Competition for Relationship Capital

Wei-Fong Pan[*]

Xiamen University

Abstract

This paper studies whether firms strategically adjust demand for relationship-oriented human capital in response to competitive pressure, using relationship banking as a setting where customer relationships and soft information are partly embodied in workers. Using vacancy postings from three Chinese recruitment platforms, we construct a shift-share-style exposure measure that combines predetermined local rival relevance with leave-one-city-out forecast innovations in rivals' relationship-banking recruitment. Greater rival pressure increases banks' relationship-banking vacancy demand, but not IT/data, funding, or back-office hiring, and has no comparable effect for geographically irrelevant rivals. Vacancy text shows stronger demand for workers with client resources, relationship experience, local market knowledge, and seniority. The findings are robust to alternative shock constructions, network weights, geographic leave-out procedures, and shock-level inference. The evidence points to a labor-market channel through which firms compete for worker-embedded relationship capital before realized worker or customer mobility occurs.

Keywords: relationship banking; labor markets; strategic hiring; soft information; bank competition

JEL codes: G21, J23, J24, L13

[*] Department of Finance, School of Economics, and Wang Yanan Institute for Studies in Economics (WISE), Xiamen University, Xiamen, China. Email: panwf@xmu.edu.cn. I am especially grateful to Yan Liu for generously sharing the China Banking Database (CBD). All errors are my own.

1. Introduction

Firms compete not only for customers and products, but also for employees who embody valuable knowledge and relationships. This margin is especially important in relationship-intensive services, where part of a firm's customer capital may reside in individual workers. In financial advisory markets, Gurun, Stoffman, and Yonker (2021) show that adviser-client relationships shape employment choices and that client assets often follow advisers across firms. Outside finance, Patault and Lenoir (2024) show that sales managers carry customer capital across employers, increasing the new employer's probability of trading with the manager's former buyers. These studies establish that economically valuable relationships can be partly worker-embedded. Yet we know much less about the demand side of this process: when a rival expands its demand for relationship-intensive workers, do competing firms strategically expand their own demand for similar human capital before workers or customers actually move?

Relationship banking provides a particularly natural setting in which to study this question. Relationship banking is usually analyzed through the contracts, information, and repeated interactions that connect a bank to its borrowers (Petersen and Rajan, 1994; Berger and Udell, 1995; Boot, 2000). Yet these relationships are produced and maintained by people. Corporate relationship managers originate business, gather soft information, monitor firms, coordinate internal credit decisions, and often span multiple products on both sides of the bank balance sheet. Evidence that individual loan officers and personal connections matter for information production and lending relationships further underscores the role of employees in this process (Hertzberg, Liberti, and Paravisini, 2010; Uchida, Udell, and Yamori, 2012; Karolyi, 2018). If these employees embody economically valuable relationship capital, rivalry in banking should extend beyond loan prices, branches, and product terms to the labor market itself. A bank that observes a locally relevant competitor expanding its demand for relationship bankers may therefore have an incentive to expand its own relationship-oriented hiring demand before any customer movement is observed.

This paper asks whether banks strategically adjust their demand for relationship-intensive human capital in response to competitive pressure. The empirical challenge is that banks often hire at the same time for reasons unrelated to strategic rivalry. National credit conditions can move several banks together; city-level economic conditions can raise local labor demand across institutions; and a bank can undertake a nationwide hiring campaign that mechanically affects many locations at once. Identifying a strategic labor-market response therefore requires variation

in rival hiring that is both unexpected and differentially relevant to the focal bank in a particular local market, while removing common bank-wide and local shocks.

We bring this question to the data using vacancy postings from three major Chinese recruitment platforms—51job, Zhaopin, and Liepin. After matching postings to banks and locations, we use job titles, descriptions, and qualification requirements to identify relationship-banking positions, with a deliberately conservative classification that distinguishes relationship origination from risk, operations, IT, general customer service, and other support functions. Our baseline measure focuses on corporate- and SME-facing relationship roles and is validated against independent human coding. The resulting data allow us to observe banks' demand for relationship-oriented labor at the bank-city-month level. Because vacancy postings capture recruitment demand rather than completed hires, we interpret the evidence as changes in hiring demand rather than realized worker mobility or employee poaching.

Our empirical design combines predetermined local competitive exposure with unexpected innovations in rivals' relationship-hiring demand. The construction follows the logic of predetermined-exposure and shift-share designs, in which pre-existing exposure maps common or entity-level shocks into heterogeneous local treatment intensity (Adão, Kolesár, and Morales, 2019; Goldsmith-Pinkham, Sorkin, and Swift, 2020; Borusyak, Hull, and Jaravel, 2022). In banking, Greenstone, Mas, and Nguyen (2020) similarly use pre-existing local bank exposure to map bank-level credit-supply shifts into local credit conditions. We adapt this logic to competition for relationship-oriented labor. Using 2014-2016 data, we construct a predetermined map of banks' geographic footprints and measure which rival banks are most relevant to each focal bank-city market. We then estimate leave-one-city-out forecasts of each rival's relationship-banking recruitment outside the focal city and use the forecast residual, net of a common monthly component, as the rival's hiring innovation. Excluding the focal city prevents its vacancies from mechanically entering both the rival-side innovation and the focal-bank outcome. We aggregate these innovations across locally present rivals using predetermined geographic-similarity weights.

The fixed-effect structure is central to identification. Bank-by-city fixed effects absorb persistent differences in a bank's local presence, market positioning, and recruiting intensity. Bank-by-month fixed effects absorb any bank-wide monthly shock, including national recruitment campaigns, balance-sheet conditions, or changes in corporate strategy that affect all cities in which the bank operates. City-by-month fixed effects absorb local credit demand, labor-market conditions, regulation, and other city-specific monthly forces common to banks in the same market.

Identification therefore comes from comparing, within the same bank and month, cities that differ in their predetermined exposure to innovations at locally relevant rivals, after removing shocks shared by all banks in each city-month. This combination of predetermined local exposure, leave-one-city-out innovations, and saturated fixed effects sharply narrows the set of common-demand and common-recruitment explanations for the estimated response, while not requiring us to claim that forecast residuals are literally randomly assigned.

The central result is that banks respond to locally relevant rival expansions by increasing their own demand for relationship-oriented labor. The coefficient on rival relationship-hiring pressure in the relationship-banking vacancy equation is 0.9444 (standard error = 0.1668). On the extensive margin, the coefficient is 0.6163 (standard error = 0.1126). Given that the standard deviation of the preferred shock is approximately 0.045, a one-standard-deviation increase in rival pressure raises the probability that a bank posts at least one relationship-banking vacancy in a city-month by approximately 2.8 percentage points, relative to an unconditional probability of 7.7 percent.

The response is highly occupation-specific and geographically targeted. Rival relationship-hiring pressure does not generate comparable increases in IT/data, funding, or back-office vacancies, and formal tests reject equality between the relationship coefficient and each comparison occupation, reducing the concern that the effect reflects a broad expansion in bank-wide hiring. A complementary locality placebo shows that hiring innovations at banks without predetermined presence in the focal market produce no comparable response, making it less likely that the results are driven by national recruitment cycles or other common bank-level shocks. The combination of occupational specificity and local relevance therefore points to a targeted response to competition in relationship banking rather than a generic increase in labor demand.

Vacancy text provides a second layer of evidence on the type of human capital banks seek. Rival pressure increases the absolute number and incidence of relationship-banking vacancies that mention client resources, prior relationship or banking experience, senior relationship-management responsibilities, local resources, and substantial prior experience. These characteristics are consistent with banks expanding demand for employees whose productivity depends on customer knowledge, relationship-building capability, and market-specific expertise. We do not interpret these measures as evidence that employees literally transfer a portable book of customers. Rather, they capture observable dimensions of relationship-oriented human capital that banks explicitly value when expanding relationship capacity.

A battery of tests addresses main alternative explanations. The estimates are stable across alternative forecasting models and predetermined network weights, reducing concern that the results are driven by a particular way of extracting rival innovations or mapping rivals into local exposure. Excluding the pandemic years shows that the findings are not specific to the exceptional labor-market and banking conditions of 2020–2022. Reconstructing the shock leave-one-province-out instead of leave-one-city-out addresses the possibility that broader provincial shocks jointly affect rival and focal-bank hiring, while excluding large rival banks and the largest exposure contributors checks that the results are not driven by a small set of dominant institutions. Future-shock placebos for the relationship-minus-nonrelationship outcome provide no evidence that focal-bank hiring systematically anticipates subsequent rival innovations. Because the exposure is generated by common rival-bank shocks, we also implement the Borusyak-Hull-Jaravel shock-level equivalent representation; inference remains similar when clustering by rival bank and calendar month, and wild-cluster bootstrap p-values are below 0.001 for both the relationship outcome and the relationship-minus-nonrelationship gap. Finally, cross-platform exercises—constructing the rival shock from one recruitment platform and the focal outcome from another—reduce concern that the findings reflect platform-specific posting algorithms, duplicate timing, or source-specific coverage. At the same time, primitive-shock predictability tests and residual persistence lead us to frame the design as a high-dimensional reduced-form exposure design rather than as literal random assignment.

We finally examine whether competitive relationship-hiring pressure is followed by changes in bank-level outcomes. Aggregating the city-month exposure measure to the bank-year level, we find that greater rival relationship-hiring pressure is followed by higher deposit funding costs and a relative shift in lending growth toward SME credit rather than consumer credit. The lending-composition result is estimated less precisely and on a smaller sample, so we treat it as suggestive. We find no detectable changes in aggregate lending, employment, profitability, or credit risk. These results indicate that competitive relationship-hiring pressure may be associated with changes in banks' funding conditions and business composition, but they should be interpreted as reduced-form downstream consequences rather than causal effects of the focal bank's own hiring response.

This paper contributes first to the relationship-banking literature by shifting attention from borrower contracts and realized lending relationships to the labor market through which relationship capacity is created. Classic work shows that repeated bank-borrower interactions mitigate information frictions and affect credit availability and terms (Petersen and Rajan, 1994;

Berger and Udell, 1995; Boot, 2000), while organizational work shows that the production and use of soft information depend on organizational structure, delegation, communication costs, and the incentives of individual bank employees (Stein, 2002; Berger et al., 2005; Liberti and Mian, 2009; Agarwal and Hauswald, 2010; Qian, Strahan, and Yang, 2015; Skrastins and Vig, 2019). A related literature shows that individual bankers matter for information production and lending relationships (Hertzberg, Liberti, and Paravisini, 2010; Uchida, Udell, and Yamori, 2012; Karolyi, 2018; Agarwal et al., 2024), while recent work continues to document the economic value of bank-customer relationships in lending and cross-selling (Berger et al., 2024; Basten and Juelsrud, 2026). Particularly close to our mechanism, Wang (2019) shows that greater interbank employee mobility facilitates de novo entry by outside banks when local information is valuable, and Chu et al. (2022) develop a theory in which credit-market competition induces competition for relationship-bank loan officers who possess borrower information. We provide a complementary empirical contribution by directly observing this competitive margin in banks' labor demand. Rather than studying realized employee mobility, bank entry, or the consequences of existing relationships, we show that banks strategically expand demand for relationship-oriented workers when locally relevant rivals increase their own demand for such labor. This response is concentrated in relationship-banking occupations and in vacancies emphasizing client resources, relationship experience, and local market knowledge, linking competitive hiring directly to the production of relationship capacity.

Second, the paper contributes to research on worker-embedded customer capital and the mobility of economically valuable human capital. A broad literature shows that knowledge and productive capabilities can move across firm boundaries with workers (Stoyanov and Zubanov, 2012). Recent research further emphasizes that human capital is an organizational resource that firms can redeploy internally, acquire through hiring, or obtain through changes in firm boundaries. Cestone et al. (2024) show how internal labor markets help firms exploit growth opportunities following competitive shocks, while Tate and Yang (2024) and Beaumont, Hebert, and Lyonnet (2025) highlight the role of worker skills and human capital in acquisitions and corporate expansion. In financial services, Jeffers (2024) shows that restrictions on labor mobility affect investment and entrepreneurial entry, Gao, Wang, and Wu (2024) quantify portable and nonportable components of human capital among M&A advisory bankers, and Huang et al. (2026) show that financial institutions systematically allocate loan-officer human capital across markets. In customer-facing settings, Gurun, Stoffman, and Yonker (2021) provide especially close

evidence that client relationships are partly attached to individual financial advisers: adviser mobility and client mobility move together when contractual restrictions are relaxed. Patault and Lenoir (2024) document a related mechanism outside finance, showing that firms hiring sales managers become more likely to trade with the managers' former customers. These papers primarily study realized mobility and its consequences. Our contribution is to move one step earlier in the adjustment process: we show that firms respond strategically on the demand side of the labor market, expanding vacancy demand for relationship-intensive workers when relevant rivals increase their own hiring pressure, even before we observe a particular employee move or customer reallocation.

Third, the paper adds a labor-market margin to research on banking competition and the transmission of bank-level shocks through pre-existing market exposures. Geographic proximity and local presence shape the production of private information and the terms of bank competition (Degryse and Ongena, 2005), while bank-level shocks propagate through pre-existing relationships and exposures to affect credit supply and real activity (Amiti and Weinstein, 2018; Greenstone, Mas, and Nguyen, 2020). Recent work continues to emphasize both dimensions. Hasan, Shen, and Yuan (2021) show that local product-market competition affects bank loan pricing, Jiang (2023) studies strategic interactions between banks and competing shadow lenders in mortgage markets, and Chava et al. (2023) document the transmission of bank funding shocks through consumer credit. Related evidence from Schneider, Strahan, and Yang (2024) shows that competitive incentives among banks can alter their willingness to cooperate in syndicated lending. We show that the same logic extends to banks' demand for relationship-oriented labor. Predetermined local overlap determines which rival hiring innovations generate competitive pressure for a focal bank, but the response we study is the bank's demand for relationship-oriented human capital rather than lending quantities, loan terms, or market entry. The results therefore identify a labor-market channel through which local banking competition operates: when economically valuable information and customer-facing capabilities are partly embodied in employees, competitive pressure can alter firms' demand for those capabilities before it becomes visible in realized worker movement, customer reallocation, or product-market outcomes.

