Demand Spillovers from Star Employees: Evidence from Intra-Firm Reassignment Across Brokerage Shops

Ruiqi Wang* Hefan Zheng†

July 2026

Abstract

Star employees are among a firm's scarcest resources. Understanding how they create value is central to determining where firms should deploy them. Research on star spillovers has emphasized supply-side channels such as knowledge transfer, peer learning, and collaboration, under which a star's value depends on interactions with coworkers. We study a different mechanism arising where demand attaches to individual employees rather than to the firm alone. Using intra-firm reassignment of agents across the shops of a large real estate brokerage, we find that shop performance rises immediately and persistently when a star arrives and reverses when she departs. Star arrival sharply expands the receiving shop's customer pool. Yet the star does not increase her own service volume; instead, coworkers absorb the additional demand, lifting their output by 11.8%. Tracing customer inflows across shops, we find the decline at the star's origin shop accounts for two-fifths of the receiving shop's gain, with no offsetting decline elsewhere in the firm. Most of the inflow is net new demand the firm would not otherwise have captured. But shop capacity is bounded: as coworkers take on more buyers, conversion does not improve and buyers search longer, with the strain sharpest where the shop is already congested. The same star therefore creates more value in shops with room to serve additional customers and the ability to convert them. These findings reframe star deployment from fostering post-arrival coworker interaction to matching scarce demand-drawing talent to the units best positioned to capitalize on the demand it attracts.

Keywords: Star employees, intra-firm talent allocation, demand-side spillovers, organizational capacity, real estate brokerage

*Department of Economics, National University of Singapore.
†East China Normal University.

1 Introduction

Star employees are among a firm's scarcest and most valuable resources (Rosen, 1981; Morris et al., 2021). The value of a star employee depends not only on who she is, but on where the firm places her. In multi-unit organizations, managers must decide where to deploy the few employees whose visibility, track record, or client following set them apart. The decision is consequential: the same star may create very different value across shops, branches, or offices. The value she creates depends on the context of the unit she joins: the coworkers around her, the customers it serves, and the conditions under which it operates. Where to place a star is therefore a central organizational problem, and answering it turns on a prior question: how a star creates value for the unit she joins.

Stars do not only produce more themselves; they also shape the performance of the coworkers around them and the units they join (Prato and Ferraro, 2018; Ray et al., 2026). Research on star employees has largely located these spillovers inside the production process. A star affects coworkers through her interactions with them: she transfers knowledge or demonstrates superior practices (Azoulay et al., 2010; Oettl, 2012), raises effort through peer pressure (Mas and Moretti, 2009; Bandiera et al., 2010), or complements their work in joint production (Arcidiacono et al., 2017). These accounts differ in their microfoundations, but share a common implication for deployment. If a star creates value by changing how her coworkers produce, the firm should deploy her where such interactions are expected to generate the greatest returns: among the coworkers who can best learn from her, raise their effort alongside her, or collaborate with her most productively. However, this view leaves the unit's demand environment largely in the background. It treats stars as improving coworkers' ability to convert existing demand into output, rather than as reshaping the demand they face.

This logic may be incomplete in settings where demand attaches to individual employees rather than to the firm alone. In these settings, customers observe, search for, and choose particular individuals (Gubler, 2019). Reputation, prior performance, ratings, and client relationships make a star visible not only inside the firm, but also in the market around it. Client attachments and external relationships can therefore be portable, following individuals when they move (Somaya et al., 2008; Raffiee, 2017). When such a star moves from one unit to another, she brings more than her own productive ability. She may redirect customer attention associated with the market she previously served and make the receiving unit more attractive to new clients. The spillover from a star then begins not with coworkers learning from her or working harder alongside her, but with a change in the demand environment they face. Under this mechanism, star deployment becomes a matching problem between demand-drawing talent and the unit best positioned to benefit from the demand she attracts.

We bring this demand-side perspective to the study of intra-firm star placement. We ask whether

stars generate spillovers by attracting demand to the units they join, how that demand is allocated within the unit, whether it reflects redistribution across units or an expansion of the firm's customer base, and how this mechanism changes the criteria for star placement. We study these questions in the internal labor market of a large real estate brokerage, a customer-facing service in which the stars are individual agents and the units are local shops. The firm routinely reassigns agents, including stars, across shops under an established internal rotation practice that promotes cooperation within local markets rather than responding to short-run shop-specific demand. This reassignment generates within-firm variation in star placement that holds firm-level factors fixed and isolates the effect of where a star is placed. We exploit this variation in a staggered difference-in-differences design built around each shop's first inflow of a star reassigned from elsewhere in the firm, comparing changes in receiving shops with contemporaneous changes in shops that do not receive a star. This contrast isolates the effect of star placement rather than differences across shops that are more likely to host stars.

We find that star placement has large and immediate effects on receiving-shop performance. When a star arrives, quarterly transactions increase by 7.8 in the event quarter, and commission revenue increases by roughly RMB 412,000; these gains persist throughout the post-event window. The pattern reverses around star outflow, a symmetry that indicates that gains are driven by the star's placement rather than unrelated shop-level trends. These shop-level gains also extend beyond the star's own output to her coworkers. Per-agent performance rises even when the star is excluded, and the effect remains positive among incumbent coworkers who were already at the shop before the star arrived. Star placement therefore generates genuine workplace spillovers.

The mechanism generating these spillovers, however, does not primarily operate through traditional supply-side channels. Star arrival sharply expands the receiving shop's customer base: buyer inflows increase by about 195 customers per quarter, while seller inflows rise more modestly. This additional demand is not absorbed by the star herself. Her own client load does not rise after the move; instead, coworkers take on the new buyers, with their per-agent buyer counts increasing by about 16%. Stars in this setting therefore lift coworker output by reshaping the demand environment in which coworkers operate. The inflow demand comes through two margins. Part of the inflow is portable: new customers who would otherwise have entered through the star's origin shop are redirected to the receiving shop, drawn by the prior-market experience and reputation she carries with her. The larger margin is net new relative to the counterfactual in which the star remains at her origin shop: it is not offset by declines at the origin or neighboring shops, but reflects additional customers the firm captures because the star becomes visible in a new local market. Star placement therefore does more than reallocate customer inflows across the firm's shops; it expands firm-wide demand.

However, demand-side spillovers are bounded by the receiving shop's capacity. As coworkers

take on more buyers, their conversion rates do not improve and buyers spend longer in search before they transact. These frictions are strongest where the receiving shop was already more congested before the star arrived. Consistent with this capacity constraint, star inflows generate substantially larger gains, roughly double, in shops with either more room to absorb additional demand or stronger pre-arrival conversion ability. The mechanism thus illustrates that placements are valuable through whether the receiving shop is already positioned to serve and convert the demand she brings. A stylized counterfactual reassigns the same stars within the same market-quarter from more congested to less congested shops and, separately, from lower-converting to higher-converting shops. Holding fixed the number and timing of moves and adding no new talent, the results suggest that aligning star placement with shops' pre-arrival operating conditions could increase aggregate transactions and commission revenue by roughly 20%.

These findings speak to three literatures. First, we contribute to research on star employees and peer spillovers by documenting a demand-side channel through which stars affect coworker output. Research on peer effects has emphasized supply-side channels such as learning, effort, and collaboration (Mas and Moretti, 2009; Azoulay et al., 2010; Arcidiacono et al., 2017), largely overlooking the role of the demand side. Prior work on employee mobility has shown demand itself is portable: client relationships travel with the individual across organizations (Somaya et al., 2008; Raffiee, 2017), but tracks the mover's own outcomes rather than the performance of coworkers at the same unit. We connect the two: star arrival expands the receiving shop's customer base by about 210 customers in the event quarter, lifting coworkers' commission revenue by 11.8%. Second, we contribute to research on talent allocation and organizational design by reframing star deployment as a matching problem between demand-generating talent and a receiving unit's operating conditions (Weller et al., 2019), extending a matching view of strategic human-capital decisions from whom firms hire (Elfenbein and Sterling, 2018) to where they deploy the talent they already employ (Minni, 2026; Ray et al., 2026). Under the supply-side view, deployment is a question of post-arrival interaction: whom a star can mentor and where collaboration is most productive. We show instead that the returns to star placement are set by conditions observable before the star ever moves: shops with more room to absorb additional demand and stronger conversion ability gain roughly twice as much from star arrival. Reallocating stars toward such shops, holding fixed the number and timing of moves, raises aggregate commission revenue by 15.4 to 26.5 percent. Third, we contribute to research on service congestion (Tan and Netessine, 2014; Berry Jaeker and Tucker, 2017) by documenting the customer-side costs of star-induced demand concentration. This literature studies how workload erodes output once demand is already given; we show where such a shock can originate, a star's arrival, and what it costs the customers who follow her. Buyers search 37 more days on average after star inflow, a penalty that grows by 15 more days when the shop was already congested beforehand.

More broadly, our results show that the internal allocation of scarce talent is a strategic lever in its own right. Firms can create substantial value not only by hiring or developing more stars, but also by placing existing stars in units that can absorb and convert the demand they attract.

2 Star Allocation and Spillover Mechanisms

Star employees are scarce organizational resources. High-performing workers generate disproportionate value, not only through their own output (Rosen, 1981; Chetty et al., 2014), but also through their effects on the people around them (Lazear et al., 2015; Hendricks et al., 2023). Yet the value a star creates depends on the context in which she works. Prior research shows that individual performance, including star performance, is only partly portable across organizations (Huckman and Pisano, 2006; Groysberg et al., 2008; Groysberg and Lee, 2009). This context dependence has direct allocation consequences for multi-unit firms, which operate multiple shops, branches, or teams. These units differ in coworker composition and operating conditions, so the same star may create different value depending on where she is placed. Star allocation is therefore an internal organizational problem: the firm must decide where to deploy its scarce star talent to create the greatest value. Making this deployment decision requires understanding how a star affects coworkers' output in the units she joins, because different mechanisms imply different criteria for star placement.

Existing research has offered several explanations for how stars affect the performance of those around them. One channel is knowledge transfer or peer learning: coworkers may learn from observing stars, adopt better practices, or draw on the expertise a star brings into the workplace (Jackson and Bruegmann, 2009; Azoulay et al., 2010; Oettl, 2012; Waldinger, 2012). The second is peer-induced effort: the presence of a high performer raises others' effort through social pressure, monitoring, or shifts in workplace norms (Mas and Moretti, 2009; Bandiera et al., 2010). A third channel is production complementarity, in which a star's contribution can raise coworkers' output in joint production, and the resulting value may reflect complementarities between the star and her collaborators rather than either party's contribution alone (Arcidiacono et al., 2017; Mindruta et al., 2025). These mechanisms differ in their microfoundations, but they share a common supply-side logic. Stars generate spillovers by changing what coworkers know, how hard they work, or how their inputs combine in production. Under this view, star allocation should prioritize units where the star's presence can most improve coworkers' productive capability.

