specialties, we examine monthly

employment histories to see whether traders move toward more specialized employers. These tests capture the slower adjustment of staffing that complements the higher-frequency allocation of trading activity.

7.3.1 Realized Trader Flows

If specialization reflects genuine expertise that is costly to build, then the labor market should respect these investments (motivated by the complementarity in 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. We test whether 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 use a monthly panel from Revelio Labs that records traders' employers, job descriptions, and seniority. We assign asset classes using the recorded positions and job descriptions, as described in Appendix IA.A.5. We follow anyone observed in a trading role at any point, including months when their title is not a trading title. The data covers 25,000 people with trading roles, 5,448 hires by banks in our sample, and 3,374 hires between covered banks. A move is a change in employer between consecutive observed employment records, which need not be consecutive calendar months.

We measure how central a trader's field is to the bank using the monthly asset-class share of trading activity defined in the postings tests. A position is *peripheral* if the trader's asset class accounts for less than the bank's median share across its asset classes that month, and *core* otherwise. The same field can therefore be peripheral at one bank and core at another.

We first test whether traders are more likely to leave banks where their field accounts for a smaller part of the business:

$$100 \times \mathbb{I}(\textit{Moved})_{i,t} = \alpha + \beta\ \textit{Origin AC Share}_{i,t} + \mathbf{X}'_{i,t}\gamma + \varepsilon_{i,t}.$$

The indicator equals one if trader i changes employers by the next observed employment record. *Origin AC Share* is the origin bank's share of trading activity in the trader's asset class, standardized within each month across bank–asset classes. Higher values mean the trader works at a bank with a larger business in that field. The controls $\mathbf{X}_{i,t}$ include fixed effects and, where indicated, seniority.

Table 12 shows that traders in peripheral positions leave more often. Our preferred column (2) includes bank and asset-class-by-month effects, comparing traders in the same field and month across banks after accounting for persistent differences in turnover across banks. Consider two traders in the same field and month whose employers' asset-class shares differ by one standard deviation. The trader whose employer has the smaller share is 0.04 pp more likely to move. 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 the sample to senior traders only and again finds that traders in less central fields are more likely to move.

Peripheral traders could leave because of pressure at their current employer or better opportunities elsewhere. Columns (5) and (6) examine the push channel using the origin bank's VaR-weighted share of desks near capacity. We take the last daily value of this measure in each month and standardize it over the trader-month panel. A one-standard-deviation increase is associated with a 0.032 pp increase in the probability of changing employers (column (6)), significant at the 10% level. Column (5) asks whether this relationship is stronger for traders in peripheral fields by interacting the bank's constraint with *Origin AC Share*. The interaction is statistically insignificant, so we do not detect disproportionately greater departures among peripheral traders when their employers become more constrained.

We examine the pull channel in two ways. First, columns (7) and (8) ask whether traders leave more often when expected demand is higher in their next observed asset class. Neither the demand level nor the demand change has a statistically significant relationship with departures. These estimates do not establish that rising demand makes traders more likely to change employers.[16]

Second, we ask whether traders who move join banks where their expertise is more central. Appendix Table IA.11 gives an initial view of destinations for 394 moves between covered banks. Of the 97 traders leaving peripheral positions, 70% join banks where their field is core. Among traders leaving core positions, 86% also join banks where their field is core.

To examine these destinations more closely, we restrict attention to traders who change banks and remain in the same asset class. We measure the change in their employer's asset-class share as the destination bank's share in the trader's field minus the origin bank's share. A positive change means that the trader joins a bank that devotes a larger fraction of its trading business to that specialty. We regress this change on whether the trader left a peripheral position:

$$\Delta AC\ Share_{i,t} = \alpha + \beta\, \mathbb{I}(Peripheral)_{i,t} + \mathbf{X}'_{i,t}\gamma + \varepsilon_{i,t}.$$

The coefficient β compares the change for peripheral movers with the change for core movers. Every observation is a move, so the test describes where traders go conditional on changing employers.

Columns (1) and (2) of Table 13 show that peripheral movers experience larger increases in their employer's specialization in their field. Among 205 moves, the increase in the employer's asset-class share is 6.9 pp greater for traders leaving peripheral positions than for those leaving core positions, after accounting for asset-class and month effects. The difference is 7.2 pp when we control for seniority. Note that part of this comparison is mechanical: peripheral traders start from lower

[16] Demand is matched to the trader's next observed asset class. Asset-class-by-month effects use the origin asset class and absorb demand changes shared by traders who remain in the same field.

shares, and that starting share enters both their classification and the measured change.

If growing customer demand raises the value of expertise, this sorting should strengthen when expected demand in the trader's field rises (motivated by the complementarity in Prediction 6). Columns (3) and (4) test this hypothesis by interacting peripheral status with *Demand Shock*:

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

Demand Shock uses the same RENTD-based measure as the postings tests, standardized separately over the trader panel. A positive interaction means that rising demand is associated with a larger increase in the new employer's specialization for peripheral movers relative to core movers.

A one-standard-deviation increase in *Demand Shock* is associated with an additional 2.83 pp increase in the employer's asset-class share for peripheral movers relative to core movers (column (3)). The difference is 3.58 pp when we control for seniority (column (4)). The comparison uses 77 moves with limit data and includes month effects but omits asset-class effects, so the test is underpowered and the evidence is suggestive.

These results are consistent with traders moving toward banks where their expertise is more valuable. Traders leave more often when their field is peripheral to their employer, and peripheral movers experience larger increases in their employer's specialization in that field, particularly when expected customer demand rises. The departure regressions do not establish whether pressure from the current employer or better opportunities elsewhere explain those moves.

7.3.2 Internal Movers

Traders can also move within a bank. These moves let us ask whether staff shifts toward the bank's larger businesses and whether expected customer demand strengthens that shift. Asset-class specialization need not prevent traders from applying some skills in another field or developing new expertise. Indeed, a bank may find it less costly to train an existing employee than to hire a specialist from another firm.

The between-bank tests follow traders who remain in the same asset class, providing the more direct test of whether existing expertise moves toward more specialized employers. Here we follow traders who remain at the same bank but change asset class. The employment records do not identify moves between desks within an asset class.

Columns (5) through (8) of Table 13 examine 832 internal moves among traders observed in the panel for more than 12 months. Of these moves, 355 have the limit data needed for the demand test. The outcome is the bank's share of trading activity in the trader's new asset class minus its share in the old one. A positive value means the trader moves into a larger part of the bank's business. As in the destination tests across banks, every observation is a move.

Traders leaving peripheral asset classes move further toward their bank's larger businesses. After

accounting for bank and month effects, the increase in asset-class share is 18.3 pp greater for these traders than for traders leaving core asset classes (column (5)). In the smaller sample with demand data, the corresponding difference is 15.0 pp at the mean of the standardized demand measure (column (7)). The starting share again enters both the peripheral classification and the measured change.

For movers from core asset classes, higher expected demand in the destination field is associated with a smaller increase in asset-class share. The interaction with peripheral status is statistically insignificant, so we do not detect a stronger relative shift of peripheral traders toward the bank's larger businesses when expected demand rises.

8 Conclusion

Banks maintain distinct trading specialties even in markets where they can trade the same instruments. Using supervisory data, we show that each bank concentrates its trading in a few asset classes and a small set of risk-factor clusters, and that these allocations persist for years. Most clusters are dominated by one or two banks. Within a bank, specialist desks earn higher risk-adjusted returns and leave a smaller share of P&L unattributed. The return result also holds within a desk as its specialization changes.

Trading capacity follows expected demand and specialization. When expected customer demand in a cluster is higher than usual, more desks become active in it over the following month. Within a bank, specialist risk factors have larger demand-based limits and operate closer to them, and a factor's limit grows more when it accounts for more of the bank's activity. As a bank runs closer to its limits, its activity shifts toward its specialist risk factors relative to its peripheral ones.

Monthly labor data gives complementary evidence on where banks seek expertise and where traders take it. When expected demand in an asset class rises, banks raise their job postings by more in the asset classes they specialize in. Traders leave slightly more often when their asset class is peripheral to their employer. Among traders who change banks and keep their asset class, those leaving a peripheral position gain more in their asset class's importance than those leaving a core one, partly by construction. The gap widens when expected demand in the asset class rises, although that comparison rests on a small sample of moves. Traders who switch asset classes inside their bank show the same pattern, but demand does not widen it.

Our model explains these patterns with two features. Specialized traders are more productive where the bank already has installed organizational capital, and the resulting expertise lets the bank manage risk more precisely and face less risk there. These organizational gains can coexist with benefits from diversification. The evidence points to an organizational source of segmentation in trading markets. How much capacity a bank can profitably offer in a risk depends on the expertise it has assembled there.

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.

Patricia Beard. *Blue Blood and Mutiny: The Fight for the Soul of Morgan Stanley*. William Morrow, 2007. ISBN 9780060881917.

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. Lucas, Jr. 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.

J. B. Rosen. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica*, 33(3):520–534, 1965. doi: 10.2307/1911749.

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 2010. doi: 10.1111/j.1467-937X.2009.00583.x.

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

9 Figures

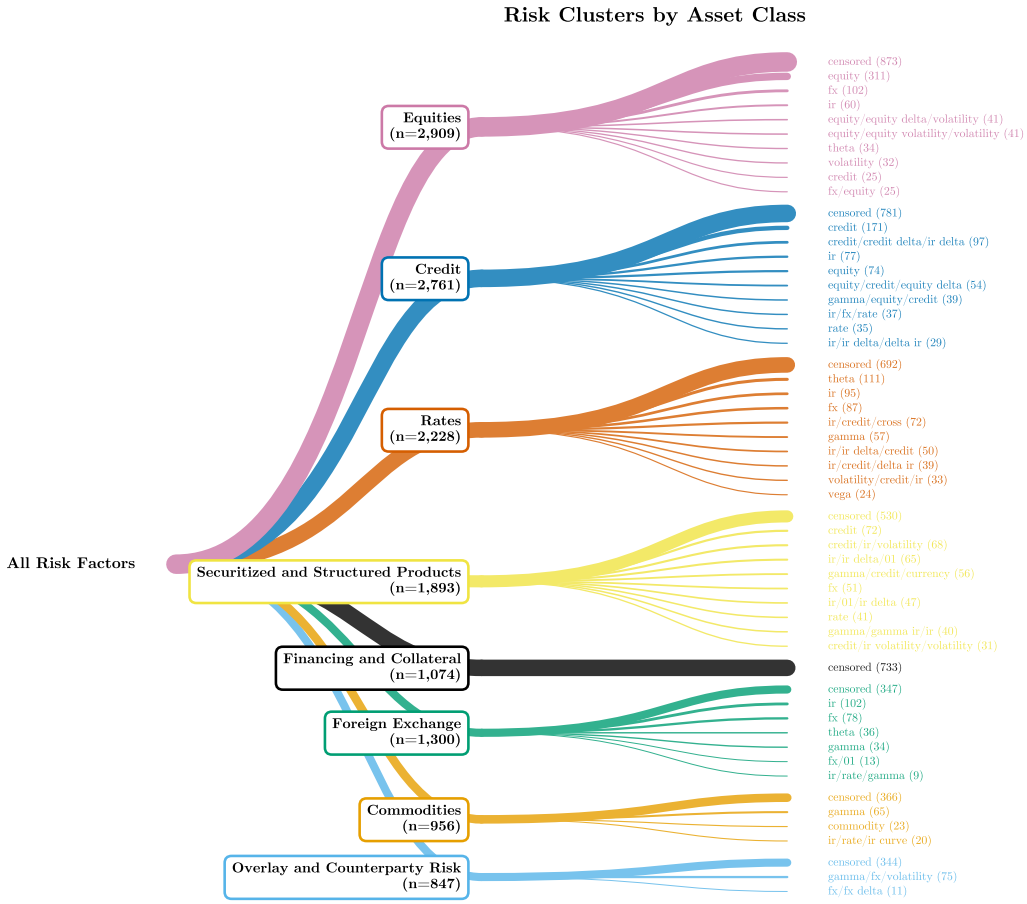


Figure 1: Risk-Factor Classification. Figure shows the largest risk-factor clusters within each asset class. Asset-class nodes report risk-factor counts. Up to ten leaves per asset class show the largest clusters, with member counts in parentheses; branch widths are proportional to those counts. Cluster labels come from constituent risk-factor descriptions. Labels are censored and asset classes omitted when too few banks contribute data. Counts pool the pre-2021 and post-2020 reporting eras, in which clusters are estimated separately. The diagram classifies risk factors into asset classes and clusters within each asset class.

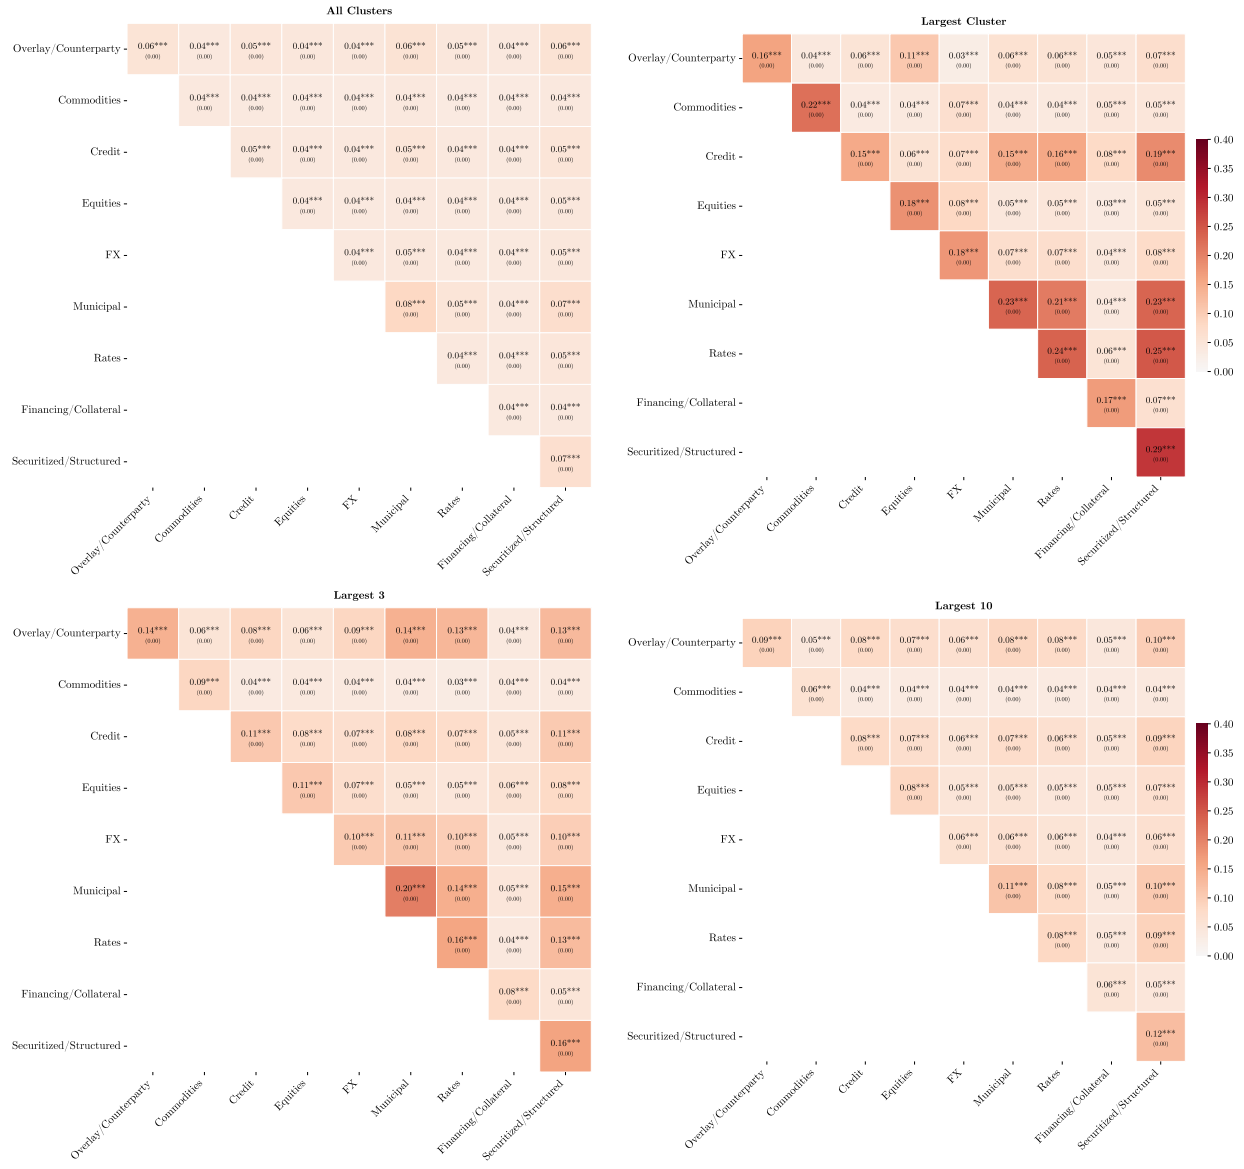


Figure 2: Risk-Factor Correlations Across Asset Classes. Figure plots the average absolute correlation of the risk-factor return proxies, which scale P&L by its trailing three-month volatility, aggregated using Fisher's z-transformation and weighting both by the inverse variance of each pairwise estimate and by risk-factor volatility. The top-left panel uses all clusters, the top-right panel the largest cluster in each asset class, the bottom-left panel the largest three clusters, and the bottom-right panel the largest ten clusters, where cluster size is defined as the sum of constituent risk factors' P&L standard deviations. p-values are in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Sample: post-2020 observations. Each cell aggregates risk-factor-pair correlations for a pair of asset classes, with within-asset-class pairs on the diagonal.

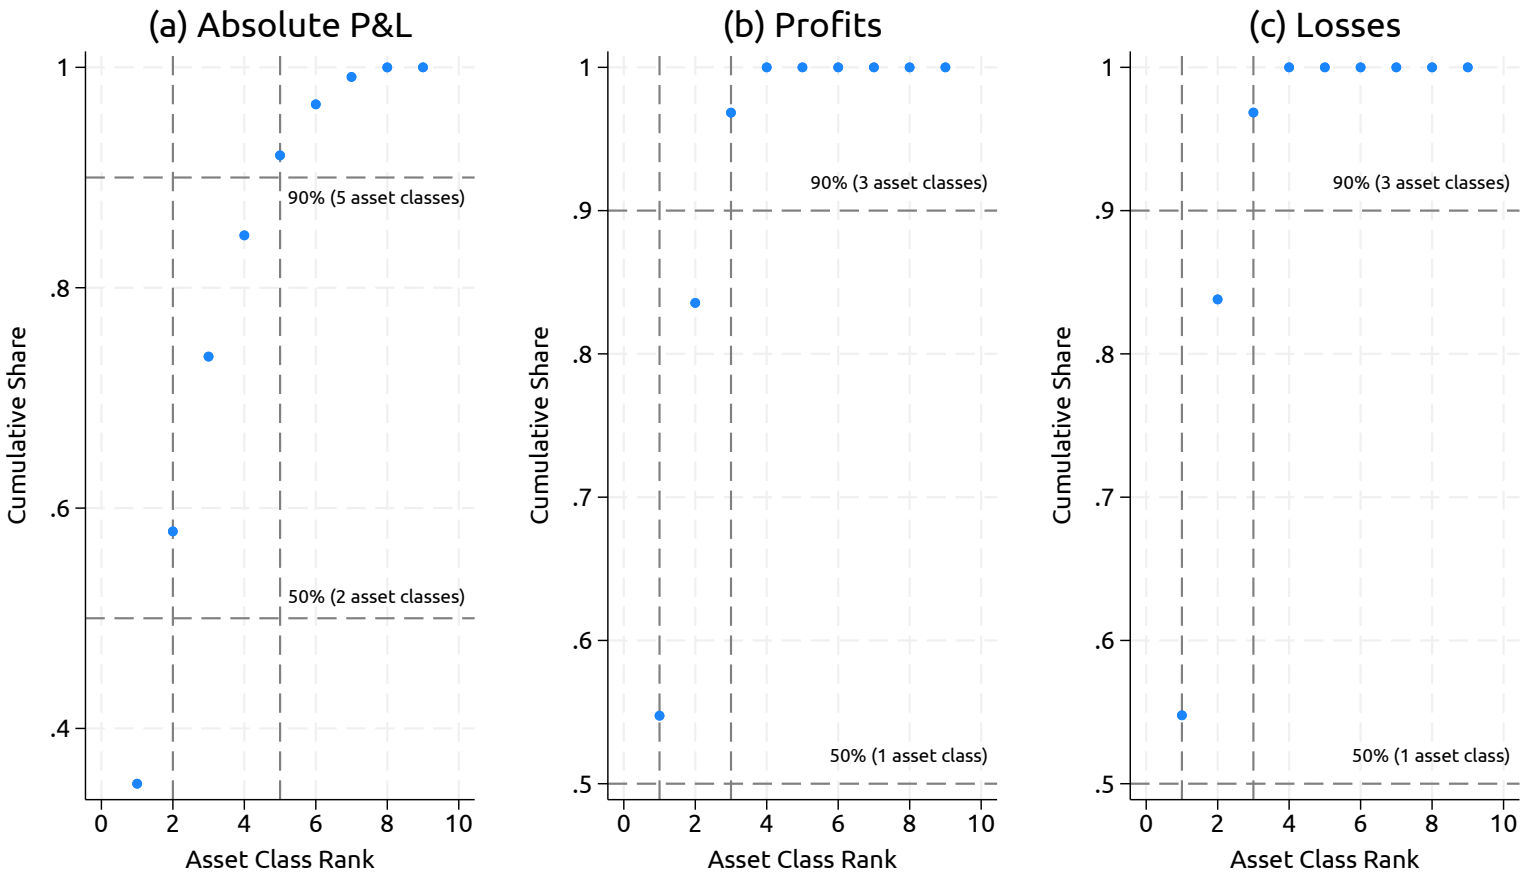


Figure 3: Concentration Across Asset Classes. Figure plots the cumulative share of bank activity accounted for by ranked asset classes. Signed risk-factor P&L is summed within each bank–asset-class–day before calculating activity. Panel (a) uses the absolute value of this net P&L, Panel (b) its positive part (profits), and Panel (c) the absolute value of its negative part (losses). Asset classes are ranked within each bank-day by their shares of the corresponding activity measure. Dashed lines mark the ranks accounting for 50% and 90% of activity. The figure uses the full sample. Rankings use bank × asset class × day observations and are summarized by the median across days within each bank, then the median across banks.