2. Data and Measurement

2.1 Vacancy data

We obtain the vacancy microdata from CnOpenData (China Open Data Platform), a commercial research database that compiles online recruitment advertisements from multiple Chinese job-search platforms. CnOpenData's online recruitment database contains daily vacancy records with information on employers, job locations, positions, recruitment quantities, compensation, experience requirements, job descriptions, education requirements, and posting dates. Its commercial-bank recruitment database further identifies vacancies posted by banks and their branches in China. Our extract contains three platform-specific sources—51job, Zhaopin, and Liepin—which we harmonize into a common posting-level dataset.

We standardize employer names, job locations, titles, descriptions, posting dates, and platform-specific fields, and match employers to banks in the China Banking Database (CBD). We exclude postings associated with staffing agencies, labor-dispatch arrangements, and outsourcing firms, and retain only direct bank vacancies for which a usable city can be identified. Importantly, the data measure vacancy postings and therefore banks' demand for labor, rather than realized hiring or worker flows.

The vacancy data span 2014–2024. We reserve 2014–2016 as a predetermined pre-period used only to define banks' geographic footprints, construct pairwise geographic similarity across banks, and estimate the preferred forecasting model. The main estimation period runs from January 2017 through December 2024. A bank-city pair enters the predetermined support if the bank posted at least three vacancies and was active in at least two months in that city during 2014–2016. We then construct a balanced monthly panel over the supported bank-city pairs. After singleton observations are removed by the high-dimensional fixed-effect estimator, the preferred relationship-hiring specification contains 147,840 bank-city-month observations.

Table 1 reports descriptive statistics for the baseline estimation sample. Vacancy posting is highly episodic: the average bank-city-month contains 2.34 postings, while the median is zero. Relationship-banking vacancies are substantially sparser, with an unconditional mean of 0.174 postings and a 7.7 percent probability of observing at least one relationship-banking vacancy in a given bank-city-month. Because unconditional quartiles are therefore dominated by zeros, Table 1 separately reports posting intensity conditional on positive activity. Among the 11,369 bank-city-months with at least one relationship-banking posting, the mean is 2.265 postings, the median is one, the 75th percentile is two, and the maximum is 56. Thus, relationship-banking recruitment

is infrequent at the monthly local level, but when banks recruit for these positions, positive cells typically contain a small number of postings rather than mechanically reflecting only very large recruitment episodes.

2.2 Measuring relationship-banking vacancy demand

Our baseline relationship category is deliberately tied to corporate relationship banking rather than to all customer-facing jobs. We classify postings using job titles, structured job-function fields, and vacancy descriptions. The classification combines high-signal job titles with contextual information from vacancy descriptions and qualification requirements. Detailed coding rules, priority rules, representative dictionary terms, and validation procedures are provided in Appendix A.[1] High-signal phrases include corporate customer manager, enterprise customer manager, institutional customer manager, micro and small-business customer manager, inclusive-finance customer manager, agricultural relationship manager, corporate finance, corporate banking, enterprise credit, supply-chain finance, trade finance, and related roles. The classifier assigns mutually exclusive banking functions with priority rules that prevent generic risk or retail language from overriding specific relationship-banking terms.

For each bank-city-month, we count relationship-banking vacancy postings and all direct vacancy postings; nonrelationship postings are the difference between the two. Our primary intensive-margin outcome is the inverse hyperbolic sine of the relationship-banking posting count, which retains zero cells and compresses the influence of large counts. We also examine an indicator for any relationship-banking posting, inverse-hyperbolic-sine counts of nonrelationship and all postings, the relationship-minus-nonrelationship difference, and the relationship share when total hiring is positive.

The distinction between postings and realized hiring is important. A posting records a bank's decision to advertise demand for a role; it does not consistently reveal the number of workers eventually hired. In particular, platform-specific fields are not reliable headcount measures across sources. We therefore interpret the outcomes as vacancy demand. This is exactly the margin needed for the paper's question—whether banks adjust recruitment demand in response to rival

[1] The classification rules are developed using a separate pilot sample and frozen before evaluation on an independent 900-vacancy holdout. Two human coders independently classify each holdout vacancy before reconciliation into a final human benchmark. For the baseline core relationship measure, precision is 95.3 percent, recall is 91.1 percent, and the F1 score is 93.2 percent. Appendix A reports the full sampling design, coding protocol, and validation results.

pressure—but it does not by itself establish worker movement, client migration, or successful poaching.

2.3 Vacancy text and relationship-capital characteristics

We use vacancy text to characterize the type of relationship labor for which banks recruit. Five pre-specified indicators flag whether a relationship-banking vacancy mentions client resources, prior relationship or banking experience, senior relationship-management responsibilities, local client or market resources, or at least five years of experience. The client-resource dictionary includes phrases referring to customer resources, enterprise or institutional customer resources, networks, channels, self-sourced clients, and related language. The experience dictionary includes prior banking, corporate banking, SME, inclusive-finance, credit, customer-management, and customer-development experience.

These text indicators are designed to capture observable features of relationship-oriented human capital, not to label a worker as owning a transferable client book. A requirement for local resources or relationship experience can reflect local knowledge, soft-information production, business development skill, internal coordination ability, or existing networks. We therefore use the broader term relationship capital: human capital that raises the productivity of originating, maintaining, and expanding client relationships.

3. Conceptual Framework

3.1 Relationship Capital and Strategic Hiring

Our conceptual framework builds on two linked ideas. First, relationship banking relies on information and capabilities accumulated through repeated interactions with customers. A large literature shows that bank-borrower relationships facilitate the production of private and soft information and affect the availability and terms of credit (Petersen and Rajan, 1994; Berger and Udell, 1995). Because soft information is difficult to codify and transmit within organizations, its production and use depend importantly on organizational structure and the employees who acquire it (Stein, 2002; Berger et al., 2005; Liberti and Petersen, 2019). Evidence from loan-officer rotation and delegated lending further shows that individual bank employees play an important role in producing, communicating, and acting on borrower-specific information (Hertzberg, Liberti, and Paravisini, 2010; Qian, Strahan, and Yang, 2015).

Chu et al. (2022) provide a particularly useful banking-specific microfoundation for why this human capital can become a margin of competition. In their model, relationship-bank loan officers possess borrower-specific information that is valuable to competing banks, so competition for borrowers can spill over into competition for the employees who hold that information. We adopt this insight but use a broader and deliberately reduced-form notion of relationship capital. Relationship bankers may embody soft information, local market knowledge, client-development skills, and networks. Our framework does not require a literal transfer of a client book or proprietary borrower file, nor does it require realized poaching. It requires only that an expansion in a rival bank's relationship-labor capacity changes the marginal value of the focal bank's own relationship capacity.

We use the term *relationship capital* to encompass these human-capital-intensive capabilities. Relationship bankers may originate and maintain business relationships, monitor borrowers, develop local client networks, coordinate products across business lines, and support deposit gathering and cross-selling. What matters for the framework is that banks use relationship-oriented employees to build and maintain customer-facing capacity, not that customers necessarily follow individual employees across banks.

Second, local banking markets are competitive, and organizational responses to competition can depend on how banks produce and use soft information. Canales and Nanda (2012), for example, show that decentralized bank structures interact with local competitive conditions in shaping lending behavior. More generally, geographic proximity affects banks' ability to collect private information and the terms on which they compete for borrowers, implying an inherently local dimension to relationship banking (Degryse and Ongena, 2005). The insight from Chu et al. (2022) adds a labor-market dimension: when economically valuable information and relationship capabilities are partly embodied in employees, competition in customer markets can increase the value of securing relationship-oriented labor. We therefore model locally relevant rival relationship capacity as a source of competitive pressure on the focal bank.

Consider bank i in local market c, suppressing time subscripts for clarity. Let h_{ic} denote its demand for relationship-oriented labor, and let relationship capacity be

$$q_{ic} = a_{ic}\, h_{ic} \tag{1}$$

where $a_{ic} > 0$ captures the productivity of relationship-oriented human capital. To make local competitive relevance explicit, define the focal bank's exposure to rivals' relationship capacity as

$$R_{ic} = \sum_{j \neq i} w_{ijc}\, q_{jc} \tag{2}$$

where $w_{ijc} \geq 0$ measures the predetermined local relevance of rival bank j to bank i in market c. The weights are normalized across locally relevant rivals, and a rival with no pre-existing local presence has weight zero. Equation (2) is a conceptual counterpart to the empirical exposure measure: empirically, we observe innovations in rivals' relationship-banking recruitment rather than realized relationship capacity and combine those innovations using predetermined local-relevance weights.

The focal bank chooses h_{ic} to maximize

$$\Pi_{ic} = V(q_{ic}, R_{ic}) - C(h_{ic}) \tag{3}$$

where $V(\cdot)$ is the value generated by the bank's relationship capacity in the local market and $C(\cdot)$ is increasing and convex. The key comparative-static condition is that the marginal value of the focal bank's relationship capacity rises when locally relevant rivals expand theirs:

$$V_{qR} > 0 \tag{4}$$

This is a strategic-complementarity condition in the spirit of Bulow, Geanakoplos, and Klemperer (1985), applied to relationship capacity. Several mechanisms can generate it. A rival expansion can intensify competition for prospective customers, increase the value of defending existing relationships, or raise the benefit of maintaining coverage in contested local markets. It can also intensify competition for relationship bankers themselves when those employees embody information or client-development capabilities that rivals value. The latter mechanism is directly motivated by Chu et al. (2022), in which loan officers' borrower information creates a link between credit-market and labor-market competition. Our empirical design does not require us to distinguish among these mechanisms; each implies that locally relevant rival expansion raises the marginal value of the focal bank's relationship capacity.

The first-order condition is

$$a_{ic} V_q(a_{ic} h_{ic}, R_{ic}) = C'(h_{ic})$$

Under the usual second-order condition,

$$\frac{\partial h_{ic}}{\partial R_{ic}} = \frac{a_{ic} V_{qR}}{C''(h_{ic}) - a_{ic}^2 V_{qq}} > 0 \tag{5}$$

Equation (5) provides the central empirical prediction: an increase in locally relevant rivals' relationship-banking labor demand increases the focal bank's desired relationship-oriented labor demand. The model is intentionally agnostic about whether the response is primarily defensive - protecting existing or prospective relationships - or offensive - expanding in response to changing competitive conditions. Both motives imply the same positive hiring response.

Our empirical counterpart to h_{ic} is vacancy posting rather than realized employment. The framework therefore predicts an adjustment in desired labor demand as revealed through vacancy creation. It does not require vacancies to be filled, workers to move between banks, or customers to migrate with particular employees. This distinction is important because the empirical analysis identifies competition for relationship capital in recruitment demand, not realized poaching.

3.2 Occupation Specificity and Relationship-Capital Intensity

The framework yields three testable implications. First, the response should be occupation-specific. Suppose the bank chooses relationship-oriented labor h_{ic} together with nonrelationship labor n_{ic}. If rival relationship expansion primarily raises the marginal value of customer-facing relationship capacity rather than the value of generic bank labor, then

$$\frac{\partial h_{ic}}{\partial R_{ic}} > \frac{\partial n_{ic}^{*}}{\partial R_{ic}} \tag{6}$$

A positive response across all occupations could simply reflect a common expansion in local banking activity or general recruitment. By contrast, a response concentrated in relationship-banking vacancies is more consistent with strategic adjustment in relationship capacity. This motivates the comparison with IT/data, back-office, and funding vacancies.

Second, Equation (2) makes geographic relevance part of the mechanism rather than merely a sample restriction. If a rival has no meaningful pre-existing presence in the focal market, then $w_{ijc} = 0$ and an innovation in that rival's relationship hiring should generate little or no competitive pressure for the focal bank in that market. Rivals with greater predetermined local overlap receive greater weight because their expansion is more likely to affect the same customers, local information environment, or pool of relationship-oriented workers. This implication motivates both the predetermined geographic-similarity weights and the never-local-rival placebo.

Third, Equation (1) treats relationship labor as a composite input, but relationship capacity can be produced by several types of relationship-oriented workers. More generally, let

$$q_{ic} = \sum_{k=1}^{K} a_{ikc}\, h_{ikc} \tag{7}$$

where k indexes different types of relationship-oriented labor. Vacancies emphasizing client resources, prior relationship or banking experience, local market resources, senior relationship-management responsibility, or substantial prior experience are natural empirical proxies for inputs that contribute to relationship capacity. An increase in R_{ic} can therefore raise the absolute demand for several such types simultaneously. Without additional restrictions on relative adjustment costs or substitution across worker types, however, the model does not imply that any one characteristic

must rise as a share of total relationship hiring. This is why the characteristic-level tests are interpreted primarily as evidence about the type of human capital being demanded, while composition within relationship hiring is treated separately.

The framework gives three empirical predictions:

Prediction 1. Greater locally relevant rival relationship-hiring pressure increases the focal bank's relationship-banking vacancy demand.

Prediction 2. The response is stronger for relationship-banking vacancies than for nonrelationship occupations and is weak or absent for rivals without predetermined local relevance.

Prediction 3. Rival relationship-hiring pressure increases the absolute demand for vacancies that explicitly emphasize client resources, relationship experience, local market resources, senior responsibility, or substantial prior experience, although these characteristics need not increase as a share of total relationship hiring.

4. Empirical Strategy

4.1 Construction of Rival Relationship-Hiring Pressure

Our empirical design combines predetermined exposure to locally relevant competing banks with time-varying innovations in those rivals' relationship-banking recruitment. The construction is related to the broader class of exposure and shift-share designs, in which predetermined exposure shares map entity-level shocks into heterogeneous local exposure (Adão, Kolesár, and Morales, 2019; Goldsmith-Pinkham, Sorkin, and Swift, 2020; Borusyak, Hull, and Jaravel, 2022). In a banking setting, Greenstone, Mas, and Nguyen (2020) similarly combine pre-existing local bank shares with bank-level credit-supply shifts to construct local credit-market shocks. We adapt this general logic to competition for relationship-oriented labor.