These supply-side accounts are compelling, but demand matters as well, especially where demand attaches to individual professionals rather than to the firm. When quality, reputation, and past performance are often visible to customers, higher-quality providers can attract more demand (Anderson and Magruder, 2012; Han and Ravid, 2020). Meanwhile, this demand can also be portable: client relationships built around an individual move with her across organizations and

locations (Broschak, 2004; Somaya et al., 2008; Mawdsley and Somaya, 2016; Raffiee, 2017). This demand-side spillover is distinct from the supply-side channels above. However, this literature has focused on the mover's own outcomes rather than on the coworkers at the receiving unit. How the demand a star brings affects coworkers is less understood, and the mechanism implies a different deployment criterion. Supply-side channels raise how productively coworkers convert a given level of demand, so placement turns on the quality of interaction the star can have with them. A demand-side channel raises the volume of demand the unit faces, which a unit with finite capacity must absorb and convert. The relevant criterion therefore shifts from what unfolds between the star and her coworkers after arrival to whether the receiving unit is already positioned to serve and convert the demand she draws before she ever arrives.

Once the focus shifts to the demand-side value of star placement, the unit becomes the first site where the demand associated with the star is processed. Customer attention creates value only if the unit can accommodate that demand. Capacity alone is not sufficient: the unit must also have the ability to convert customer opportunities into completed transactions. When the demand-to-output process is constrained, the consequences appear not only in lower realized output but also in the customer experience during service (KC and Terwiesch, 2009; Anand et al., 2011; Tan and Netessine, 2014; Berry Jaeker and Tucker, 2017). Accordingly, the returns to star placement should be greater in units with more available capacity and stronger pre-arrival conversion ability.

Supply-side and demand-side mechanisms could both produce gains from star placement, but they imply different paths through which value arises and different criteria for allocating stars. Prior research on supply-side star spillovers has greatly advanced our understanding of why stars matter and how they affect those around them. Yet the relative lack of attention to demand-side mechanisms leaves incomplete our understanding of stars as deployable organizational resources. The analysis below therefore examines whether star inflow affects not only overall shop performance, but also the coworkers and customers connected to the shop. By tracing these coworker-side and customer-side consequences, we assess whether a demand-side mechanism operates in this setting and what it implies for deploying scarce star talent across units.

3 Study Context and Empirical Strategy

3.1 Organizational Setting

We study a leading real estate brokerage firm in China that operates a large network of local brokerage shops together with online housing platforms. The firm is organized as a multi-unit service organization. Agents are affiliated with local shops, and shops serve as the primary offline units for customer reception, daily operations, and local market coverage. Shops are typically located

in residential communities and commercial districts, giving them access to local housing demand and neighborhood-level information. The firm's online platforms complement this offline network by allowing customers to browse available listings, search housing information, and find agents. This combination of offline shops and online platforms makes both shops and individual agents visible to potential customers.

Real estate brokerage is a customer-facing professional service in which agents work directly with clients, largely one-to-one. Customers reach a shop by browsing listings on the platform, which links properties to nearby shops, by contacting an agent directly, or by walking in. The shop manager assigns each customer to a specific agent based on the customer's preferences and agents' current availability. The platform facilitates the initial property search but does not direct customers to specific agents. Throughout the service process, a customer is typically served by a single agent. On the buyer side, agents interpret housing needs, recommend listings, arrange and conduct showings, and help buyers navigate the purchase process. Seller-side agents help manage listings, coordinate showings, and maintain relationships with property owners (Han and Strange, 2015). The firm's operations rely heavily on its agents, whose income comes primarily from commissions on the deals they close. Performance evaluation, however, takes place in part at the shop level rather than the individual level, giving agents in the same shop a strong incentive to share listings and leads. Cross-shop showings are common and even encouraged, so buyers can work with an agent regardless of where her shop sits. Sellers are more geographically anchored because the property is fixed in place, and nearby shops are better positioned to maintain listings and arrange showings.

Brokerage relies on individualized service, so an agent's value to clients depends heavily on her own experience, market knowledge, and track record (Fang and Hayunga, 2024; Gilbukh and Goldsmith-Pinkham, 2024). These attributes are unusually visible to customers. The shop and firm's platform display the completed transactions and ratings of each agent, and the reputation also circulates among clients within a neighborhood. Clients use such visible signals to select particular agents rather than simply the firm (Gubler, 2019). Within this organization, some agents become especially visible to customers in this way. We refer to these high-performing, highly visible agents as *stars*. This definition is consistent with research conceptualizing star employees as workers who combine sustained high performance with visibility and relevant social capital (Call et al., 2015). We defer a precise empirical definition of stars to Section 3.2.

A defining feature of our setting is that the firm routinely reassigns agents across its shops. Such moves are common: 31.16 percent of agents change shops at least once during their time at the firm. These transfers follow an established internal rotation practice[1] and are predominantly organized

[1]Institutional details in this section and in Section 5.2 draw on semi-structured field interviews with one firm-level manager, two shop managers, and five agents, three of whom were star agents under the definition in Section 3.2, conducted in 2025. Participation was voluntary, and we do not identify individual respondents. Appendix EC.1 summarizes the institutional evidence.

within local operating regions. The firm's Beijing operations are organized into ten operating markets based on administrative districts. In our data, 90.79% of rotations occur within the same operating market, and rotation frequencies are relatively balanced across markets. These rotations are intended to strengthen cooperation among agents. By working in different shops within the same market over time, agents become familiar with nearby neighborhoods, local listings, and the working routines of other agents, which supports cross-agent collaboration in matching buyers and sellers, arranging property showings, and coordinating transactions. These routine reassignments move agents across existing shops and are in practice largely reallocative, not expansionary. Around the reassignment quarter, the average net change in shop headcount is close to zero and statistically indistinguishable from zero. [2] A separate motive for transfers is to staff newly opened shops, but such cases account for about 10.9 percent of rotation activity. Transfers to new shops may be directly tied to shop expansion and local growth opportunities, so our analysis of star reassignment excludes them through the construction of the event. These features make routine reassignment a useful source of within-firm variation in the placement of star agents across shops. We use this variation in Section 3.3 to construct the star inflow event and estimate the effect of star placement.

3.2 Data and Key Measures

Our data come from the real estate brokerage firm described above. The firm operates nationally, but we focus on its Beijing operations. As the largest brokerage in Beijing's second-hand residential housing market, the firm captures a substantial share of local buyer and seller activity, so our customer-flow measures reflect most of the relevant demand rather than a fragment that leaks to unobserved competitors. Restricting attention to a single city also holds constant many facets of the regulatory and macroeconomic conditions that vary widely across Chinese housing markets and that would otherwise confound comparisons across shops. Finally, Beijing is the firm's most extensive and mature operation, providing a dense network of shops within well-defined operating markets and enough star reassignment events to estimate dynamic effects with precision. Each operating market contains multiple business circles, finer commercial zones we use to measure geographic proximity among shops (Section 5.3) and to characterize buyers' search area. Over the sample period (July 1, 2021 to December 31, 2023), the Beijing dataset covers 23,044 agents and 2,244 shops. A novel feature is that for each agent, we observe a complete work record, including her shop affiliation, customer relationships, managed listings, showing activities, and completed transactions. Linking these agent records to shops lets us place reassignments within the organizational context

[2] We consider shop-quarters in which at least one agent is routinely reassigned into or out of the shop, excluding shop opening quarters (4,036 shop-quarters). The mean change in the shop's agent count from the quarter before to the quarter of the move is 0.024 agents (median 0), statistically indistinguishable from zero ($p = 0.485$). 60.6% of these shop-quarters change by at most one agent.

of shop operations and performance.

The dataset also includes the property visit records of 1,997,387 buyers who were actively searching during the sample period. Among them, 117,679 completed a transaction observed in the firm's records. The seller-side records cover 822,794 listed properties on the platform. For each property in the sample, we observe its list price, location, property size, number of living rooms, and number of bedrooms. Since buyer visits are recorded at the property level, we can link each visited property to its listing characteristics, allowing us to characterize buyers' search profiles from the types of properties they consider before purchase.

The complete agent work histories allow us to identify high-performing agents in the firm, whom we refer to as *stars*. Our empirical definition follows the firm's own performance-based recognition system, which highlights agents ranked in the top 10 percent of business performance. For each calendar year, we compute an agent's annual income from commissions earned on completed transactions. Because agents are compensated primarily through transaction-based commissions, this income measure captures their realized sales performance. Agents are ranked within the year by this annual income measure. We then carry star status forward by one year: if agent i ranks in the top 10% of agents in year $y-1$, she is classified as a star in all months of year y. This lagged definition ensures that star status is predetermined relative to the outcomes we study in year y. Because star status is defined using commission income, her own output could mechanically drive any shop-level gain we document; Section 4 addresses this by isolating effects on coworkers, excluding the star's own contribution. Although our measure is performance-based, performance and visibility are tightly linked in this setting: the platform publicly displays each agent's transactions and ratings, making top-performing agents especially observable to customers (Kehoe et al., 2018).

Our outcome variables are measured at two levels. At the shop-quarter level, we measure performance using the number of completed transactions and total commission revenue. We measure customer inflows using the number of buyer and seller customers whose first recorded interaction with the firm is associated with a given shop in that quarter. This measure captures newly arriving customers in the firm's observed records rather than continuing service relationships with existing customers. [3] At the buyer level, we measure the search experience of buyers who complete a purchase. These outcomes include search duration, defined as the number of days between the first observed property visit and purchase; the number of properties visited before purchase; and search density, defined as visits per day. Table 1 reports summary statistics for these measures. Additional measures used in the mechanism and heterogeneity analyses are defined when introduced in the corresponding sections, with construction details provided in the Appendix EC.2.

[3] In our setting, customers typically maintain an ongoing relationship with a fixed agent during the service process. Existing customers are therefore treated as continuing relationships and are not mechanically counted as new inflows when an agent moves across shops.

Table 1. Summary Statistics

	Mean	Std	Median	Observations
Panel A: Shop-quarter outcomes				
Number of Transactions	21.049	13.587	19	16,177
Commission Revenue (RMB)	936,676	612,973	845,461	16,177
Buyer Customers	521.174	311.368	476	16,177
Seller Customers	45.275	29.536	40	16,177
Panel B: Buyer-level search outcomes				
Search Duration (days)	92.032	133.190	34	117,679
Number of Properties Visited	16.282	14.504	12	117,679
Search Density	0.884	1.306	0.390	117,679

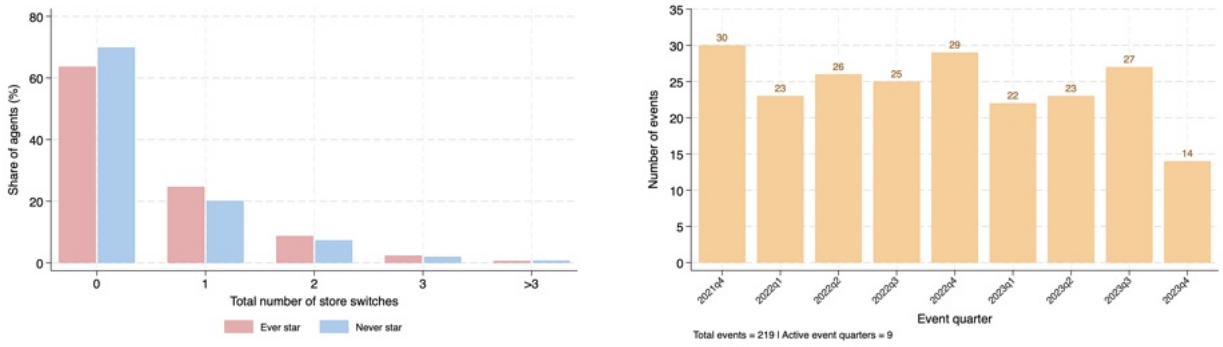
3.3 Empirical Strategy

Our analysis revolves around the causal effect of star employees on the shops they join and the coworkers within them. Stars move across shops within the firm, and we exploit these naturally occurring reassignments to identify the effect of where a star is placed. We define a star inflow event at the shop level. A shop is treated in the quarter when it receives a star agent reassigned from another shop and transitions from having no star agents to having at least one star agent. The event quarter is the first such zero-to-one transition observed during the sample period. We refer to the shop that receives the star as the receiving shop and to the shop the star leaves as the origin shop. To distinguish star reassignment from the staffing of newly opened shops, we require receiving shops to have been operating for at least one quarter prior to the inflow.