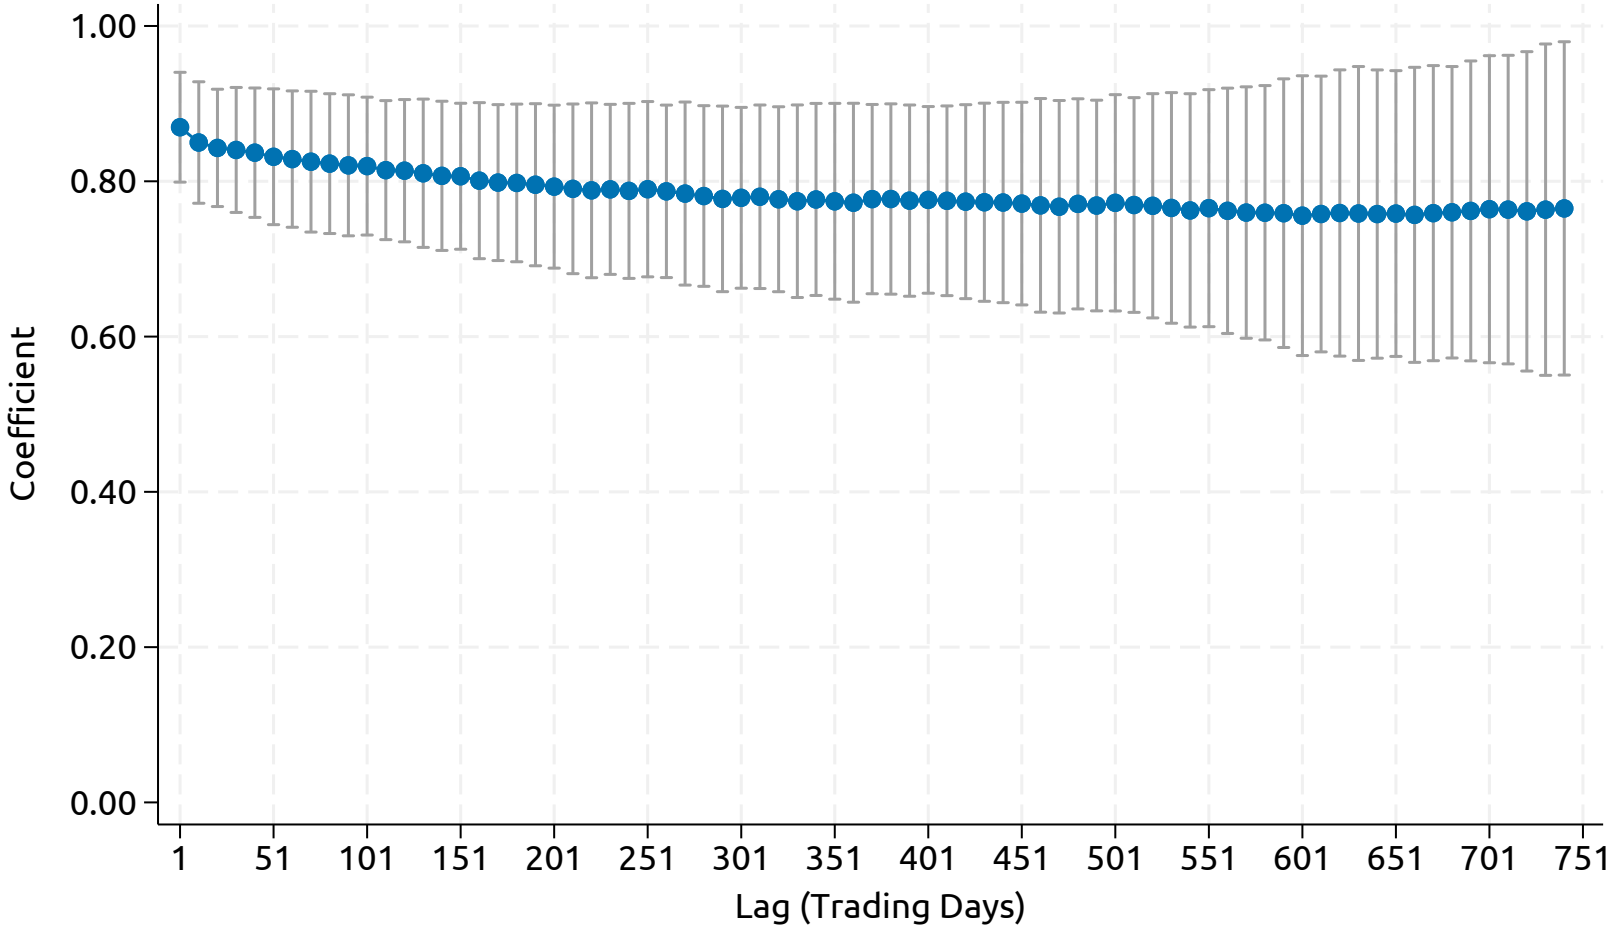


Figure 4: Asset-Class Share Persistence. Figure plots the coefficient from the regression $Share_{b,a,t} = \alpha + \beta\, Share_{b,a,t-\tau} + \varepsilon_{b,a,t}$, estimated separately at each lag τ, where $Share_{b,a,t}$ is the sum of absolute risk-factor P&L in asset class a as a percentage of bank b's total absolute risk-factor P&L on date t. Lags are measured in trading days. The regressions absorb no fixed effects and standard errors are clustered by bank and asset class; capped bars are 95% confidence intervals. Sample: the post-2020 sample. Observations are bank × asset class × trading day.

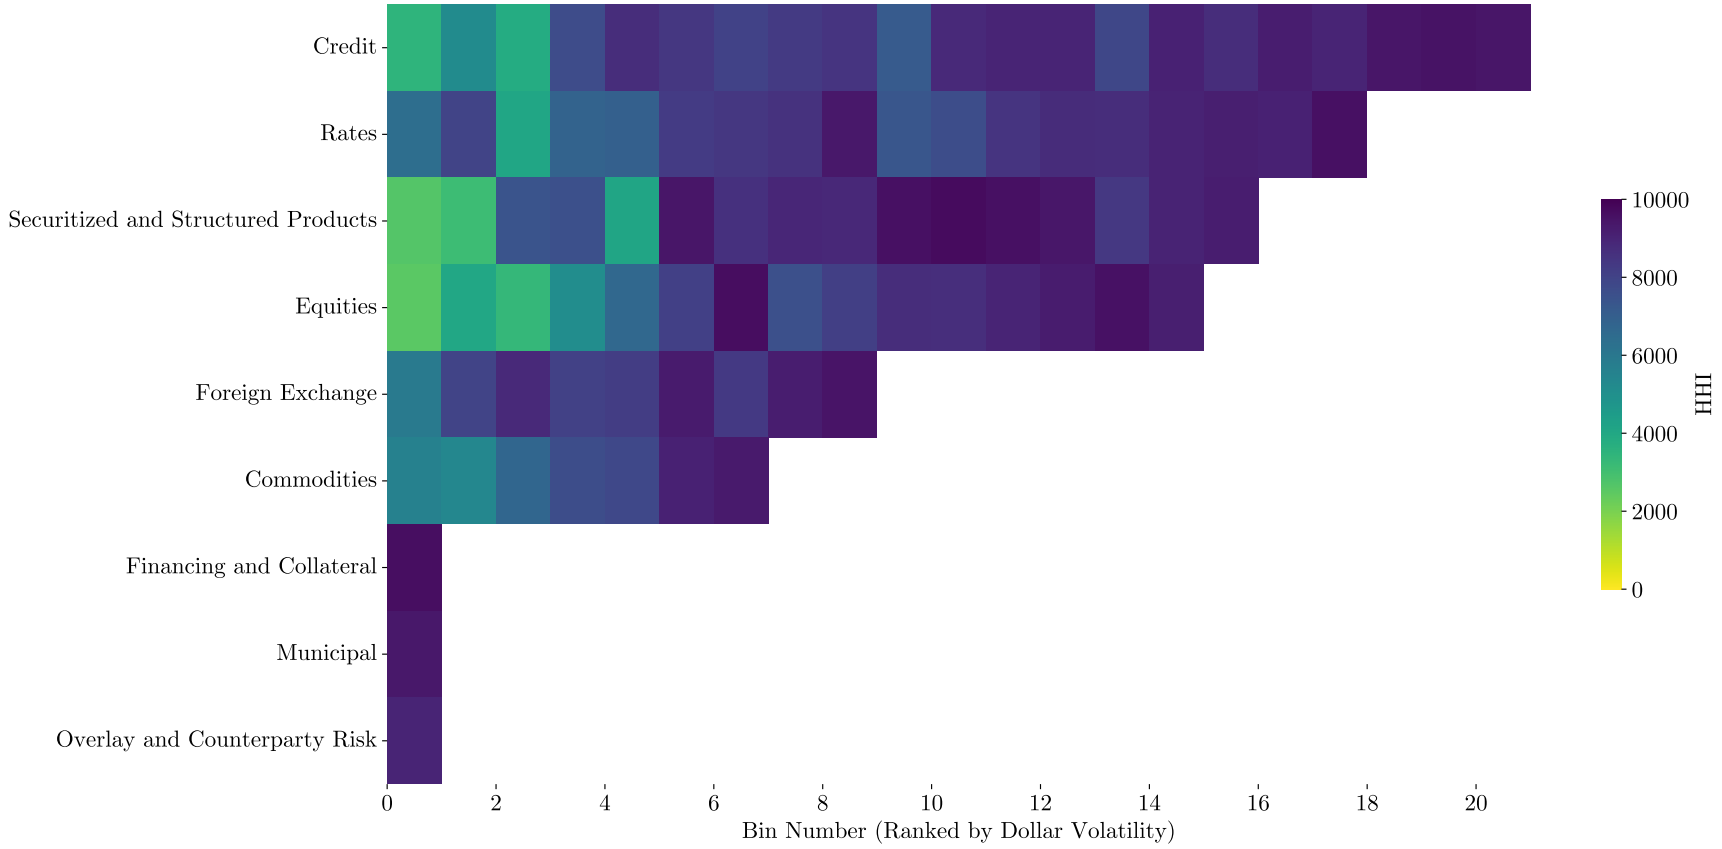


Figure 5: Cluster Concentration and Size. Figure plots a heatmap of concentration across banks within risk-factor clusters. Rows are asset classes. Within each row, clusters are ranked from highest to lowest average monthly dollar P&L volatility and grouped into adjacent bins containing observations from all sample banks. Color shows the average HHI of bank P&L-volatility shares, on a 0–10,000 scale, with darker cells denoting higher values. Each cell summarizes time-averaged daily cluster HHI for a bin of clusters within an asset class.

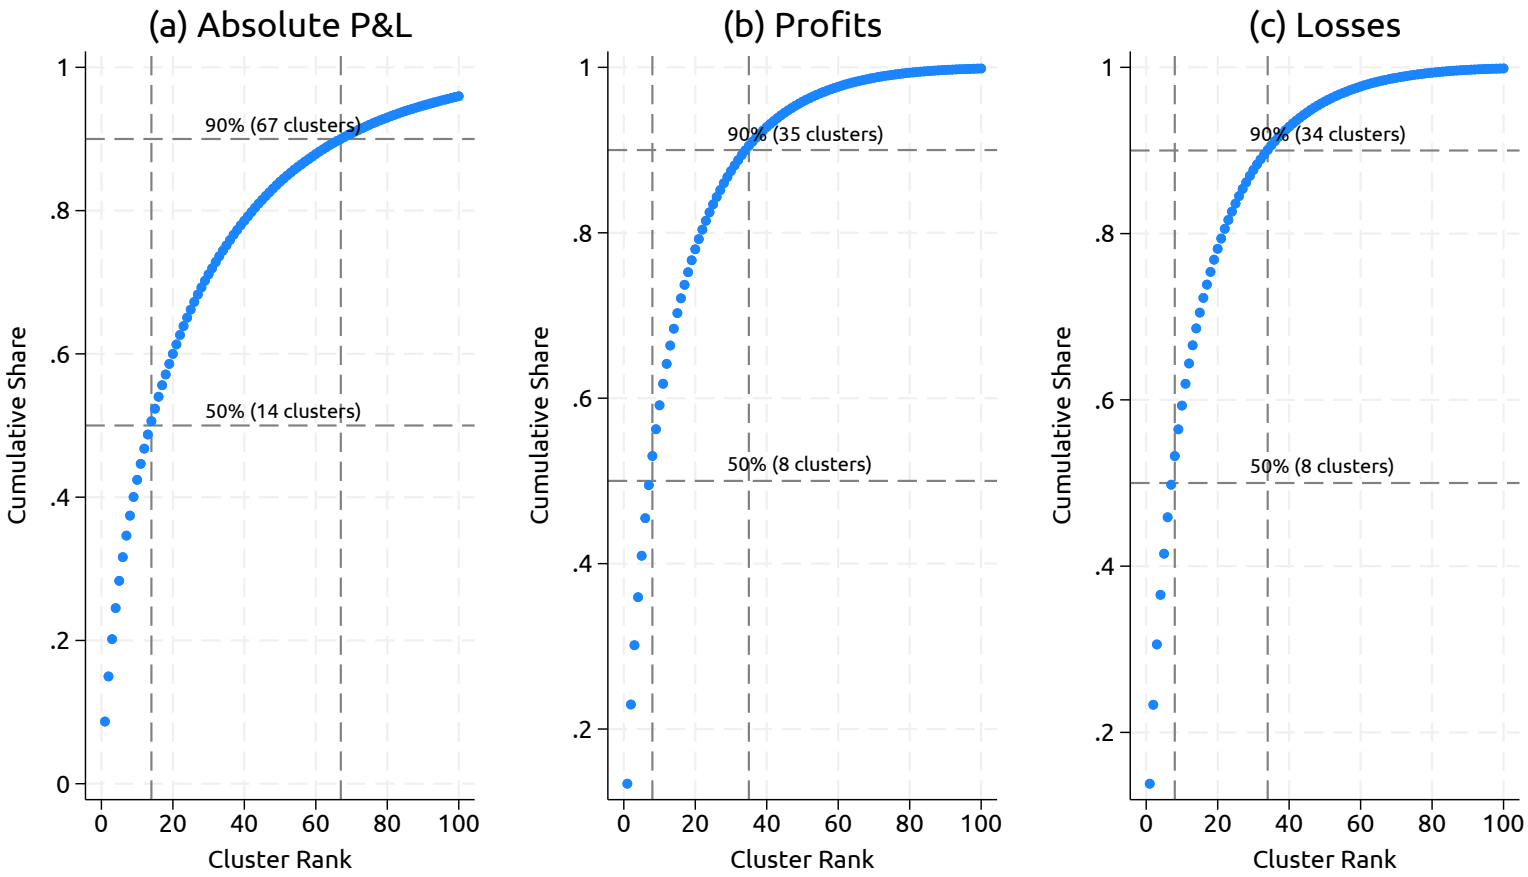


Figure 6: Concentration Across Clusters. Figure plots the cumulative share of bank activity accounted for by the top 100 ranked clusters. Signed risk-factor P&L is summed within each bank–cluster–day before calculating activity. Panel (a) uses the absolute value of this net P&L, Panel (b) its positive part (profits), and Panel (c) the absolute value of its negative part (losses). Clusters are ranked within each bank-day by their shares of the corresponding activity measure. Dashed lines mark the ranks accounting for 50% and 90% of activity. The figure uses the full sample, pooling the pre-2021 and post-2020 reporting eras in which clusters are estimated separately. Rankings use bank × cluster × day observations and are summarized by the median across days within each bank, then the median across banks.

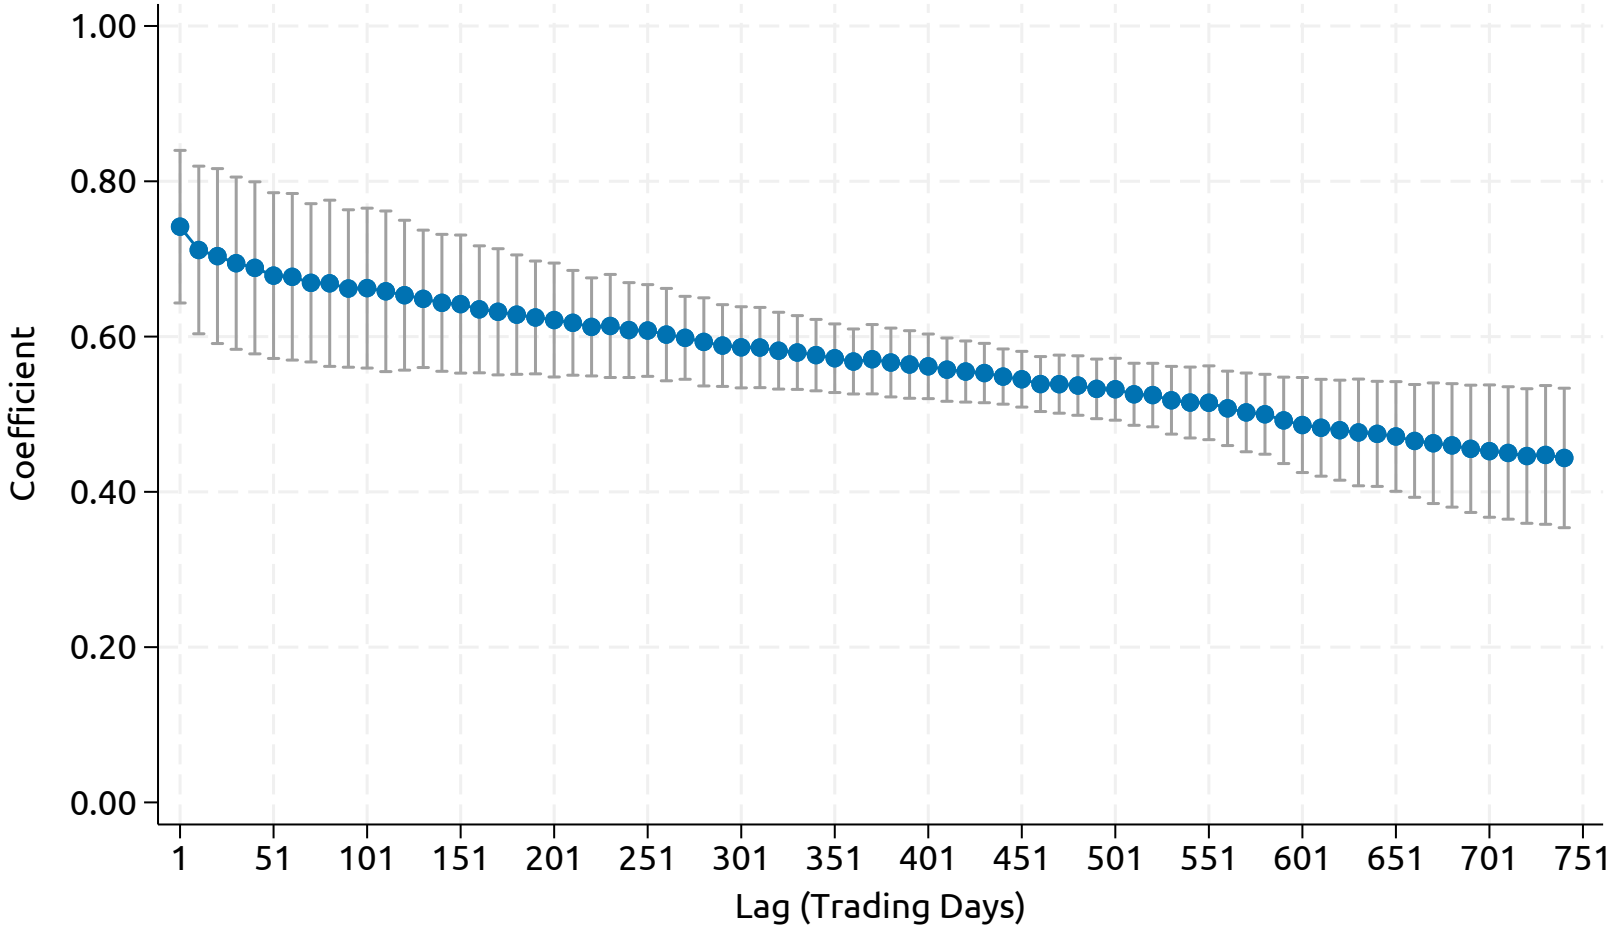


Figure 7: Cluster Share Persistence. Figure plots the coefficient from the regression $Share_{b,c,t} = \alpha + \beta\, Share_{b,c,t-\tau} + \varepsilon_{b,c,t}$, estimated separately at each lag τ, where $Share_{b,c,t}$ is the sum of absolute risk-factor P&L in cluster c as a percentage of bank b's total absolute P&L across clustered risk factors on date t. Lags are measured in trading days. The regressions absorb no fixed effects and standard errors are clustered by bank and cluster; capped bars are 95% confidence intervals. Sample: the post-2020 sample. Observations are bank $\times$ cluster $\times$ trading day.

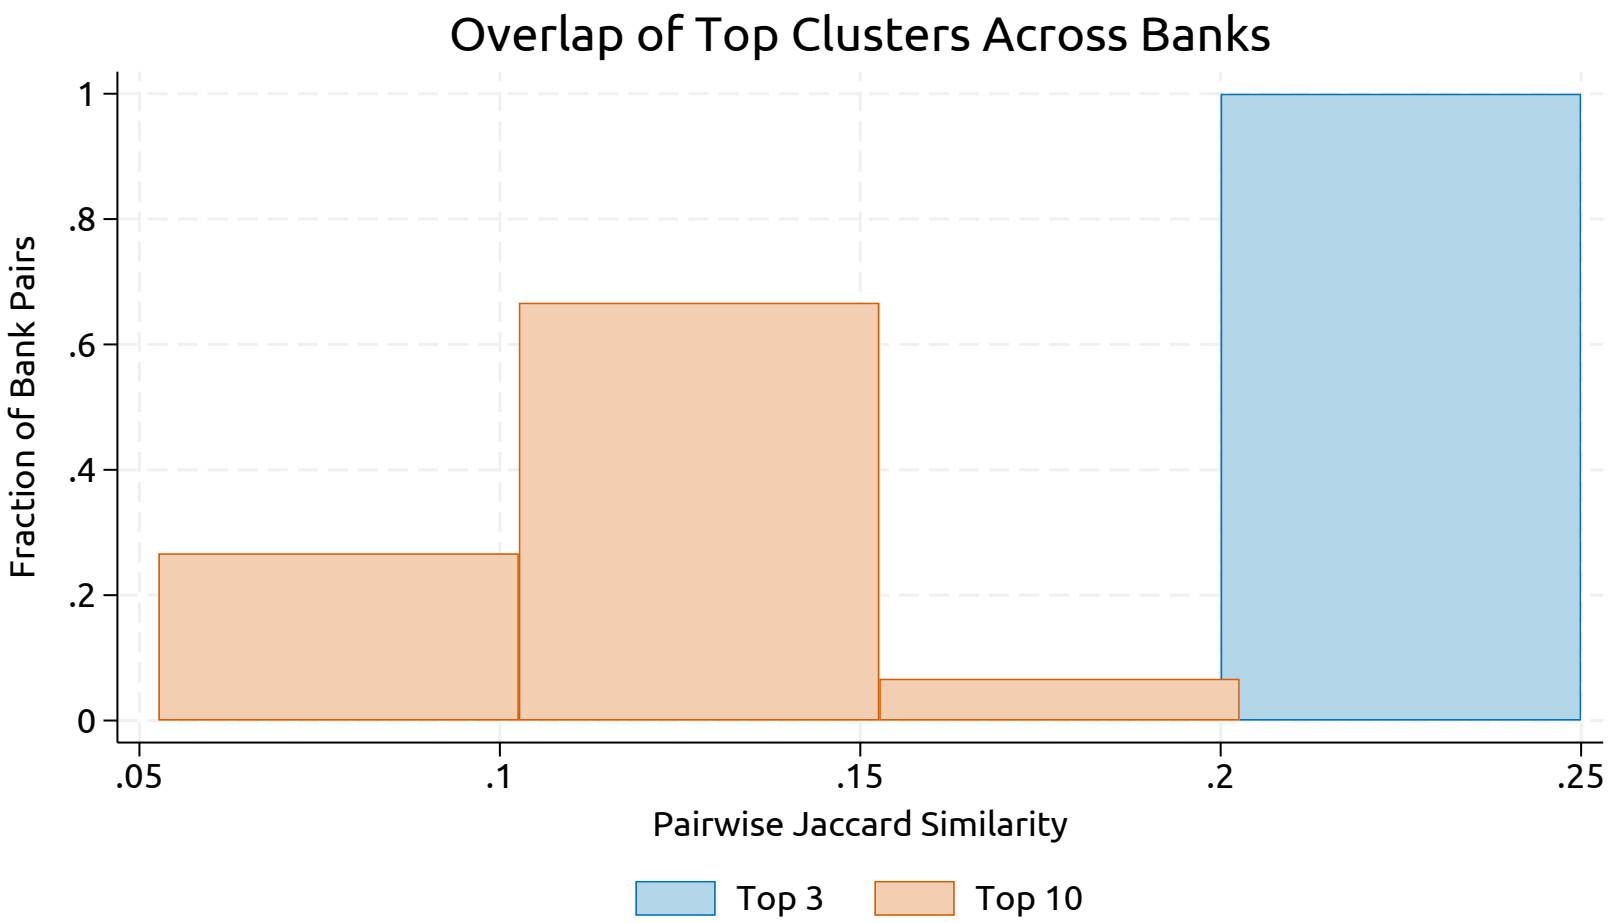


Figure 8: Overlap of Banks' Top Clusters. Figure plots histograms of banks' top-cluster overlap for the top three and top ten clusters. For each bank-month, we rank clusters by their share of the bank's total absolute risk-factor P&L and select the top K. The plotted index is $I/(2K - I)$, where I is the number of shared clusters; it equals Jaccard similarity when both banks have K clusters. Months with no shared top-ranked cluster are excluded. Sample: the full sample period. Each histogram summarizes bank-pair medians across months with positive overlap.