We first determine which rival banks are locally relevant using information measured entirely before the estimation period. Using vacancy data from 2014–2016, we define an eligible bank-city presence when a bank posts at least three vacancies and is active in at least two months in that city. For each bank, we form a geographic-footprint vector across cities whose elements equal log(1 + pre-period postings). We operationalize the similarity of two banks' predetermined geographic footprints using the cosine similarity of these vectors, with the diagonal set to zero. For a focal bank-city pair (j,c), the relevant rival set consists of banks $b \neq j$ that also have eligible pre-period presence in city c. Our preferred exposure weight is the geographic similarity between the focal and rival banks, normalized within each focal bank-city:

$$s_{jcb}^{pre} = \frac{\text{Similarity}_{jb}{}^{pre}}{\sum_{k \in B_c \setminus \{j\}} \text{Similarity}_{jk}{}^{pre}}$$

All network weights are therefore fixed using information observed before 2017. Cosine similarity is our preferred way of translating predetermined geographic overlap into rival relevance, rather than an assumption inherited from the shift-share literature. We accordingly show that the results are similar under equal weights, rival-footprint weights, and similarity-by-footprint weights (Table 6).

The second component isolates unexpected changes in each rival bank's relationship-hiring mix outside the focal city. For rival bank b, city c, and month t, we aggregate the bank's national vacancy postings and subtract postings in city c. We then calculate the relationship-banking share of these leave-one-city-out postings, setting the share to zero when no postings remain outside the focal city. Excluding city c ensures that vacancies posted there cannot mechanically enter both the rival-side innovation and the focal-city outcome. This construction follows the general logic of leave-one-out procedures used in shift-share settings to reduce mechanical dependence between estimated shocks and local outcomes (Borusyak, Hull, and Jaravel, 2022).

We forecast the leave-one-city-out relationship share using information from the 2014–2016 pre-period. The preferred model absorbs the bank-by-excluded-city mean and month-of-year seasonality and includes pooled autoregressive terms at lags 1, 2, 3, and 12. The raw innovation is the realized relationship share minus its forecast. We then remove an equal-bank-weight common monthly component to isolate bank-specific rather than aggregate recruitment variation. The resulting innovation equals the realized share minus its forecast and the common monthly component:

$$g_{b,-c,t} = r_{b,-c,t} - \widehat{r_{b,-c,t}} - \overline{\varepsilon_t}$$

This step is motivated by the broader banking-shock literature, which seeks to isolate idiosyncratic bank-level variation from common aggregate movements (Amiti and Weinstein, 2018). Our exact forecasting specification is tailored to the monthly recruitment data rather than taken directly from any single prior study. To assess sensitivity to the forecasting procedure, we construct two alternatives. One fixed model includes all monthly lags from 1 through 12. A recursive specification fixes the pre-period means and seasonal components but re-estimates the AR(1,2,3,12) coefficients each month using only information available through t−1. Table 6, Panel A, shows that the estimated relationship-hiring response is nearly unchanged across these forecasting families.

Finally, we combine the predetermined local network weights with the time-varying rival innovations to construct focal bank j's relationship-hiring pressure in city c and month t:

$$Z_{\mathrm{jct}} = \sum_{b \in B_c \setminus \{\mathrm{j}\}} s_{\mathrm{jcb}}^{\mathrm{pre}} \, g_{\mathrm{b,-c,t}}$$

This variable has the familiar shift-share form in which predetermined exposure shares weight time-varying shifts (Adão, Kolesár, and Morales, 2019; Goldsmith-Pinkham, Sorkin, and Swift, 2020; Borusyak, Hull, and Jaravel, 2022). In our setting, the shares capture the predetermined relevance of competing banks to a focal bank-city market, while the shifts capture innovations in those rivals' relationship-oriented recruitment outside the focal city. Importantly, we interpret the measure as a reduced-form indicator of competitive relationship-hiring pressure, not as a structural instrument for the focal bank's own hiring. Nor does a forecasting residual become mechanically exogenous simply because predictable components have been removed. Identification is therefore assessed through the predetermined exposure structure, the leave-one-city-out construction, the fixed effects in the baseline specification, and the placebo, predictability, alternative-shock, and shock-level inference exercises reported below.

4.2 Baseline Specification and Identifications

We estimate the response of focal-bank vacancy demand to rival relationship-hiring pressure at the bank-city-month level. The baseline specification is

$$Y_{jct} = \beta Z_{jct} + \alpha_{jc} + \gamma_{jt} + \delta_{ct} + \varepsilon_{jct}. \tag{8}$$

Here, Y_{jct} denotes a hiring outcome for focal bank j in city c and month t, and Z_{jct} is the rival relationship-hiring pressure measure constructed above. The preferred outcome is the inverse hyperbolic sine of the number of relationship-banking vacancy postings. We also examine an indicator for any relationship-banking vacancy, the inverse hyperbolic sine of nonrelationship and total vacancy postings, a relationship-minus-nonrelationship hiring gap, and the relationship-banking share among bank-city-months with positive total vacancy posting. These complementary outcomes distinguish the extensive margin, occupation specificity, and composition of the hiring response.

The specification includes bank-by-city fixed effects (α_{jc}), bank-by-month fixed effects (γ_{jt}), and city-by-month fixed effects (δ_{ct}). Bank-by-city fixed effects absorb persistent differences in a bank's local market presence and recruiting intensity. Bank-by-month fixed effects absorb bank-wide monthly recruitment shocks, including national expansions or contractions in the focal bank's

hiring. City-by-month fixed effects absorb local labor-market conditions and other monthly shocks common to banks operating in the same city. The coefficient β is therefore identified from differential exposure to rival innovations across bank-city pairs after absorbing both bank-wide and city-wide monthly variation.

The identifying assumption is deliberately narrower than a claim that the primitive forecast innovations are as-good-as-random. Conditional on the fixed effects and the predetermined network, innovations in a rival bank's relationship-hiring share outside the focal city should not be systematically related to unobserved determinants of the focal bank's relative relationship-hiring demand in that city, except through competitive pressure associated with locally relevant rivals. The predetermined exposure weights and the leave-one-city-out construction remove important mechanical links, while the saturated fixed effects absorb bank-wide and city-wide monthly shocks. Forecast residualization, however, does not by itself make the primitive innovations exogenous. We therefore interpret the design as a quasi-experimental reduced-form exposure design and evaluate the remaining identifying assumptions with locality and timing placebos, alternative shock constructions, and primitive-shock predictability diagnostics in the empirical results and the Online Appendix.

Baseline standard errors are two-way clustered by bank and city, allowing for arbitrary dependence in recruitment outcomes within a bank and within a local market over time. Because is a shift-share-style exposure generated from rival-bank innovations, we also report shock-level equivalent inference following Adão, Kolesár, and Morales (2019) and Borusyak, Hull, and Jaravel (2022), clustering at the rival-bank and calendar-month levels. Wild-cluster bootstrap inference provides an additional finite-sample check in the Online Appendix. These procedures address inference concerns associated with the exposure design without changing the economic estimand of the baseline reduced-form specification.Z_{jct}

5. Results

5.1 Baseline results

Table 2 reports the baseline estimates of the response of focal-bank hiring to rival relationship-hiring pressure. Column (1) uses the inverse hyperbolic sine of relationship-banking vacancy postings as the dependent variable. The coefficient on rival relationship-hiring pressure is 0.9444, with a standard error of 0.1668, and is statistically significant at the 1 percent level. Column (2) examines the extensive margin using an indicator for whether the focal bank posts any

relationship-banking vacancy in a given city-month. The estimated coefficient is 0.6163, with a standard error of 0.1126, and is again significant at the 1 percent level. The two specifications therefore point to the same pattern: greater relationship-hiring pressure from locally relevant rivals is followed by an increase in both the incidence and intensity of the focal bank's relationship-oriented recruitment.

The extensive-margin estimate provides a particularly transparent economic magnitude. The standard deviation of the preferred rival-pressure measure is approximately 0.045. Multiplying this standard deviation by the coefficient in column (2) implies that a one-standard-deviation increase in rival pressure raises the probability that a focal bank posts at least one relationship-banking vacancy in a city-month by approximately 0.028, or 2.8 percentage points. Relative to the unconditional probability of 7.7 percent reported in Table 1, this represents an economically meaningful increase. The extensive-margin result is also useful given the sparsity of relationship hiring documented in the summary statistics: because the dependent variable treats a city-month with one relationship vacancy the same as one with many vacancies, the finding cannot be explained solely by a small number of unusually large posting episodes.

The remaining columns examine whether this response reflects a broad increase in recruitment or is concentrated in relationship-oriented positions. In column (3), the coefficient for nonrelationship vacancies is -0.1850 (s.e. = 0.1088) and is significant only at the 10 percent level. Column (4) shows a positive response in total vacancy postings, with a coefficient of 0.3672 (s.e. = 0.1472), but this aggregate response is considerably smaller than the relationship-hiring response. More directly, column (5) uses the difference between relationship and nonrelationship hiring as the dependent variable. The coefficient is 1.1294 (s.e. = 0.1837), significant at the 1 percent level. Column (6) further shows a significant increase in the relationship-banking share among bank-city-months with positive total hiring. Taken together, these estimates indicate that the response to rival pressure is not a uniform expansion across occupations. Rather, focal banks disproportionately expand vacancy demand for relationship-banking positions, producing a clear shift in the composition of recruitment toward relationship-oriented labor.

Figure 1 examines the timing of this response. We augment the baseline specification with a set of lead and lag measures of rival relationship-hiring pressure. Specifically, we divide the exposure into seven mutually exclusive event-time bins: 7–12 months before the rival innovation, 4–6 months before, 1–3 months before, the contemporaneous month, 1–3 months after, 4–6 months after, and 7–12 months after. Within each multi-month window, we average the corresponding

monthly exposure measures, and all seven terms are included jointly in the regression. The specification retains the same bank-by-city, bank-by-month, and city-by-month fixed effects as the baseline regression, with standard errors two-way clustered by bank and city. Figure 1 plots the resulting coefficients together with 95 percent confidence intervals. Panel A uses relationship-banking hiring as the outcome, while Panel B uses the relationship-minus-nonrelationship hiring gap. Because the construction requires up to twelve months of leads and lags, the dynamic estimation sample contains 110,880 bank-city-month observations.

The lead coefficients provide a timing diagnostic. A coefficient on a future rival innovation asks whether the focal bank's current hiring already moves before that rival innovation occurs, whereas coefficients on lagged innovations trace the persistence of the response after the shock. In Panel A, the relationship-hiring response is concentrated around the month of the rival innovation. The contemporaneous coefficient is 1.0787 (s.e. = 0.1943) and is significant at the 1 percent level. Relationship hiring remains elevated over the following one to three months, with a coefficient of 0.3722 (s.e. = 0.1201), before attenuating at longer horizons. The 4–6 and 7–12 month lag coefficients are small and statistically insignificant.

The pre-shock pattern is generally much weaker than the contemporaneous response, although the 1–3 month lead in Panel A is positive and marginally significant at the 10 percent level. We therefore do not interpret the relationship-level series as exhibiting a perfectly flat pre-trend. Panel B provides a cleaner timing pattern. All three lead coefficients for the relationship-minus-nonrelationship gap are statistically insignificant, whereas the contemporaneous coefficient rises sharply to 1.2208 (s.e. = 0.2322). The subsequent coefficients are statistically indistinguishable from zero. Thus, the most occupation-specific measure of the response emerges sharply in the month of the rival innovation rather than gradually building in the preceding months.

Overall, Table 2 and Figure 1 deliver two complementary findings. The cross-sectional estimates show that banks exposed to stronger rival relationship-hiring pressure expand their own relationship-oriented vacancy demand relative to other occupations. The dynamic estimates show that this differential response is concentrated around the timing of the rival innovation. These patterns are consistent with a rapid strategic adjustment in banks' demand for relationship-oriented human capital.

5.2 Local Relevance and Placebo Evidence

A central feature of our interpretation is that the rival-hiring exposure captures local competitive pressure rather than generic comovement in banks' recruitment activity. This distinction matters because relationship-banking hiring may be correlated across banks for reasons unrelated to local competition. For example, banks may respond simultaneously to nationwide changes in credit demand, regulation, macroeconomic conditions, or aggregate recruitment cycles. Large banks may also expand hiring across many regions at the same time. If the baseline results merely reflected such common bank-industry forces, hiring innovations at geographically irrelevant rivals could predict the focal bank's recruitment almost as strongly as innovations at rivals that were already active in the same local market.

Table 3 provides a direct falsification test of this concern. We decompose rival innovations according to whether the rival had a predetermined presence in the focal city during the 2014–2016 pre-period. The local-rival innovation is constructed from banks that were already present in the focal market, whereas the never-local-rival innovation is constructed from banks without predetermined local presence in that city. Importantly, both variables are built from rival-bank hiring innovations using the same underlying forecasting logic. The key difference is therefore the rival's pre-existing local relevance. If the estimated response reflects local competition for relationship-oriented capacity, the local component should matter while the never-local component should have little predictive power.

The results show a sharp contrast. In column (1), the coefficient on the local-rival innovation is 0.9650 with a standard error of 0.1695 and is significant at the 1 percent level. By comparison, the coefficient on the never-local-rival innovation is -0.0425 with a standard error of 0.0457 and is economically small and statistically insignificant. The same pattern appears for the relative and compositional outcomes. For the relationship-minus-nonrelationship hiring gap, the local-rival coefficient is 1.1294, whereas the never-local coefficient is 0.0254. For the relationship share, the corresponding estimates are 3.1961 and 0.1344. Thus, the focal bank's response is concentrated in innovations originating from rivals with an established local connection.