Agent reassignment provides the key within-firm variation in star placement. Because these moves take place within a single firm, a star's reassignment changes her location within the organization while holding fixed the firm's brand, platform, commission structure, and operating practices. This helps isolate the effect of placement from the firm-specific human-capital requirements and organizational-context changes that typically accompany inter-firm mobility (Lazear, 2009; Bidwell, 2011). Reassignment activity is also frequent and dispersed, providing sufficient variation for empirical design. Figure 1 summarizes the frequency of agent shop switching and the timing of first star inflow events. About 31 percent of agents switch shops at least once during the sample period, and stars switch at a similar rate to non-stars, consistent with star mobility being part of the firm's routine rotation rather than a distinct, selectively managed process. We identify 219 first star inflow events spread across nine active quarters, broadly balanced across calendar quarters. The frequency and dispersion of these events provide substantial within-firm variation in star placement and reduce the concern that the estimates are driven by aggregate shocks coinciding with a single

wave of reassignments.

Figure 1. Agent Shop-Switching Frequency and Timing of Star Inflow Events



(a) Agent Shop-Switching Frequency

(b) Timing of Star Inflow Events

The identifying comparison is between changes in outcomes at receiving shops around the first star inflow and contemporaneous changes at shops that do not receive a star inflow during the sample period. These control shops include shops that never have a star agent and shops in which a star is developed internally rather than reassigned from another shop. Including the latter sharpens the interpretation: the treatment is the arrival of a reassigned star, not the mere presence of a star, which may itself be endogenous. We also exclude from the control group shops that are directly exposed to the star reassignment, such as the star's origin shop. We return to this issue in Section 5.3, where we examine which shops are directly affected by the spillover induced by star reassignment. Our design therefore identifies the effect of star placement across shops, rather than the broader association between having a high-performing agent and shop performance.

Identification rests on a parallel-trends assumption: absent the star inflow, treated and control shops would have followed similar paths. That is, after accounting for permanent differences across shops and shocks common to all shops in a given quarter, the two groups would have experienced similar changes in outcomes over time. The lead coefficients of the event-study specification provide a direct diagnostic, since any pre-arrival divergence between receiving and control shops would appear as nonzero leads. In the results below, these coefficients are close to zero and show no systematic pre-trend. However, flat pre-trends do not by themselves rule out selective reassignment. We therefore assess the identifying assumption along three dimensions through which selection could operate: which shops receive stars, which stars are moved, and when the moves occur.

We first examine whether reassignment selects the shops that receive stars. The concern is that star placement may be correlated with receiving shops' pre-event trajectories. The firm's rotation practice mitigates this concern, since reassignments are organized primarily around cooperation within local markets rather than the short-run demand of particular shops, a description that

managers and agents gave consistently in our interviews (Appendix EC.1). We evaluate it directly by comparing treated and control shops within market-quarter cells on pre-event shop characteristics and local market conditions. Table 2 Panel A shows no statistically significant differences across these dimensions. Treated and control shops have similar staffing levels, shop age, performance, and surrounding market conditions. In particular, treated and control shops are similar in pre-event commission revenue and transaction growth, suggesting that star inflows are not targeted to shops already performing unusually well or trending upward. We further test whether recent changes in shop staffing and customer demand predict subsequent star inflow. We find no such relationship in Appendix EC.3.1. We then ask whether the stars who move differ from those who stay. Table 2 Panel B compares reassigned and non-reassigned stars on observable characteristics and pre-move performance. As in Panel A, we test equality of means via two group t-tests. No difference is significant on any individual dimension, indicating that reassignment does not select stars on observable productivity or background. The timing of moves is also unlikely to drive the estimates, since star inflows are dispersed across quarters rather than concentrated in a single wave (Figure 1, Panel (b)). Together, these checks rule out selection along the dimensions we can observe and test directly: pre-event shop trajectories, star characteristics, and predictors of move timing. They cannot rule out a time-varying, shop-specific shock that is anticipated by the firm but not yet reflected in observed outcomes, such as private information about an upcoming demand increase or a concurrent organizational initiative that both triggers star reassignment and raises subsequent performance. Our interpretation therefore relies on the assumption that, conditional on the fixed effects and observed pre-event dynamics, no such unobserved shock differentially affects receiving shops at the time of reassignment. We view the flat pre-trends, the balance tests, and the evidence that pre-event shop characteristics do not predict the timing of star inflow as jointly reducing this concern, but not eliminating it.

Because first star inflows occur in different calendar quarters, treatment timing is staggered across receiving shops. These events arise from the firm's recurring rotation process rather than from a single reassignment wave. We therefore estimate the effect of star placement using a staggered difference-in-differences event-study design. The design compares receiving shops with control shops observed in the same calendar quarter, while organizing the estimates by quarters relative to the star's arrival. This allows us to exploit variation in the timing of star inflows across shops rather than relying on a single before-after comparison around a common event date. It also traces the dynamic path of outcomes before and after star arrival, so that the post-event coefficients capture the evolution of treatment effects and the lead coefficients provide evidence on differential pre-event trends.

Table 2. Balance Tests

	Group 1	Group 2	Diff	p-value
Panel A: Shop-level (treated shop vs control shop)				
Agent Headcount	8.478	9.204	-0.727	0.880
Commission Revenue (RMB)	654,617	673,114	-18,497	0.101
Transaction Growth	-0.044	-0.018	-0.025	0.243
Shop Age (quarter)	13.079	12.863	0.216	0.274
Nearby Customer Pool	5,908	6,655	-747	0.593
Nearby Shop Count	11.818	13.025	-1.208	0.336
Panel B: Star Agent-level (reassigned stars vs non-reassigned stars)				
Female	0.315	0.270	0.045	0.142
Bachelor's Degree or Above	0.409	0.419	-0.009	0.781
More than 5 years' experience	0.623	0.614	0.009	0.910
Transaction Number	3.217	3.163	0.054	0.629
Commission Revenue (RMB)	178,061	169,191	8,870	0.229
Customer Number	72.578	69.473	3.105	0.269

Notes: All variables are measured before the star inflow event. In Panel A, Group 1 consists of treated shops, and Group 2 consists of control shops in the same market-quarter cell. In Panel B, Group 1 consists of reassigned stars whose moves constitute inflow events, and Group 2 consists of non-reassigned stars aligned in calendar time. The p-values test equality of means between treated and control shops. Agent headcount is the number of agents affiliated with the shop. Nearby customer pool and nearby shop count are measured within the same business circle, excluding the shop's own.

We estimate the dynamic effect of receiving a reassigned star at the shop-quarter level using the following specification:

$$Y_{st} = \sum_{k \neq -1} \beta_k D_{st}^{(k)} + \alpha_s + \tau_t + \varepsilon_{st},$$

where Y_{st} is an outcome for shop s in quarter t. The indicator $D_{st}^{(k)}$ equals one if shop s is k quarters away from its first star inflow in quarter t. We omit $k = -1$ as the baseline, so the coefficients β_k measure outcome differences relative to the quarter immediately before star arrival. The shop fixed effects α_s absorb time-invariant differences across shops, and the quarter fixed effects τ_t absorb market-wide shocks common to all shops in a given quarter. Standard errors are clustered at the shop level. In the baseline specification, we report leads and lags from seven quarters before to seven quarters after the event. The lag coefficients β_k for $k \geq 0$ trace the dynamic effect of star inflow after the event, while the lead coefficients β_k for $k < 0$ test for differential pre-trends. We use this specification for the shop-quarter outcomes and an analogous design at the buyer level for buyers' search outcomes.

We complement the main inflow design with two additional validation exercises that probe the causal interpretation and the robustness of our estimates. First, we examine the reverse event

of star outflow. A star outflow event occurs when a shop loses its last star agent and transitions from having at least one star to having no star agents. If the inflow effects are caused by the star's arrival, outcomes should move in the opposite direction when the star departs; if they instead reflect confounding shop-level trends, no such reversal should occur. Second, the baseline estimates use a two-way fixed effects (TWFE) event-study specification. Because our treatment timing is staggered, TWFE estimates may be sensitive to treatment-effect heterogeneity because staggered treatment timing allows already-treated shops to serve as controls for later-treated shops. These comparisons, often termed forbidden comparisons, need not recover a clean untreated counterfactual when effects vary across cohorts or over event time. To assess this concern, we report a Goodman-Bacon (2021) decomposition to characterize the identifying variation in the TWFE estimates and verify the main findings using the imputation estimator of Borusyak et al. (2024) and the group-time average treatment effect on the treated (ATT) estimator of Callaway and Sant'Anna (2021).

4 Performance Effects of Star Placement

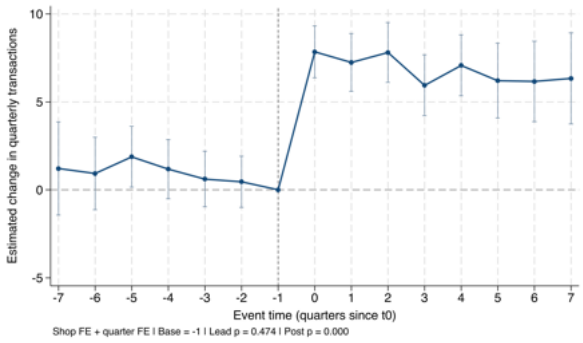
We begin by examining the performance effects of star placement on the receiving shop. In our setting, performance can be observed directly through shop-level transactions and commission revenue, allowing us to trace how outcomes change around the first inflow of star agents.

Using the staggered DID specification discussed in Section 3.3, we find that star agent inflow leads to an immediate and persistent increase in shop performance. Figure 2, Panel (a) and (b), show that both transactions and total commission revenue rise sharply in the event quarter and remain elevated thereafter. Quarterly transactions increase by 7.8 at the event quarter and stabilize around 6.3 by the end of the window, while quarterly commission revenue rises by roughly RMB 411,855 on impact and remains above RMB 321,332 seven quarters later. Prior to the event quarter, the pre-event coefficients are small and display no systematic pattern. Joint tests fail to reject the null that the pre-event leads are jointly equal to zero ($p = 0.474$ and 0.229), providing support for the parallel-trends assumption. Consistent with the robustness checks described in Section 3.3, a Goodman-Bacon decomposition attributes over 94% of the identifying variation to clean comparisons between treated and never-treated shops, and our main results are robust to the imputation estimator of Borusyak et al. (2024) and the group-time ATT estimator of Callaway and Sant'Anna (2021). (See Appendix EC.3.2 for additional details.)