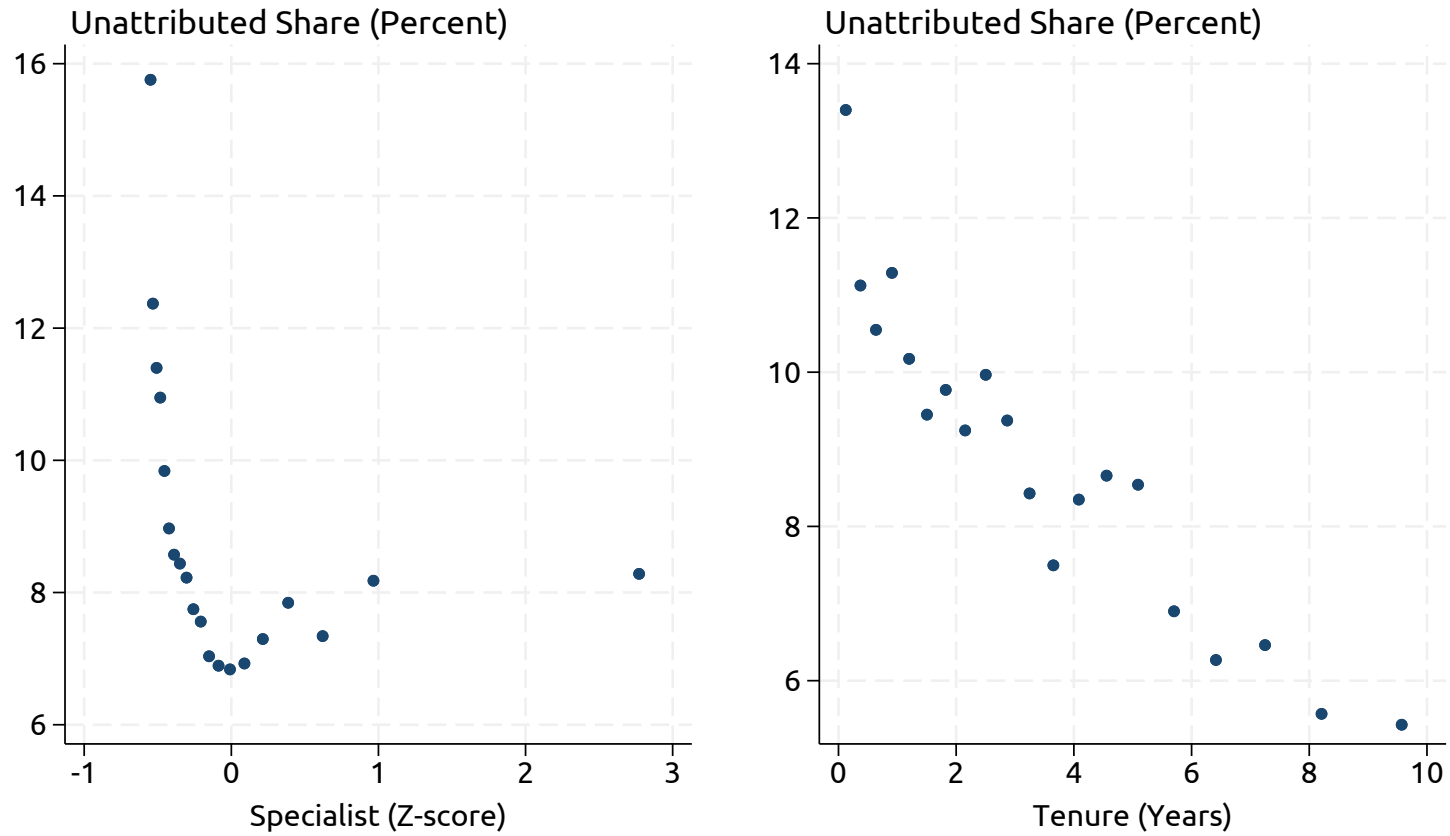


Figure 9: Specialization and Unattributed P&L. Figure plots binned means of the unattributable share of P&L, in %, against desk specialization and tenure. The unattributable share is unexplained P&L (Item 10 under the FR VV-1 reporting requirements) divided by the sum of the P&L from existing positions' risk factors, carry, and the unattributed component, with all components entering in absolute value. The left panel plots binned means of the unattributable share against the desk-level specialization measure, the desk's share of the bank's total absolute desk P&L, averaged over the 90 calendar days ending the day before the observation. The right panel plots binned means against the desk's cumulative number of active trading years; its sample excludes desks that were already active on the bank's first reporting date. Sample: the full sample period. Points are binned means of bank × desk × day observations.

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 desk counts, profitability, and VaR by asset class for the daily sample from 2014 to 2025. *Count* is the number of distinct desks assigned to each asset class over the sample; *Desk–Day N* counts daily desk observations; and *Avg. Desks/Day* is the average number of active desks per day. Profitability and VaR are in millions of dollars. Counts are at the desk level, daily desk counts are aggregated by asset class, and profitability and VaR moments use bank × desk × day observations.

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 Limit Types	4,451,870	19.6	25.0	0.0	0.0	8.2	32.9	58.0
Aging	423,099	24.3	33.8	0.0	0.0	0.2	48.1	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	0.3	12.0	36.8	62.5	80.3
Scenario	1,528,811	8.0	18.1	0.0	0.0	0.0	3.8	33.5
Sensitivity	1,566,027	23.6	23.4	0.1	4.3	16.5	36.9	59.1
VaR	779,949	27.8	24.4	0.0	6.6	23.7	43.1	61.9
All Limit Sources	4,451,870	19.6	25.0	0.0	0.0	8.2	32.9	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	0.4	6.5	20.6	46.0	72.5
Risk Reducing	410	20.5	17.0	0.0	6.3	23.0	30.8	35.8
Other	1,438,585	23.8	26.9	0.0	0.0	15.2	38.4	63.7
Panel C: Proximity Thresholds (%)	Share >50	>80	>90	>95	>100			
All Limit 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	13.9	3.0	1.4	0.9	0.2			
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 Limit 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. Table reports risk limits from form FR VV-1 for dates after December 31, 2020. Panel A reports the number of limits per desk-day. Panel B reports proximity overall, by limit type (aging, other, position, scenario, sensitivity, and VaR), and by source (RENTD, risk appetite, risk reducing, and other). Proximity is 0 at maximum slack, 100 at the allowable bound, and above 100 for a breach; it is bounded below at zero and winsorized at the 99.9th percentile (Appendix IA.A.4). Panel C reports the percentage of limit-days above each proximity threshold. Panel D reports the length in days of spells during which a limit's proximity remains above 80 on consecutive days. Panel E reports the frequency of limit increases and decreases, the median absolute percentage change conditional on a change, and the median days between changes, by limit source and type. Both medians use limit-days with a nonzero change. Observations are bank × desk × day in Panel A, bank × desk × limit × day in Panels B, C, and E, and spells for a bank's desk-level limit in Panel D.

Panel A: Characteristics	Avg. Factors/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 reports risk-factor and cluster characteristics by asset class over the full sample. Panel A reports the average number of risk factors per desk-day, counts of distinct risk-factor–desk combinations before clustering, and profitability in millions of dollars. Panel B reports counts of classified factors and clusters and the percentage of clusters containing a single factor. Clusters are estimated separately for the pre-2021 and post-2020 reporting eras using the method in Appendix IA.A.3. Panel A averages factor counts over bank × desk × day observations and calculates profitability moments over bank × desk × risk factor × day observations; factor counts in both panels and cluster statistics in Panel B are aggregated by asset class and reporting era.

Panel A: Asset-Class Concentration	N	Mean	StDev	Min	Max
Asset Class HHI	25,281	3,063.00	1,317.07	1,687.43	8,779.85
$\ln(1 + \lvert\text{P\&L}\rvert)$	25,281	17.82	1.12	13.08	20.92
Correlation			-0.346^{***}		
Panel B: Cluster Concentration	N	Mean	StDev	Min	Max
Cluster HHI	2,930,779	8,870.57	1,909.71	1,758.26	10,000.00
$\ln(1 + \lvert\text{P\&L}\rvert)$	2,930,779	10.53	4.02	0.00	19.89
Correlation			-0.295^{***}		

Table 4: Asset-Class and Cluster Concentration. Table reports summary statistics for bank concentration and log absolute P&L and the pairwise correlation between them, for asset classes over the full sample in Panel A and for clusters in Panel B. Both HHIs are on a 0–10,000 scale. *Asset Class HHI* is computed from banks' shares of VaR within each asset class; a bank's share is the sum of its desks' VaR in the asset class divided by the same sum across all banks on that day. *Cluster HHI* is computed from banks' shares within a cluster; a bank's share is the sum of its risk factors' trailing 90-day P&L standard deviations within the cluster divided by the same sum across all banks. $\ln(1 + |\text{P\&L}|)$ is the log of one plus the sum of absolute comprehensive desk P&L in the asset class (Panel A) or of absolute risk-factor P&L in the cluster on that day (Panel B), measured in dollars. Panel A excludes asset-class-days with zero total VaR, and Panel B excludes cluster-days with a zero sum of P&L standard deviations. Clusters are estimated separately in the pre-2021 and post-2020 reporting eras. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Observations are asset class $\times$ business day in Panel A and cluster $\times$ day in Panel B.

	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)	(6.91)	(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 5: Specialist Returns. Table reports regressions of desk returns on specialization. The dependent variable is the desk's daily comprehensive P&L divided by its VaR on the previous business day, winsorized at the 5th and 95th percentiles. *Specialist* is the desk's share of the bank's total absolute daily P&L, averaged over the 90 calendar days ending the day before date t, and standardized over the full desk panel. *Excess Specialization*, added in column (7), is the bank's share of its own VaR in the desk's asset class minus that asset class's share of VaR summed across banks; it is measured at the bank × asset-class × day level, is contemporaneous, and is not standardized. Constant omitted. Within R^2 reported. Standard errors are clustered by date and bank–asset class; t-statistics in parentheses, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Observations are bank × desk × day.

	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 6: Specialization and Unattributed P&L. Table reports regressions of the desk's unattributable P&L share on specialization. In columns (1)–(3), the dependent variable is absolute unexplained P&L (Item 10 of form FR VV-1) divided by the sum of absolute P&L attributed to existing positions' risk factors, carry, and the unattributable component. Columns (4)–(6) add absolute other attributable P&L to the denominator. *Specialist* is the desk's share of the bank's total absolute daily P&L, averaged over the 90 calendar days ending the day before date t, and standardized over the full desk panel. Constant omitted. Within R^2 reported. Standard errors are clustered by date and bank–asset class; t-statistics in parentheses, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Observations are bank × desk × day.

	# 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	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Asset Class × Date FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
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 7: Demand and Desk Participation. Table reports regressions of subsequent desk participation on cluster demand measures. The dependent variable is the maximum number of distinct desks, across all banks, active in the cluster over the following 30 calendar days ("1m") or 90 calendar days ("3m"). Four demand proxies enter one at a time. *Activity Level* uses the trailing 90-day mean of the cluster's total gross risk-factor P&L, summed in absolute value across all banks in the cluster. *Activity Volatility* uses the trailing 90-day standard deviation of the same series. *Average RENTD Limit* and *Total RENTD Limit* are built from banks' internal risk limits justified by RENTD analysis. For each U.S. dollar-denominated RENTD sensitivity limit we take the maximum absolute value of the upper and lower bounds as a measure of demand capacity, aggregate across a bank's limits within a bank-cluster-date by the mean (Average) or the sum (Total), and then average across banks to reach the cluster-date. Each proxy is log-transformed and standardized over the cluster-day panel before selecting the monthly observations in columns (9)–(10). Constant omitted. Within R^2 reported. Standard errors are clustered by date and cluster in columns (1)–(8) and by cluster alone in columns (9)–(10); t-statistics in parentheses, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. RENTD regressions use post-2020 observations with cluster-level limit data. Observations are cluster × day in columns (1)–(8) and cluster × month in columns (9)–(10).

	Number of Postings				$\mathbb{I}$(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 8: Demand and Trading Job Postings. Table reports regressions of trading-related job postings on specialization, demand shocks, and their interaction. The dependent variable is the number of the bank's trading-related job postings in the asset class during the month in columns (1)–(4), and an indicator equal to one if the bank posts any opening in the asset class that month in columns (5)–(8). Columns (1)–(2) and (5)–(6) use *AC Share*, which is the monthly mean of the daily share of the bank's total absolute desk P&L that sits in the asset class, in %, standardized within each month across bank–asset classes and lagged one month. Columns (3)–(4) and (7)–(8) use *Specialist*, which is the trailing three-month mean of the same monthly share, standardized over the full monthly bank–asset-class panel. *Demand Shock* is the log of the market-wide mean RENTD limit in the asset class minus its trailing three-month moving average, which includes the current month, standardized over the postings panel. Constant omitted. Within R^2 reported. Standard errors are clustered by month and bank–asset class; t-statistics in parentheses, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Sample: job postings from July 2014 onward for the covered banks matched to the postings data. Observations are bank × asset class × month.

	RENTD Limit		Proximity		$\mathbb{I}$(Near Capacity)		$\mathbb{I}$(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 9: Specialists and Risk Limits. Table reports regressions of risk-limit capacity and usage on specialization. The dependent variables are the log of total RENTD limit capacity mapped to the risk factor in columns (1)–(2); the standardized maximum proximity across those limits in columns (3)–(4); an indicator for proximity above 80 in columns (5)–(6); and an indicator for proximity above 80 and at or below 100 in columns (7)–(8). *Proximity* measures how close current usage is to the allowable bound (0 = maximum slack, 100 = at the limit, and values above 100 = breach) and enters columns (3)–(4) as a z-score taken over the limit panel. *Specialist* is the factor's share of the bank's total absolute risk-factor P&L that day, averaged over the 90 calendar days ending the day before date t and standardized over the full risk-factor panel. Constant omitted. Within R^2 reported. Standard errors are clustered by date and risk limit, the bank × desk × risk-factor limit panel; t-statistics in parentheses, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Sample: U.S. dollar-denominated RENTD sensitivity limits that map to risk factors, which are reported only in the post-2020 era. Observations are bank × desk × risk factor × day.

	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)
σ(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 10: Specialization and Subsequent Limit Growth. Table reports regressions of subsequent RENTD limit growth on specialization and limit proximity. The dependent variable is $100 \times [\ln(\overline{Limit}_h^{\text{fwd}} + 1) - \ln(Limit_t + 1)]$, where the forward limit is the mean limit capacity over the following $h = 14$, 30, or 90 calendar days; it is winsorized at the 1st and 99th percentiles. *Specialist* is the factor's absolute P&L divided by the sum of absolute risk-factor P&L across the bank's factors that day, averaged over the 90 calendar days ending the day before date t and standardized over the full risk-factor panel. *Proximity* is the maximum across the factor's limits and measures how close current usage is to the allowable bound (0 = maximum slack, 100 = at the limit, and values above 100 = breach) and enters as its z-score, standardized over the limit panel. Controls are the log of current limit size and the trailing 90-day standard deviation of the log limit. Columns (4)–(6) add *Proximity* and its interaction with *Specialist* using the same sample as columns (1)–(3). Constant omitted. Within R^2 reported. Standard errors are clustered by date and risk limit; t-statistics in parentheses, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Sample: U.S. dollar-denominated RENTD sensitivity limits that map to risk factors, which are reported only in the post-2020 era. Observations are bank × desk × risk factor × day.

	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 11: Constraints and Risk-Factor Activity. Table reports regressions of changes in risk-factor activity on specialization and its interactions with bank constraints. The dependent variable is $100 \times [\ln(\overline{Activity}^{\text{fwd}} + 1) - \ln(\overline{Activity}^{\text{trail}} + 1)]$, where *Activity* is the risk factor's absolute P&L, the forward mean is taken over the following 1, 3, or 5 calendar days and the trailing mean over the previous 90 calendar days. *Specialist* is the factor's absolute P&L divided by the bank's sum of absolute risk-factor P&L that day, averaged over the 90 calendar days ending the day before date t and standardized over the full risk-factor panel. *Near Capacity* is the lagged-VaR-weighted share of the bank's desks whose maximum limit proximity exceeds 80 on that day. *Utilization* is constructed from the lagged-VaR-weighted average across the bank's desks of each desk's maximum limit proximity. Both constraint measures are unstandardized. Odd columns use *Near Capacity*; even columns use *Utilization*. Every column includes bank-by-date and cluster-by-date fixed effects. The sample uses dates with bank constraint data in the post-2020 reporting era. Constant omitted. Within R^2 reported. Standard errors are clustered by date and bank–cluster; t-statistics in parentheses, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Observations are bank × desk × risk factor × day.

	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 12: Trader Employer Changes. Table reports regressions of employer changes on origin specialization, bank constraints, and asset-class demand. The dependent variable is 100× an indicator that the trader reports a different employer, at any firm, in the next observed employment record. *Origin AC Share* is the origin bank's monthly share of its absolute desk P&L in the trader's asset class, standardized within each month across bank–asset classes. *Bank Constraint* is the bank's lagged-VaR-weighted share of desks whose maximum limit proximity exceeds 80, taken as the last daily value in the month and standardized over the trader-month panel. *Demand* is the log market-wide RENTD limit in the asset class of the trader's next observed employment record, measured in the origin month, and *Demand Shock* is that series less its trailing three-month moving average, including the current month; both are standardized over the trader panel. *Seniority* is the reported seniority level. Constant omitted. Within R^2 reported. Standard errors are clustered by month and bank–asset class; t-statistics in parentheses, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Sample: trader-months at covered banks from July 2014, excluding each trader's final observation; column (4) restricts the sample to senior traders (seniority ≥ 6) and columns (7)–(8) to the post-2020 era in which RENTD limits are reported. Observations are trader × month.

	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 13: Trader Moves and Asset-Class Shares. Table reports regressions of changes in asset-class shares for traders who move. Columns (1)–(4) are movers between covered banks whose asset class is the same at origin and destination; columns (5)–(8) are internal movers, traders who stay at the same bank, change asset class, and are observed more than 12 months after their first retained panel record. The dependent variable is $\Delta(AC\ Share)$, the destination asset-class share minus the origin asset-class share, in pp. Asset-class shares are monthly averages of the share of each bank's total absolute desk P&L in the asset class. *Peripheral* is an indicator that the origin asset-class share is below the bank's within-month median across its asset classes. *Demand Shock* is the log market-wide mean RENTD limit in the destination asset class in the origin month minus its trailing three-month moving average, including the current month, standardized over the trader panel. *Seniority* is the reported seniority level. Constant omitted. Within R^2 reported. Standard errors are clustered by month; t-statistics in parentheses, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Observations are trader moves.

Appendix to

Expertise and Specialization in Bank Trading

Chase P. Ross, Sharon Y. Ross, Kaushik Vasudevan, Martin M. Waibel

The Appendix consists of three sections. Section IA.A provides additional details about the data, including the construction of risk-factor clusters. Section IA.B presents proofs for the model. Section IA.C presents additional tables and figures.

IA.A Data Cleaning and Construction

We construct our analysis sample from regulatory filings submitted by banks under the Federal Reserve's Volcker Rule trading desk reporting requirements. The data comes in two distinct reporting formats: the pre-2021 format and the post-2020 format, which we merge to create a consistent time series. The reporting instructions changed in December 2020 to provide additional granularity on certain fields, particularly limit characteristics and desk classifications. We clean the data so observations under both formats are directly comparable.

IA.A.1 Sample Cleaning

We retain banks that consistently report throughout the sample and convert all monetary values to U.S. dollars using the bank-reported currency conversion rates. We exclude observations missing a bank, desk, or risk-factor identifier. We also remove desks and risk factors that never report activity. A desk is excluded if both its comprehensive P&L and its VaR are missing or zero for all dates; a risk factor within a desk is excluded if its risk-move P&L is missing or zero for all dates.

We limit the influence of outliers by winsorizing risk-factor P&L and desk-level variables (comprehensive P&L, VaR, and P&L components) 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.2 Desk Classifications

We sort desks into ten asset-class groups using their names and descriptions. The classification combines keyword matching with supervised machine learning, trained on a subset of desks for which we have high confidence in the asset-class assignment.

1. Credit: desks that trade corporate debt or credit derivatives, or that 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., collateralized loan obligations, collateralized mortgage obligations, and collateralized debt obligations), 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 or collateral swaps, provide financing for internal counterparties or prime brokerage clients, or engage in securities financing.

9. XVA: desks that conduct valuation adjustment trading, including valuation adjustments to hedge counterparty risk and hedges to funding conditions; desks that focus on centralized hedging of other desks' net positions; or desks that 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.

We manually construct a keyword list for each asset class from trading desk naming conventions. The lists include standard terminology (e.g., "equity," "stock," and "ipo" for the equity class) and specialized jargon (e.g., "sonia," "sofr," and "eurodollar" for the rates class). Some keywords unambiguously signal an asset class (e.g., "xva," "cva," and "counterparty credit risk" for XVA desks). We designate these keywords as "dominant" and assign the asset class whenever one appears in a desk's name or description, without further disambiguation.

For the remaining desks, we match keywords to the combined desk name and description. A match to only one asset class determines the assignment. When the text matches several asset classes, we use TF-IDF similarity to disambiguate. In the rare case a desk has no name or description, we exclude it from the classification rather than setting it to unclassified.

IA.A.3 Risk-Factor Clusters

We sort risk factors into clusters so we can compare what banks trade without relying on the names banks give their risk factors. Two factors at different banks belong to the same cluster when their returns comove, whatever each bank calls them. We cluster within an asset class and across banks: if one bank has two equities desks and another has one, the equities clustering uses the risk factors of all three desks. This keeps the comparison between a bank's equities factors and other banks' equities factors and lets a widely populated cluster stand in for a risk the market trades in common.