This comparison addresses an important alternative explanation for the baseline estimates. A nationwide recruitment cycle, a common bank-level expansion shock, or a generic tendency for banks to hire at the same time should not depend sharply on whether a rival happened to have pre-period local presence in the focal city. By contrast, a competition-based interpretation does make this prediction: rival hiring should matter more where the rival was already economically relevant

to the local market. The absence of a comparable response to never-local rivals therefore supports the interpretation of Z_{jct} as a measure of locally relevant competitive pressure, rather than a generic national bank-hiring factor.

At the same time, we view this exercise as a falsification test rather than a complete identification argument. The never-local comparison does not establish that the local-rival innovation is fully exogenous. Its value is that it isolates one specific implication of the proposed mechanism—local relevance should matter—and the data strongly conform to that prediction.

5.3 Robustness and Inference

We next examine whether the relationship-hiring response depends on particular sample periods, geographic definitions, influential rival banks, forecasting choices, or inference procedures. These exercises address several distinct concerns. The baseline estimate could be driven by the unusual recruitment environment surrounding the COVID-19 period, by correlated shocks elsewhere in the same province, by a small number of large or highly exposed rival banks, or by specific choices in constructing the forecast innovation and the predetermined rival network. In addition, because the exposure variable aggregates shocks from multiple rival banks, conventional location-level standard errors may not fully capture dependence induced by common underlying shocks. Figure 2 and Table 4 address these concerns in turn.

Alternative Samples and Shock Definitions. Figure 2 reports the coefficient on rival relationship-hiring pressure and its 95 percent confidence interval across five specifications. The figure deliberately reports the raw coefficient rather than rescaling the estimates by the standard deviation of the exposure, allowing the stability of the estimated slope to be compared directly across specifications. The baseline coefficient is 0.9444. Excluding 2020–2022 yields a coefficient of 0.8276, indicating that the main finding is not driven by the exceptional labor-market and banking conditions surrounding the pandemic.

The second geographic robustness exercise reconstructs the primitive rival innovation using a leave-one-province-out rather than leave-one-city-out procedure. Thus, when measuring a rival bank's hiring innovation relevant to a focal bank-city, we remove all of the rival's vacancy postings in the focal province before constructing the shock. This is a substantially more demanding geographic exclusion because it removes not only the focal city but also other cities potentially exposed to common provincial economic, regulatory, or labor-market conditions. The resulting coefficient is 0.9548, nearly identical to the baseline estimate. This result reduces the concern that

the baseline relationship is generated by broader regional shocks that simultaneously affect rival and focal-bank hiring within the same province.

Figure 2 also examines whether the estimate is dominated by a small set of influential competitors. Excluding large rival banks produces a coefficient of 0.8693, while excluding the ten rivals that account for the largest exposure contributions yields an estimate of 0.8571. Both remain positive and statistically significant. These exercises are particularly useful in a network exposure design, where a small number of large institutions could otherwise contribute disproportionately to measured competitive pressure. The stability of the coefficient after removing such rivals indicates that the result is not attributable to one or a few dominant shock contributors. Online Appendix Figure B1 repeats the same exercises for the relationship-minus-nonrelationship outcome and yields the same qualitative conclusion.

Exposure Construction, Inference, and Timing. Table 4 examines the construction of the exposure measure more systematically. Panel A varies the model used to forecast rival relationship-hiring shares. The preferred AR(1,2,3,12) specification produces a relationship-hiring coefficient of 0.9444. Expanding the forecasting model to include all twelve monthly lags yields 0.9404, while the recursive AR(1,2,3,12) model produces 0.9173. The corresponding relationship-minus-nonrelationship estimates range only from 1.1067 to 1.1294. Thus, the estimated response is largely insensitive to whether the predictable component of rival hiring is modeled using a parsimonious fixed autoregression, a richer lag structure, or a recursively updated forecasting procedure.

Panel B varies how predetermined rival relevance is translated into exposure weights. Our preferred specification weights rivals according to the similarity of their pre-period geographic footprints, but this choice is not essential to the result. Equal weighting produces a coefficient of 1.2718, rival-footprint weighting produces 1.1765, and similarity-by-footprint weighting produces 0.8992, compared with 0.9444 under the preferred similarity weights. All four estimates are positive and statistically significant at the 1 percent level. The finding therefore does not depend on a particular functional form for mapping the predetermined competitive network into rival exposure.

Panel C addresses a different issue: inference in the presence of a shift-share-style exposure structure. The baseline regressions cluster standard errors by focal bank and city, whereas the underlying exposure is generated by innovations at rival banks that can affect many focal

observations simultaneously. We therefore implement the Borusyak, Hull, and Jaravel (2022) shock-level equivalent representation of the reduced-form regression. This transformation leaves the economic coefficient unchanged but permits inference at the level at which the primitive shocks are generated. For relationship hiring, the standard error is 0.1326 when clustered by rival bank and 0.1335 when two-way clustered by rival bank and calendar month. The corresponding two-way standard error for the relationship-minus-nonrelationship outcome is 0.2346. Both coefficients remain statistically significant at the 1 percent level. The shock-level analysis therefore indicates that the baseline statistical significance is not an artifact of treating many focal bank-city observations exposed to common rival shocks as independent. Wild-cluster bootstrap inference reported in the Online Appendix leads to the same conclusion.

Finally, Panel D examines whether future rival innovations predict the current relationship-minus-nonrelationship hiring gap. We conduct this placebo separately under each of the three forecasting models considered in Panel A. Across specifications, coefficients on innovations occurring one to three months and four to six months in the future are small and statistically insignificant. Joint tests of the future terms yield p-values of 0.8266, 0.8012, and 0.8304, respectively. By contrast, the contemporaneous rival innovation remains large and significant in every forecasting specification. This exercise complements the dynamic evidence in Figure 1 by showing that the absence of systematic future-shock predictability is not specific to the preferred forecasting model.

Across these exercises, the estimated relationship-hiring response remains remarkably stable to changes in the sample period, geographic leave-out definition, set of rival banks, forecasting model, and predetermined network weights. The result also remains statistically strong when inference is reorganized around the underlying rival-bank shocks rather than the focal bank-city observations. These checks address different potential sources of bias—regional comovement, influential rivals, model dependence, and the correlation structure induced by common shocks—and none materially alters the main finding. They do not imply that the rival innovations are randomly assigned, but they make it less likely that the estimated response is driven by a particular construction choice or by conventional inference applied to a high-dimensional exposure design.

5.4 Relationship-Capital Content of the Hiring Response

The preceding results show that banks expand relationship-banking vacancies when locally relevant rivals experience unexpected increases in relationship hiring. We next ask whether this

response has the occupational and human-capital content implied by a relationship-capital interpretation. This question matters because a positive relationship-hiring coefficient alone could still reflect a broader expansion in bank recruitment or imperfections in how vacancies are classified. We therefore examine whether the response is specific to relationship-oriented occupations and whether the additional vacancies display characteristics associated with customer-facing, information-intensive banking work.[2]

This interpretation builds on a large literature emphasizing that relationship banking relies on repeated customer interaction, soft information, and the human capital of bank employees (Petersen and Rajan, 1994; Berger and Udell, 1995; Stein, 2002). Loan officers and other customer-facing employees play an important role in acquiring and using borrower-specific information (Hertzberg, Liberti, and Paravisini, 2010; Uchida, Udell, and Yamori, 2012), while related evidence shows that customer relationships can be partly embodied in individual financial professionals and other customer-facing workers (Gurun, Stoffman, and Yonker, 2021; Patault and Lenoir, 2024). These findings motivate our use of occupation and vacancy-text characteristics as observable proxies for the type of relationship-oriented human capital banks seek to acquire.

Occupation specificity. Table 5 first asks whether rival relationship-hiring pressure generates a general recruitment response or one concentrated in relationship-oriented jobs. We re-estimate the baseline specification using the same shock, sample, fixed effects, and clustering structure, but replace the dependent variable with vacancy postings in alternative occupational functions. If the baseline result simply reflected bank-wide expansion, an unobserved local recruitment boom, or a mechanical increase in posting activity, we would expect other occupational categories to respond in the same direction.

The evidence instead shows a pronounced occupational contrast. The coefficient for relationship-banking vacancies is 0.9444, whereas the corresponding coefficients are -0.0699 for IT and data positions, -0.1685 for back-office positions, and -0.0145 for funding positions. The

[2] We also assess whether the relationship-hiring result could reflect platform-specific posting practices or measurement noise. When relationship hiring is measured separately on 51job, Liepin, and Zhaopin, the estimated coefficient is positive and statistically significant on each platform (Internet Appendix Table A8). In a more demanding exercise, we construct the rival innovation using one recruitment platform and measure the focal bank's relationship-hiring outcome on a different platform, excluding same-platform shock-outcome pairs by construction. All six cross-platform estimates are positive and four are statistically significant (Internet Appendix Table A14). Because the rival shock and focal outcome are drawn from different recruitment platforms in these specifications, the results are unlikely to be explained solely by platform-specific posting algorithms, duplicate timing, or source-specific coverage.

latter two categories are economically distinct from relationship banking, while IT and back-office hiring should also respond to a broad organizational expansion if banks were simply increasing recruitment across the board. Formal tests reject equality between the relationship-hiring coefficient and each of the three comparison coefficients at \(p<0.001\). The response is therefore concentrated in the occupation directly related to customer-facing relationship activity rather than reflecting a generic increase in banks' demand for labor.

Relationship-capital characteristics. Table 6 goes one step further and examines what kinds of relationship-banking positions account for the response. We use information in vacancy text and job requirements to identify five observable characteristics that plausibly load more heavily on relationship-oriented human capital: explicit references to **client resources**, prior **relationship experience**, **senior relationship-manager** positions, **local resources**, and requirements of at least **five years of prior experience**. These characteristics are not intended to identify a worker's portable client book. Rather, they capture observable dimensions of vacancies for which customer knowledge, prior banking experience, local familiarity, or relationship-management capability are particularly likely to matter.

Panel A measures the inverse hyperbolic sine of the number of relationship vacancies satisfying each characteristic. Rival relationship-hiring pressure significantly increases vacancies mentioning client resources (0.2933), relationship experience (0.1190), senior relationship-manager responsibilities (0.1279), local resources (0.0331), and at least five years of experience (0.2755). Panel B examines the extensive margin—whether the focal bank posts at least one relationship vacancy with each characteristic—and yields the same qualitative pattern: every coefficient is positive and statistically significant at conventional levels. Thus, the competitive response is visible not only in the aggregate relationship-banking category but also among vacancies whose descriptions explicitly emphasize customer-facing resources, prior relationship expertise, locality, or senior experience.

Taken together, the occupation and vacancy-text evidence sharpens the interpretation of the baseline result. Rival relationship-hiring pressure does not generate a comparable increase across unrelated bank occupations, while the relationship vacancies that do increase include positions explicitly requiring customer resources, prior relationship experience, local resources, senior responsibilities, and substantial experience.

5.5 Bank-Level Consequences

The preceding analysis establishes that banks adjust their relationship-oriented labor demand when locally relevant rivals expand relationship hiring. We next ask whether the same competitive pressure is followed by changes in banks' realized business outcomes. We next ask whether competitive relationship-hiring pressure is also reflected in banks' realized business outcomes. Whereas the vacancy regressions capture adjustments in labor demand, the bank-level analysis provides evidence on whether the same competitive pressure has consequences beyond recruitment. In particular, relationship banking is closely tied to the production of information and the provision of credit to informationally intensive customers, especially small and medium-sized firms (Petersen and Rajan, 1994; Berger and Udell, 1995; Boot, 2000). Consistent with this view, local banks can be particularly difficult to replace as lenders to SMEs when other local banks become distressed (Hasan et al., 2021). Competitive pressure in relationship banking may therefore affect the composition of lending as well as the cost of maintaining the funding base that supports bank activity.

To examine these downstream outcomes, we merge the recruitment-based exposure measure with the China Banking Database (CBD), a hand-collected bank-level database constructed from banks' public information disclosures. CBD organizes standardized bank-year information on Chinese banks and contains detailed modules on bank operations, risk, and governance. The operating module includes balance-sheet and income-statement items, loan and deposit information, and interest income and expense, while the risk module contains measures of credit risk, capital, liquidity, and regulatory indicators. The underlying information is collected from sources such as annual reports, audited financial statements, bank websites, and regulatory disclosure platforms, and is standardized across banks in terms of institutional identity, accounting conventions, and reporting units.

For each bank, we aggregate the monthly city-level rival-pressure measure using predetermined 2014–2016 city weights among supported cities. We require complete monthly exposure within a calendar year and average the resulting monthly bank-level pressure measure over the year. We then relate competitive pressure in year t to changes in bank outcomes from t to t+1. Table 7 standardizes the pressure measure so that each coefficient corresponds to a one-standard-deviation increase in annual rival relationship-hiring pressure. Continuous outcomes and predetermined controls are winsorized at the 1st and 99th percentiles. All specifications include

bank and year fixed effects, control for lagged bank size, nonperforming loans, capital adequacy, loans-to-assets, and deposits-to-assets, and cluster standard errors by bank.

Funding conditions. The clearest financial response appears on the liability side. A one-standard-deviation increase in rival relationship-hiring pressure is associated with an 8.84 basis point increase in the subsequent change in deposit funding cost, with a standard error of 3.59 basis points. Deposit funding cost is measured as deposit interest expense divided by deposits and should therefore be interpreted as an accounting measure of the average cost of deposit funding rather than a contractual deposit rate. The result is consistent with greater competitive pressure in relationship banking being accompanied by more costly deposit funding. The data do not allow us to distinguish whether this reflects higher offered deposit rates, a shift toward more expensive forms of deposits, the loss of lower-cost funding, or other adjustments in deposit composition. We therefore interpret the estimate as evidence of **funding pressure**, rather than as a direct estimate of deposit-price competition.