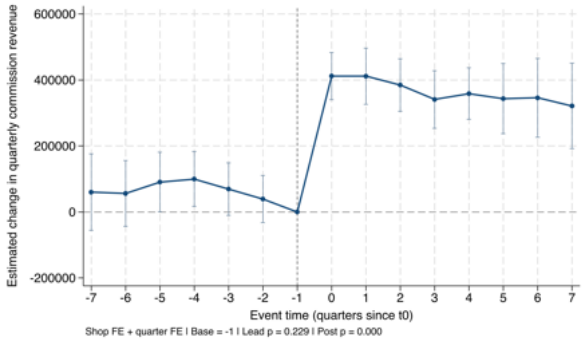
To further assess whether these gains are indeed driven by the presence of star agents, we turn to the reverse experiment of star outflow. If the performance gains documented above are causally attributable to stars rather than to contemporaneous changes in shop conditions, we should observe the opposite pattern when a shop loses its star. As shown in Figure 2 Panel (c) and (d), following star outflow, transactions and commission revenue decline substantially. This symmetric pattern

strengthens the interpretation that the performance gains at inflow shops are driven by star placement itself.

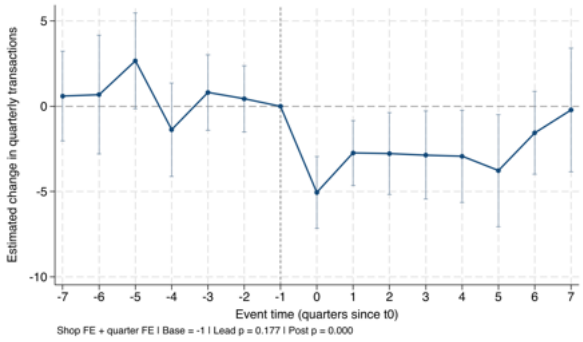
Figure 2. Transactions and Total Commission Revenue



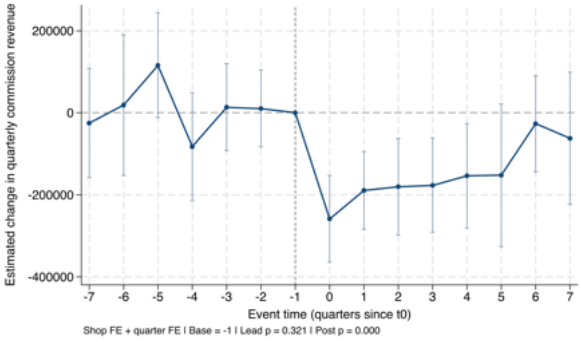
(a) Star Inflow: Transactions



(b) Star Inflow: Commission Revenue



(c) Star Outflow: Transactions



(d) Star Outflow: Commission Revenue

We examine per-agent performance along three progressively restrictive dimensions, reported in Table 3, each successively ruling out an alternative source of the shop-level gain and isolating the spillover to pre-existing coworkers. Columns (1) and (4) report output per agent at the shop level. Per-agent transactions and commission revenue both rise on impact, implying that the aggregate shop gains are not driven purely by scale expansion. In columns (2) and (5), the per-agent estimates remain positive when the inflowing star is excluded, with especially clear evidence for commission revenue, indicating that the improvement in shop-level performance cannot be attributed solely to the star's own output. More importantly, these specifications reveal that star arrival generates immediate and economically meaningful spillovers to coworkers. Excluding the inflowing star, per-agent transactions increase by 0.157 in the event quarter, or roughly 9.91% relative to the pre-event treated shop mean, and per-agent commission revenue rises by about 11.84%. Columns (3) and (6) further restrict attention to incumbent agents, those already at the shop before the star arrived. This holds the workforce fixed at its pre-event composition, so the per-agent effect reflects a genuine change

among the same coworkers rather than a shift in who staffs the shop. The persistence of positive per-agent effects in this setting provides the cleanest evidence that the benefits of star arrival extend to pre-existing coworkers.

Table 3. Per-Agent Performance

	Transactions			Commission Revenue		
	Y/N	Y/N excl star	Incumbent Y/N	Y/N	Y/N excl star	Incumbent Y/N
	(1)	(2)	(3)	(4)	(5)	(6)
Event Quarter ($t = 0$)	0.270***	0.157**	0.188**	15,471***	8,616**	10,777**
Post Average ($t = 0$ to $t = 3$)	0.224***	0.108	0.145*	14,581***	7,130**	9,193**
Shop FE	✓	✓	✓	✓	✓	✓
Quarter FE	✓	✓	✓	✓	✓	✓
Joint lead p-value	0.483	0.547	0.778	0.853	0.938	0.661
Observations	11,825	11,818	10,636	11,825	11,818	10,636

Notes: Columns (1) and (4) report per-agent transactions and commission revenue. Columns (2) and (5) exclude star agents, i.e., $(Y - Y^{star})/(N - N^{star})$. Columns (3) and (6) report per-agent outcomes for incumbent agents (agents present before the star inflow), i.e., $Y^{Incumbent}/N^{Incumbent}$. The unit of observation is a shop-quarter. Standard errors are clustered at the shop level.

Taken together, these patterns establish that star inflow produces genuine shop gains beyond a mechanical scale effect, and that a meaningful portion of the aggregate improvement accrues to coworkers rather than to the star alone. Yet the way these gains unfold does not fit within the mechanisms the existing literature has emphasized. Prior work has largely attributed such spillovers to supply-side channels, such as knowledge transfer, peer learning, collaboration externalities and effort responses to social pressure. The first three share a common empirical signature: effects that accumulate gradually as coworkers absorb skills, and persist after the source of the spillover leaves. However, our results display the opposite pattern: the effects are immediate, flat throughout the post-event window and symmetrically reversed upon star outflow. Effort responses to social pressure can in principle generate such immediate and reversible effects, but are typically difficult to reconcile with the sizable per-agent gains we observe. Beyond this magnitude, the coworker responses documented in Section 6.1 weigh more directly against this channel. These discrepancies suggest that a different kind of mechanism may be at work. We therefore turn in Section 5 to an alternative mechanism: stars as portable sources of demand, reshaping the work environment of their coworkers upon arrival.

5 Demand-Side Mechanism of Star Spillovers

In professional service industries such as real estate brokerage, insurance, financial advisory, health care and law, stars are high-productivity professionals whose work is largely individualized, yet

whose reputation and professional standing are observable to potential customers. This structure limits the scope for traditional supply-side mechanisms, which rely on peer interaction, observation, or collaboration. Two well-documented features of these industries instead make demand-side channels a natural place to look: high-productivity professionals attract disproportionate customer demand and client relationships tend to move with professionals when they switch firms. When a star relocates, a portion of her client base and reputation-driven customer inflow may relocate with her, expanding the pool of demand at the receiving shop.

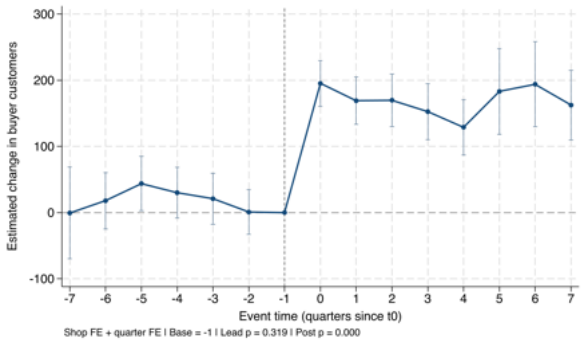
We advance this as the explanation for the spillovers documented in Section 4. Stars affect coworkers not primarily through what they know or how they cooperate, but through the demand environment those coworkers face. This account predicts the features we observe: effects are immediate (demand flows upon arrival), reversible (demand follows the star upon departure), and concentrated among coworkers (since the star's own capacity is bounded). We develop and test this argument in the remainder of this section.

5.1 Star Inflow Expands the Demand Pool

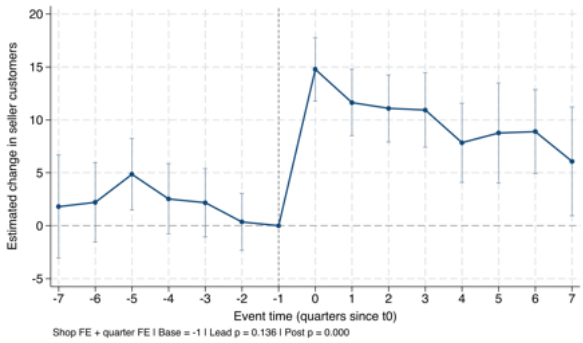
We first examine customer inflow around star arrival: if stars act as portable sources of demand, their arrival should expand the receiving shop's customer base.

As depicted in Figure 3 Panel (a), buyer inflows rise sharply in the event quarter, with the coefficient increasing from near zero in the pre-event period to approximately 195 additional buyers per quarter, and remain at this elevated level throughout the post-event window. Panel (b) reports that seller inflows show a similar pattern, rising by roughly 15 per quarter, though smaller in magnitude given the lower baseline level. Both show that star arrival substantially expands the demand pool at the receiving shop, with particularly strong effects on buyer inflows. This asymmetry between buyer and seller inflows reflects a basic feature of the housing market. Buyers search across neighborhoods and readily engage an agent regardless of the agent's shop location, making buyer-to-star attachment highly portable across workplaces. By contrast, sellers typically list with a nearby shop because tasks such as showings and property management favor geographic proximity, limiting the extent to which seller attachments can follow a star across locations. At the same time, a mirror-image prediction follows: if star arrival brings demand with it, star departure should take demand away. Figure 3 Panels (c) and (d) support this prediction. Among shops that experience a star outflow and transition from having a star to having no star, buyer inflows fall by approximately 106 per quarter and remain persistently depressed throughout the post-event window, roughly symmetric to the effect at receiving shops. The pattern also holds when we restrict attention to the origin shops of the stars that appear in our inflow events. A direct link between the demand gain at receiving shops and the demand loss at the star's origin shop is provided in Section 5.3.

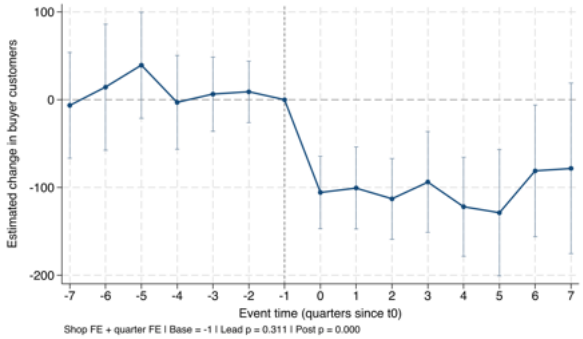
Figure 3. Customer Inflows



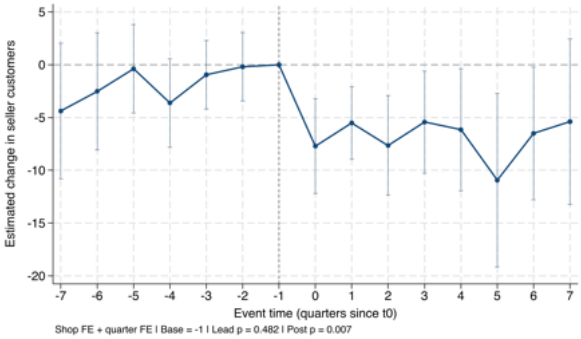
(a) Star Inflow: Buyer



(b) Star Inflow: Seller



(c) Star Outflow: Buyer



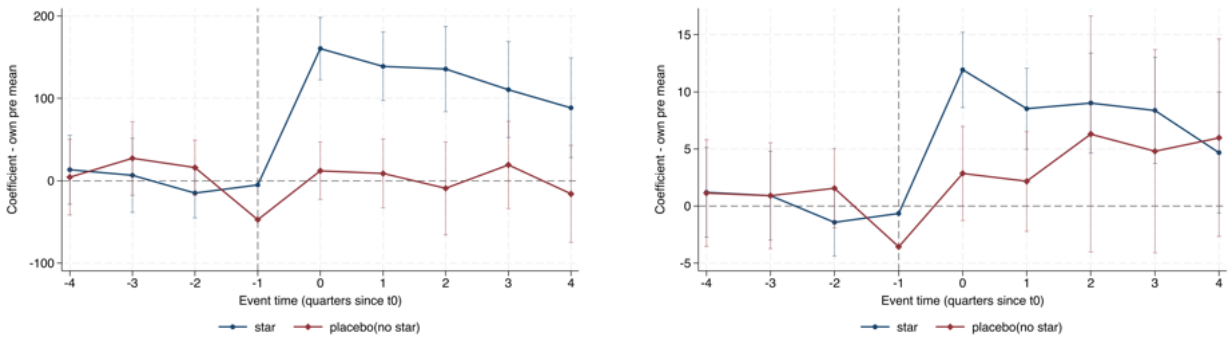
(d) Star Outflow: Seller

One alternative explanation is that customer inflows respond to any agent transfer, rather than specifically to star arrival. To address this possibility, we conduct a placebo test using shops that receive non-star agent inflows during the same period, estimated under a common control group specification that holds the control group fixed across the two treatments (details in Appendix EC.4). As shown in Figure 4, placebo shops display no comparable increase in either buyer or seller inflows following non-star agent inflow, providing supportive evidence that the demand response documented above is specific to star arrival rather than a general feature of personnel movement.