We use absolute correlations of signed daily risk-factor return proxies because we care whether a desk trades a given risk, not whether it is long or short. Desks switch between long and short expressions of the same risk as they hedge, make markets, manage inventory, or absorb client flow. A signed-correlation rule would separate desks that are active in the same risk but sit on opposite sides of it.

Two-Stage Construction Trading strategies differ widely in how long they last. We therefore cluster daily return proxies, scaled by trailing three-month P&L volatility, in two stages: long-lived factors determine the correlation structure, and shorter-lived factors join the resulting clusters. In the first stage we keep risk factors with at least a year of data and compute a pairwise correlation only when two factors have at least a year of overlapping observations. Pairs that fail the overlap requirement get zero similarity, so sparse overlap cannot pull them together. We then apply hierarchical agglomerative clustering with average linkage and a distance of one minus the absolute correlation, motivated by Jensen et al. (2023). We cut the resulting dendrogram at distance 0.90, corresponding to an average-linkage absolute-correlation threshold of 0.10.

In the second stage we assign short-lived risk factors, defined as factors with at least ten daily observations, to those clusters. A short-lived factor need not overlap a cluster's centroid, so we use a best-overlap rule: within each candidate cluster we find the member with which the factor has the longest temporal overlap, compute the correlation over that overlap, and assign the factor when the correlation is significant under a Fisher z-test against the null of zero and at least 0.10 in magnitude. Factors that match no cluster are labeled unclassified—idiosyncratic or transient risks. The two stages cover a wide set of factors while estimating the correlation structure from the factors whose correlations are best measured.

We estimate the clustering separately for the two reporting formats. Banks did not keep identical risk-factor identifiers when the data collection changed at the end of 2020, so a factor's return series splits across the pre-2021 and post-2020 samples and cluster identities are specific to an era.

Cluster labels describe the assignments but do not determine them. For exposition we apply TF-IDF to the names and descriptions of a cluster's member risk factors, which yields labels such as "sovereign/credit/european." The labels are poor descriptors because many of the names and descriptions banks report are too broad or general to separate one factor from another. We therefore cluster on realized profits rather than on text. Figure 1 shows the asset-class groupings and their ten largest clusters.

The Correlation Threshold Section 3 explains the choice of $\rho = 0.10$ and reports how the number of equities clusters rises with the threshold. Those counts come from the pre-2021 columns of Appendix Table IA.1; Appendix Figure IA.1 plots the post-2020 clustering.

Correlation Separation and Stability Correlations near 0.10 could still group factors that are only spuriously related. Appendix Figure IA.9 plots the distribution of pairwise correlations

within an asset class, separating pairs in the same cluster from pairs in different clusters. The between-cluster mass spikes just above zero—because we take absolute values, uncorrelated pairs have a small positive average correlation—while the within-cluster mass sits well to the right, with a sizable share of pairs far above the threshold rather than bunched at it. The mean correlation within a cluster is 0.17, with a 75th percentile of 0.30. Pairs in the same asset class but different clusters average 0.04, with a 75th percentile of 0.09.

Appendix Table IA.12 measures this separation and asks whether it is stable across subsamples of days. We summarize the separation with the area under the curve (AUC), $\Pr(|\rho|_{\text{within}} > |\rho|_{\text{between}})$, which equals 0.5 when the two distributions coincide and rises as high-correlation pairs concentrate within clusters. The average AUC across asset classes is 0.81 in the full sample, ranging from 0.76 to 0.86. Because we formed the clusters from correlations, a higher average within-cluster correlation is partly mechanical.

We recompute the correlations on alternating-day subsamples—odd trading days and even trading days—holding the full-sample cluster assignments fixed. The separation persists: the AUC runs from 0.68 to 0.78 across asset classes in both subsamples. Because both subsamples contributed to cluster formation, selection can contribute to this persistence, and the exercise is a stability check on the correlations rather than a held-out test of the clustering itself.

The table also reports hit rates, the share of pairs whose correlation clears 0.10 when estimated on a subsample. Within-cluster pairs clear the threshold in both subsamples 48% of the time, against 16% for between-cluster pairs, so a within-cluster pair is about three times as likely to hold an above-threshold relationship in both halves of the data.

Appendix Table IA.4 compares these correlations across banks as well as across clusters and asset classes. It regresses the absolute correlation of a risk-factor pair on indicators for the two factors being at the same bank, in the same asset class, and in the same cluster. Section 4 reads the coefficients.

Absolute versus Signed Correlations Appendix Figures IA.2 and IA.3 compare the correlation structure of risk factors across banks, asset classes, and clusters under clustering based on absolute and signed correlations of the same signed return proxies. Appendix Figure IA.10 gives the cluster counts across thresholds under signed correlations. Appendix Table IA.4 uses clusters formed from absolute correlations of signed return proxies, while Appendix Table IA.5 uses clusters formed from signed correlations of the same proxies. The dependent variables are absolute and signed correlations, respectively. The two tables therefore use different cluster assignments.

Aggregate Indices and Risk-Factor Returns Within-cluster comovement is not a repackaging of a common aggregate factor in trading returns. Appendix Table IA.13 follows Fama and French (1993) and regresses each risk factor's return proxy on aggregate return indices built from the risk factors themselves, weighted either equally or by the standard deviation of the previous month's

P&L. The first columns report the moments of those indices. The aggregate value-weighted index has a daily return proxy of 0.005, meaning the average P&L is 0.005 times its volatility over the past three months, with a standard deviation of 0.067. Because we do not observe portfolio sizes, the indices are unitless and not comparable to percentages.

The remaining columns show that aggregate measures explain little of the variation across risk factors. Rates and foreign exchange factors have the largest median betas to the aggregate value-weighted index, 0.38 in both cases, but the median R^2 is 0.25%—not 25%—far below the 70% to 90% Fama and French (1993) report for equity portfolios. The typical risk factor has no significant relationship with the index. Indices built only from factors inside the same asset class do no better: betas are mostly below 0.2 and R^2 stays well below 1%. The standard Fama–French market index does worse still, with betas near zero even for risk factors at equity desks.

The low explanatory power of aggregate indices is consistent with the low average absolute correlations within and across asset classes in Figure 2.

IA.A.4 Measuring Risk Limit Utilization

We measure risk-limit utilization from the usage and bounds reported on form FR VV-1. We first compute proximity for each limit-date observation, then aggregate to the desk and bank. The scale runs from 0 for maximum slack to 100 for usage exactly at the limit; values above 100 indicate a breach. To apply the same measure throughout the sample, we harmonize the bounds across reporting formats.

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 is a symmetric constraint (where exposure must stay within $\pm|\textit{Limit Size}|$) 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 is 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 $\textit{Usage}/\textit{Limit Size}$ is negative, but a bank using an absolute-value convention will report a positive limit percent used equal to $|\textit{Usage}|/|\textit{Limit Size}|$. 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 about equal to the absolute value of that ratio (within a tolerance of 1 pp). 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 $+|\textit{Limit Size}|$ and a lower bound of $-|\textit{Limit Size}|$. Remaining limits are treated as directional, with positive limit sizes assigned as upper bounds and negative limit sizes as lower bounds.

With harmonized bounds, headroom is the distance from current usage to the nearest bound. For two-sided limits with upper bound u, lower bound l, and usage x, headroom 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. The resulting proximity score equals 0 at the midpoint of the allowable range, 100 at either boundary, and exceeds 100 in a breach.

For one-sided limits, usage on the opposite side of the constraint can produce negative proximity, indicating that the desk is moving away from the relevant bound. We clamp proximity at zero from below, so that observations far from their relevant constraint register as unconstrained. We winsorize the right tail at the 99.9th percentile.

A desk's proximity on a date is the maximum across its limits, 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.

The bank-date constraint measure aggregates desk proximity using VaR 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 at the bank has a breach (proximity above 100) or is near capacity (proximity above 80) on a given day.

IA.A.5 Trader Employment Data

We construct a monthly employment panel from Revelio Labs. The records link workers to employers and report job titles, descriptions, start and end dates, and seniority. We begin with employment histories for people who worked at banks in our search universe. An initial screen retains positions whose titles match a broad set of keywords, including trading terms, financial products, and seniority terms. A worker enters the screened sample if at least one such position at a bank in the search universe has a start or end date in 2014 or later.

We identify traders within this screened sample using a narrower title search. A position is a trading role if its title contains a term such as *trader*, *quant*, *market maker*, *sales trading*, or *trading desk*. We retain all screened positions for anyone observed in a trading role, including positions at other employers and positions without a trading title. Non-trading positions must still pass the initial title screen. We combine the bank-specific files and remove duplicate position identifiers so that a job appearing in several files enters once.

We assign asset classes from each position's title, description, and Revelio role category. We search the combined text for asset-class keywords, ignoring capitalization and requiring word boundaries so that a keyword does not match part of a longer word. Each keyword contributes at most one match, regardless of how often it appears. The asset class with the most matches becomes the position's primary assignment; ties are broken alphabetically. We retain up to three additional asset classes in order of their match counts, but the labor tests use the primary assignment. Positions with no keyword matches remain unclassified.

We expand the employment spells into a trader $\times$ month panel from January 2014 through June 2025. We truncate spells to this window and treat missing end dates as ongoing through the

panel's end. We drop spells without a usable start date or a positive duration within the window. Each remaining spell contributes observations from its starting month through its ending month, including both endpoints.

We retain one position per trader-month when spells overlap. Trading roles take priority over non-trading roles. Within each group, we give priority to positions at a designated set of large banks. Any remaining tie is resolved by keeping the first record. We carry the selected position's employer, asset class, and seniority into the monthly panel and use Revelio's numerical seniority score in the regressions.

The mobility tests use trader-months with a primary asset-class assignment from July 2014 onward. We drop unclassified observations before identifying moves and match the remaining records to the regulatory data by employer, asset class, and month. The employer histories extend beyond the banks covered by the regulatory data, so we can observe departures to other firms. Comparing specialization at the origin and destination requires both employers to have the relevant trading data.

A move is a change in employer between consecutive retained employment records. We assign the departure to the origin observation and exclude each trader's final observation from the departure regressions because no later employment record is available. Consecutive records need not fall in consecutive calendar months. The destination tests across banks follow traders who remain in the same asset class. Internal moves are changes in asset class at the same bank, measured more than 12 months after the trader's first retained observation. This restriction measures time since entering the panel, not tenure at the current bank. The records do not identify moves between desks within an asset class.

IA.A.6 Job Postings Data

Dewey LinkUp supplies bank job postings scraped daily. Each posting records the hiring bank, job title, description, and the dates on which the posting is first observed and later removed.

We classify each posting into an asset class using a large language model (LLM) applied to the posting text description. For the 26% of postings still unclassified after the LLM classification, we use a keyword search over job titles. We drop postings that we cannot assign to any asset class after these steps (about 22% of postings).

We restrict the sample to trading roles and exclude internships, summer positions, and industrial-placement programs, identified from the job title.

Re-listed vacancies can appear as several consecutive records in the daily scrape. We combine these records to avoid double-counting the vacancy. 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 securitized products—and reshape the data to the posting $\times$ asset-class level. We attribute each posting 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 get the flow of new postings.

IA.B Model Derivations

IA.B.1 Model Preliminaries

We summarize anticipated customer demand in cluster c by its scaled second moment, $M_c \equiv a^2 \sum_{f \in c} \mathbb{E}[d_f^2]$, which enters the hiring condition in Equation (6).

For predictions about binding constraints, we hold hiring fixed and assume that one bank has a binding risk budget while the other banks have slack budgets.

IA.B.2 Trading and Hiring Equilibrium

We solve by backward induction for every $\ell \geq 0$. We first characterize the unique trading Nash equilibrium after every hiring profile and realization of customer demand. Substituting the resulting prices and positions into banks' expected utility gives their hiring objectives. We establish existence of a hiring Nash equilibrium using concavity and a bound on hiring, and uniqueness by comparing hiring optimality conditions across two assumed equilibria. The unique hiring equilibrium and the trading equilibrium after every hiring profile and demand realization together give the unique subgame perfect equilibrium.

For comparisons at sufficiently small positive ℓ, we additionally assume that the marginal value of the first hire at $\ell = 0$ differs from the wage:

$$M_c O_{b,c} \neq 4\gamma w[1 + \kappa^2(\tau_0 + O_{b,c})]^2$$

for every bank b and cluster c. This excludes equality between the marginal value of hiring at zero and the wage: each bank is strictly above or below the threshold for starting to hire. This restriction provides differentiable continuation of hiring and its demand response from $\ell = 0$. Existence and uniqueness below hold for every $\ell \geq 0$ without this additional restriction.

Trading equilibrium. We hold all banks' hiring fixed. For a positive position in factor f, bank b's first order condition gives $z_{b,f} = p_f/(\ell + \gamma v_{b,c})$. We write $h_{b,c} = (\ell + \gamma v_{b,c})^{-1}$ for the coefficient on p_f, so $z_{b,f} = h_{b,c} p_f$. Greater expertise lowers $v_{b,c}$ and therefore gives bank b a larger position at a given intermediation price.

Bank b's objective in Equation (3) has second derivative $-2\ell - \gamma v_{b,c} < 0$ with respect to $z_{b,f}$, so its first order condition is sufficient for its optimal position given rivals' positions.

For $d_f > 0$, all banks' equilibrium positions are positive, and Equation (5) is the unique solution to their first order conditions and customer inverse demand and therefore gives the unique trading Nash equilibrium. If $d_f = 0$, any positive position yields negative trading utility, so every bank chooses zero. Thus the trading equilibrium exists and is unique after every hiring profile and demand realization.

Hiring objective. Holding rivals' hiring fixed, we define bank b's expected trading utility in cluster c before wage costs as

$$F_{b,c}(n_{b,c}) \equiv \mathbb{E}\left[\sum_{f\in c}\left\{z_{b,f}p_f - \frac{\gamma}{2}v_{b,c}z_{b,f}^2\right\}\right].$$

For each hiring choice $n_{b,c}$ and realization of customer demand, p_f and $z_{b,f}$ are the resulting trading equilibrium price and position, and $v_{b,c}$ is the residual variance given by Equation (2). We substitute the equilibrium prices and positions from Equation (5) to obtain

$$F_{b,c}(n_{b,c}) = \frac{M_c h_{b,c}(n_{b,c})[1 + \ell h_{b,c}(n_{b,c})]}{2[1 + \ell\sum_{j\neq b} h_{j,c} + \ell h_{b,c}(n_{b,c})]^2}. \tag{15}$$

Bank b chooses hiring in cluster c to maximize $F_{b,c}(n_{b,c}) - wn_{b,c}$ over $n_{b,c} \geq 0$, holding rivals' hiring fixed and anticipating the trading equilibrium after each hiring choice.

Concavity in hiring. We show that $F_{b,c}$ is strictly concave in bank b's own hiring, holding rivals' hiring fixed. Strict concavity makes the hiring first order condition and the inequality for zero hiring sufficient for an optimum. We also use it to establish equilibrium existence below.

Bank b's expertise is $E_{b,c} = O_{b,c}\sqrt{1 + n_{b,c}}$. We define its residual payoff precision and differentiate with respect to expertise:

$$\begin{aligned}
\rho_{b,c}(E_{b,c}) &= \frac{1}{v_{b,c}} = \frac{\tau_0 + E_{b,c}}{1 + \kappa^2(\tau_0 + E_{b,c})},\\
\frac{\partial\rho_{b,c}}{\partial E_{b,c}} &= \frac{1}{[1 + \kappa^2(\tau_0 + E_{b,c})]^2} > 0,\\
\frac{\partial^2\rho_{b,c}}{\partial E_{b,c}^2} &= -\frac{2\kappa^2}{[1 + \kappa^2(\tau_0 + E_{b,c})]^3} < 0.
\end{aligned}$$

We write $T_{b,c} = 1 + \ell\sum_{j\neq b} h_{j,c} \geq 1$ for the part of the price denominator that does not vary with bank b's own hiring. Since $h_{b,c} = \rho_{b,c}/(\gamma + \ell\rho_{b,c})$, we express bank b's expected trading utility as a function of its residual payoff precision:

$$\Phi_{b,c}(\rho_{b,c}) = \frac{M_c}{2}\frac{\rho_{b,c}(\gamma + 2\ell\rho_{b,c})}{[T_{b,c}\gamma + \ell(T_{b,c} + 1)\rho_{b,c}]^2}.$$

We hold rivals' hiring and hence $T_{b,c}$ fixed and differentiate with respect to $\rho_{b,c}$ to obtain

$$\Phi'_{b,c}(\rho_{b,c}) = \frac{M_c\gamma}{2}\frac{T_{b,c}\gamma + \ell(3T_{b,c}-1)\rho_{b,c}}{[T_{b,c}\gamma + \ell(T_{b,c}+1)\rho_{b,c}]^3} > 0,$$
$$\Phi''_{b,c}(\rho_{b,c}) = -M_c\gamma\ell\frac{2T_{b,c}\gamma + \ell(T_{b,c}+1)(3T_{b,c}-1)\rho_{b,c}}{[T_{b,c}\gamma + \ell(T_{b,c}+1)\rho_{b,c}]^4} \leq 0.$$

Bank b's expertise increases at a decreasing rate with hiring:

$$\frac{\partial E_{b,c}}{\partial n_{b,c}} = \frac{O_{b,c}}{2\sqrt{1+n_{b,c}}} > 0,$$
$$\frac{\partial^2 E_{b,c}}{\partial n_{b,c}^2} = -\frac{O_{b,c}}{4(1+n_{b,c})^{3/2}} < 0.$$

We apply the chain rule using the derivatives of precision with respect to expertise, $\partial\rho_{b,c}/\partial E_{b,c} > 0$ and $\partial^2\rho_{b,c}/\partial E_{b,c}^2 < 0$, calculated above:

$$\frac{\partial \rho_{b,c}}{\partial n_{b,c}} = \frac{\partial \rho_{b,c}}{\partial E_{b,c}}\frac{\partial E_{b,c}}{\partial n_{b,c}} > 0,$$
$$\frac{\partial^2 \rho_{b,c}}{\partial n_{b,c}^2} = \frac{\partial^2 \rho_{b,c}}{\partial E_{b,c}^2}\left(\frac{\partial E_{b,c}}{\partial n_{b,c}}\right)^2 + \frac{\partial \rho_{b,c}}{\partial E_{b,c}}\frac{\partial^2 E_{b,c}}{\partial n_{b,c}^2} < 0.$$

We apply the chain rule to $F_{b,c}(n_{b,c}) = \Phi_{b,c}(\rho_{b,c}(E_{b,c}(n_{b,c})))$ to obtain

$$F''_{b,c}(n_{b,c}) = \Phi''_{b,c}(\rho_{b,c})\left(\frac{\partial \rho_{b,c}}{\partial n_{b,c}}\right)^2 + \Phi'_{b,c}(\rho_{b,c})\frac{\partial^2 \rho_{b,c}}{\partial n_{b,c}^2} < 0$$

for every $\ell \geq 0$. Subtracting the linear wage bill $wn_{b,c}$ preserves this strict concavity.

Existence of a hiring equilibrium. The bounds $\Phi'_{b,c}(\rho_{b,c}) \leq \Phi'_{b,c}(0) = M_c/(2T_{b,c}^2\gamma) \leq M_c/(2\gamma)$ and $\partial\rho_{b,c}/\partial E_{b,c} \leq 1$ give

$$F'_{b,c}(n_{b,c}) \leq \frac{M_c O_{b,c}}{4\gamma\sqrt{1+n_{b,c}}}.$$

Above the following bound, marginal trading utility is strictly below the wage w, so reducing hiring raises the bank's objective. Every hiring best response therefore satisfies

$$0 \leq n_{b,c} \leq \left[\left(\frac{M_c O_{b,c}}{4\gamma w}\right)^2 - 1\right]_+. \tag{16}$$

The bound on marginal trading utility also gives $F'_{b,c}(n_{b,c}) \to 0$ as $n_{b,c} \to \infty$. By strict concavity, bank b has a unique hiring best response in cluster c: zero if $F'_{b,c}(0) \leq w$, and otherwise the positive solution to $F'_{b,c}(n_{b,c}) = w$.

Payoffs in the hiring game are continuous in all banks' hiring choices: residual variances are bounded below by $\kappa^2 > 0$, so the trading equilibrium formulas have positive denominators. Using

Equation (16), we restrict hiring to compact, convex strategy sets without excluding any best response. Since each bank's payoff is concave in its own hiring, a pure strategy hiring Nash equilibrium exists.[17]

We combine this hiring equilibrium with the unique trading equilibrium after every hiring profile to establish existence of a subgame perfect equilibrium. We next establish uniqueness of hiring.

Uniqueness of hiring. At $\ell = 0$, the intermediation price is $p_f = ad_f$, so each bank's hiring problem is independent of rivals' hiring. Each bank's unique hiring best response therefore determines the unique hiring equilibrium.

For $\ell > 0$, fix cluster c. We prove uniqueness by contradiction. In particular, assume by contradiction that two distinct hiring equilibria exist. Write bank j's hiring and position coefficient in equilibrium $k \in \{1, 2\}$ as $n_{j,c}^{(k)}$ and $h_{j,c}^{(k)}$, respectively. The common price denominator in equilibrium k is

$$H_c^{(k)} = 1 + \ell \sum_j h_{j,c}^{(k)}.$$

Label the two assumed equilibria so that $H_c^{(2)} \geq H_c^{(1)}$. Since $h_{j,c}$ increases strictly with $n_{j,c}$ for each bank j, at least one coefficient differs across the two equilibria. As their sum is weakly greater in equilibrium 2, some bank b must have $h_{b,c}^{(2)} > h_{b,c}^{(1)}$ and hence $n_{b,c}^{(2)} > n_{b,c}^{(1)} \geq 0$.

We differentiate Equation (15) with respect to bank b's hiring, holding rivals' hiring fixed. Since $\partial H_c / \partial n_{b,c} = \ell \, \partial h_{b,c} / \partial n_{b,c}$, marginal trading utility is

$$F'_{b,c}(n_{b,c}) = \frac{M_c}{2H_c^3} \, h_{b,c} \frac{\partial h_{b,c}}{\partial n_{b,c}} \left[\frac{H_c}{h_{b,c}} + 2\ell(H_c - 1) - 2\ell^2 h_{b,c} \right].$$

We compare this expression as a function of $n_{b,c}$ and H_c on $H_c \geq 1 + \ell h_{b,c}$. For given customer demand, holding H_c fixed holds the price $p_f = ad_f / H_c$ fixed.

Holding H_c fixed, we first vary bank b's hiring. Hiring raises expertise $E_{b,c}$ and the coefficient $h_{b,c}$. Substituting the expertise and variance formulas gives

$$h_{b,c} \frac{\partial h_{b,c}}{\partial n_{b,c}} = \frac{\gamma O_{b,c}^2}{2} \, \frac{1 + \tau_0 / E_{b,c}}{[\gamma + (\ell + \gamma\kappa^2)(\tau_0 + E_{b,c})]^3} > 0.$$

The numerator $1 + \tau_0 / E_{b,c}$ decreases with expertise, and the denominator increases. Thus the product $h_{b,c} \, \partial h_{b,c} / \partial n_{b,c}$ decreases with hiring.