Lending composition. The credit results point more toward a change in the composition of lending than toward aggregate credit expansion. Total loan growth is essentially unchanged. Enterprise-loan growth is also statistically insignificant. SME-loan growth is positive, while consumer-loan growth is negative, but neither component is individually estimated with sufficient precision to reject zero. The more direct comparison is the SME-minus-consumer growth differential, which rises by 46.33 points in the $100 \times$ log-growth metric and is statistically significant at the 5 percent level. This specification is estimated on 141 bank-year observations from 44 banks, so the magnitude is large and relatively imprecise. We therefore treat it as suggestive evidence that relationship-hiring pressure is associated with a shift in credit activity toward the more relationship-intensive SME segment relative to retail consumer lending, rather than as evidence that SME lending itself increases significantly. This interpretation is consistent with the established role of relationship lending in overcoming informational frictions faced by smaller firms (Petersen and Rajan, 1994; Berger and Udell, 1995).

Other bank outcomes. We find little evidence that the competitive pressure produces a broad expansion or deterioration in bank balance sheets over the following year. Employee growth and next-year log employment are both statistically insignificant. This distinction is informative because the main analysis measures vacancy creation rather than realized hiring: an increase in recruitment demand need not translate one-for-one into the annual stock of bank employees. Likewise, total loan growth shows no significant response, and changes in net interest income

relative to assets and the nonperforming-loan ratio are small and statistically insignificant. The observed consequences therefore appear concentrated in **funding conditions and business composition**, rather than in aggregate lending, employment, profitability, or near-term credit risk.

6. Conclusion

This paper documents a labor-market margin of competition in banking. Using high-frequency vacancy postings and a predetermined map of local rival relevance, we show that banks expand relationship-banking vacancy demand when locally relevant rivals experience positive innovations in relationship-oriented hiring outside the focal city. The response appears on both the intensive and extensive margins, is substantially stronger for relationship banking than for nonrelationship occupations, and is absent for rivals without predetermined local relevance. Vacancy text further shows that the response includes positions requiring client resources, prior relationship experience, local market resources, senior relationship-management responsibility, and substantial prior experience.

Our empirical design combines predetermined local exposure with leave-one-city-out rival hiring innovations and saturated bank-by-city, bank-by-month, and city-by-month fixed effects. Because the underlying innovations are forecast residuals rather than randomized treatments, we do not interpret them as mechanically exogenous. Instead, the identifying case rests on a set of complementary tests of the design's observable implications. The estimated response is concentrated among locally relevant rivals, occurs around the timing of the rival innovation, is stable across alternative forecasting models, geographic leave-out definitions, and network weights, and remains statistically strong under shock-level inference. The evidence points to a rapid, occupation-specific adjustment in banks' local demand for relationship-oriented labor.

These findings broaden the conventional view of relationship banking as an informational and contracting technology. Producing and maintaining relationships also requires human capital. Banks appear to respond to competitive pressure not only through lending terms, branch networks, or product-market choices, but also by adjusting the people and capabilities deployed to originate, serve, and maintain customer relationships. Consistent with this interpretation, the bank-level evidence suggests that stronger competitive relationship-hiring pressure is followed by higher deposit funding costs and, more tentatively, a reallocation of lending growth toward SME credit relative to consumer credit. We view these downstream estimates as suggestive rather than as causal effects of the focal bank's own hiring.

Vacancy data stop short of showing what happens after recruitment demand changes. We observe banks' demand for labor, not whether individual vacancies are filled, where newly hired employees come from, whether workers bring customer relationships with them, or whether clients subsequently move across banks. Linking vacancy demand to matched employer-employee histories, branch-level personnel records, or customer-level banking relationships would make it possible to trace the next step in this process: whether competition for relationship-oriented labor generates realized worker mobility and customer reallocation. Such evidence would help quantify how much of competition in relationship banking operates through the market for human capital.

References

Adão, R., Kolesár, M., & Morales, E. (2019). Shift-share designs: Theory and inference. *Quarterly Journal of Economics, 134*(4), 1949–2010.

Agarwal, S., & Hauswald, R. (2010). Distance and private information in lending. *Review of Financial Studies, 23*(7), 2757–2788.

Agarwal, S., Lin, Y., Zhang, Y., & Zhang, Z. (2024). Labor mobility and loan origination. *Journal of Financial and Quantitative Analysis*, *59*(5), 2099-2132.

Amiti, M., & Weinstein, D. E. (2018). How much do idiosyncratic bank shocks affect investment? Evidence from matched bank-firm loan data. *Journal of Political Economy, 126*(2), 525–587.

Basten, C., & Juelsrud, R. (2026). Cross-selling in bank-household relationships: Mechanisms and implications for pricing. *Review of Financial Studies*, 39(6), 1751–1784.

Beaumont, P., Hebert, C., & Lyonnet, V. (2025). Build or buy? Human capital and corporate diversification. *The Review of Financial Studies*, *38*(5), 1333-1367.

Berger, A. N., Bouwman, C. H. S., Norden, L., Roman, R. A., Udell, G. F., & Wang, T. (2024). Piercing through opacity: Relationships and credit card lending to consumers and small businesses during normal times and the COVID-19 crisis. *Journal of Political Economy*, 132(2), 484–551.

Berger, A. N., Miller, N. H., Petersen, M. A., Rajan, R. G., & Stein, J. C. (2005). Does function follow organizational form? Evidence from the lending practices of large and small banks. *Journal of Financial Economics*, *76*(2), 237-269.

Berger, A. N., & Udell, G. F. (1995). Relationship lending and lines of credit in small firm finance. *Journal of Business, 68*(3), 351–381.

Bharath, S. T., Dahiya, S., Saunders, A., & Srinivasan, A. (2007). So what do I get? The bank's view of lending relationships. *Journal of Financial Economics, 85*(2), 368–419.

Boot, A. W. A. (2000). Relationship banking: What do we know? *Journal of Financial Intermediation, 9*(1), 7–25.

Borusyak, K., Hull, P., & Jaravel, X. (2022). Quasi-experimental shift-share research designs. *Review of Economic Studies, 89*(1), 181–213.

Bulow, J. I., Geanakoplos, J. D., & Klemperer, P. D. (1985). Multimarket oligopoly: Strategic substitutes and complements. *Journal of Political Economy, 93*(3), 488–511.

Canales, R., & Nanda, R. (2012). A darker side to decentralized banks: Market power and credit rationing in SME lending. *Journal of Financial Economics*, *105*(2), 353-366.

Cestone, G., Fumagalli, C., Kramarz, F., & Pica, G. (2024). Exploiting growth opportunities: The role of internal labour markets. *Review of Economic Studies*, *91*(5), 2676-2716.

Chava, S., Ganduri, R., Paradkar, N., & Zeng, L. (2023). Shocked by bank funding shocks: Evidence from consumer credit cards. *Review of Financial Studies*, 36(10), 3906–3952.

Chu, Y., Li, Z., Wei, J., & Wu, W. (2022). A tale of two markets: Labor market mobility and bank information sharing. *Journal of Economic Dynamics and Control*, *141*, 104381.

Degryse, H., & Ongena, S. (2005). Distance, lending relationships, and competition. *Journal of Finance, 60*(1), 231–266.

Gao, J., Wang, W., & Wu, Y. (2024). Human capital portability and careers in finance. *Review of Financial Studies*, 37(9), 2732–2778.

Goldsmith-Pinkham, P., Sorkin, I., & Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review, 110*(8), 2586–2624.

Greenstone, M., Mas, A., & Nguyen, H.-L. (2020). Do credit market shocks affect the real economy? Quasi-experimental evidence from the Great Recession and "normal" economic times. *American Economic Journal: Economic Policy, 12*(1), 200–225.

Gurun, U. G., Stoffman, N., & Yonker, S. E. (2021). Unlocking clients: The importance of relationships in the financial advisory industry. *Journal of Financial Economics, 141*(3), 1218–1243.

Hasan, I., Jackowicz, K., Jagiełło, R., Kowalewski, O., & Kozłowski, Ł. (2021). Local banks as difficult-to-replace SME lenders: Evidence from bank corrective programs. *Journal of Banking & Finance*, *123*, 106029.

Hasan, I., Shen, Y., & Yuan, X. (2021). Local product market competition and bank loans. *Journal of Corporate Finance*, *70*, 102054.

Heider, F., & Inderst, R. (2012). Loan prospecting. *Review of Financial Studies, 25*(8), 2381–2415.

Hertzberg, A., Liberti, J. M., & Paravisini, D. (2010). Information and incentives inside the firm: Evidence from loan officer rotation. *Journal of Finance, 65*(3), 795–828.

Huang, R., Linck, J. S., Mayer, E. J., & Parsons, C. A. (2026). Can human capital explain income-based disparities in financial services?. *The Review of Financial Studies*, *39*(6), 1785-1822.

Jeffers, J. S. (2024). The impact of restricting labor mobility on corporate investment and entrepreneurship. *Review of Financial Studies*, 37(1), 1–44.

Jiang, E. X. (2023). Financing competitors: Shadow banks' funding and mortgage market competition. *Review of Financial Studies*, 36(10), 3861–3905.

Karolyi, S. A. (2018). Personal lending relationships. *Journal of Finance, 73*(1), 5–49.

Liberti, J. M., & Mian, A. R. (2009). Estimating the effect of hierarchies on information use. *Review of Financial Studies, 22*(10), 4057–4090.

Liberti, J. M., & Petersen, M. A. (2019). Information: Hard and soft. *Review of Corporate Finance Studies*, *8*(1), 1-41.

Patault, B., & Lenoir, C. (2024). Customer capital spillovers: Evidence from sales managers in international markets. *American Economic Journal: Applied Economics, 16*(4), 404–443.

Petersen, M. A., & Rajan, R. G. (1994). The benefits of lending relationships: Evidence from small business data. *Journal of Finance, 49*(1), 3–37.

Qian, J., Strahan, P. E., & Yang, Z. (2015). The impact of incentives and communication costs on information production and use: Evidence from bank lending. *The Journal of Finance*, *70*(4), 1457-1493.

Schneider, T., Strahan, P. E., & Yang, J. (2024). Syndicated lending, competition, and relative performance evaluation. *Review of Financial Studies*, 37(12), 3802–3834.

Skrastins, J., & Vig, V. (2019). How organizational hierarchy affects information production. *The Review of Financial Studies*, *32*(2), 564-604.

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

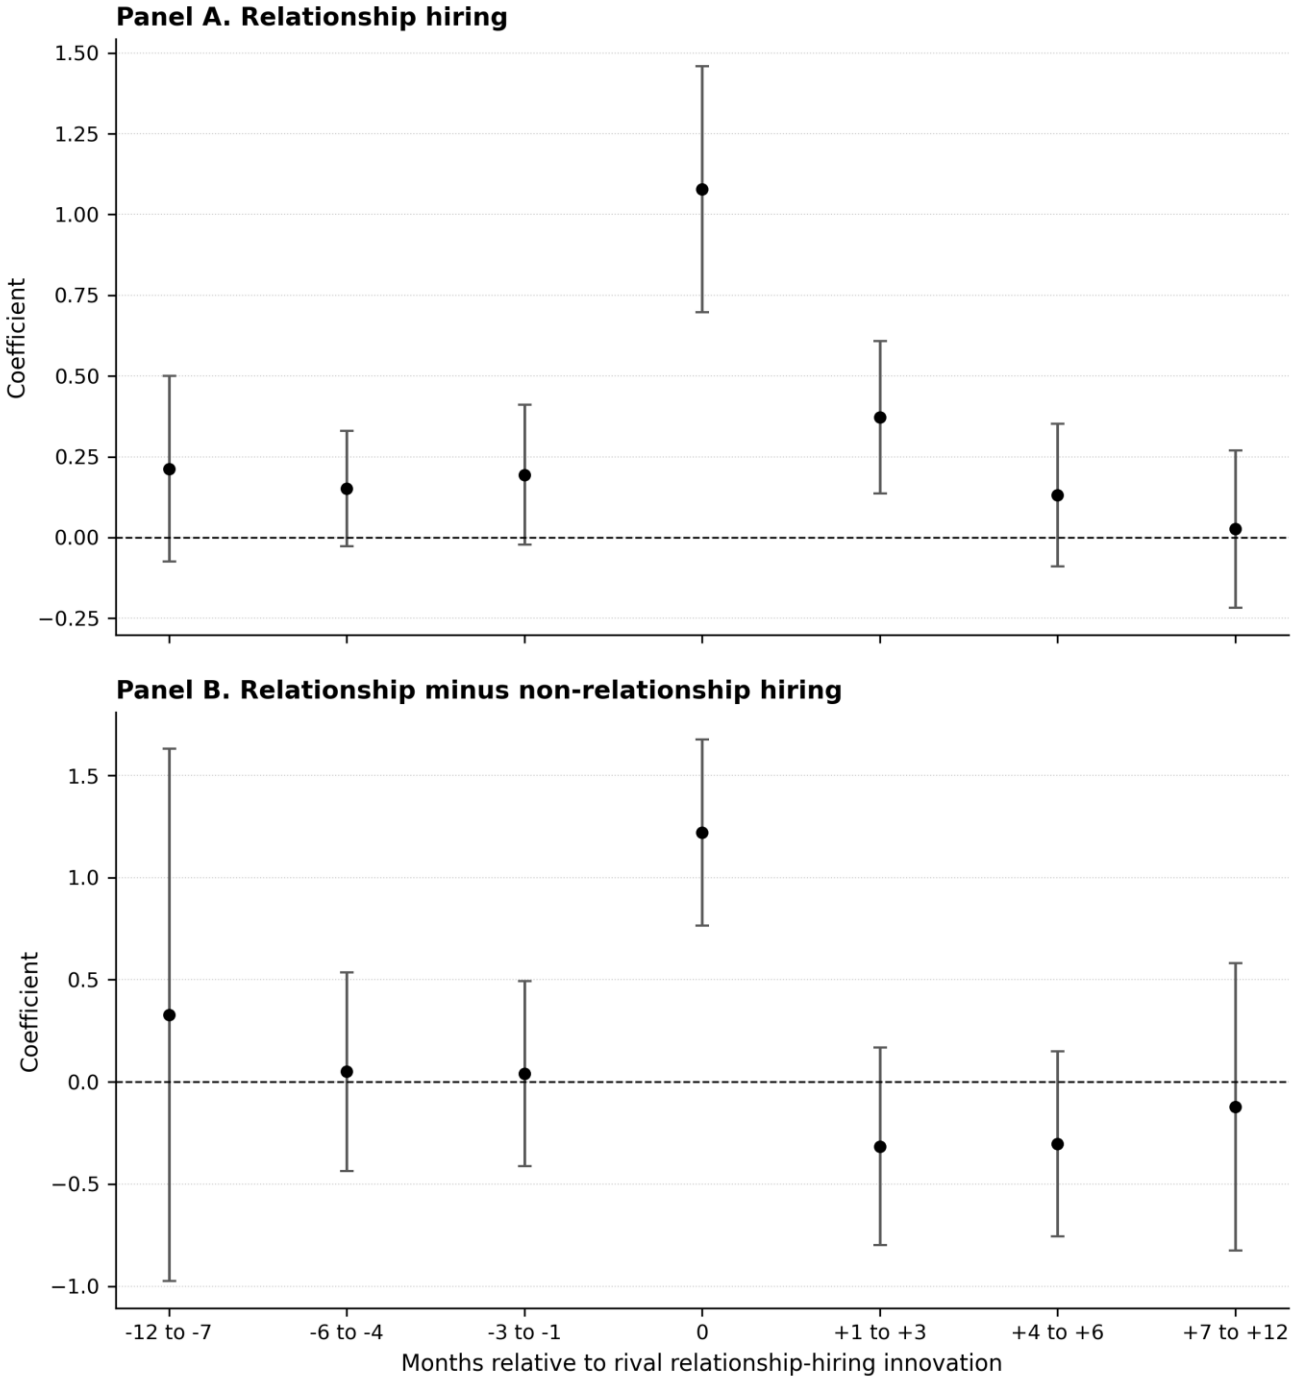
Stoyanov, A., & Zubanov, N. (2012). Productivity spillovers across firms through worker mobility. *American Economic Journal: Applied Economics, 4*(2), 168–198.