5.2 Demand Spillovers to Coworkers

An important implication of the demand-side mechanism is that the star need not be the primary beneficiary of the additional inflow. Stars in professional service settings often maintain a stable client load, whether due to capacity constraints or to preserve service quality. The incremental demand would therefore be reallocated to coworkers within the receiving shop. We test this implication by examining the star's own customer load. If she absorbs the incremental demand herself, her per-quarter customer count should rise after the move; if instead the demand flows to her

Figure 4. Placebo Treatment: Non-star Agent Inflow



(a) Buyer **(b)** Seller

Notes: This placebo comparison is estimated using a shared-control DID design in which star-treated and placebo-treated shops share the same never-inflow control group. Coefficients are normalized relative to each group's own pre-event mean. The event window is [-4, +4].

coworkers, her load should remain unchanged.

The star's own customer load remains broadly stable after the move. Both her buyer-customer and seller-customer counts fluctuate around zero throughout the event window, with no discernible change at the time of transfer or thereafter (Appendix EC.5). This pattern stands in sharp contrast to the substantial increase in total customer inflow documented in Section 5.1, indicating that the additional demand is not absorbed primarily through an expansion in the star's own activity. The natural implication is that the incremental demand generated by the star inflow is absorbed by coworkers within the receiving shop. We test this directly by examining per-agent customer counts among non-star agents at the receiving shop. Table 4 supports this prediction: buyer customers per agent rise significantly among coworkers, whereas the corresponding effects for seller customers are small and statistically insignificant. Excluding the inflowing star, buyer customers per agent in the receiving shop increase by 6.1 in the event quarter, equivalent to a 16.31% rise relative to the pre-event mean, and remain 3.9 higher on average over the first four post-event quarters. This increase is especially notable because star inflow also expands shop headcount: the receiving shop adds roughly two non-star agents around the event, yet even with demand spread across a larger workforce, each non-star coworker still serves substantially more buyers than before. The pattern persists when we restrict attention to incumbent coworkers. Their per-agent buyer load rises by 4.5 in the event quarter, or 11.19% relative to baseline, and remains 2.4 higher on average over the post-event window. This result rules out a purely compositional explanation based on newly added agents. These existing coworkers themselves serve more buyer customers afterward, consistent with a demand redistribution mechanism within the shop. For sellers, by contrast, per-agent counts show no comparable response. This asymmetry mirrors the buyer-seller divergence in Section 5.1: seller demand is anchored to neighborhood geography, limiting both its responsiveness to star arrival and

its redistribution across coworkers within the shop.

Table 4. Non-star Agent Customer Counts

	Buyer		Seller	
	Y/N excl star	Incumbent Y/N	Y/N excl star	Incumbent Y/N
	(1)	(2)	(3)	(4)
Event Quarter ($t = 0$)	6.122***	4.456***	0.192	0.070
Post Average ($t = 0$ to $t = 3$)	3.906***	2.434*	0.018	0.044
Shop FE	✓	✓	✓	✓
Quarter FE	✓	✓	✓	✓
Joint lead p-value	0.655	0.208	0.444	0.074
Observations	11,818	10,636	11,818	10,636

Notes: Columns (1) and (3) exclude star agents, i.e., $(Y - Y^{star})/(N - N^{star})$. Columns (2) and (4) report per-agent outcomes for incumbent agents (agents present before the star inflow), i.e., $Y^{Incumbent}/N^{Incumbent}$. The unit of observation is a shop-quarter. Standard errors are clustered at the shop level.

These patterns reflect the institutional structure of the setting. A star's track record and reputation raise the attractiveness of the shop she joins, drawing demand that arrives at the shop level. Because agents serve clients largely one-to-one, the star cannot absorb a large inflow herself without compromising service quality. The shop manager balances the allocation and assigns overflow customers to available coworkers (see Appendix EC.1), a staffing response consistent with evidence that client-related demands and the availability of internal human capital shape collaboration among professionals (Mawdsley et al., 2022). Both sides make this redistribution sustainable. For the star, performance evaluation depends in part on shop-level outcomes, so retaining excess buyers within the receiving shop can serve her own interest. For customers, the star's arrival can enhance the shop's visibility and credibility, making them willing to remain with the shop even when served by affiliated coworkers. The additional demand is therefore absorbed by coworkers while the star maintains her own client load.

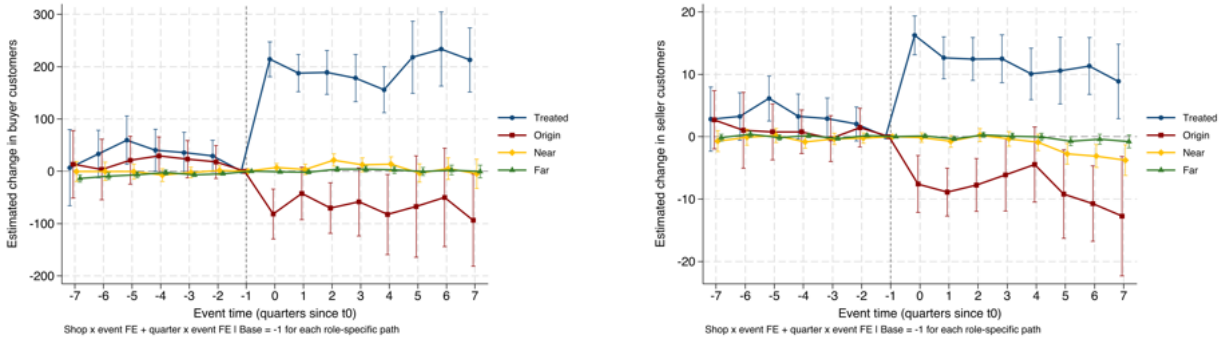
5.3 Where Does the Additional Demand Come From?

Coworkers absorb the customers a star draws into the receiving shop. Where those customers come from determines whether star placement merely reallocates demand within the firm or expands the firm's overall customer base. The demand-side mechanism points to two channels. The first is a portable demand channel. Because a star has accumulated reputation, experience, and client relationships in her prior market, some newly arriving customers who would otherwise have entered the firm through her origin shop may instead enter through the receiving shop after she moves. The

second is a local visibility channel. Once the star is placed in the receiving shop, her presence may make that shop more visible and attractive to buyers in the surrounding market through storefront promotion and neighborhood marketing.

We begin by tracing the source of the additional demand across shops, estimating demand changes at the receiving shop, the star's origin shop, and surrounding shops grouped by their proximity to the receiving shop. Figure 5 shows that buyer inflow rises at the receiving shop and falls at the origin shop, while nearby and more distant shops show no comparable change. Consistent with the portable demand channel, part of the receiving shop's gain reflects customer-flow reallocation from the star's origin shop rather than displacement from neighboring shops. At the same time, the gain exceeds the origin shop's loss. In the ring specification, the contemporaneous decline at the origin accounts for only about two-fifths of the receiving shop's event-quarter increase, and the remainder is not offset by declines at any other shop. This larger part is therefore net new demand the star brings to the firm rather than demand reallocated across its shops. Star placement thus does more than redistribute customer flow across existing shops; it expands the demand captured by the firm. This localized, neighbor-neutral spillover also explains a feature of our design: because the reassignment directly affects the origin shop but not surrounding shops in the same way, excluding only the origin shop from the control group mitigates the main source of treatment-induced contamination.

Figure 5. Demand Changes Across Receiving, Origin, and Nearby Shops



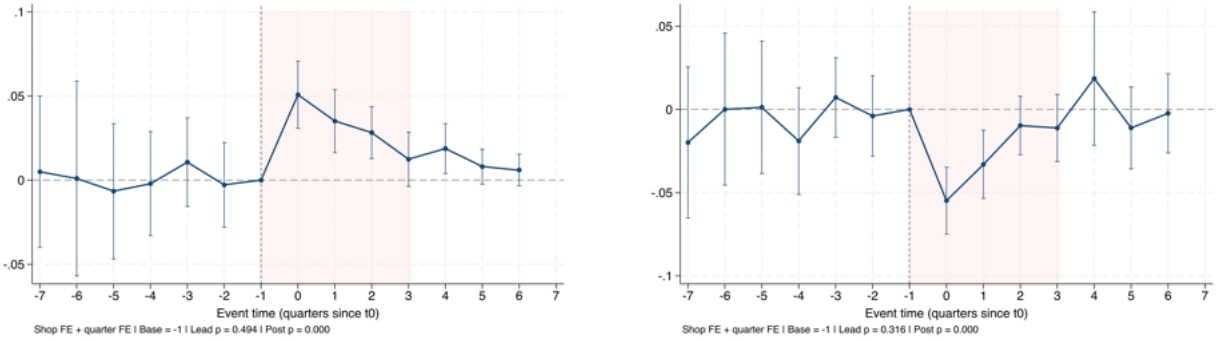
(a) Star Inflow: Buyer **(b)** Star Inflow: Seller

Notes: This figure adapts the baseline event-study design to a stacked ring specification around each star inflow event. Each event stack retains the receiving shop, the star's origin shop, and surrounding shops, and estimates separate event-time paths for each group. "Treated" denotes the receiving shop; "Origin" denotes the shop the star leaves; "Near" denotes shops in the same business circle as the receiving shop; and "Far" denotes shops in the same district but outside the receiving shop's business circle.