The term $H_c / h_{b,c} + 2\ell(H_c - 1) - 2\ell^2 h_{b,c}$ is positive: using $H_c \geq 1 + \ell h_{b,c}$ gives a lower bound of $1/h_{b,c} + \ell > 0$. Its derivative with respect to $h_{b,c}$ at fixed H_c is

$$\frac{\partial}{\partial h_{b,c}} \left[\frac{H_c}{h_{b,c}} + 2\ell(H_c - 1) - 2\ell^2 h_{b,c} \right] = -\frac{H_c}{h_{b,c}^2} - 2\ell^2 < 0.$$

[17] See Rosen (1965).

Since $h_{b,c}$ rises with hiring, this positive term decreases as bank b hires more. Both positive expressions multiplying $M_c/(2H_c^3)$ therefore decrease with hiring, while $M_c/(2H_c^3)$ stays fixed. Marginal trading utility falls with bank b's hiring at fixed H_c.

Next, hold bank b's hiring fixed. Then $M_c h_{b,c}(\partial h_{b,c}/\partial n_{b,c})/2$ is a positive constant, so the sign of the change in marginal trading utility with H_c follows from

$$\begin{aligned}
&\frac{\partial}{\partial H_c}\left[\frac{H_c/h_{b,c} + 2\ell(H_c-1) - 2\ell^2 h_{b,c}}{H_c^3}\right] \\
&\quad = \frac{-2H_c(1+2\ell h_{b,c}) + 6\ell h_{b,c}(1+\ell h_{b,c})}{h_{b,c}H_c^4} \\
&\quad \leq -\frac{2(1+\ell h_{b,c})(1-\ell h_{b,c})}{h_{b,c}H_c^4} < 0.
\end{aligned}$$

The weak inequality uses $H_c \geq 1+\ell h_{b,c}$. The strict inequality follows from $h_{b,c} = (\ell+\gamma v_{b,c})^{-1} < 1/\ell$. Thus a larger H_c, which lowers the price at given customer demand, also lowers bank b's marginal trading utility when its hiring is fixed.

Starting from equilibrium 1, increase H_c from $H_c^{(1)}$ to $H_c^{(2)}$ at fixed $n_{b,c}^{(1)}$, then increase $n_{b,c}$ from $n_{b,c}^{(1)}$ to $n_{b,c}^{(2)}$ at fixed $H_c^{(2)}$. The first comparison weakly lowers marginal trading utility and the second lowers it strictly. Both comparisons stay in the required domain because $H_c^{(2)} \geq 1+\ell h_{b,c}^{(2)}$ and $h_{b,c}$ rises with hiring.

Bank b's marginal trading utility is therefore strictly lower in equilibrium 2 than in equilibrium 1. Since $n_{b,c}^{(2)} > 0$, marginal trading utility equals w in equilibrium 2 and must therefore exceed w in equilibrium 1. This contradicts the hiring optimality condition in equilibrium 1, which requires marginal trading utility to be at most w. Hence two distinct hiring equilibria cannot exist in cluster c.

Hiring objectives separate across clusters, so uniqueness in every cluster gives a unique hiring equilibrium. Together with the unique trading equilibrium after every hiring profile and demand realization, this gives a unique subgame perfect equilibrium.

Hiring benchmark at $\ell = 0$. For the hiring comparisons below, we use the benchmark $\ell = 0$, where the intermediation price $p_f = ad_f$ does not depend on banks' hiring. We define

$$A_0 \equiv 1 + \kappa^2\tau_0 = \frac{\kappa^2 + 1/\tau_0}{1/\tau_0},$$

which measures total payoff variance relative to unexplained variance at the hypothetical benchmark of zero expertise, and is common across banks and clusters and remains fixed as expertise changes.

We define bank b's expertise multiplier from hiring in cluster c as $t_{b,c} = \sqrt{1 + n_{b,c}}$. Then

$$\begin{aligned} F_{b,c}(n_{b,c}) &= \frac{M_c}{2\gamma} \frac{\tau_0 + O_{b,c} t_{b,c}}{1 + \kappa^2(\tau_0 + O_{b,c} t_{b,c})}, \\ F'_{b,c}(n_{b,c}) &= \frac{M_c O_{b,c}}{4\gamma t_{b,c}(A_0 + \kappa^2 O_{b,c} t_{b,c})^2}, \\ F''_{b,c}(n_{b,c}) &= -\frac{M_c O_{b,c}(A_0 + 3\kappa^2 O_{b,c} t_{b,c})}{8\gamma t_{b,c}^3 (A_0 + \kappa^2 O_{b,c} t_{b,c})^3} < 0. \end{aligned} \tag{17}$$

At zero hiring, marginal trading utility is $F'_{b,c}(0) = M_c O_{b,c}/[4\gamma(A_0 + \kappa^2 O_{b,c})^2]$. The bank hires zero if this is at or below w. Otherwise, its unique optimum satisfies Equation (6). We substitute $t_{b,c} = E_{b,c}/O_{b,c}$ to obtain

$$E_{b,c}(A_0 + \kappa^2 E_{b,c})^2 = \frac{M_c O_{b,c}^2}{4\gamma w}, \qquad n_{b,c} = \frac{E_{b,c}^2}{O_{b,c}^2} - 1 > 0. \tag{18}$$

The left hand side of Equation (18) is strictly increasing in expertise $E_{b,c}$. Together with the condition for zero hiring, this determines bank b's unique hiring choice in cluster c at $\ell = 0$.

Hiring comparisons for small positive ℓ. At $\ell = 0$, fix the set of bank and cluster pairs with positive equilibrium hiring, and hold the remaining hiring choices at zero. For each pair with positive hiring, the first order condition is $F'_{b,c}(n_{b,c}) = w$. Banks' hiring problems are independent at $\ell = 0$, so the Jacobian of these conditions with respect to the positive hiring choices is diagonal. Its diagonal entries are the strictly negative second derivatives in Equation (17), so it is invertible.

All denominators in the trading utility formulas are strictly positive at $\ell = 0$, so the formulas extend smoothly to an open neighborhood in hiring, ℓ, anticipated demand M_c, and organizational capital. The implicit function theorem therefore gives smooth hiring choices for the bank and cluster pairs that hire at $\ell = 0$, for sufficiently small changes in these parameters. Their hiring remains positive by continuity.

For each bank and cluster with zero hiring at $\ell = 0$, the threshold condition at the beginning of this section makes marginal trading utility at zero strictly less than w. By continuity, this inequality remains strict at the continued hiring choices. Thus zero remains optimal for these bank and cluster pairs. If no bank hires at $\ell = 0$, the same strict inequalities keep every bank's hiring at zero for sufficiently small parameter changes.

The continued hiring choices satisfy every bank's hiring optimality conditions. Strict concavity makes these conditions sufficient, and uniqueness established above identifies these choices as the unique hiring equilibrium. Equilibrium hiring and its derivatives with respect to M_c and organizational capital vary continuously with ℓ. Any fixed strict comparison of hiring or its demand response established at $\ell = 0$ therefore persists for sufficiently small positive ℓ. The same argument applies locally in M_c and organizational capital while the inequalities at zero hiring remain strict.

IA.B.3 Organizational Capital, Hiring, and Residual Risk

Before presenting the derivations of Predictions 1, 2, and 3, we present Lemma 1, which shows that, when anticipated customer demand is equal, greater organizational capital gives greater expertise and lower unexplained and total residual risk. The comparison applies both across a bank's clusters and across banks in the same cluster. When expertise is low enough that organizational capital raises the marginal value of hiring, greater organizational capital also gives weakly greater hiring.

The result holds at zero price impact. It extends to sufficiently small positive price impact when, at $\ell = 0$, no bank is exactly at the threshold for starting to hire. We then use this result in the proofs of Predictions 1, 2, and 3 to compare expected trading activity, risk-adjusted returns, and the expected unattributed share of P&L, respectively, and in the proof of Prediction 5 and the competitor-constraint extension in Section IA.B.10 to compare residual risk.

We compare bank b in cluster c with bank j in cluster d under the following conditions for Lemma 1.

1. Anticipated customer demand is equal, $M_c = M_d$, and organizational capital is greater for bank b in cluster c, $O_{b,c} > O_{j,d} > 0$. The comparison covers two clusters at the same bank ($j = b$, $d \neq c$) and two banks in the same cluster ($d = c$, $j \neq b$).

2. For the hiring comparison, bank b's expertise in cluster c is low enough that the complementarity between organizational capital and hiring outweighs the diminishing benefit of reducing residual risk. Specifically, at $\ell = 0$,

$$\kappa^2 E_{b,c} < A_0 = 1 + \kappa^2 \tau_0.$$

3. For comparisons with sufficiently small positive price impact, $\ell > 0$, we impose the strict threshold condition in Section IA.B.2: at $\ell = 0$, every bank's marginal value of the first hire in each cluster differs from the wage. This excludes a bank being exactly at the threshold for starting to hire and keeps the set of positive hiring choices unchanged under small changes in ℓ.

Lemma 1 (Organizational capital, hiring, and residual risk)**.** *Under condition 1, at $\ell = 0$, expertise is greater for bank b in cluster c, and unexplained and total residual variances are lower:*

$$E_{b,c} > E_{j,d}, \qquad \sigma^2_{u,b,c} < \sigma^2_{u,j,d}, \qquad v_{b,c} < v_{j,d}.$$

Under condition 2 as well, hiring is weakly greater, $n_{b,c} \geq n_{j,d}$, with strict inequality if $n_{b,c} > 0$. With condition 3, these rankings also hold for sufficiently small positive ℓ.

Proof. At $\ell = 0$, bank b's hiring problem in cluster c depends on the bank and cluster only through organizational capital $O_{b,c}$ and anticipated demand M_c. We denote the unique positive solution of

the hiring condition in Equation (18) as $E^*_{b,c}$:

$$E^*_{b,c}(A_0 + \kappa^2 E^*_{b,c})^2 = \frac{M_c O^2_{b,c}}{4\gamma w}.$$

The left hand side increases strictly and without bound in $E^*_{b,c}$ from zero, and the right hand side increases strictly in $O_{b,c}$. Since $M_c = M_d$ and $O_{b,c} > O_{j,d}$, this gives $E^*_{b,c} > E^*_{j,d}$. Marginal trading utility at zero hiring, $F'_{b,c}(0) = M_c O_{b,c}/[4\gamma(A_0 + \kappa^2 O_{b,c})^2]$, exceeds w if and only if $O_{b,c}(A_0 + \kappa^2 O_{b,c})^2 < M_c O^2_{b,c}/(4\gamma w)$, or equivalently $E^*_{b,c} > O_{b,c}$. By Equation (18) and the condition for zero hiring, expertise is therefore

$$E_{b,c} = \max\{O_{b,c},\, E^*_{b,c}\}.$$

Both terms are greater for bank b in cluster c, so $E_{b,c} > E_{j,d}$. The formulas $\sigma^2_{u,b,c} = 1/(\tau_0 + E_{b,c})$ and $v_{b,c} = \kappa^2 + \sigma^2_{u,b,c}$ then give both variance rankings.

Hiring. By Equation (18) and the condition for zero hiring, bank b's hiring in cluster c is

$$n_{b,c} = \max\left\{0,\, \left(\frac{E^*_{b,c}}{O_{b,c}}\right)^2 - 1\right\}.$$

We vary bank b's organizational capital in cluster c between the two capital levels, holding M_c fixed. We denote a capital level in this range as $\widetilde{O}_{b,c} \in [O_{j,d}, O_{b,c}]$ and the corresponding solution of the hiring condition as $\widetilde{E}^*_{b,c}$. At the lower endpoint, $M_c = M_d$ gives $\widetilde{E}^*_{b,c} = E^*_{j,d}$. We differentiate the hiring condition with respect to organizational capital to obtain

$$\frac{\partial \log \widetilde{E}^*_{b,c}}{\partial \log \widetilde{O}_{b,c}} = \frac{2(A_0 + \kappa^2 \widetilde{E}^*_{b,c})}{A_0 + 3\kappa^2 \widetilde{E}^*_{b,c}} > 0,$$
$$\frac{\partial \log(\widetilde{E}^*_{b,c}/\widetilde{O}_{b,c})}{\partial \log \widetilde{O}_{b,c}} = \frac{A_0 - \kappa^2 \widetilde{E}^*_{b,c}}{A_0 + 3\kappa^2 \widetilde{E}^*_{b,c}} > 0.$$

The first inequality gives $\widetilde{E}^*_{b,c} \le E^*_{b,c} \le E_{b,c}$ at every capital level in this range, so condition 2 gives $\kappa^2 \widetilde{E}^*_{b,c} < A_0$ and hence the second inequality. Integrating from $O_{j,d}$ to $O_{b,c}$ gives $E^*_{b,c}/O_{b,c} > E^*_{j,d}/O_{j,d}$. Hence $n_{b,c} \ge n_{j,d}$, with strict inequality if $n_{b,c} > 0$.

At zero hiring, marginal trading utility at capital level $\widetilde{O}_{b,c}$ is $M_c \widetilde{O}_{b,c}/[4\gamma(A_0 + \kappa^2 \widetilde{O}_{b,c})^2]$. Its derivative with respect to $\widetilde{O}_{b,c}$ is

$$\frac{\partial}{\partial \widetilde{O}_{b,c}}\left[\frac{M_c \widetilde{O}_{b,c}}{4\gamma(A_0 + \kappa^2 \widetilde{O}_{b,c})^2}\right] = \frac{M_c(A_0 - \kappa^2 \widetilde{O}_{b,c})}{4\gamma(A_0 + \kappa^2 \widetilde{O}_{b,c})^3} > 0,$$

where the inequality follows from $\kappa^2 \widetilde{O}_{b,c} \le \kappa^2 O_{b,c} \le \kappa^2 E_{b,c} < A_0$. Thus greater organizational

capital also raises the marginal value of the first hire.

Small positive ℓ. The continuation argument in Section IA.B.2 preserves the strict hiring, expertise, and variance rankings for sufficiently small positive ℓ. If neither bank and cluster pair hires at $\ell = 0$, the strict threshold condition keeps both hiring choices at zero for sufficiently small positive ℓ, so the weak hiring ranking is also preserved. □

IA.B.4 Derivation of Prediction 1

We impose conditions 1–3 in Section IA.B.3 for clusters c and d at bank b, taking $\ell > 0$ sufficiently small. For a strict hiring comparison, we assume positive hiring in cluster c at $\ell = 0$.

For trading activity, we compare a core factor f in cluster c with a peripheral factor g in cluster d, with equal positive realized demand $d_f = d_g > 0$. We evaluate positions after banks choose their equilibrium hiring.

For comparisons of cluster totals, the clusters contain the same number of factors, paired to have equal positive realized demand. Trading activity is absolute attributed P&L, $|z_{b,f}\kappa R_f|$. As in the model, positions are chosen before the standard normal payoff shock R_f arrives.

Hiring. Lemma 1 gives $n_{b,c} > n_{b,d}$ and $E_{b,c} > E_{b,d}$ for sufficiently small positive ℓ.

In particular, at $\ell = 0$, hiring occurs only in cluster c when

$$\frac{M_d O_{b,d}}{4\gamma(A_0 + \kappa^2 O_{b,d})^2} < w < \frac{M_c O_{b,c}}{4\gamma(A_0 + \kappa^2 O_{b,c})^2}.$$

The interval is nonempty because anticipated demand is equal and the marginal value of the first hire increases with organizational capital, as shown in the proof of Lemma 1. The continuation argument in Section IA.B.2 keeps hiring zero in cluster d and positive in cluster c for sufficiently small positive ℓ.

Trading activity. Within a factor, position shares are $z_{b,f}/\sum_j z_{j,f} = h_{b,c}/\sum_j h_{j,c}$, which rise with the bank's expertise holding rivals' expertise fixed. Expected absolute attributed P&L is

$$\mathbb{E}[|z_{b,f}\kappa R_f|] = \sqrt{\frac{2}{\pi}}\,\kappa z_{b,f}. \tag{19}$$

Thus expected trading activity is a common positive multiple of the position. At $\ell = 0$, Lemma 1 gives $v_{b,c} < v_{b,d}$. With equal positive demand, Equation (5) therefore gives

$$z_{b,f} = \frac{a d_f}{\gamma v_{b,c}} > \frac{a d_g}{\gamma v_{b,d}} = z_{b,g}.$$

Section IA.B.2 establishes that equilibrium hiring is continuous at $\ell = 0$. Residual variances and equilibrium positions are therefore continuous as well, so the strict position ranking persists for

sufficiently small positive ℓ. Bank b then has greater expected trading activity in its core factor. We divide by the same bank total to obtain the same ranking for shares of expected activity.

Summing the inequalities for the paired factors gives the same ranking for cluster activity.

IA.B.5 Derivation of Prediction 2

We impose conditions 1 and 3 in Section IA.B.3, taking $\ell > 0$ sufficiently small. Within bank b, these conditions apply to clusters c and d. Across banks, they apply to two banks in the same cluster. For the return comparison within bank b, we compare a core factor f in cluster c and a peripheral factor g in cluster d, with equal positive realized demand $d_f = d_g > 0$.

We compare desks with the same number of independent factors. Within each desk, factors have the same realized demand, positions, and residual variance.

For expected P&L divided by lagged reported VaR, reported VaR is a common multiple of P&L standard deviation for the same horizon. We hold demand, positions, and variances fixed from the lagged risk measurement through the return interval.

Lemma 1 gives greater expertise and lower total residual variance in the core, $E_{b,c} > E_{b,d}$ and $v_{b,c} < v_{b,d}$. For a given realization of customer demand and fixed hiring, bank b's position $z_{b,f}$ and intermediation price p_f are fixed. Its residual payoff $\varepsilon_{b,f}$ has mean zero and variance $v_{b,c}$, so

$$\mathbb{E}[\pi_{b,f}] = z_{b,f} p_f, \qquad \mathrm{Var}(\pi_{b,f}) = z_{b,f}^2 v_{b,c}.$$

We measure the bank's risk-adjusted trading return before wage costs as

$$\frac{\mathbb{E}[\pi_{b,f}]}{\sqrt{\mathrm{Var}(\pi_{b,f})}} = \frac{p_f}{\sqrt{v_{b,c}}}. \tag{20}$$

Across banks in the same factor, the price is common and greater expertise gives lower variance and higher risk-adjusted return. By Lemma 1, specialist banks have greater expertise in the factor's cluster than non-specialist banks.

Within bank b, at $\ell = 0$ the intermediation prices are equal, $p_f = a d_f = a d_g = p_g$. Lower core variance therefore gives

$$\frac{p_f}{\sqrt{v_{b,c}}} > \frac{p_g}{\sqrt{v_{b,d}}}.$$

By Section IA.B.2, equilibrium hiring and hence residual variances and intermediation prices are continuous at $\ell = 0$. The strict return ranking therefore persists for sufficiently small positive ℓ.

Desk returns. Expected P&L scales with the factor count and its standard deviation with the square root of that count. The desk's return ratio therefore equals the ratio for one factor times the square root of the count, preserving the core and peripheral ranking for equal counts. Dividing expected P&L by lagged reported VaR rescales both ratios by the same constant, so the ranking is

unchanged.

IA.B.6 Derivation of Prediction 3

We impose conditions 1 and 3 in Section IA.B.3, taking $\ell > 0$ sufficiently small. Within bank b, these conditions apply to clusters c and d. Across banks, they apply to two banks in the same cluster. We compare positive positions in a core factor f in cluster c and a peripheral factor g in cluster d. Positions and payoff loadings are fixed before returns arrive.

Lemma 1 gives $E_{b,c} > E_{b,d}$ and $\sigma_{u,b,c} < \sigma_{u,b,d}$.

Unattributed share. Bank b's attributed and unexplained P&L from factor f are

$$A_{b,f} = z_{b,f}\kappa R_f, \qquad U_{b,f} = z_{b,f}u_{b,f}.$$

These components decompose bank b's P&L from factor payoffs, $z_{b,f}\varepsilon_{b,f}$; its revenue from intermediation, $z_{b,f}p_f$, is outside this share. Bank b observes the unexplained P&L $U_{b,f}$ afterward.

We define bank b's standardized unexplained payoff shocks as $\xi_{b,f} = u_{b,f}/\sigma_{u,b,c}$ for the core factor and $\xi_{b,g} = u_{b,g}/\sigma_{u,b,d}$ for the peripheral factor. Both are standard normal and independent of their respective common modeled shocks. For a positive position, bank b's unattributed share from factor f is

$$\mathcal{U}_{b,f} = \frac{\sigma_{u,b,c}|\xi_{b,f}|}{\kappa|R_f| + \sigma_{u,b,c}|\xi_{b,f}|}.$$

Position size cancels. We hold standardized shock realizations fixed and differentiate with respect to bank b's standard deviation of unexplained risk to obtain, almost surely,

$$\frac{\partial \mathcal{U}_{b,f}}{\partial \sigma_{u,b,c}} = \frac{|\xi_{b,f}|\kappa|R_f|}{(\kappa|R_f| + \sigma_{u,b,c}|\xi_{b,f}|)^2} > 0.$$

The pairs $(R_f, \xi_{b,f})$ and $(R_g, \xi_{b,g})$ have the same joint distribution. We evaluate the shares at common standardized realizations and take expectations to obtain $\mathbb{E}[\mathcal{U}_{b,f}] < \mathbb{E}[\mathcal{U}_{b,g}]$, since $\sigma_{u,b,c} < \sigma_{u,b,d}$. The shares are bounded by one, so expectations exist. Greater expertise lowers the expected unattributed share even when bank b takes a larger position.

Across banks, we compare positive positions of banks b and j in a common factor f in cluster c, with greater organizational capital at bank b. Lemma 1 gives $\sigma_{u,b,c} < \sigma_{u,j,c}$. With $\xi_{j,f} = u_{j,f}/\sigma_{u,j,c}$, the pairs $(R_f, \xi_{b,f})$ and $(R_f, \xi_{j,f})$ have the same joint distribution, and the same calculation gives $\mathbb{E}[\mathcal{U}_{b,f}] < \mathbb{E}[\mathcal{U}_{j,f}]$.

IA.B.7 Derivation of Prediction 4

For sufficiently small positive ℓ, we consider a local anticipated increase in customer demand in cluster c. The increase multiplies every d_f in this cluster by $(1+\delta)$ in every realization, with a held fixed, and is anticipated by all banks before hiring.

We maintain the strict threshold condition in Section IA.B.2: at $\ell = 0$, every bank's marginal value of hiring at zero differs from the wage in each cluster. We compare banks b and j in cluster c, with greater organizational capital at bank b, $O_{b,c} > O_{j,c} > 0$, and positive hiring there at $\ell = 0$.

We assume that bank b's expertise at $\ell = 0$ satisfies

$$\kappa^2 E_{b,c} < \frac{A_0}{\sqrt{3}}.$$

This restricts expertise to the range where greater expertise raises the sensitivity of hiring to anticipated demand.

Benchmark at $\ell = 0$. Expected trading utility depends on the demand distribution only through M_c. For bank b, which hires a positive amount in cluster c, we differentiate Equation (18) to obtain

$$\frac{\partial E_{b,c}}{\partial M_c} = \frac{O_{b,c}^2}{4\gamma w(A_0 + \kappa^2 E_{b,c})(A_0 + 3\kappa^2 E_{b,c})} > 0.$$

Since $n_{b,c} = E_{b,c}^2 / O_{b,c}^2 - 1$, bank b's hiring response is

$$\frac{\partial n_{b,c}}{\partial M_c} = \frac{E_{b,c}}{2\gamma w(A_0 + \kappa^2 E_{b,c})(A_0 + 3\kappa^2 E_{b,c})} > 0. \tag{21}$$

The same calculation applies to each bank with positive hiring in cluster c, so each has a positive response to M_c.