Tate, G., & Yang, L. (2024). The human factor in acquisitions: Cross-industry labor mobility and corporate diversification. *The Review of Financial Studies*, *37*(1), 45-88.

Uchida, H., Udell, G. F., & Yamori, N. (2012). Loan officers and relationship lending to SMEs. *Journal of Financial Intermediation, 21*(1), 97–122.

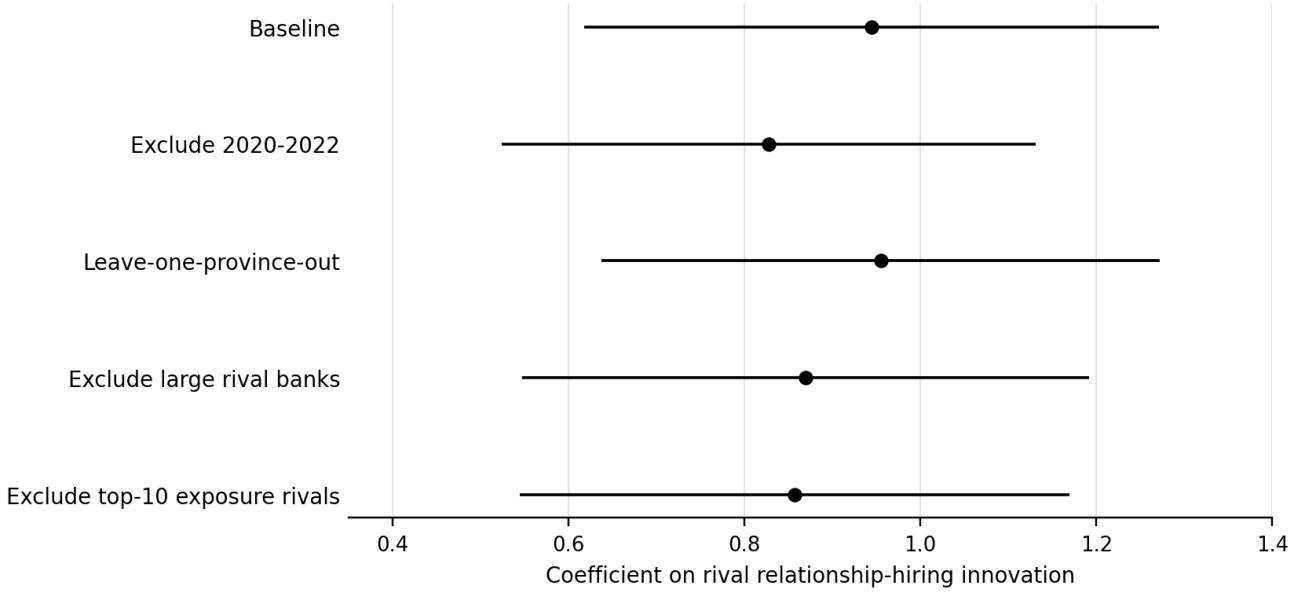
Wang, T. (2019). To build or to buy? The role of local information in credit market development. *Management Science*, *65*(12), 5838-5860.

Figure 1. Dynamic response to unexpected rival relationship-hiring expansion



Notes: The figure plots coefficients and 95% confidence intervals from the dynamic specification. The shock is the similarity-weighted leave-one-city-out rival relationship-hiring forecast residual. The forecasting model is estimated in the 2014-2016 pre-period. Regressions absorb bank x city, bank x month, and city x month fixed effects; standard errors are two-way clustered by city and bank.

Figure 2. Robustness of the relationship-hiring response



Notes: The figure plots raw coefficient estimates and 95% confidence intervals. The baseline uses the fixed pre-period AR(1,2,3,12) forecast residual and predetermined similarity weights. Alternative specifications exclude 2020-2022, rebuild the shock leave-one-province-out, exclude large rival banks, or exclude the top-10 exposure rivals. All regressions absorb bank x city, bank x month, and city x month fixed effects; standard errors are two-way clustered by city and bank. The large-focal-bank exclusion is treated as heterogeneity rather than a robustness specification and is not plotted.

Table 1. Summary statistics

Panel A. Full baseline estimation sample

	N	Mean	SD	Min	P25	Median	P75	Max
All vacancy postings	147,840	2.340	11.459	0.000	0.000	0.000	1.000	1,040.000
Relationship-banking vacancy postings	147,840	0.174	1.012	0.000	0.000	0.000	0.000	56.000
Non-relationship vacancy postings	147,840	2.166	11.110	0.000	0.000	0.000	1.000	1,030.000
Any vacancy posting	147,840	0.324	0.468	0.000	0.000	0.000	1.000	1.000
Any relationship-banking vacancy posting	147,840	0.077	0.266	0.000	0.000	0.000	0.000	1.000
Any non-relationship vacancy posting	147,840	0.309	0.462	0.000	0.000	0.000	1.000	1.000
Relationship-banking share	47,913	0.104	0.245	0.000	0.000	0.000	0.000	1.000
Rival relationship-hiring innovation exposure	147,840	-0.016	0.045	-0.383	-0.039	-0.019	0.002	0.986
Pre-period bank city footprint	147,840	54.103	40.447	2.000	14.000	53.000	85.000	139.000
Pre-period bank vacancy postings	147,840	3,398.652	4,643.572	6.000	440.000	1,876.000	3,594.000	17,712.000
Large-footprint focal bank	147,840	0.649	0.477	0.000	0.000	1.000	1.000	1.000

Panel B. Posting intensity conditional on positive activity

	N	Mean	SD	Min	P25	Median	P75	Max
All vacancy postings, conditional on >0	47,913	7.220	19.234	1.000	1.000	3.000	6.000	1,040.000
Relationship-banking postings, conditional on >0	11,369	2.265	2.928	1.000	1.000	1.000	2.000	56.000
Non-relationship postings, conditional on >0	45,649	7.014	19.125	1.000	1.000	2.000	6.000	1,030.000

Notes: The unit of observation is bank-city-month, 2017-2024, restricted to the preferred baseline relationship-hiring estimation sample after singleton removal. Panel A reports unconditional statistics. Relationship-banking share is defined only for bank-city-months with positive total vacancy postings; the separate share regression in Table 2 can have fewer observations after its own fixed-effect singleton removal. Panel B reports posting intensity conditional on the corresponding posting-count variable being strictly positive. Vacancy variables measure vacancy postings rather than realized worker flows or consistently observed headcounts.

Table 2. Unexpected rival relationship-hiring expansion and local hiring demand

	(1)	(2)	(3)	(4)	(5)	(6)
	Relationship (asinh)	**Relationship (any)**	**Non-relationship (asinh)**	**All hiring (asinh)**	**Relationship - non-relationship**	**Relationship share**
Rival relationship-hiring innovation	0.9444***	0.6163***	-0.1850*	0.3672**	1.1294***	3.1836***
	(0.1668)	(0.1126)	(0.1088)	(0.1472)	(0.1837)	(0.4714)
Observations	147,840	147,840	147,840	147,840	147,840	42,085
R-squared	0.3759	0.3457	0.5949	0.5938	0.5491	0.5179

Notes: The sample covers 2017-2024 bank-city-month observations in the preferred baseline estimation sample after singleton removal. The baseline shock uses the fixed pre-period AR(1,2,3,12) forecasting model and predetermined geographic-similarity weights. Hiring outcomes are constructed from vacancy postings rather than realized hires. All specifications absorb bank x city, bank x month, and city x month fixed effects. Standard errors, in parentheses, are two-way clustered by city and bank. The relationship-share specification is estimated only when total vacancy postings are positive. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table 3. Local competitive network versus never-local rival innovations

	(1)	(2)	(3)
	Relationship hiring	**Relationship - non-relationship**	**Relationship share**
Local-rival innovation	0.9650*** (0.1695)	1.1294*** (0.1901)	3.1961*** (0.4749)
Never-local-rival innovation	-0.0425 (0.0457)	0.0254 (0.0885)	0.1344 (0.1458)
Observations	145,536	145,536	41,983
R-squared	0.3742	0.5484	0.5175

Notes: The local-rival shock uses only rivals with predetermined local presence. The never-local measure is constructed from rival banks without predetermined local presence and serves as a placebo. All specifications use the baseline fixed effects and two-way clustering by city and bank described in Table 2. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table 4. Identification and robustness of the relationship-hiring response

Panel A. Forecast-model family

	AR(1,2,3,12)	AR(1,...,12)	Recursive AR(1,2,3,12)
Relationship hiring	0.9444*** (0.1668)	0.9404*** (0.1658)	0.9173*** (0.1610)
Relationship - non-relationship	1.1294*** (0.1837)	1.1090*** (0.1801)	1.1067*** (0.1835)

Panel B. Predetermined network-weight family

	Equal weight	Similarity	Footprint	Similarity x footprint
Relationship hiring	1.2718*** (0.2463)	0.9444*** (0.1668)	1.1765*** (0.2194)	0.8992*** (0.1550)
Observations	147,840	147,840	147,840	147,840
R-squared	0.3769	0.3759	0.3766	0.3757

Panel C. BHJ shock-level equivalent inference

	Heteroskedasticity robust	Cluster rival bank	Two-way: rival bank and month
Relationship hiring	0.9444*** (0.0604)	0.9444*** (0.1326)	0.9444*** (0.1335)
Relationship - non-relationship	1.1294*** (0.1009)	1.1294*** (0.2399)	1.1294*** (0.2346)
Rival-bank clusters	150	150	150
Month clusters	96	96	96

Panel D. Future-shock placebo across forecast models

	AR(1,2,3,12)	AR(1,...,12)	Recursive
Future 1-3 months	0.0699 (0.2482)	0.0089 (0.2399)	0.0959 (0.3076)
Future 4-6 months	0.1423 (0.2386)	0.1268 (0.2319)	0.1538 (0.2692)
Current shock	1.1191*** (0.1992)	1.0836*** (0.2009)	1.1219*** (0.1954)
Joint p-value: future shocks	0.8266	0.8012	0.8304

Notes: Panels A and B use the location-level specification in Table 2. Panel C reports the Borusyak-Hull-Jaravel shock-level equivalent representation of the shift-share reduced form; it is used for inference and is not a separate structural IV estimand. The preferred shock-level standard errors cluster by rival bank and calendar month. Panel D reports future-shock placebos for the relationship-minus-non-relationship outcome. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table 5. Occupation-specific responses to rival relationship-hiring innovations

	(1) Relationship	(2) IT / Data	(3) Back office	(4) Funding
Rival relationship-hiring innovation	0.9444***	-0.0699	-0.1685**	-0.0145
	(0.1668)	(0.0425)	(0.0699)	(0.0229)
Observations	147,840	147,840	147,840	147,840
R-squared	0.3759	0.4908	0.5117	0.5752
Equality-test p-value	—	0.0005	0.0088	0.0100

Notes: Each column uses the same preferred rival relationship-hiring shock, but a different vacancy-function outcome. Formal occupation-interacted tests reject equality of the relationship coefficient with the IT/Data, back-office, and funding coefficients ($p = 0.0005$, 0.0088, and 0.0100, respectively). Fixed effects and clustering follow Table 2. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table 6. Relationship-capital characteristics of the hiring response

Panel A. Intensive margin: asinh number of flagged relationship vacancies

	(1) Client resources	(2) Relationship experience	(3) Senior RM	(4) Local resources	(5) 5+ years experience
Rival relationship-hiring innovation	0.2933***	0.1190***	0.1279***	0.0331**	0.2755***
	(0.0594)	(0.0289)	(0.0245)	(0.0136)	(0.0510)
Observations	147,840	147,840	147,840	147,840	147,840
R-squared	0.2611	0.2616	0.2647	0.2194	0.2695