The composition of the incoming buyers confirms that part of the gain is portable. If part of the additional demand is linked to the star's prior shop, the buyers arriving at the receiving shop should become more similar to the clientele of the star's origin shop after the move. To test this prediction, we compare incoming buyers at the receiving shop with buyers at the star's origin shop along two

dimensions: search area and budget. We construct two similarity metrics: Sim_to_Outflow$_{st}$, which captures how closely the receiving shop's incoming buyers resemble the star's origin-shop clientele in their search areas, and Price_Gap$_{st}$, which captures their similarity in search budgets. Appendix EC.2.1 provides the exact construction of these measures. Figure 6 shows that both measures move toward the origin shop's clientele profile immediately after the star arrives. Search-area similarity rises and the budget gap narrows, with both effects peaking around the event quarter. The incoming buyers are therefore not simply an arbitrary draw of local demand. In the immediate post-event period, they resemble the clientele associated with the star's prior market. Together with the contemporaneous decline at the origin shop documented above, this pattern suggests that part of the receiving shop's gain reflects origin-linked portable demand tied to the star's established, market-specific reputation and experience.

Figure 6. Similarity of Incoming Buyers to Star's Origin Shop Clientele



(a) Search Area Similarity **(b)** Search Budget Gap

Notes: Event-study coefficients estimated on the sample of treated shops, with shop and quarter fixed effects. The sample is restricted to treated shops, since both outcome measures are defined relative to the inflowing star's origin shop.

The resemblance to the origin shop gradually fades, but the buyer-inflow effect remains elevated. This combination is informative about how the source of star-generated demand changes over time. The portable component is most visible immediately after reassignment, when the star's prior market-specific reputation and experience are still current and most relevant to customers from her former service segment. As the star spends more time in the receiving shop, this origin-linked component weakens as her origin-market experience depreciates over time. The persistence of buyer inflow, meanwhile, points to the local visibility channel and is consistent with the net new demand identified in the ring analysis: as the star builds visibility and local expertise at the receiving shop, locally generated demand replaces the fading portable component and sustains the elevated inflow. Thus, the receiving shop's demand gain initially reflects a portable, origin-linked component, while its persistence increasingly comes from demand the star captures through her visibility and

experience in the new local market.

These results show that star reassignment does not simply move a fixed pool of customers across shops. Part of the inflow is portable: customers are redirected from the star's origin shop, drawn by the market-specific reputation and experience she built in her prior market. The larger part, by contrast, is not pulled from neighboring shops. It represents additional market demand that the firm captures through the star's visibility at the receiving shop. The mix of demand then shifts over time. Because the star's prior expertise is geographically bounded and depreciates, the origin-linked component recedes. Meanwhile, demand sustained by her local visibility and newly accumulated market experience comes to dominate. Star placement therefore reshapes the geographic composition of the firm's customer demand: it redirects some origin-linked demand toward the receiving shop while also helping the firm capture new demand in the market where the star is placed.

Wherever it originates, the demand drawn by the star reaches the receiving shop and is passed on to coworkers. The spillover from a star to her coworkers therefore runs through the demand they face rather than through what they learn from her or how hard they work alongside her. Whether that demand becomes output is a separate question, one that turns on whether the receiving shop can serve and convert the buyers she draws in.

6 Capacity Constraints and Star Allocation

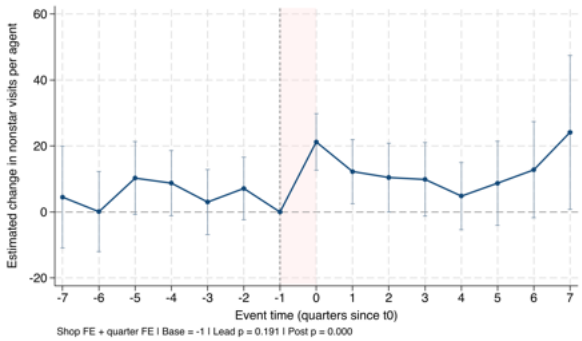
The results in Section 5 show that star inflow expands buyer demand and redistributes that demand to coworkers within the receiving shop. A star is not simply a highly productive individual whose contribution can be evaluated in isolation, but a scarce organizational resource whose placement changes the demand faced by an entire unit. The focus of star deployment therefore shifts away from post-arrival interactions between the star and her coworkers, such as collaboration or learning, and toward the receiving shop's pre-event operating conditions. What matters is whether the shop, before the star arrives, is positioned to absorb the additional demand and convert it into transactions. When this capacity is limited, the demand a star brings may not fully translate into output.

6.1 Coworker Capacity and Congestion Effects

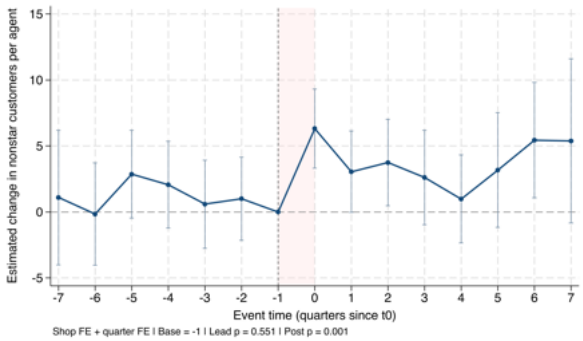
Coworkers absorb a substantial share of the additional buyer inflow that star arrival brings. We next ask whether the resulting rise in their activity reflects a peer-effort response or a demand-driven increase in workload. Under the first, a supply-side channel, the star's presence raises coworkers' effort through peer pressure. A larger workload instead reflects a passive response: coworkers serve more customers because the star draws them into the shop, not because they work harder. Panel

(a) of Figure 7 shows that showings per non-star agent rise after the event, but this could reflect either more customers per coworker or more service per customer. Panels (b) and (c) distinguish the two: customers per coworker rise, while showings per customer remain flat throughout the event window. Coworkers therefore conduct more showings only because they serve more customers, not because they devote more attention to each one. The added activity loads on caseload rather than effort. A peer-effort channel would instead surface on the intensive margin, with more service per buyer and higher conversion. We find neither. Per-customer service intensity does not move, and conversion does not improve around star inflow (Figure 7, Panel (d)), contrary to what learning or effort channels would predict. The composition of the inflow discussed in Section 5.3 reinforces this interpretation. Incoming buyers tilt toward the star's origin-shop clientele, a pattern coworker effort cannot produce. Coworkers therefore do not raise output by intensifying effort for each customer after a star arrives; they absorb the larger pool of buyers drawn into the shop.

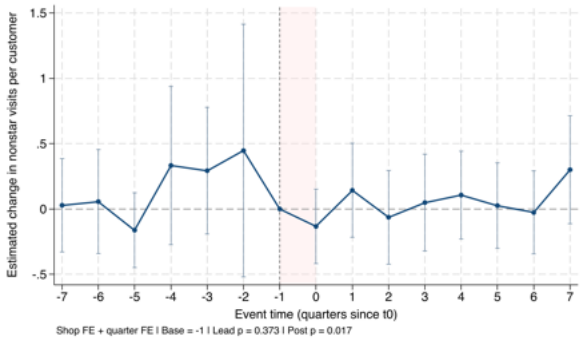
Figure 7. Service Burden, Service Intensity, and Conversion among Non-Star Coworkers



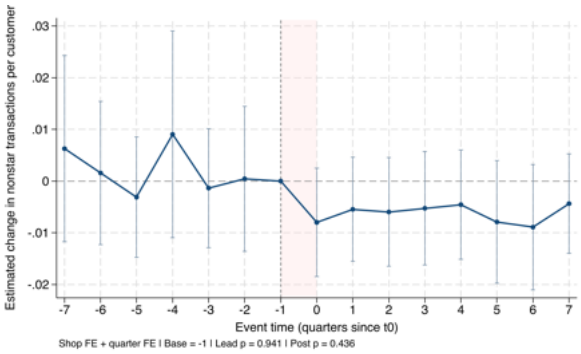
(a) Showings per Non-Star Coworker



(b) Customers per Non-Star Coworker



(c) Showings per Customer



(d) Conversion Rate

Higher caseloads, unchanged per-buyer intensity, and no improvement in conversion together indicate that coworkers' ability to absorb and convert buyer inflow is bounded: the shop can take on more customers, but not without strain. If congestion is rising inside the shop, it should not be

confined to the agents themselves. Buyers served by coworkers operating closer to capacity should also face a deterioration in their search conditions. We measure the buyer's search experience along three dimensions: how many days a buyer spends in the search process (search_dur), how many properties a buyer visits (n_search), and how intensively a buyer searches per day (search_density). Table 5 reveals exactly this pattern. After star inflow, buyers transacting through treated shops spend roughly 37 additional days in the search process and conduct about 2.1 more searches before signing. At the same time, search density declines, indicating that buyers search longer and more extensively, but not more intensively. They are not exerting more search effort in pursuit of a better match; rather, the process itself becomes slower and less efficient. As agents become busier serving the additional demand, buyers must wait longer for agents to arrange property showings. These results are robust across a range of alternative specifications, winsorizing search outcomes at the 1st and 99th percentiles, re-estimating in logs, applying Lee (2009) bounds for selection into the transacting sample (Appendix EC.6).

Table 5. Buyer Search Frictions and Shop Congestion

	Average Effects			With Pre-Event Shop Congestion		
	search_dur	n_search	search_density	search_dur	n_search	search_density
	(1)	(2)	(3)	(4)	(5)	(6)
PostStar$_i$	36.925***	2.141***	-0.170***	42.721***	2.318***	-0.207***
PostStar$_i$× ln($\overline{Cong}_{s(i)}$)				21.287***	0.636	-0.133
X$_i$	✓	✓	✓	✓	✓	✓
Shop FE	✓	✓	✓	✓	✓	✓
Biz FE	✓	✓	✓	✓	✓	✓
Start-Quarter FE	✓	✓	✓	✓	✓	✓
Observations	87,853	87,853	87,853	87,187	87,187	87,187

Notes: $PostStar_i$ equals one for buyers signing at treated shops after the star inflow. Shop's congestion $\overline{Cong}_{s(i)}$ denotes the average ratio of customers to agent headcount over the $[-4,-1]$ pre-event quarters for treated shops and over all observed quarters for control shops. We also control for buyer characteristics through X_i = (ln(First search price), ln(First search size)). Standard errors clustered at the shop level. All specifications absorb buyer search business circle fixed effects and search start quarter fixed effects and shop fixed effects.

A natural question is whether this deterioration is specifically driven by congestion at the receiving shop, or whether it instead reflects other changes coinciding with star arrival. To distinguish between these possibilities, we examine whether the friction is more severe at shops that were already operating under congested conditions before the event. For each shop, we measure pre-event congestion as the average ratio of customers to agent headcount over the four quarters preceding star inflow, and test whether the post-star effect on buyer search outcomes is amplified at shops with higher pre-event congestion. Columns (4)-(6) of Table 5 report the interaction estimates.

Pre-event congestion substantially amplifies the effect of star inflow on search duration: doubling a shop's pre-event congestion adds roughly 15 days to the post-star increase in search duration, equivalent to about 35% of the average post-star effect on duration. By contrast, the interaction terms for the number of searches and for the search density are smaller and statistically insignificant. This asymmetry is informative. Among the three outcomes, search duration most directly captures the time buyers wait while the shop serves them; the other two reflect the buyers' own search behavior more directly. The fact that pre-event congestion amplifies the duration effect but not the behavioral margins indicates that the friction documented above is driven by capacity pressure at the receiving shop rather than by other shop-level changes coinciding with star arrival.

The evidence therefore shows that the demand expansion documented in Section 5 is absorbed at a cost. Agents take on more customers without converting them more effectively, buyers wait longer to complete their search, and both costs are sharper where the shop's pre-event capacity was already tight. This carries a direct implication for star deployment: the value of placing a star at any given shop depends not only on how much demand she attracts, but also on the shop's capacity to absorb and convert that demand into transactions. Where stars should be deployed therefore depends on the receiving shop's pre-event conditions.