Hiring in other clusters is unchanged because hiring objectives separate across clusters.

Differences across banks. When both banks hire at $\ell = 0$, Equation (18) gives $E_{b,c} > E_{j,c}$ because its right hand side increases with organizational capital and its left hand side increases strictly with expertise. We denote an expertise level between the two banks' levels as $\widetilde{E}_{b,c} \in [E_{j,c}, E_{b,c}]$. Then

$$\begin{aligned}
&\frac{\partial}{\partial \widetilde{E}_{b,c}} \left[\frac{\widetilde{E}_{b,c}}{(A_0 + \kappa^2 \widetilde{E}_{b,c})(A_0 + 3\kappa^2 \widetilde{E}_{b,c})} \right] \\
&\qquad = \frac{A_0^2 - 3\kappa^4 \widetilde{E}_{b,c}^2}{(A_0 + \kappa^2 \widetilde{E}_{b,c})^2 (A_0 + 3\kappa^2 \widetilde{E}_{b,c})^2} > 0,
\end{aligned}$$

where the inequality follows from $\kappa^2 \widetilde{E}_{b,c} \leq \kappa^2 E_{b,c} < A_0/\sqrt{3}$. Equation (21) therefore gives $\partial n_{b,c}/\partial M_c > \partial n_{j,c}/\partial M_c$.

If bank j hires zero in cluster c with marginal trading utility at zero hiring strictly below w, its hiring there stays zero after a sufficiently small demand increase. Bank b already hires and has a positive response by Equation (21).

Positive price impact. Section IA.B.2 establishes that equilibrium hiring and its derivatives with respect to M_c vary continuously with ℓ under the strict threshold condition. Every bank that

hires at $\ell = 0$ has a strictly positive demand response there, and bank b's response is strictly greater than bank j's. Thus, for some $\bar{\ell} > 0$ depending on the parameter values being compared, these strict comparisons hold for every $0 < \ell < \bar{\ell}$. In particular,

$$\frac{\partial n_{b,c}}{\partial M_c} > \frac{\partial n_{j,c}}{\partial M_c} \geq 0.$$

These are derivatives of the joint hiring equilibrium: all banks adjust their hiring and anticipate the resulting trading equilibrium. Banks that hire zero at $\ell = 0$ continue to hire zero locally, by the strict threshold condition.

Proportional demand changes. Anticipated demand changes to $(1 + \delta)^2 M_c$, so bank b's equilibrium hiring derivative with respect to δ at $\delta = 0$ is $2M_c \, \partial n_{b,c}/\partial M_c$. The multiplier $2M_c$ is common across banks, so the same ranking applies to responses to the proportional demand increase. Integrating these continuous derivative inequalities over a sufficiently small demand increase gives the same ranking for hiring changes.

Clusters within a bank. The same comparison applies to a core cluster c and a peripheral cluster d at bank b, $O_{b,c} > O_{b,d}$, with $M_c = M_d$ and the same proportional increase in anticipated demand in each cluster. Hiring objectives separate across clusters, and Equation (21) depends on the cluster only through expertise, so the steps above apply with cluster d at bank b in place of cluster c at bank j.

IA.B.8 Derivation of Prediction 5

We consider bank b with a binding risk budget in Equation (8) and hold customer demand and expertise fixed. All competitors have slack budgets. We compare a core factor f in cluster c and a peripheral factor g in cluster d. We impose conditions 1 and 3 in Section IA.B.3 for clusters c and d at bank b, taking $\ell > 0$ sufficiently small, so Lemma 1 gives $v_{b,c} < v_{b,d}$. We assume equal positive realized demand $d_f = d_g$. We compare percentage changes in positions and expected trading activity while both positions remain positive.

Positions. We denote the Lagrange multiplier on bank b's risk budget as $\lambda_b \geq 0$. It measures bank b's marginal value of an additional unit of risk budget, holding the other banks' positions fixed. Bank b's first order condition for its factor f position is $p_f = (\ell + \gamma v_{b,c}) z_{b,f} + \lambda_b \sqrt{v_{b,c}}$ when $z_{b,f} > 0$. Each rival has a slack budget and satisfies $z_{j,f} = h_{j,c} p_f$. We combine these conditions with inverse demand to obtain

$$z_{b,f}(\lambda_b) = \frac{h_{b,c}}{1 + \ell \sum_j h_{j,c}} \left[a d_f - \lambda_b \sqrt{v_{b,c}} (1 + \ell \sum_{j \neq b} h_{j,c}) \right]_+ . \tag{22}$$

The positive part sets bank b's factor f position to zero when $\lambda_b\sqrt{v_{b,c}}(1+\ell\sum_{j\neq b}h_{j,c})\geq ad_f$. Each positive position decreases with λ_b. Budget usage is their sum weighted by the fixed positive risk charges $\sqrt{v_{b,c}}$, so a lower binding positive budget $\overline{R}_b$ requires a larger λ_b.

While bank b's factor f position remains positive,

$$\frac{\partial \log z_{b,f}}{\partial \lambda_b} = -\frac{1}{ad_f/[\sqrt{v_{b,c}}(1+\ell\sum_{j\neq b}h_{j,c})]-\lambda_b}.$$

At $\ell = 0$, the ratio $ad_f/[\sqrt{v_{b,c}}(1+\ell\sum_{j\neq b}h_{j,c})]$ equals $ad_f/\sqrt{v_{b,c}}$, so equal demand and lower core variance make it strictly greater in the core. With $d_f = d_g$, this ranking is equivalent to $\sqrt{v_{b,c}}(1+\ell\sum_{j\neq b}h_{j,c}) < \sqrt{v_{b,d}}(1+\ell\sum_{j\neq b}h_{j,d})$, which does not depend on λ_b or the level of demand. By Section IA.B.2, equilibrium hiring and hence residual variances and the coefficients $h_{j,c}$ and $h_{j,d}$ are continuous at $\ell = 0$, so the inequality holds for sufficiently small positive ℓ. At the same λ_b, the derivative of the log position is therefore less negative in the core. The peripheral position also reaches zero first, since each position reaches zero when λ_b equals this ratio. More generally, both comparisons hold whenever

$$\frac{ad_f}{\sqrt{v_{b,c}}(1+\ell\sum_{j\neq b}h_{j,c})} > \frac{ad_g}{\sqrt{v_{b,d}}(1+\ell\sum_{j\neq b}h_{j,d})}. \tag{23}$$

Trading activity. We integrate the derivatives of log positions over the change in λ_b to obtain the same ranking of finite percentage contractions while both positions remain positive. With hiring fixed, Equation (19) makes bank b's expected trading activity in each factor a fixed multiple of its position. Its percentage contraction therefore equals the percentage contraction in that position and is larger in the peripheral factor.

IA.B.9 Activity Responses with Fixed Hiring

We use Lemma 2 in the competitor-constraint extension in Section IA.B.10 to compare the other banks' responses when rival bank r reduces its position in factor f in cluster c and raises the intermediation price. The lemma shows that greater expertise in cluster c gives larger increases in positions and expected trading activity.

Lemma 2 (Activity responses to a common price increase)**.** *We hold all banks' hiring fixed and consider banks with slack risk budgets trading factor f in cluster c. An intermediation price increase from $p_f > 0$ to $p'_f > p_f$ gives banks with greater expertise larger increases in positions and expected trading activity.*

Proof. With hiring fixed, bank b's variance $v_{b,c}$ and hence $h_{b,c} = (\ell+\gamma v_{b,c})^{-1}$ are unchanged. Its factor f position therefore changes by

$$z'_{b,f} - z_{b,f} = h_{b,c}(p'_f - p_f).$$

Greater expertise lowers $v_{b,c}$ and raises $h_{b,c}$, giving a larger position increase.

By Equation (19), bank b's expected trading activity increases by

$$\mathbb{E}[|z'_{b,f}\kappa R_f|] - \mathbb{E}[|z_{b,f}\kappa R_f|] = \sqrt{\frac{2}{\pi}}\,\kappa h_{b,c}(p'_f - p_f).$$

The coefficient $\sqrt{2/\pi}\,\kappa$ is common across banks, so expected activity has the same ranking of increases as positions. □

The lemma also applies to an increase in realized demand with hiring fixed. Equation (5) gives

$$p_f = \frac{ad_f}{1+\ell\sum_j h_{j,c}}, \qquad z_{b,f}(d'_f) - z_{b,f}(d_f) = \frac{ah_{b,c}(d'_f - d_f)}{1+\ell\sum_j h_{j,c}}.$$

For $d'_f > d_f > 0$, each bank's position share remains $h_{b,c}/\sum_j h_{j,c}$.

IA.B.10 Competitor Constraints and Specialist Expansion

We impose a binding risk budget from Equation (8) on rival bank r. We hold customer demand $ad_f > 0$ and expertise fixed, and compare the other banks trading factor f. These banks have slack budgets. Under conditions 1 and 3 in Section IA.B.3 for two banks in cluster c, Lemma 1 gives specialist banks greater expertise in cluster c than non-specialist banks for sufficiently small positive ℓ.

We first consider a local tightening while bank r remains active in factor f, then its exit. We assume $\ell > 0$, since at $\ell = 0$ its withdrawal leaves the price and other banks' positions unchanged.

Tightening rival bank r's budget. We denote the Lagrange multiplier on rival bank r's risk budget as $\lambda_r \geq 0$. It measures bank r's marginal value of an additional unit of risk budget, holding the other banks' positions fixed. Rival bank r's first order condition for a positive position in factor f is

$$p_f = (\ell + \gamma v_{r,c})z_{r,f} + \lambda_r\sqrt{v_{r,c}}.$$

Each responding bank $j \neq r$ has a slack budget, so its first order condition for its factor f position gives $z_{j,f} = h_{j,c}p_f$. We combine these conditions with bank r's first order condition and inverse demand to obtain

$$p_f(\lambda_r) = \frac{ad_f + \ell h_{r,c}\lambda_r\sqrt{v_{r,c}}}{1+\ell\sum_j h_{j,c}}, \qquad z_{j,f} = h_{j,c}p_f(\lambda_r) \quad (j \neq r). \tag{24}$$

By Equation (22), each positive position at bank r falls with λ_r, so a tighter binding positive budget raises λ_r. The price in Equation (24) then rises because the coefficient on λ_r is positive for $\ell > 0$. Each responding bank's position changes by

$$\frac{\partial z_{j,f}}{\partial \lambda_r} = \frac{\ell h_{j,c}h_{r,c}\sqrt{v_{r,c}}}{1+\ell\sum_b h_{b,c}} > 0.$$

All terms except $h_{j,c}$ are common across responding banks. Greater expertise raises $h_{j,c}$ and hence the position increase.

To compare market shares, we define total factor f exposure as $S_f = \sum_b z_{b,f} = (ad_f - p_f)/\ell$. Bank j's share satisfies

$$\frac{\partial}{\partial \lambda_r}\left(\frac{z_{j,f}}{S_f}\right) = h_{j,c}\frac{\ell a d_f}{(ad_f - p_f)^2}\frac{\partial p_f}{\partial \lambda_r} > 0.$$

The share derivative is positive and proportional to $h_{j,c}$, so greater expertise gives a larger market share gain.

Rival bank r's exit. We next consider a previously unconstrained bank r exiting factor f and hold demand and surviving banks' hiring fixed. We remove $h_{r,c}$ from Equation (5). The price rises from $ad_f/(1 + \ell \sum_b h_{b,c})$ to $ad_f/(1 + \ell \sum_{b \neq r} h_{b,c})$. Bank j's position increases by $h_{j,c}$ times this common price increase, and its market share gain is

$$\frac{h_{j,c}h_{r,c}}{(\sum_b h_{b,c})(\sum_{b \neq r} h_{b,c})} > 0. \tag{25}$$

The denominators are common across surviving banks, so specialist banks gain more market share. Tightening bank r's budget until its factor f position reaches zero gives the same factor f equilibrium and share gains relative to the equilibrium before tightening.

Trading activity. Lemma 2 shows that the common price increase gives banks with greater expertise larger increases in expected trading activity. Market shares here refer to positions.

IA.B.11 Derivation of Prediction 6

We compare banks b and j in cluster c, with greater organizational capital at bank b, $O_{b,c} > O_{j,c} > 0$. We hold all banks' hiring fixed, with the same number of specialists at both banks, $n_{b,c} = n_{j,c} = n$, so both banks have the same expertise multiplier, $t_{b,c} = t_{j,c} = \sqrt{1+n}$, which we denote as t. We compare the banks' marginal trading utilities at this common staffing level, $F'_{b,c}(n) - F'_{j,c}(n)$.

As in condition 2, we assume that bank b's expertise with n specialists satisfies

$$\kappa^2 E_{b,c} = \kappa^2 O_{b,c} t < A_0.$$

This restricts expertise to the range where organizational capital raises the marginal value of hiring.

Benchmark at $\ell = 0$. By Equation (17), bank b's marginal trading utility is

$$F'_{b,c}(n) = \frac{M_c O_{b,c}}{4\gamma t(A_0 + \kappa^2 O_{b,c} t)^2}.$$

We denote a capital level between the two banks' levels as $\widetilde{O}_{b,c} \in [O_{j,c}, O_{b,c}]$, holding t fixed. Then

$$\frac{\partial}{\partial \widetilde{O}_{b,c}} \left[\frac{M_c \widetilde{O}_{b,c}}{4\gamma t (A_0 + \kappa^2 \widetilde{O}_{b,c} t)^2} \right] = \frac{M_c (A_0 - \kappa^2 \widetilde{O}_{b,c} t)}{4\gamma t (A_0 + \kappa^2 \widetilde{O}_{b,c} t)^3} > 0,$$

where the inequality follows from $\kappa^2 \widetilde{O}_{b,c} t \leq \kappa^2 O_{b,c} t < A_0$. Integrating from $O_{j,c}$ to $O_{b,c}$ gives $F'_{b,c}(n) > F'_{j,c}(n)$. This extends the comparison of the marginal value of the first hire in the proof of Lemma 1 to any common number of specialists that satisfies this bound.

Anticipated demand. With hiring fixed, Equation (15) makes each bank's expected trading utility, and hence its marginal trading utility, proportional to M_c. At $\ell = 0$, the difference in marginal trading utilities is

$$F'_{b,c}(n) - F'_{j,c}(n) = \frac{M_c}{4\gamma t} \left[\frac{O_{b,c}}{(A_0 + \kappa^2 O_{b,c} t)^2} - \frac{O_{j,c}}{(A_0 + \kappa^2 O_{j,c} t)^2} \right],$$

where the bracketed term is positive and does not depend on M_c. An increase in anticipated customer demand in cluster c therefore raises the difference in marginal trading utilities. An increase that multiplies every d_f in cluster c by $(1 + \delta)$ in every realization changes anticipated demand to $(1 + \delta)^2 M_c$ and multiplies the difference in marginal trading utilities by $(1 + \delta)^2$.

Positive price impact. With hiring fixed, all denominators in Equation (15) are strictly positive, so each bank's marginal trading utility varies continuously with ℓ. Thus, for some $\bar{\ell} > 0$ depending on the parameter values being compared, $F'_{b,c}(n) > F'_{j,c}(n)$ for every $0 < \ell < \bar{\ell}$. Marginal trading utility remains proportional to M_c, so an increase in anticipated customer demand also raises the difference in marginal trading utilities for these values of ℓ.

Core and peripheral clusters. Suppose cluster c is core at bank b and peripheral at bank j relative to another cluster d, $O_{b,c} > O_{b,d}$ and $O_{j,c} < O_{j,d}$, and bank b has at least as much organizational capital in cluster d as bank j, $O_{b,d} \geq O_{j,d}$. Then $O_{b,c} > O_{b,d} \geq O_{j,d} > O_{j,c}$, so the comparisons above apply to comparing marginal trading utilities at the common staffing level.

IA.C Figures and Tables



Figure IA.1: Clusters by Absolute Correlation Threshold. Figure plots cluster counts (solid lines) and singleton-cluster counts (dashed lines) as the absolute-correlation threshold varies. Singleton clusters contain one risk factor. The panels show the aggregate and nine asset classes. Correlations use the risk-factor return proxy based on trailing 90-day P&L volatility. Clustering uses the post-2020 sample. Counts are reported at each correlation threshold for the aggregate panel and for each asset class.

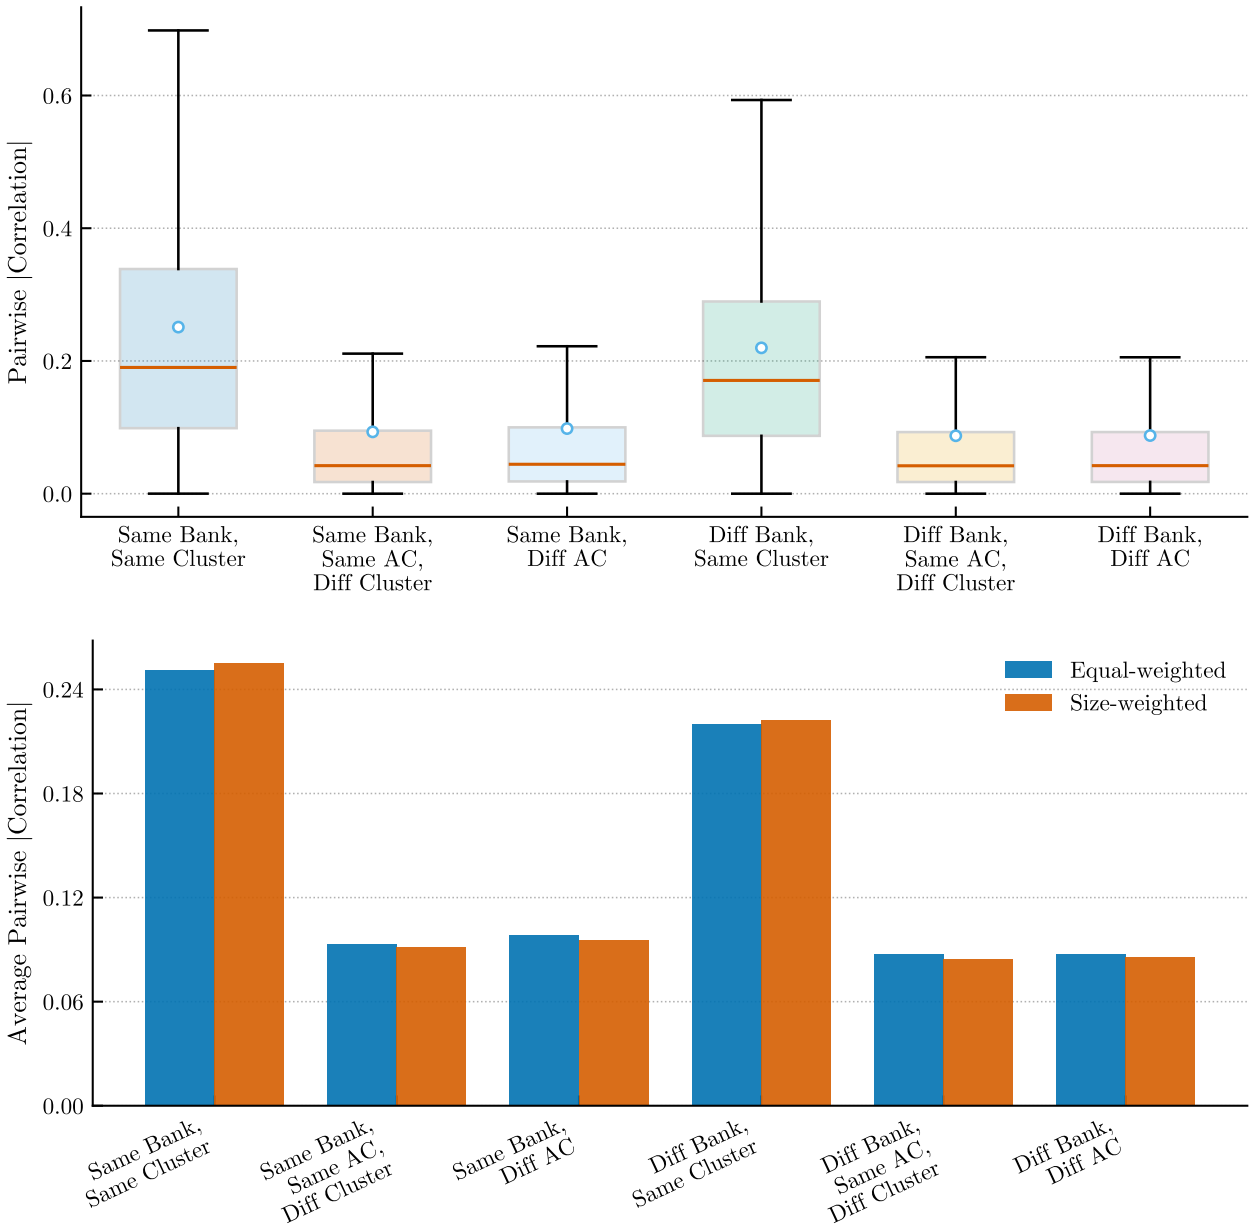


Figure IA.2: Absolute Risk-Factor Correlations. Figure plots absolute risk-factor correlations by shared bank, cluster, and asset-class membership. The top panel shows box plots; the bottom panel shows equal-weighted and size-weighted means. Cluster assignments follow the benchmark based on absolute correlations of risk-factor return proxies scaled by trailing 90-day P&L volatility. Observations are risk-factor pairs grouped by shared bank, cluster, and asset-class membership.