Panel B. Extensive margin: any flagged relationship vacancy

	(1) Client resources	(2) Relationship experience	(3) Senior RM	(4) Local resources	(5) 5+ years experience
Rival relationship-hiring innovation	0.2207***	0.1000***	0.1079***	0.0285**	0.2238***
	(0.0456)	(0.0233)	(0.0205)	(0.0112)	(0.0432)
Observations	147,840	147,840	147,840	147,840	147,840
R-squared	0.2543	0.2514	0.2330	0.2165	0.2589

Notes: Outcomes are constructed from relationship-banking vacancy text. Panel A uses the inverse-hyperbolic-sine number of flagged vacancies; Panel B uses an indicator for whether at least one such vacancy is posted. These regressions capture absolute demand for relationship-capital-related positions, not their share within relationship hiring. Fixed effects and clustering follow Table 2. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table 7. Bank-level consequences of rival relationship-hiring pressure

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Employee growth	**Log employment t+1**	**Total loan growth**	**Enterprise loan growth**	**SME loan growth**	**Consumer loan growth**	**SME - consumer growth**	**Change in deposit funding cost (bp)**	**Change in NII/assets (pp)**	**Change in NPL**
Rival relationship-hiring pressure	1.4963	0.0097	-0.2532	-0.0874	12.3434	-7.6703	46.3304**	8.8449**	-0.0091	-0.0283
	(1.2708)	(0.0082)	(0.8681)	(1.4739)	(12.7743)	(8.3557)	(22.9533)	(3.5934)	(0.0330)	(0.0407)
Observations	503	526	673	600	169	417	141	608	658	648
R-squared	0.4406	0.9978	0.5417	0.4376	0.3111	0.3364	0.3974	0.2180	0.2163	0.2905
Banks	112	117	142	133	52	97	44	133	141	139

Notes: All continuous outcomes and lagged controls are winsorized at the 1st and 99th percentiles. All specifications include bank and year fixed effects; standard errors are clustered by bank. Rival relationship-hiring pressure is standardized to one standard deviation. Growth outcomes are 100 times one-year log changes; Column (2) reports log employment at t+1. Enterprise and SME lending are relationship-oriented outcomes, while consumer lending is the retail benchmark; the SME-minus-consumer outcome is the formal relationship-versus-retail comparison. Deposit funding cost is deposit interest expense divided by deposits, and Column (8) reports its one-year change in basis points. Column (9) reports the one-year change in NII/assets in percentage points, and Column (10) reports the one-year change in the reported NPL ratio. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Appendix A. Detailed Construction and Validation of Bank Vacancy Measures

A.1. Vacancy Data and Sample Construction

Our vacancy data are obtained from CnOpenData, which consolidates online job postings from major Chinese recruitment platforms, including 51job, Zhaopin, and Liepin. The raw records contain information on the employer, job title, job description, salary, required work experience, education, location, and posting date. Our sample covers 2014–2024.

We first standardize employer names and match recruiting entities to our bank master list. We retain only vacancies posted directly by banks or their branches. Vacancies posted by staffing agencies, labor-dispatch companies, outsourcing companies, and other third-party recruiters are excluded. This restriction is important because the identity of the ultimate employing bank and the location of the position are substantially less reliable for third-party postings.

We standardize city names and map each vacancy to a bank-city-month. Observations without a valid bank, city, posting month, or job title are excluded. We further remove exact duplicate records and repeated scrapes of the same vacancy. The resulting direct-bank sample contains 511,084 vacancies from 309 banks operating across 444 cities.

Our empirical analysis is therefore based on recruiting decisions made directly by banks rather than vacancies posted on their behalf by external intermediaries.

A.2. Text Processing and Functional Classification

For each vacancy, we use the job title as the primary source for identifying the function of the position and use the job description and qualification requirements to resolve ambiguous titles. Before classification, we normalize punctuation, whitespace, English capitalization, and common textual variants.

We classify vacancies into mutually interpretable functional groups, including relationship-oriented business positions, credit approval, risk and compliance, funding and financial markets, information technology and data, operations and back-office functions, retail service positions, wealth and private banking, and other positions.

The distinction between relationship-oriented positions and other bank jobs is central to the analysis. Conceptually, a relationship-oriented employee is one whose primary function involves acquiring, developing, or maintaining external customer relationships. Such employees may embody relationship capital because part of their productivity depends on knowledge of customers, established client connections, and the ability to originate new business.

Importantly, the classification is not based simply on whether the word "customer" (客户) appears in a vacancy. Customer-service, teller, operational, risk-control, and administrative positions may interact with customers without performing an economically meaningful relationship-origination function. We therefore employ a deliberately conservative classification.

A.3. Identification of Relationship-Oriented Vacancies

Our baseline measure focuses on positions associated with corporate, SME, inclusive-finance, and institutional customer relationships.

A vacancy is classified as a **core relationship vacancy** when either (i) its job title directly identifies a relationship-management position, or (ii) the job description jointly indicates a relevant banking business line and responsibilities involving customer acquisition, marketing, development, or relationship maintenance.

The principal Chinese terms used in the classification are reported below. The table reports representative word stems and expressions; the actual search also accommodates common textual variants.

Table A1. Dictionary for Relationship-Capital Measures

Construct	Representative terms
Direct relationship-job titles	对公客户经理; 公司客户经理; 企业客户经理; 小微客户经理; 普惠客户经理; 机构客户经理; 三农客户经理; 科技金融客户经理; corporate relationship manager; corporate banking manager; MNC relationship manager; business banking relationship manager; corporate coverage
Corporate/relationship business lines	对公业务; 公司业务; 公司金融; 企业金融; 企业信贷; 小微金融; 普惠金融; 供应链金融; 贸易融资; 项目融资; 科技金融; 三农金融
Relationship-building activities	客户开发; 客户拓展; 客户维护; 客户营销; 客户关系; 客户关系维护; 客户拜访; 业务拓展; 市场拓展; 营销拓展; 渠道拓展
Broader relationship-job terms	客户经理; account manager; relationship manager; retail relationship manager; Premier Relationship Manager; 财富经理; 理财经理; 私人银行; wealth manager; private banker
Existing client/network resources	客户资源; 企业客户资源; 机构客户资源; 政府资源; 人脉资源; 社会资源; 营销资源; 渠道资源; 丰富客户资源; 有客户资源者优先; 有资源者优先; 自带资源及流量; 业务网络
Prior banking/client experience	银行工作经验; 银行从业经验; 同业经验; 金融机构工作经历; 信贷/贷款经验; 对公/公司业务经验; 小微/普惠经验; 客户营销经验; wealth management experience; commercial/corporate banking experience; client-facing sales/marketing experience

The baseline definition gives greater weight to precise job titles. For example, “对公客户经理” is classified directly as a relationship vacancy. By contrast, a generic title such as “业务经理” is classified as relationship-oriented only when the vacancy description also indicates both a relevant banking business line and client-development responsibilities.

Terms referring only to a banking product or business area are not sufficient by themselves. For instance, a vacancy mentioning “供应链金融” is classified as a relationship vacancy only when the employee is expected to originate, develop, market to, or maintain customers rather than perform risk control, product design, approval, operations, or technical support.

Likewise, positions whose primary functions are credit approval, risk management, compliance, auditing, legal affairs, accounting, human resources, administration, information technology, data management, operations, teller services, or general customer service are not classified as relationship vacancies solely because their descriptions mention customers.

We separately construct a **broader relationship measure** that additionally includes generic customer-manager positions and wealth/private-banking positions. We use this broader definition in robustness tests rather than in the baseline specification.

A.4. Explicit Demand for Relationship Capital

The relationship-job classification captures the nature of the position. We construct a separate measure to capture whether the bank explicitly seeks workers who already possess potentially transferable relationship capital.

For each vacancy, we search the qualification and job-description fields for expressions indicating that the applicant already possesses customer, network, marketing, channel, institutional, government, or other business resources. When a posting contains an identifiable qualification section, that section receives priority in the resource classification. Examples include “客户资源,” “企业客户资源,” “机构客户资源,” “人脉资源,” “社会资源,” “渠道资源,” “政府资源,” “丰富客户资源,” “有客户资源者优先,” and equivalent English expressions.

We define:

ClientResource_j = 1

when vacancy (j) explicitly values such pre-existing resources.

The coding is intentionally conservative. Negated expressions such as “无需客户资源” or “不要求客户资源” are coded as zero. We also distinguish possession of pre-existing resources

from job duties that ask the employee to build resources after hiring. Thus, phrases such as “开发客户资源,” “拓展新的客户资源,” “积累优质客户资源,” or “维护客户资源” do not, by themselves, imply that the applicant brings portable relationship capital. Likewise, resources supplied by the employer, such as a bank-provided customer list, are not coded as applicant resources.

This measure is intentionally distinct from the relationship-job indicator. A bank may hire a corporate relationship manager without explicitly requiring the applicant to bring an existing client network. Conversely, an explicit client-resource requirement provides particularly direct evidence that the bank values relationship capital embodied in the worker.

We do not interpret this variable as evidence that an employee literally transfers a book of customers from a previous employer. Rather, it captures employers' revealed demand for workers possessing established networks and customer relationships that may be valuable across employers.

A.5. Prior Banking and Client-Marketing Experience

We additionally identify vacancies that explicitly require or prefer prior banking, financial-institution, credit/lending, corporate-banking, financial-market or wealth-management, or relevant client-facing sales/marketing experience. Representative expressions include “银行工作经验,” “银行从业经验,” “同业经验,” “金融机构工作经历,” “信贷经验,” “对公业务经验,” “客户营销经验,” “商业银行工作经验,” “小微/普惠业务经验,” “wealth management experience,” “commercial banking experience,” “corporate banking experience,” and “sales/marketing experience” when the surrounding text makes the relevant client-facing context clear.

These requirements provide a complementary measure of the type of human capital banks seek when competition for relationship-intensive employees increases. Generic statements about work experience are not sufficient unless the required background is relevant to banking, finance, credit/lending, or client-facing business development. Expressions such as “经验不限,” “有无经验均可,” or equivalent language are coded as zero. Unlike the client-resource variable, however, prior banking/client experience may reflect industry-specific skills as well as relationship capital. We therefore treat the experience requirement as a separate outcome rather than equating it with relationship capital itself.

A.6. Human Validation

We validate the four text-based measures using a two-stage procedure that separates dictionary development from out-of-sample evaluation. The pilot observations are used only

to refine the written coding rules and classifier. The subsequent 900-observation holdout is drawn after the classifier is frozen and excludes every pilot observation.

Pilot coding and classifier freezing

We first construct a 300-vacancy pilot sample with equal representation from three diagnostic strata: 100 algorithm-positive vacancies, 100 hard negatives, and 100 clean negatives. A hard negative is a vacancy not classified as a core relationship position but containing plausibly relevant language (for example, terms related to customers, marketing, business development, credit, corporate banking, or similar activities). Clean negatives are the remaining algorithm-negative vacancies. The pilot is manually reviewed under a written protocol.

The pilot is used to identify ambiguous Chinese and English titles, overly broad keywords, negations, support and product-management positions, and the distinction between possessing pre-existing customer resources and being asked to develop customers after hiring. Based on this exercise, we finalize the classification rules for the core relationship, broad relationship, explicit client-resource, and prior banking/client-experience measures. The resulting classifier is then frozen. No pilot observation is eligible for the formal holdout sample.

Independent 900-vacancy holdout

Using the frozen classifier, we rebuild the vacancy corpus, remove the pilot observations, and partition the remaining vacancies into three mutually exclusive strata. We then draw 300 observations without replacement from each stratum using a fixed random seed, producing a 900-vacancy holdout sample:

1. Algorithm positive: 300 vacancies classified as core relationship positions by the frozen baseline algorithm.

2. Hard negative: 300 vacancies with rel_core = 0 but containing plausibly relationship-relevant language.

3. Clean negative: 300 vacancies with rel_core = 0 that do not meet the hard-negative screen.

The 900 observations are randomly reordered and assigned anonymous validation IDs. The algorithmic labels, validation stratum, and sampling weights are stored in a separate secret key and are not included in the initial coding file. The reviewer initially sees only the job title, job description, and qualification requirements; treatment/exposure variables and regression results are not part of the coding materials.

Manual coding protocol

For each vacancy, the reviewer records the following variables:

1. human_relationship = 1 if the employee's primary function involves acquiring, developing, marketing to, or maintaining external client relationships.

2. human_corporate_relationship = 1 if the relationship role primarily serves corporate, SME, inclusive-finance, institutional, or similar business clients.

3. human_client_resource = 1 only when the posting explicitly requires or prefers pre-existing customer, social, network, marketing, channel, institutional, government, or comparable business resources.

4. human_bank_experience = 1 when the posting explicitly requires or prefers relevant prior banking, financial-industry, credit/lending, wealth/financial-market, or client-facing sales/marketing experience.

5. human_function assigns the vacancy to one primary category: corporate relationship, retail relationship, wealth/private banking, risk/credit, operations, IT/data, administration, or other.

The core algorithm is benchmarked against human_corporate_relationship; the broad relationship algorithm against human_relationship; the client-resource algorithm against human_client_resource; and the experience algorithm against human_bank_experience.

Independent coding, reconciliation, and consistency checks

Each vacancy in the holdout sample is independently labeled by two human coders using the same written coding protocol. The coders separately review the job title, job description, and qualification requirements and assign the relevant binary indicators and primary functional category without observing the algorithmic classification.

After the independent coding is completed, the two sets of manual labels are compared and reconciled by the senior researcher to produce a single unified human classification for each vacancy. Cases in which the two coders reach different judgments are reviewed against the full vacancy text and the written coding rules. Particular attention is given to distinctions that require greater judgment, such as separating core corporate-relationship positions from adjacent product, management, operational, or support roles and identifying explicit requirements for prior banking or client-marketing experience.

Before computing the validation statistics, we conduct an additional internal consistency check across the binary indicators and the primary functional classifications. Any remaining inconsistencies are reviewed using the same coding protocol. The resulting unified human labels constitute the final manual benchmark against which the algorithmic classifications are evaluated.