6.2 Heterogeneous Returns to Star Placement

The congestion costs documented in Section 6.1 are not uniform across shops. They are sharper when the receiving shop is less well positioned to handle the additional buyer demand before the star arrives. This heterogeneity has a corresponding implication for the benefits of star inflow. If stars generate spillovers by attracting demand that coworkers then absorb and convert, the value of star placement should depend on the receiving shop's capacity to absorb and convert that demand. Star placement therefore becomes an internal allocation problem: placing the same scarce star in shops with different operating conditions can generate different value, and the firm's challenge is to direct each star toward a shop with the room and conversion ability to handle the demand she brings.

We examine two pre-arrival dimensions of the receiving shop's capacity to absorb and convert demand. The first is the shop's demand environment, which captures how busy its customer flow already is, in turn, how much room remains to accommodate additional buyers. The second is the shop's conversion ability, which captures how effectively its agents turn customer flow into transactions and therefore whether additional inflow is likely to translate into output. Together, these dimensions summarize the receiving shop's capacity to absorb and convert the demand a star brings. Empirically, the demand environment is measured by the pre-event customer-to-agent ratio, with higher ratios indicating busier shops with less room to absorb additional buyers. Conversion ability is measured by a pre-event conversion-ability index: a shop's staffing residual after partialling out

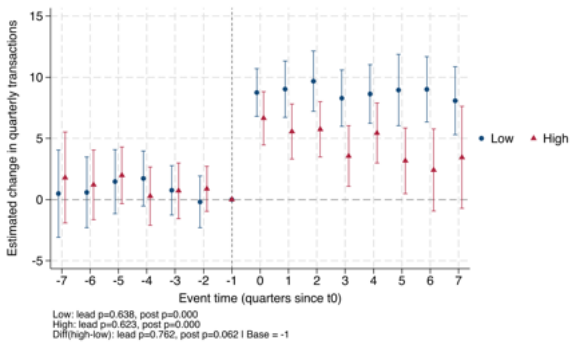
its demand level, so that shops serving a given demand with fewer agents reveal higher conversion ability. Because the conversion-ability index nets out demand, it is conceptually distinct from the customer-to-agent ratio rather than a transformation of it, and the two are only weakly correlated (Pearson 0.16, Spearman 0.16). The two measures therefore capture separate dimensions of a shop's capacity. Construction details and a balance check are in the Appendix EC.2.2.

We test these implications by splitting treated shops at the median of each pre-arrival measure and estimating event study coefficients separately for the corresponding Low and High groups. Panels (a) and (b) of Figure 8 examine heterogeneity by the shop's demand environment. Returns to star inflow are larger at shops with less congested pre-event customer flow, where agents have more room to accommodate additional buyers. At low-congestion shops, star arrival increases quarterly transactions by 8.8 on average across the post-event window, roughly double the gain of 4.5 transactions per quarter at high-congestion shops. The contrast for commission revenue is similarly pronounced: RMB 492,400 per quarter at low-congestion shops, compared with RMB 237,600 at high-congestion shops. Pre-trends are flat in both groups for both outcomes. And the post-event revenue gap is statistically significant (post gap $p = 0.025$), while the transactions gap is economically large and marginally significant (post gap $p = 0.062$). The pattern indicates that the same star creates substantially more value when she joins a shop where agents have greater bandwidth to serve the additional buyers she draws.

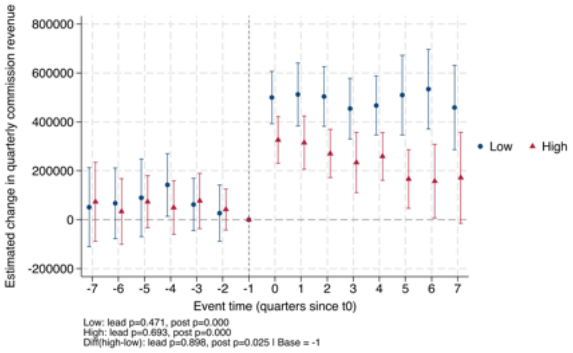
Panels (c) and (d) turn to conversion ability. Shops that serve a given level of demand with fewer agents than predicted, classified as having higher conversion ability, gain substantially more from star arrival. At high-conversion shops, star inflow increases quarterly transactions by 9.1 on average across the post-event window, more than double the gain of 4.2 transactions per quarter at low-conversion shops. Commission revenue follows the same pattern, with average post-event gains of RMB 491,600 versus RMB 239,200 per quarter at high-conversion and low-conversion shops, respectively. The two groups diverge significantly after star arrival, with post-event gap tests yielding $p = 0.025$ for transactions and $p = 0.002$ for commission revenue. This second dimension shows that receiving-shop capacity is not only about having enough bandwidth to serve additional buyers. It also depends on whether agents can turn that buyer flow into completed transactions. The demand a star brings generates the largest output gains when it is routed to shops with stronger pre-event conversion ability.

These results reveal a deployment principle that distinguishes our setting from the supply-side spillover literature. When star spillovers operate through skill transfer, peer learning, or collaboration, the value of star placement depends primarily on post-arrival interactions: whom the star can mentor and where collaboration is most productive. In our setting, spillovers operate through demand attraction and coworker absorption. The relevant deployment criterion is therefore the operating state of the receiving shop before the star arrives: the demand environment its agents already face

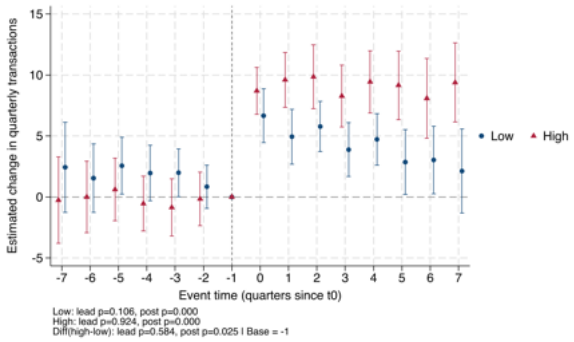
Figure 8. Heterogeneous Returns to Star Placement by Pre-Event Demand Environment and Conversion Ability



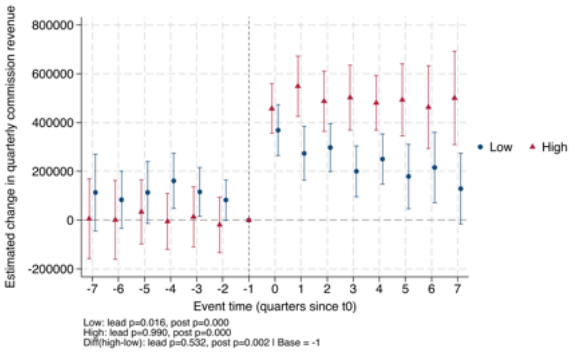
(a) Demand Environment: Transactions



(b) Demand Environment: Commission Revenue



(c) Conversion Ability: Transactions



(d) Conversion Ability: Commission Revenue

Notes: The figure shows event-study estimates of star inflow effects on quarterly shop performance, split by pre-event shop conditions. Shops are classified as 'Low' or 'High' based on the median of each pre-event measure.

and the efficiency with which the shop converts customer flow into output. A star creates more value when she joins a shop whose operating conditions allow the additional buyers she attracts to be served, processed, and converted into transactions. Whether actual star allocations align with these conditions remains an empirical question. The next subsection quantifies the aggregate value that could be gained from a counterfactual allocation guided by this deployment principle.

6.3 Counterfactual Star Deployment

The heterogeneity results in Section 6.2 show that the returns to star inflow depend systematically on the receiving shop's pre-arrival operating conditions. Stars generate larger performance gains when they join shops with more room to serve additional buyers and stronger ability to convert customer flow into transactions. This raises a natural managerial question: how much value could the firm realize if star deployment were guided by these conditions?

We conduct a counterfactual reallocation exercise to quantify this. The analysis is not intended to

recover the firm's globally optimal assignment of stars, which would require observing all constraints faced by managers when moving agents across shops. Instead, it provides a transparent estimate of the potential gains from aligning star placement more closely with the deployment principle identified above. We compare the realized allocation of star inflow events with counterfactual allocations that reassign the same stars across eligible shops within the same market-period. We consider two simple deployment rules based on the pre-arrival conditions identified in Section 6.2. Each rule reallocates stars from shops that are unfavorable on one capacity dimension to eligible shops that are favorable on the same dimension, while leaving unchanged stars already in favorable shops. The low-congestion rule targets the demand environment, reallocating stars from high-congestion to low-congestion shops. The high-conversion rule targets conversion ability, reallocating stars from low-conversion to high-conversion shops.

Table 6 reports the predicted cumulative outcomes over the first four post-inflow quarters. Both rules yield substantial aggregate gains relative to the realized allocation. Directing star inflows to low-congestion shops is predicted to raise aggregate transactions by about 17.2% and commission revenue by an additional RMB 30.1 million, a gain of roughly 15.4%. The high-conversion rule yields even larger improvements: transactions are about 22.4% higher and commission revenue rises by an additional RMB 51.7 million, or roughly 26.5%. Gains of this magnitude indicate that the receiving shop's operating conditions are a first-order determinant of the value stars create at the firm level. Because the counterfactuals hold fixed the number and timing of star inflows, the improvement comes entirely from changing where the same scarce talent is deployed within the same market-period. The firm could realize these gains without hiring or retaining additional stars. The consistency of the gains across two distinct capacity dimensions reinforces the deployment principle established in Section 6.2: a star's value depends as much on the operating conditions of the shop she joins as on the demand she brings. These figures are model-based predictions, computed by applying the Section 6.2 heterogeneity estimates to a reassignment of the same stars, and the assumptions they rest on warrant attention.

Several qualifications apply. First, the heterogeneity estimates in Section 6.2 are identified from stars the firm actually placed in low-congestion or high-conversion shops. Reassigning stars to shops the firm chose not to target requires the same returns to hold off the support of observed placements; if placements are themselves selective, the returns to a newly designated target shop need not match those estimated on currently favorable shops. Second, the exercise abstracts from operational constraints that may limit the firm's flexibility in reallocating stars in practice, such as ongoing client relationships and managerial constraints. Third, the gains are computed using heterogeneity estimates from Section 6.2, which capture cross-sectional differences under the actual deployment pattern. Large-scale reallocation could in principle alter these returns through general equilibrium effects. The counterfactual gains should therefore be read as an illustrative comparative static under

Table 6. The Value of Matching Stars to Shop Operating Conditions

Allocation Rule	Aggregate Transactions	Transactions Gain	Aggregate Commission	Commission Gain
Actual Allocation	3,661	-	194,978,016	-
Low-congestion rule	4,290	+17.176%	225,047,760	+15.422%
High-conversion rule	4,482	+22.416%	246,706,128	+26.530%

Notes: The table reports predicted cumulative gains over the first four post-inflow quarters. The actual allocation uses the observed receiving shops of star inflows. In the counterfactual analysis, each star inflow is reassigned to eligible shops within the same market-quarter according to pre-event shop operating conditions. The low-congestion rule assigns stars to shops with lower pre-event congestion, while the high-conversion rule assigns stars to shops with higher pre-event conversion ability. Congestion and conversion ability are defined as in the preceding section. Gain is defined as the percentage change in predicted cumulative gains relative to the actual allocation.

explicit assumptions rather than as the gain the firm would necessarily realize in practice. Even with these qualifications, the magnitudes are informative. The counterfactual analysis underscores a central managerial implication: when star employees function as scarce demand magnets, firms create the most aggregate value by assigning them to units that are operationally prepared to serve and convert the demand they attract.