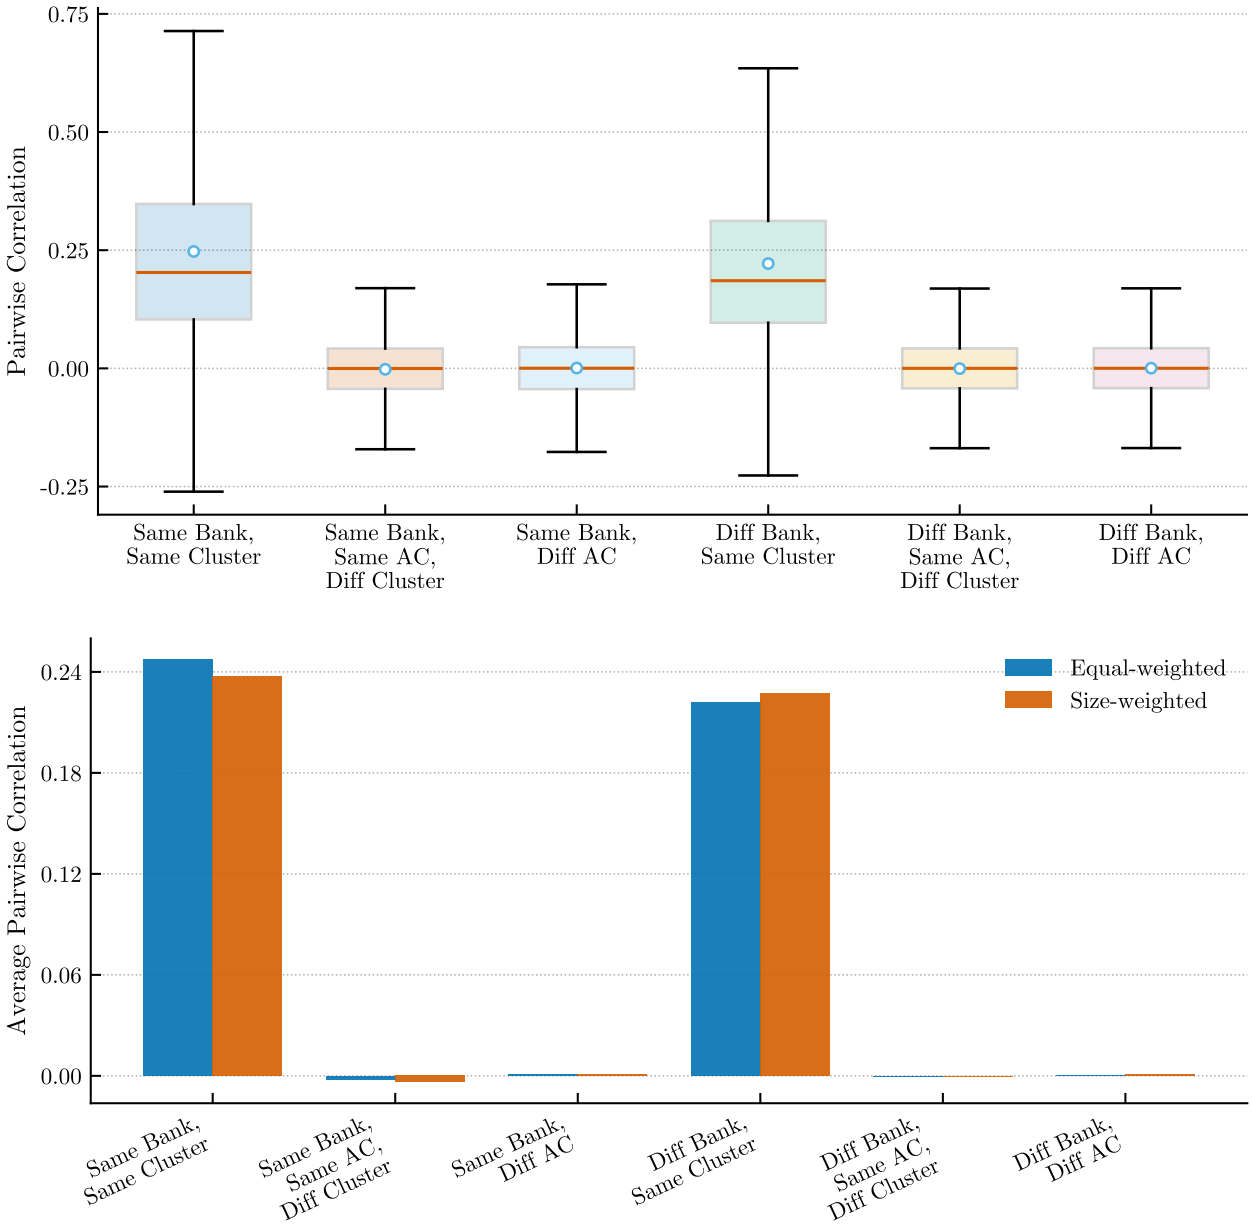


Figure IA.3: Signed Risk-Factor Correlations. Figure plots signed risk-factor correlations by shared bank, cluster, and asset-class membership. The top panel shows box plots; the bottom panel shows equal-weighted and size-weighted means. Correlations use risk-factor return proxies scaled by trailing 90-day P&L volatility. Observations are risk-factor pairs grouped by shared bank, cluster, and asset-class membership.

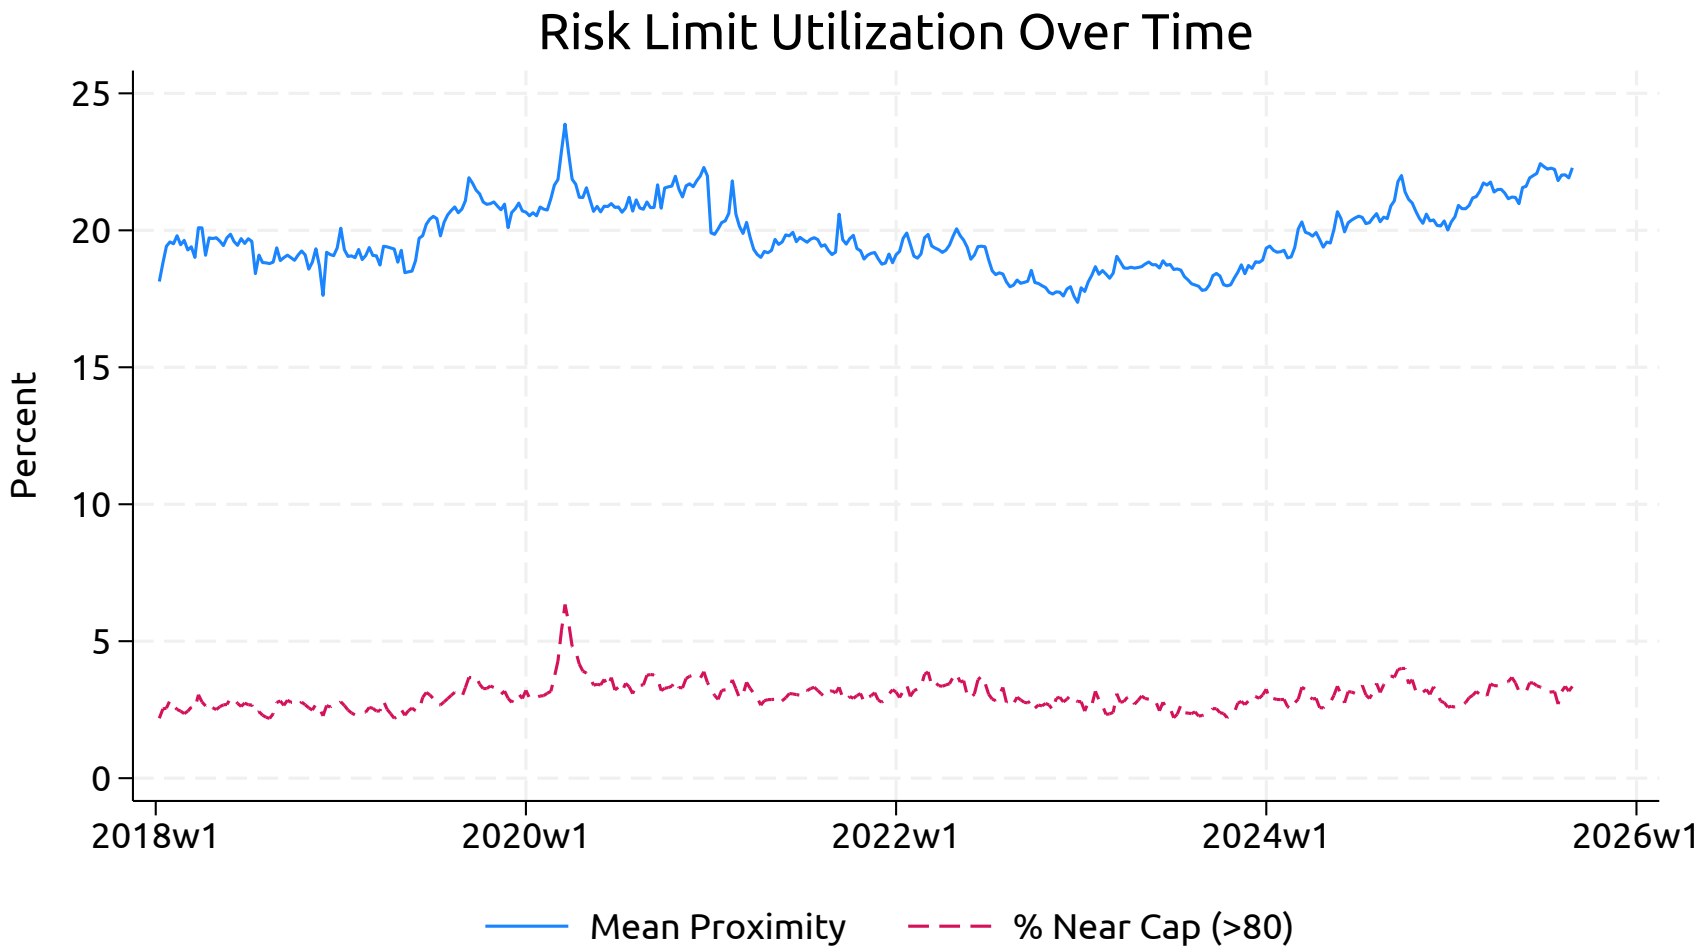


Figure IA.4: Risk Limit Utilization Over Time. Figure plots weekly averages of risk limit proximity (solid) and the weekly share of limit-date observations near capacity, with proximity above 80 (dashed). Proximity measures how close current usage is to the allowable bound, on a scale where 0 indicates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

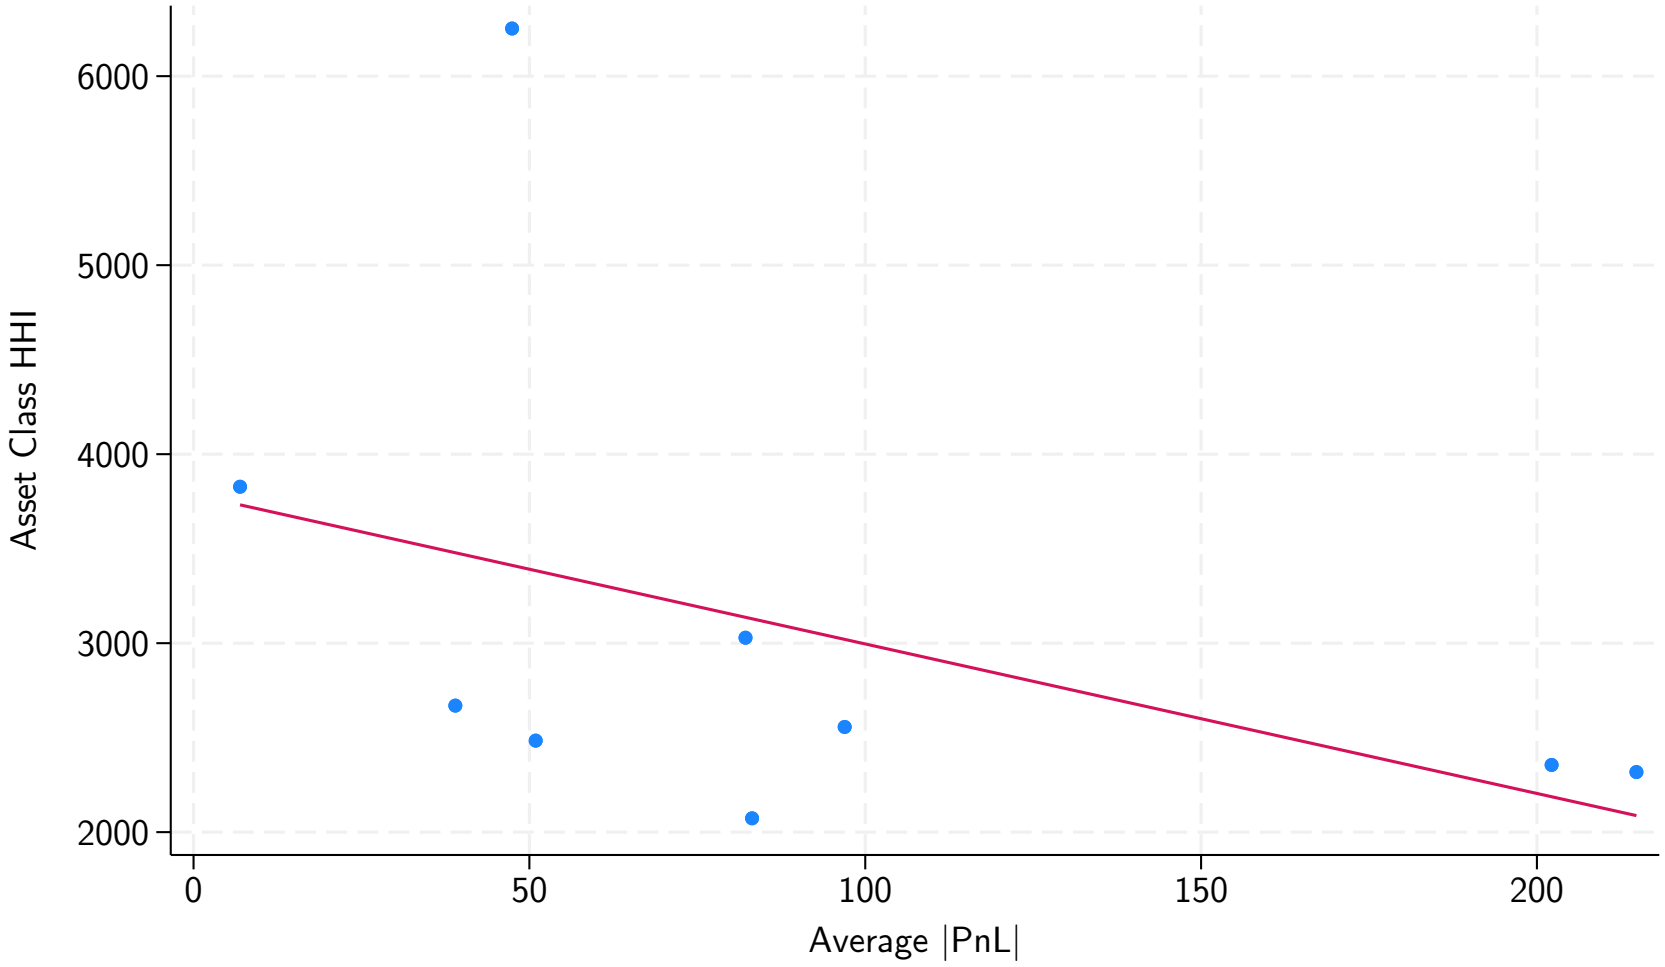


Figure IA.5: Asset-Class Concentration and Activity. Figure plots average asset-class HHI against average daily absolute comprehensive desk P&L, in millions of dollars, with a fitted line. HHI is computed from banks' shares of total desk VaR within the asset class, on a 0–10,000 scale. Absolute P&L is summed across desks before averaging over days. The figure uses the full sample. Each of the nine points is an asset class, with both variables averaged over business days.

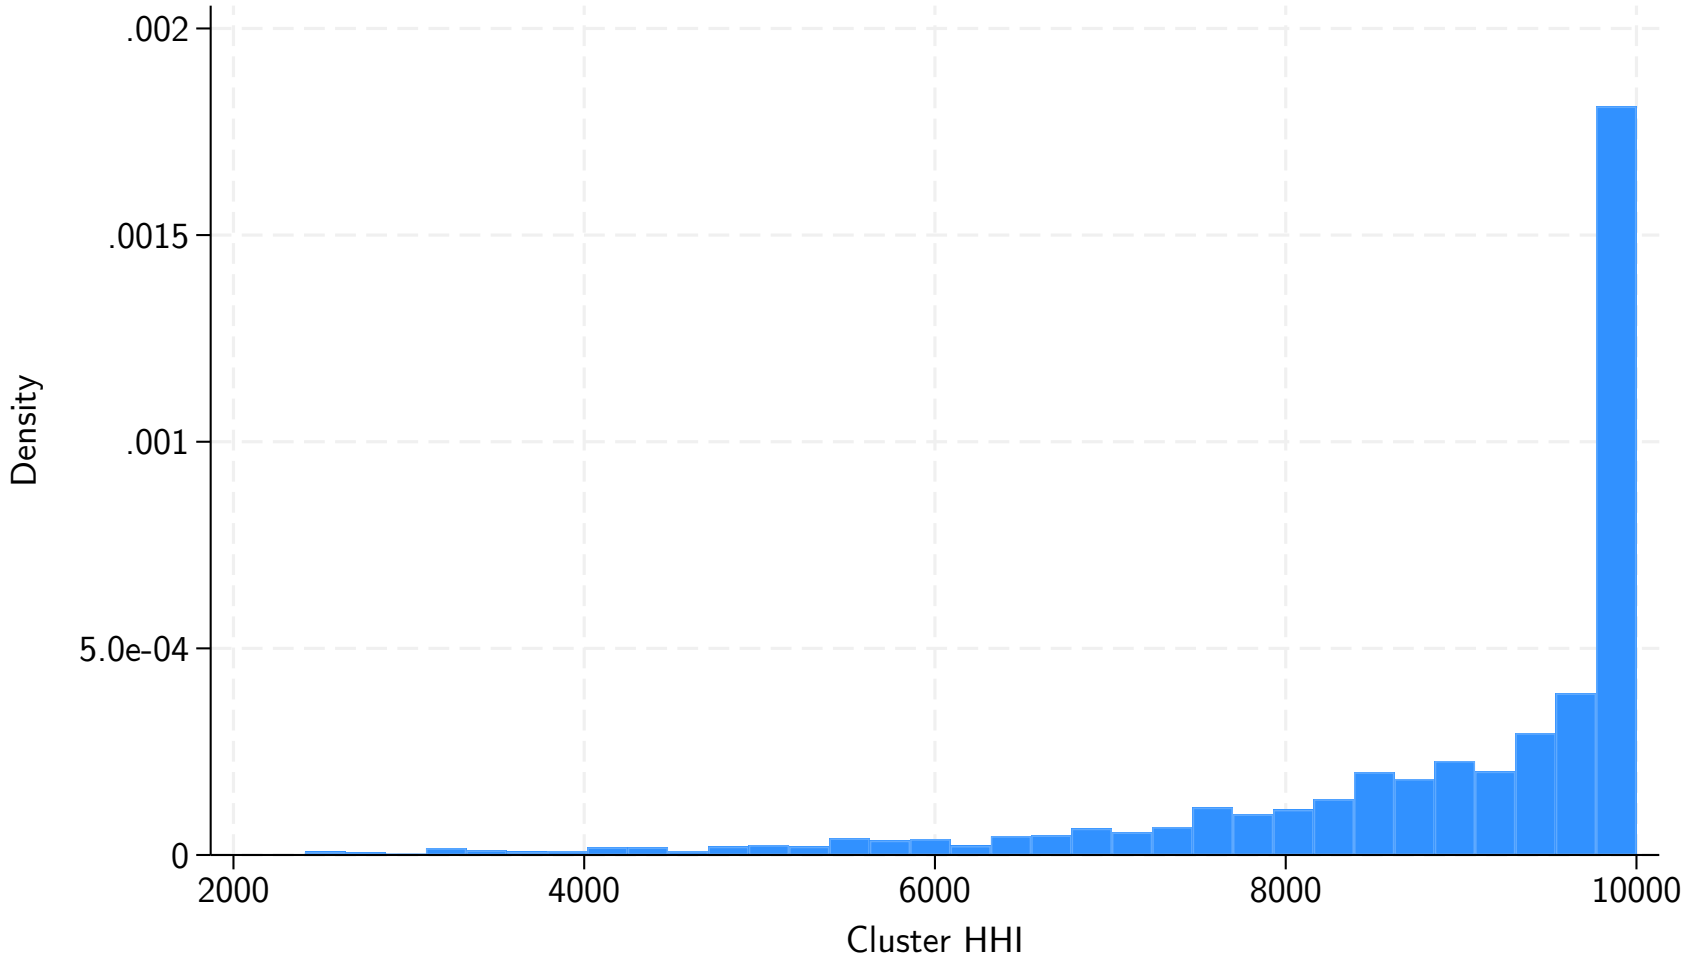


Figure IA.6: Distribution of Cluster Concentration. Figure plots a histogram of average cluster HHI. The HHI uses each bank's share of the sum of member risk factors' trailing 90-day P&L standard deviations in the cluster, on a 0–10,000 scale. The sample pools the pre-2021 and post-2020 reporting eras, in which clusters are estimated separately; factors with no cluster assignment are pooled into a residual group. Observations are clusters, with daily HHI averaged within each cluster before binning.

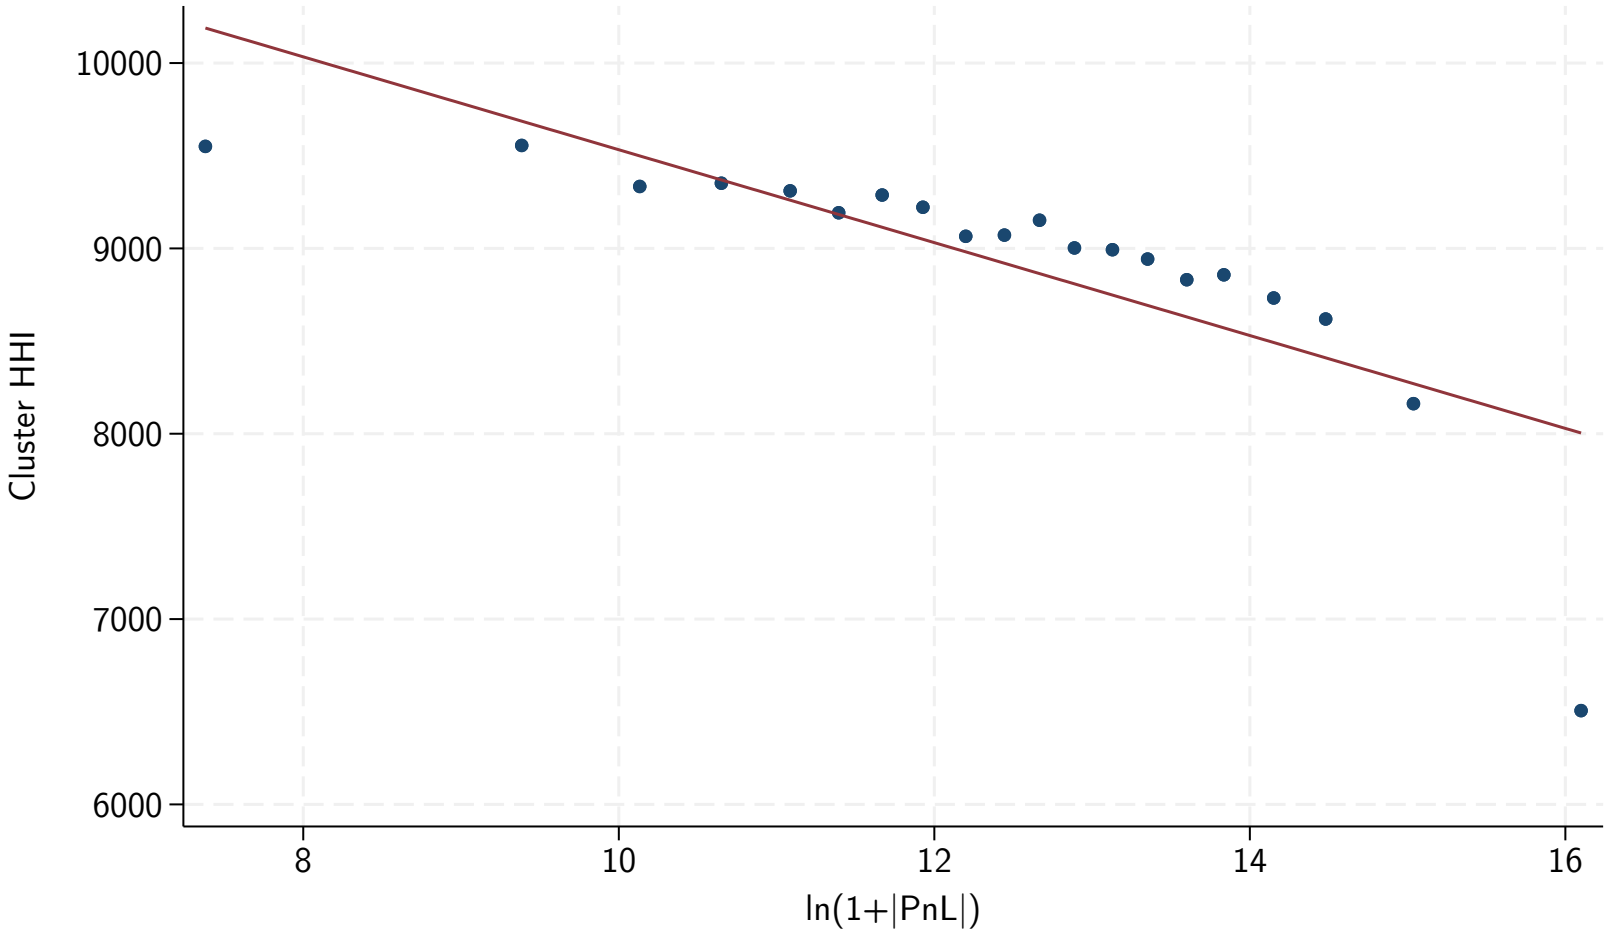


Figure IA.7: Cluster Concentration and Activity. Figure plots a binned scatter of average cluster HHI against the log of one plus average daily cluster activity, with a fitted line. Daily activity is the sum of absolute risk-factor P&L across all banks in a cluster, measured in dollars; the log is taken after averaging over days. HHI is computed from banks' shares of the sum of member risk factors' trailing 90-day P&L standard deviations, on a 0–10,000 scale. The sample pools the pre-2021 and post-2020 reporting eras, with separate cluster identifiers for each era. Points are binned means across clusters, after averaging daily HHI and daily activity within each cluster.