Importantly, the classification algorithm is frozen before the holdout validation exercise. Neither the dictionary nor the classification rules are modified using information from the holdout sample.

A.7. Validation Statistics

We compare the frozen algorithmic classifications with the final unified human benchmark. For each measure, we report precision, recall, specificity, accuracy, the F1 score, and Cohen's kappa. Precision is the fraction of algorithm-positive observations confirmed by the human benchmark; recall is the fraction of human-positive observations recovered by the algorithm; specificity is the fraction of human-negative observations correctly rejected; and the F1 score

is the harmonic mean of precision and recall. Cohen's kappa summarizes agreement between the algorithm and the human benchmark after accounting for agreement expected by chance.

Table A2 reports the unweighted diagnostic performance on the 900-vacancy holdout sample. Because the holdout deliberately contains equal numbers of algorithm-positive observations, hard negatives, and clean negatives, these statistics are designed to provide a demanding assessment of classification performance rather than to reproduce the population prevalence of each vacancy type. The validation data also retain stratum-specific sampling weights, defined as the population size of each validation stratum divided by the number of sampled observations in that stratum, which we use for population-reweighted sensitivity analyses.

Table A2. Human Validation of Vacancy Classification

Measure	Precision	Recall	Specificity	Accuracy	F1 Score	Cohen's κ	Validation N
Core relationship	95.3%	91.1%	97.6%	95.3%	93.2%	0.896	900
Broad relationship	92.4%	91.5%	89.2%	90.6%	91.9%	0.806	900
Client-resource requirement	98.3%	83.6%	99.7%	97.2%	90.3%	0.887	900
Bank/client experience	97.0%	89.6%	95.8%	92.0%	93.1%	0.836	900

Notes: The classifier is fixed before the holdout is drawn. The table uses the final adjudicated benchmark. The holdout contains 300 algorithm-positive, 300 hard-negative, and 300 clean-negative vacancies. Cohen's κ measures agreement beyond chance.

The validation results indicate strong agreement between the text classifier and the adjudicated benchmark. For the baseline core relationship measure, precision is 95.3%, recall is 91.1%, and the F1 score is 93.2%. The broader relationship measure has an F1 score of 91.9%. The explicit client-resource indicator is especially conservative: 98.3% of algorithm-positive observations are confirmed by the benchmark, while recall is 83.6%. The prior banking/client-experience measure achieves 97.0% precision, 89.6% recall, and an F1 score of 93.1%. Cohen's kappa ranges from 0.806 to 0.896 across the four measures.

Together, these checks show that the baseline classifier identifies high-confidence corporate relationship hiring rather than generic customer-facing or support positions, while the client-resource measure isolates explicit demand for pre-existing networks rather than incidental customer-related vocabulary. The broader relationship and experience measures provide complementary evidence using less restrictive definitions.

Appendix B. Additional Tables and Figures

Table B1. Shift-share shock concentration diagnostics

Diagnostic	Preferred design
Actual rival banks	150
Effective rival banks	27.15
Top-1 rival-bank weight share	8.35%
Top-5 rival-bank weight share	32.55%
Top-10 rival-bank weight share	54.78%
Actual rival-bank x excluded-city shock pairs	1,587
Effective shock pairs	1,527.43
Shock-time cells	152,352
Effective shock-time cells	146,633.65

Notes: The table summarizes concentration in the predetermined normalized similarity weights used in the preferred shift-share design. Effective counts are inverse-Herfindahl quantities based on the relevant normalized weights.

Table B2. Predictability of the primitive rival-bank shock

	Bank pre-characteristics	Excluded-city pre-characteristics	Combined
Rival-bank pre-period relationship share	-0.0025 (0.0076)		-0.0054 (0.0075)
Rival-bank pre-period total hiring (log)	-0.0077 (0.0161)		-0.0094 (0.0157)
Rival-bank pre-period number of cities (log)	-0.0062 (0.0143)		-0.0048 (0.0141)
Rival bank x excluded-city pre-period relationship share		0.0010 (0.0043)	0.0012 (0.0041)
Rival bank x excluded-city pre-period total hiring (log)		-0.0044 (0.0028)	-0.0008 (0.0028)
Rival bank x excluded-city pre-period relationship hiring (log)		0.0050 (0.0061)	0.0034 (0.0049)
Shock observations	152,352	152,352	152,352
R-squared	0.0162	0.0086	0.0168
Joint p-value	0.0504	0.3021	0.0010

Notes: The dependent variable is the baseline primitive leave-one-city-out rival-bank forecast residual. All specifications absorb calendar-month fixed effects and cluster standard errors at the rival-bank level. The joint p-value tests the displayed predetermined characteristics in each column.

Table B3. Serial predictability of forecast residuals

	AR(1,2,3,12)	AR(1,...,12)	Recursive
Lag 1	0.0747*** (0.0123)	0.0854*** (0.0124)	-0.0065 (0.0134)
Lag 2	0.1035*** (0.0182)	0.1092*** (0.0178)	0.0205 (0.0201)
Lag 3	0.0627*** (0.0120)	0.0764*** (0.0119)	-0.0156 (0.0125)
Lag 4	0.0865*** (0.0151)	0.1359*** (0.0153)	0.0814*** (0.0153)
Lag 5	0.0786*** (0.0163)	0.0864*** (0.0163)	0.0828*** (0.0155)
Lag 6	0.0809*** (0.0146)	0.1015*** (0.0142)	0.0804*** (0.0150)
Lag 12	0.1653*** (0.0187)	0.1491*** (0.0193)	0.0241 (0.0195)
Shock observations	133,308	133,308	133,308
R-squared	0.1341	0.2068	0.0301
Joint p-value: all lags	0.0000	0.0000	0.0000

Notes: Each column regresses the indicated forecast residual on its own lags 1-6 and 12, absorbs calendar-month fixed effects, and clusters standard errors by rival bank. This table is a persistence diagnostic; the design does not require i.i.d. shocks.

Table B4. Dynamic response of non-relationship hiring

	Non-relationship hiring
Lead 7-12 months	-0.1145 (0.6984)
Lead 4-6 months	0.1010 (0.2809)
Lead 1-3 months	0.1541 (0.2710)
Current shock	-0.1421 (0.1276)
Lag 1-3 months	0.6872*** (0.2479)
Lag 4-6 months	0.4329* (0.2402)
Lag 7-12 months	0.1488 (0.3483)
Observations	110,880
R-squared	0.6086

Notes: The table reports the non-relationship hiring response from the same dynamic specification used in Figure 1 of the main text. Lead rows correspond to future rival innovations relative to the focal-bank outcome month, while lag rows correspond to past rival innovations. Fixed effects and clustering follow Table 2 of the main text.

Table B5. Wild-cluster bootstrap inference for shock-level estimates

	Relationship hiring	Relationship - non-relationship
Coefficient	0.9444***	1.1294***
Two-way clustered standard error	(0.1335)	(0.2346)
Wild-cluster bootstrap p-value	<0.001	<0.001
Wild-bootstrap 95% confidence interval	[0.6877, 1.2083]	[0.6704, 1.5939]
Observations	152,352	152,352
Rival-bank clusters	150	150
Month clusters	96	96
Bootstrap replications	9,999	9,999

Notes: The table reports inference for the Borusyak-Hull-Jaravel shock-level equivalent representation of the shift-share reduced form. Coefficients are identical to the corresponding location-level reduced-form estimates. Conventional standard errors are two-way clustered by rival bank and calendar month. Wild-cluster bootstrap inference is computed at the rival-bank level with the null imposed and 9,999 replications. Bootstrap p-values reported as 0.0000 in Stata are displayed as <0.001. The bootstrap exercise is used as an inference robustness check and is not interpreted as a separate structural IV estimand. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table B6. Relationship-hiring response by recruitment platform

	51job	Liepin	Zhaopin
Rival relationship-hiring innovation	0.3608***	0.3807***	0.3286***
	(0.0695)	(0.0793)	(0.0765)
Observations	147,840	147,840	147,840
R-squared	0.2851	0.3635	0.2892

Notes: The source order follows the platform ordering in the estimation script (51job, Liepin, Zhaopin). Each column measures relationship-banking vacancy postings on one platform using the common baseline rival shock. Cross-platform shock-outcome specifications are reported separately in Table B12.

Table B7. Salary-disclosed relationship vacancies: intensive margin

	Salary disclosed
Rival relationship-hiring innovation	0.5644***
	(0.1157)
Observations	147,840
R-squared	0.3455

Notes: The dependent variable is the inverse-hyperbolic-sine number of relationship vacancies with advertised salary information. Fixed effects and clustering follow Table 2 of the main text.

Table B8. Salary-disclosed relationship vacancies: extensive margin

	Salary disclosed
Rival relationship-hiring innovation	0.4027***
	(0.0833)
Observations	147,840
R-squared	0.3226

Notes: The dependent variable equals one if the bank-city-month contains at least one relationship vacancy with advertised salary information. Fixed effects and clustering follow Table 2 of the main text.

Table B9. Within-relationship composition: flagged versus other relationship vacancies

	Client resources	Relationship experience	Senior RM	Local resources	5+ years experience
Rival relationship-hiring innovation	-0.4512*** (0.1006)	-0.7507*** (0.1365)	-0.7453*** (0.1360)	-0.8883*** (0.1626)	-0.4829*** (0.1025)
Observations	147,840	147,840	147,840	147,840	147,840
R-squared	0.3090	0.3260	0.3394	0.3663	0.3124

Notes: For characteristic k, the outcome is asinh(flagged relationship vacancies) minus asinh(other relationship vacancies). The outcome is defined even when total relationship hiring is zero. Negative coefficients therefore indicate that other relationship vacancies expand more strongly than the flagged category, not that flagged vacancies decline in absolute terms.

Table B10. Advertised relationship salary: conditional and selected

	Log advertised relationship salary
Rival relationship-hiring innovation	-4.0602** (1.8060)
Observed salary cells	3,113
R-squared	0.8216

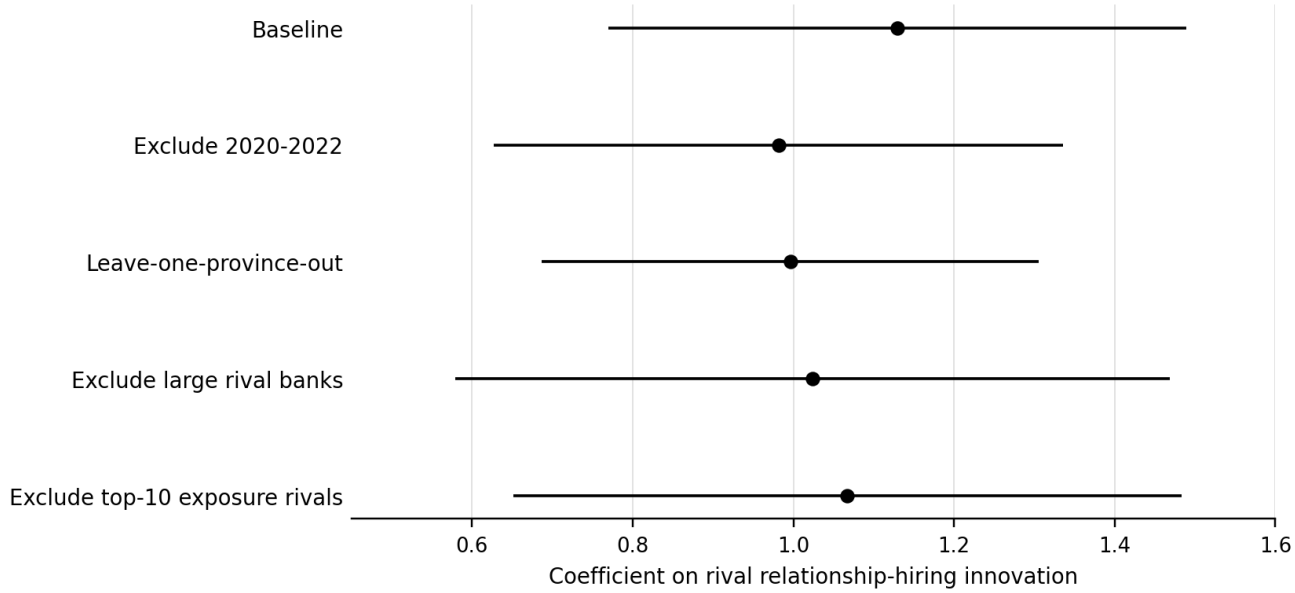
Notes: The dependent variable is log mean advertised salary among relationship vacancies in cells with observed salary information. Salary disclosure itself responds positively to the rival shock, so this conditional salary regression is selected and is not interpreted as a causal compensation effect.

Table B11. Cross-platform rival shocks and relationship-hiring outcomes

	51job -> Liepin	51job -> Zhaopin	Liepin -> 51job	Liepin -> Zhaopin	Zhaopin -> 51job	Zhaopin -> Liepin
Cross-platform rival innovation	0.1089**	0.1009***	0.0614***	0.0401	0.0595***	0.0377
	(0.0433)	(0.0280)	(0.0162)	(0.0242)	(0.0226)	(0.0277)
Observations	147,840	147,840	147,840	147,840	147,840	147,840
R-squared	0.3628	0.2885	0.2843	0.2885	0.2843	0.3627

Notes: Each column constructs the rival shock from the first-named platform and the focal-bank relationship-hiring outcome from the second-named platform. Same-platform pairs are excluded by construction. Fixed effects and clustering follow Table 2 of the main text.

Figure B1. Robustness of relationship minus non-relationship hiring



Notes: The figure plots raw coefficient estimates and 95% confidence intervals for the relationship-minus-non-relationship hiring outcome. The specifications mirror Figure 2: baseline, exclusion of 2020-2022, leave-one-province-out shock construction, exclusion of large rival banks, and exclusion of the top-10 exposure rivals. All regressions absorb bank x city, bank x month, and city x month fixed effects; standard errors are two-way clustered by city and bank.