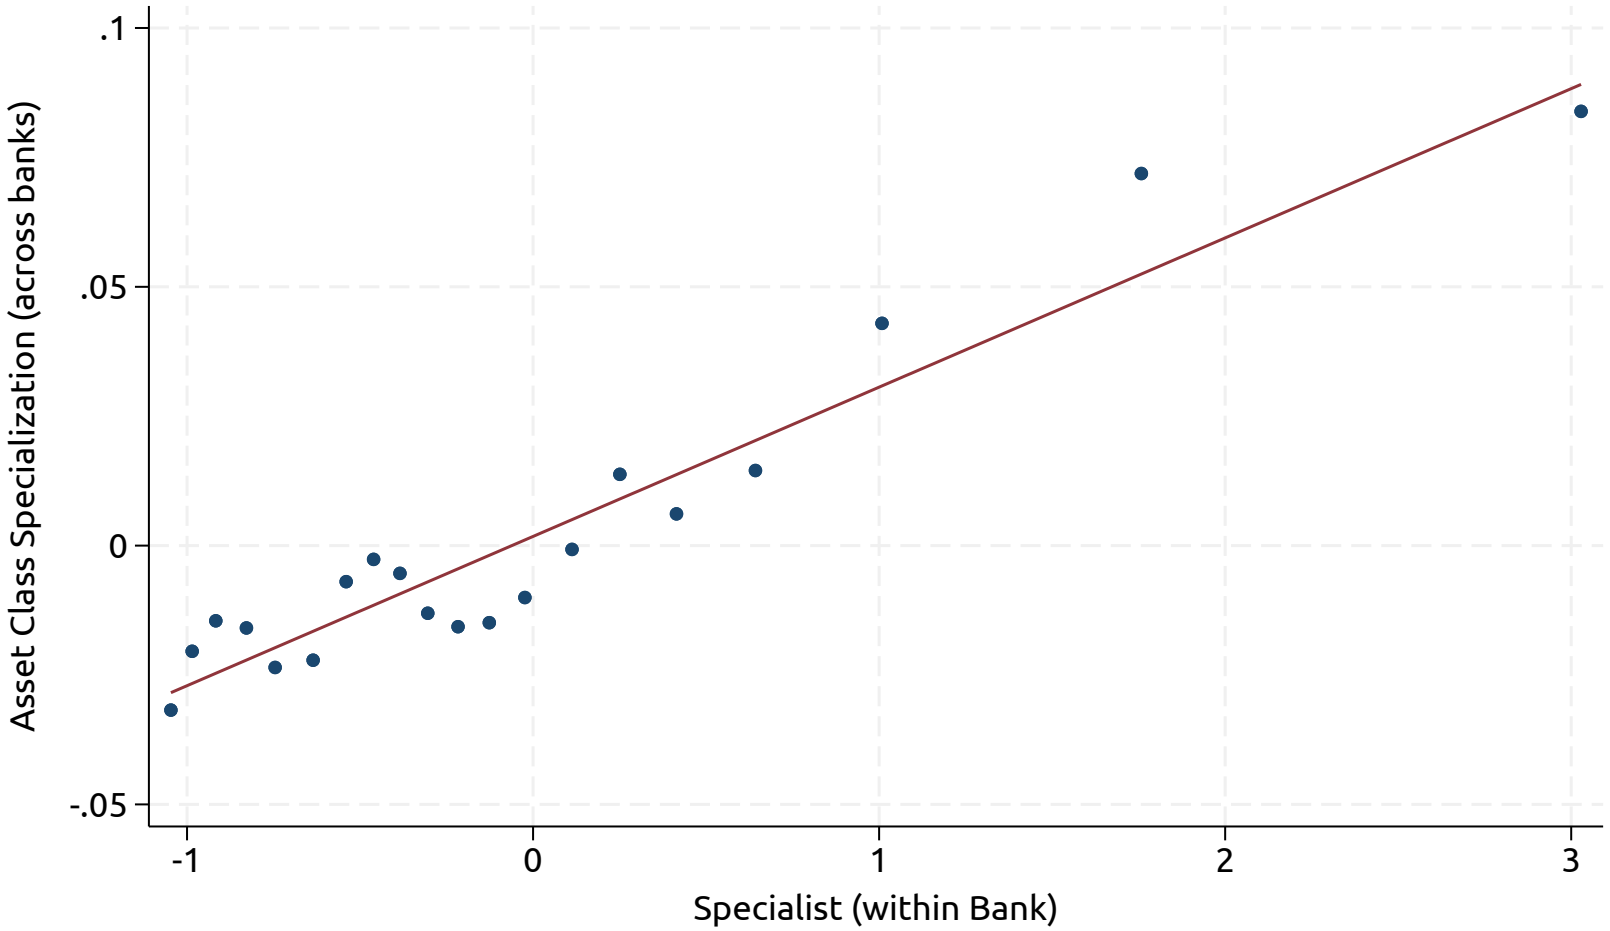


Figure IA.8: Asset Class Specialization. Figure plots a binned scatter of two asset-class specialization measures, with a fitted line. The y-axis is a contemporaneous, unstandardized measure constructed analogously to Blickle et al. (2025): the bank's share of its own VaR in the asset class minus that asset class's share of VaR summed across all banks. The x-axis is *Specialist*: the asset class's share of the bank's total absolute desk P&L, averaged over the 90 calendar days ending the day before the date and standardized over the full daily bank–asset-class panel. Sample: the full sample period. Points are binned means of bank × asset class × day observations.

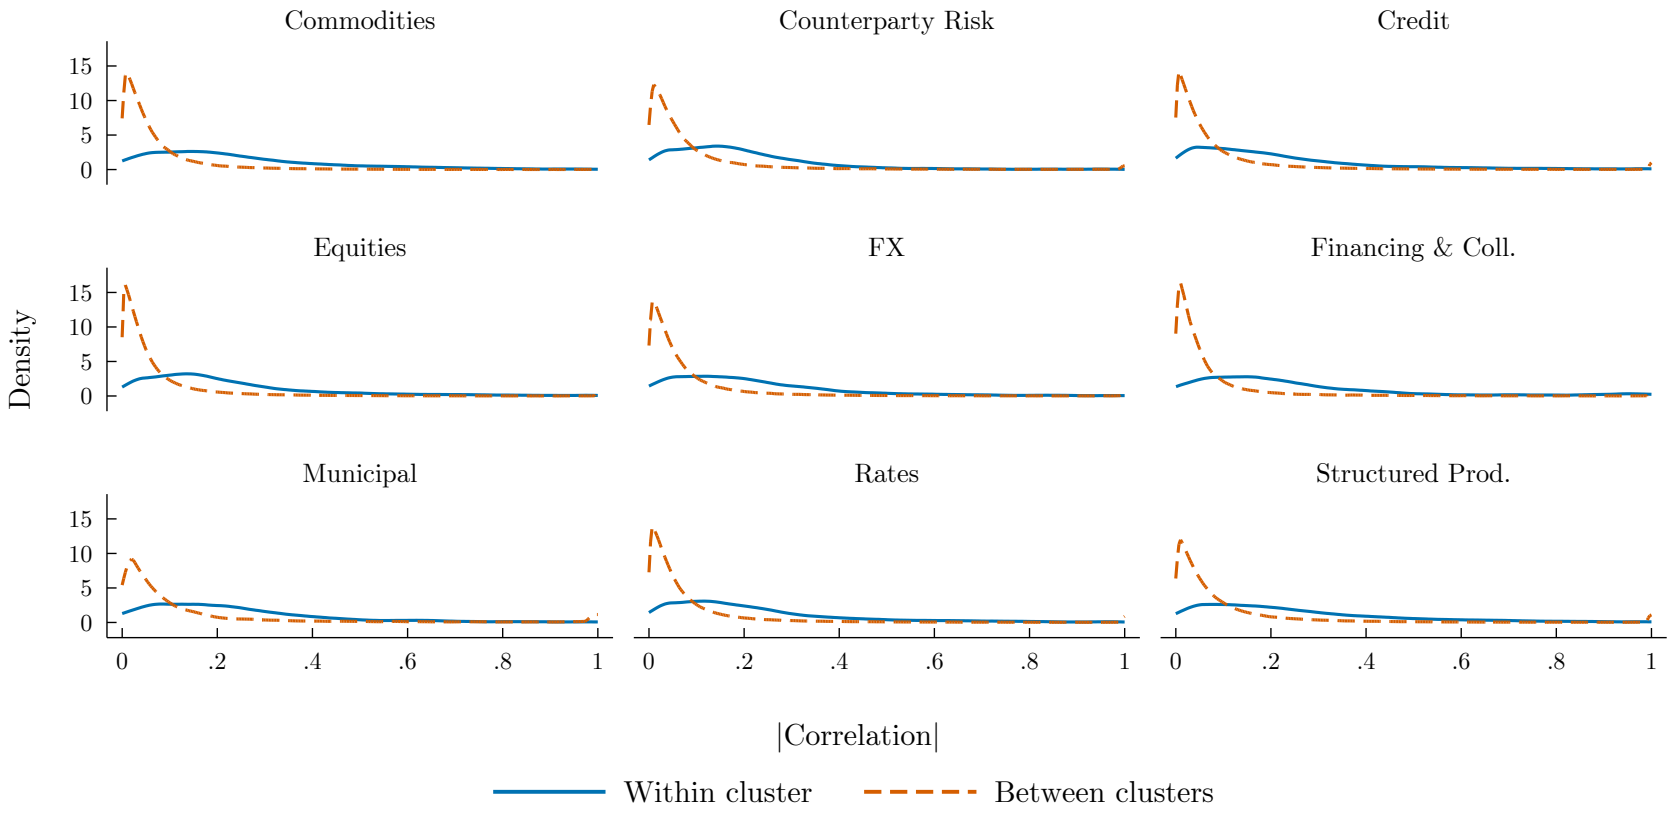


Figure IA.9: Risk-Factor Correlation Distributions. Figure plots densities of absolute correlations between risk-factor return proxies, separately by asset class. Solid curves show pairs in the same cluster; dashed curves show pairs in different clusters. Each pair must have at least one year of overlapping data. Distributions use risk-factor pairs within each asset class.



Figure IA.10: Clusters by Signed Correlation Threshold. Figure plots cluster counts (solid lines) and singleton-cluster counts (dashed lines) as the signed-correlation threshold varies. Singleton clusters contain one risk factor. The panels show the aggregate and nine asset classes. Correlations use the risk-factor return proxy based on trailing 90-day P&L volatility. Clustering uses the post-2020 sample. Counts are reported at each correlation threshold for the aggregate panel and for each asset class.

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 IA.1: Clusters by Correlation Threshold. Table reports cluster counts as the absolute-correlation threshold for average-linkage clustering varies across columns. Risk factors are clustered within asset classes using absolute correlations of their return proxies based on trailing three-month P&L volatility. The benchmark used elsewhere in the paper is 0.10. Clusters are estimated separately in the pre-2021 and post-2020 reporting eras. Sample: the full sample period. Each cell counts clusters at the asset class × reporting era × correlation threshold level.

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.2: Concentration by Asset Class. Table reports average concentration and absolute desk P&L by asset class. *HHI* is the average of the daily asset-class HHI defined in Table 4, based on banks' shares of total desk VaR in each asset class, and |P&L| is the average daily total absolute comprehensive P&L of the asset class's desks, in dollars. Rows are sorted by *HHI*. Asset-class labels follow the raw data groupings: Repo, Muni, Equity, and Securitized Products correspond to Financing and Collateral, Municipal, Equities, and Securitized and Structured Products elsewhere in the paper. Sample: the full sample period. Each row is an asset class, with daily HHI and absolute P&L averaged over business days.

	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			Yes	Yes
Date FE		Yes		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.3: Desk VaR and P&L Volatility. Table reports regressions of desk VaR on the standard deviation of comprehensive desk P&L over the previous 90 calendar days. Both measures are in dollars. Constant omitted. Within R^2 reported. Standard errors are clustered by date and desk; t-statistics in parentheses, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Sample: the full sample period. Observations are bank $\times$ desk $\times$ day.

	$\lvert\rho_{ij}\rvert$					
	(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 IA.4: Absolute Risk-Factor Correlations. Table reports regressions of absolute correlations between daily risk-factor return proxies, multiplied by 100. Cluster assignments use absolute correlations of signed daily risk-factor return proxies. The regressors are indicators for whether the two factors belong to the same bank, the same asset class, or the same cluster, together with interactions of the same-bank indicator with the other two. Columns (1)–(3) include the same-bank, same-asset-class, and same-cluster indicators separately; columns (4)–(6) combine the indicators and interactions as shown. All columns include a constant and no fixed effects. Sample: all pairs with at least one year of overlapping daily observations. Standard errors are two-way clustered on the two factors in the pair; t-statistics in parentheses, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Observations are risk-factor pairs.

	ρ_{ij}					
	(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: Signed Correlations of Risk Factors. Table reports regressions of signed correlations between daily risk-factor return proxies, $100 \times \rho_{ij}$. Cluster assignments use signed correlations of the same signed daily risk-factor return proxies used in Appendix Table IA.4, where assignments use absolute correlations. Columns (1)–(3) include the same-bank, same-asset-class, and same-cluster indicators separately. Columns (4)–(6) combine these indicators and their same-bank interactions as shown. All columns include a constant and no fixed effects. Sample: all pairs with at least one year of overlapping daily observations. Standard errors are two-way clustered on the two factors in the pair; t-statistics in parentheses, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Observations are risk-factor pairs.

	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.6: Unattributed P&L and Excess Specialization. Table reports regressions of the unattributable P&L share on desk specialization and excess specialization. Under the main definition, columns (1)–(3), the outcome is absolute unexplained P&L (Item 10 under the FR VV-1 reporting requirements) divided by the sum of the P&L attributed to existing positions' risk factors, carry, and the unattributable component, with all components entering in absolute value. Columns (4)–(6) use a broader numerator as well as a broader denominator: absolute other attributable P&L enters both. In Table 6, columns (4)–(6) instead add this component only to the denominator. *Specialist* is the desk-level measure of Table 6: the desk's share of the bank's total absolute daily P&L, averaged over the 90 calendar days ending the day before date t and standardized over the full desk panel. *Excess Specialization* is the bank's share of its own VaR in the desk's asset class minus that asset class's share of VaR summed across banks, measured at the bank × asset-class × day level, contemporaneous and not standardized. Constant omitted. Within R^2 reported. Standard errors are clustered by date and bank–asset class; t-statistics in parentheses, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Observations are bank × desk × day.

	log(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)	
log(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 IA.7: Specialization and Reporting Granularity. Table reports regressions of risk-factor counts and unattributable P&L shares. In columns (1)–(3) the dependent variable is the log of the number of distinct risk factors to which the desk attributes nonzero P&L on that day. In columns (4)–(6) it is the unattributable share under the main definition of Table 6, absolute unattributed P&L divided by the sum of absolute unattributed P&L, carry, and P&L attributed to existing positions' risk factors. *Specialist* is the desk's share of the bank's total absolute daily P&L, averaged over the 90 calendar days ending the day before date t and standardized over the full desk panel. *log(Num. of Risk Factors)* is that same log count entering as a regressor in columns (5)–(6). Columns (1)–(4) include *Specialist* alone, column (5) includes *Specialist* and the log factor count, and column (6) includes the log factor count alone. Constant omitted. Within R^2 reported. Standard errors are clustered by date and bank–asset class; t-statistics in parentheses, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Observations are bank × desk × day.

	# 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	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Asset Class × Date FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
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.8: Demand and Bank Participation. Table reports regressions of subsequent bank participation on cluster demand measures. The dependent variable is the maximum number of distinct banks active in the cluster over the following 30 calendar days ("1m") or 90 calendar days ("3m"). Four demand proxies enter one at a time. *Activity Level* uses the trailing 90-day mean of the cluster's total gross risk-factor P&L, summed in absolute value across all banks in the cluster. *Activity Volatility* uses the trailing 90-day standard deviation of the same series. *Average RENTD Limit* and *Total RENTD Limit* are built from banks' internal risk limits justified by RENTD analysis. For each U.S. dollar-denominated RENTD sensitivity limit we take the maximum absolute value of the upper and lower bounds as a measure of demand capacity, aggregate across a bank's limits within a bank-cluster-date by the mean (Average) or the sum (Total), and then average across banks to reach the cluster-date. Each proxy is log-transformed and standardized over the cluster-day panel before selecting the monthly observations in columns (9)–(10). Constant omitted. Within R^2 reported. Standard errors are clustered by date and cluster in columns (1)–(8) and by cluster alone in columns (9)–(10); t-statistics in parentheses, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. RENTD regressions use post-2020 observations with cluster-level limit data. Observations are cluster × day in columns (1)–(8) and cluster × month in columns (9)–(10).

	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 IA.9: Constraints and Activity: Bank Fixed Effects. Table reports regressions of changes in risk-factor activity on specialization and its interactions with bank constraints. The dependent variable is $100 \times [\ln(\overline{Activity}^{\text{fwd}} + 1) - \ln(\overline{Activity}^{\text{trail}} + 1)]$, where *Activity* is the risk factor's absolute P&L, the forward mean is taken over the following 1, 3, or 5 calendar days and the trailing mean over the previous 90 calendar days. *Specialist* is the factor's absolute P&L divided by the bank's sum of absolute risk-factor P&L that day, averaged over the 90 calendar days ending the day before date t and standardized over the full risk-factor panel. *Near Capacity* is the lagged-VaR-weighted share of the bank's desks whose maximum limit proximity exceeds 80 on that day. *Utilization* is constructed from the lagged-VaR-weighted average across the bank's desks of each desk's maximum limit proximity. Both constraint measures are unstandardized. Odd columns use *Near Capacity*; even columns use *Utilization*. Every column includes bank and cluster-by-date fixed effects. The regressions include each constraint measure, *Specialist*, and their interaction. The sample uses dates with bank constraint data in the post-2020 reporting era. Constant omitted. Within R^2 reported. Standard errors are clustered by date and bank–cluster; t-statistics in parentheses, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Observations are bank × desk × risk factor × day.

	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 IA.10: Competitor Constraints and Cluster Activity. Table reports regressions of bank-cluster activity changes on specialization interacted with other banks' constraints. The dependent variable is $100 \times [\ln(\overline{Activity}^{\text{fwd}} + 1) - \ln(\overline{Activity}^{\text{trail}} + 1)]$, where *Activity* is a measure of the bank's absolute risk-factor P&L in the cluster, the forward mean is taken over the following 1, 3, or 5 calendar days and the trailing mean over the previous 90 calendar days. *Specialist* is measured at the cluster level: the cluster's sum of absolute risk-factor P&L divided by the bank's sum across all its clustered factors, averaged over the 90 calendar days ending the day before date t and standardized over the full bank-cluster panel. *Other Banks Proximity* is the leave-one-out mean, across the other banks active in the cluster, of each bank's maximum limit proximity in that cluster-day, weighted by each bank's mean member-factor trailing 90-day P&L standard deviation in the cluster; *Other Banks Near Capacity* is the same weighted mean of an indicator that the other bank's maximum proximity in the cluster exceeds 80. Both constraint measures are unstandardized. Odd columns use *Other Banks Proximity*; even columns use *Other Banks Near Capacity*. The main effects of *Specialist* and the constraint measures are estimated but omitted from the table. Every column includes bank-by-cluster, bank-by-date, and cluster-by-date fixed effects. Constant omitted. Within R^2 reported. Standard errors are clustered by date and cluster; t-statistics in parentheses, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Sample: bank-cluster-days with at least one other bank active in the cluster. Observations are bank × cluster × day.

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

Table IA.11: Core and Peripheral Trader Moves. Table reports the number of moves classified by whether the trader's asset class was peripheral or core at the origin bank and at the destination bank. An asset class is core (peripheral) for a bank in a month if the bank's monthly share of absolute desk P&L in that asset class is at or above (below) the bank's within-month median across its reported asset classes. Bank–asset-class–months with no reported activity are subsequently filled with zero and classified as peripheral. The final column reports the share of movers landing at a core destination. Sample: all moves between covered banks from July 2014 onward for which origin and destination specialization are both available. Observations are trader moves between covered banks.

	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 IA.12: Within-Cluster Comovement. Table reports pair counts, AUC statistics, and correlation hit rates by asset class. The first two columns report the number of factor pairs classified as within the same cluster ("Within") versus in different clusters ("Between"). Columns *AUC* report the statistic $\Pr(|\rho|_{\text{within}} > |\rho|_{\text{between}})$ computed using absolute correlations estimated on the full sample, on odd trading days only, and on 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. Correlations are calculated for risk-factor pairs within an asset class, and the statistics are aggregated by asset class.

	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 IA.13: Risk-Factor Variation Regressions. Table reports moments of risk-factor return indices and time-series regressions of risk-factor returns on indices. The first four columns report the daily mean and standard deviation of a value-weighted and an equal-weighted index of risk-factor return proxies, computed across all risk factors within an asset class and, in the *Aggregate* row, across all risk factors. The value-weighted index uses the standard deviation of a risk factor's P&L in the previous month as weights and keeps factors with more than 15 days of data in the month. The remaining columns report median slopes and R^2 across factors in each asset class, alongside the reported t-statistics, from the time-series regressions $\textit{Risk Factor}_t^i = \alpha + \beta \textit{Index}_t + \varepsilon_t$, where $\textit{Index}_t$ is the value-weighted aggregate risk-factor index, the value-weighted index computed across risk factors in the same asset class, or the Fama–French equity market excess return. R^2 columns are in %. Sample: daily from August 2014. The first four columns summarize daily asset-class and aggregate indices; the remaining columns summarize risk-factor time-series regressions by asset class.