7 Discussion

7.1 Rethinking Star Spillovers in Employee-linked Demand Settings

The literature on star employees has shown that high-performing workers can have large effects on their coworkers. These effects may be positive or negative, but the mechanisms emphasized in prior work typically operate on the supply side of production. In research and scientific work, star scientists affect colleagues through collaboration and knowledge spillovers. In settings where workers operate in close proximity, high-performing peers can shape coworkers' effort through social pressure, monitoring, or social incentives. In team-production settings, high performers can also affect coworkers through complementarities in production, where one worker's contribution changes the marginal value of others'. These settings differ in important ways, but they share a common structure: the star affects coworker output by changing the production side of coworkers' work: what coworkers know, how much effort they exert, or how their inputs combine with the star's.

Our findings point to a broader class of settings, such as real estate brokerage, financial advisory, health care, and law, in which this production-side view is incomplete. Many of these settings fall within or adjacent to the broad domain of professional service firms, whose organization depends heavily on knowledge-intensive and professionalized labor, yet whose expertise and client relationships vary across professional sectors (Von Nordenflycht, 2010; Malhotra and Morris, 2009).

In these settings, service delivery is relatively individualized, and customers can identify, evaluate, or select particular employees. Demand is therefore shaped not only by the firm or unit, but also by the visibility, reputation, track record, and prior customer ties of particular workers. We refer to these contexts as employee-linked demand settings. This concept highlights a demand-side dimension that is not central to existing classifications of professional service firms: the extent to which customer attention and demand attach to particular employees. This structure creates a demand-side route for star spillovers. A star can affect coworkers even without first changing their skills, effort, or collaboration. She may instead redirect customer attention toward the unit she joins, and the unit then becomes the organizational channel through which this demand is allocated, served, and converted. Such demand-side effects may also exist in the settings emphasized by prior work, but they are often difficult to observe or separate from changes in workers' own productivity, collaboration, or learning. In employee-linked demand settings, by contrast, the demand channel is more salient and empirically separable because reputation and client appeal can operate outside the workplace production process. Our evidence identifies this mechanism in real estate brokerage: star movement changes the customer flow of the receiving shop, and coworkers benefit because the set of customer opportunities available to them expands.

Recognizing this mechanism reframes how we think about star spillovers. The point is not that the supply-side mechanisms documented in prior work are absent or unimportant. Rather, those mechanisms are most visible in settings where production is collaborative, knowledge-intensive, or jointly performed. Employee-linked demand settings point to a different possibility: star spillovers operate by reshaping the demand environment faced by coworkers. The spillover should therefore be understood not only from worker to worker, but from the star's visibility and client base to the shop's customer flow, which in turn increases coworker output. The mechanism through which stars matter depends on the production environment, and different mechanisms generate different organizational implications.

7.2 Implications for Star Deployment

The demand-side mechanism documented in this paper shifts the focus of the deployment problem. Under supply-side mechanisms, stars affect their organizations through interactions with coworkers, so the effect of a star depends on what unfolds between the star and her coworkers: what coworkers learn from her, how her presence shapes peer effort, or where production complementarities are strongest. Under the demand-side mechanism, the star's primary effect operates less through changes in coworkers' productivity, effort, or modes of collaboration than through changes in the demand environment of the workplace. The placement decision therefore turns on whether the receiving unit can capitalize on that changed demand environment: whether it has room to serve the customers she

draws and the conversion ability to turn that demand into output.

This distinction implies that star deployment is best understood as a matching problem between demand-drawing talent and the receiving shop's ability to absorb and convert customer inflows (Weller et al., 2019; Schmidt et al., 2024). This match is tractable because the relevant shop operating conditions are observable before the star moves. Such observability allows the firm to compare the expected value of scarce talent across alternative internal deployments, a central challenge in resource redeployment (Levinthal and Wu, 2025). Using operational data on pre-period demand pressure and conversion ability, the firm can identify which shops are best positioned to serve the demand a star would bring and turn it into output. The firm therefore does not need to wait to observe the outcome of star placement or rely only on predictions about how the star will interact with future coworkers. Placement becomes a lever the firm can pull deliberately, rather than an outcome it discovers after the fact.

A more counterintuitive implication follows from this matching logic. The natural instinct is to send stars to a firm's busiest or most prominent units, where demand is already high and the star's visibility may appear especially valuable. But this instinct misreads where the value comes from: a unit already straining to convert its existing demand gains little from more, since the additional customers that a star attracts worsen the congestion, and translate into fewer additional transactions. A unit with room to spare turns that same demand into completed transactions and gains that spread across its agents. Stars create the most value where the demand they bring can be most fully absorbed.

The counterfactual exercise in Section 6.3 quantifies what is at stake in getting this match right. Reallocating the same stars toward units with more favorable pre-arrival conditions is predicted to raise aggregate transactions and commission revenue by roughly 20 percent, holding fixed the number of stars and the timing of their moves. These gains come not from acquiring more talent or extracting more effort from existing stars, but from deploying the same scarce talent where it can be most productively absorbed. For multi-unit firms in customer-facing settings, the broader implication is that star deployment is a strategic lever for improving the productivity of existing human capital. Firms can create value not only by relying on stars to transmit expertise or enhance collaboration, but also by placing them in units where the demand they generate can be absorbed, converted, and translated into aggregate output.

8 Conclusion

Multi-unit firms must decide where to deploy their most talented employees, yet the returns to that decision depend on how stars actually generate value for the units they join. This paper studies the intra-firm reassignment of real estate agents across brokerage shops, a setting in which stars

move within a single organization and their effect on a receiving unit can be isolated. Tracking shop performance, customer flows, and buyer search behavior around the shop's first arrival of a star, we document a mechanism that differs from the supply-side channels emphasized in prior work. Rather than operating primarily through knowledge transfer, peer learning, collaboration, or peer-induced effort responses, stars in this setting create value by reshaping the demand environment of the receiving unit.

The evidence assembles into a single account. Star arrival raises shop transactions and commission revenue immediately and persistently, and the gains reverse symmetrically when a star departs. Behind this pattern is a demand-side shock: the star expands the receiving shop's pool of incoming buyers, drawn by her reputation and recent market-specific experience. The expanded demand does not pass through the star herself: her own customer load remains stable after the move, while the demand is absorbed by non-star agents at the receiving shop. Per-agent buyer counts rise among these coworkers, and the increase remains visible even when we restrict attention to incumbent agents who were already at the shop before the star arrived. Part of the inflow is redirected from the star's origin shop; the larger part reflects net new demand that the firm would not otherwise have captured. The mix shifts over time: origin-linked buyers are most evident immediately after the move, when her prior-market experience is still current, while later inflows are sustained increasingly by her visibility and accumulated experience in the new location. The spillover from the star to her coworkers therefore runs through the customer opportunities they face, rather than through changes in what they know, how hard they work, or how they collaborate with her.

But demand becomes output only when the receiving shop has the capacity to serve and convert it. As coworkers take on more buyers, conversion does not improve, and buyers face a longer and less efficient search process. These pressures are sharpest in shops that were already near capacity before the star arrived. The same star therefore creates different value across shops, depending on how fully each can take on the buyers she draws and turn them into transactions. A counterfactual reallocation makes the stakes concrete: directing stars toward less congested, higher-converting shops is predicted to raise aggregate transactions and commission revenue by roughly 20 percent, without changing the number or timing of moves.

These findings reframe the managerial problem of star deployment. The dominant view treats star spillovers as a matter of fostering productive interaction after the star arrives: mentoring, collaboration, skill diffusion, or effort responses. Our evidence is not easily reconciled with these supply-side channels alone: when service is delivered relatively independently and reputation or client relationships attach to individual employees, the pattern of results instead points to a channel running from the star's visibility and client base to the unit's customer flow, and from there to coworker output. The relevant deployment criterion therefore shifts from post-arrival interaction to

pre-arrival operating conditions: the demand its agents already face and the efficiency with which they turn customer flow into output.

However, several limitations qualify the scope of the findings. The evidence comes from one firm in one industry, and the magnitude of demand-side spillovers may differ in other settings where demand attaches to individual professionals. The counterfactual gains are best read as an illustrative comparative static rather than a guaranteed return, since the exercise abstracts from operational constraints, ongoing client relationships, and broader equilibrium adjustments that may limit large-scale reallocation in practice. These limitations point to natural extensions. One is to examine settings in which demand-side and supply-side spillovers coexist, and to ask how their relative strength shapes optimal deployment. A second concerns incentives: when a star generates demand that her coworkers convert, she produces value she does not personally capture, which raises questions about how firms should compensate and retain talent whose contribution is partly external to her own measured output.

Nonetheless, the central implication is clear. For multi-unit firms in employee-linked demand settings, the internal allocation of scarce talent is itself a strategic lever. Firms create value not only by relying on stars to transmit expertise, induce effort, or strengthen collaboration, but by placing them where the demand they attract can be absorbed and converted. Crucially, these gains require no new hiring or added talent; the firm realizes them simply by deploying the stars it already has more effectively. Identifying this demand-side mechanism gives firms a principled basis for these placement decisions. Because the way a star creates value depends on the setting, recognizing the new mechanism at work is what lets a firm place her where that value is most fully realized.

References

Anand KS, Paç MF, Veeraraghavan S (2011) Quality-speed conundrum: Trade-offs in customer-intensive services. *Management Science* 57(1):40–56.

Anderson M, Magruder J (2012) Learning from the crowd: Regression discontinuity estimates of the effects of an online review database. *The Economic Journal* 122(563):957–989.

Arcidiacono P, Kinsler J, Price J (2017) Productivity spillovers in team production: Evidence from professional basketball. *Journal of Labor Economics* 35(1):191–225.

Azoulay P, Graff Zivin JS, Wang J (2010) Superstar extinction. *The Quarterly Journal of Economics* 125(2):549–589.

Bandiera O, Barankay I, Rasul I (2010) Social incentives in the workplace. *The Review of Economic Studies* 77(2):417–458.

Berry Jaeker JA, Tucker AL (2017) Past the point of speeding up: The negative effects of workload saturation on efficiency and patient severity. *Management Science* 63(4):1042–1062.

Bidwell M (2011) Paying more to get less: The effects of external hiring versus internal mobility. *Administrative Science Quarterly* 56(3):369–407.

Borusyak K, Jaravel X, Spiess J (2024) Revisiting event-study designs: robust and efficient estimation. *Review of Economic Studies* 91(6):3253–3285.

Broschak JP (2004) Managers' mobility and market interface: The effect of managers' career mobility on the dissolution of market ties. *Administrative Science Quarterly* 49(4):608–640.

Call ML, Nyberg AJ, Thatcher S (2015) Stargazing: an integrative conceptual review, theoretical reconciliation, and extension for star employee research. *Journal of Applied Psychology* 100(3):623.

Callaway B, Sant'Anna PHC (2021) Difference-in-differences with multiple time periods. *Journal of Econometrics* 225(2):200–230.

Chetty R, Friedman JN, Rockoff JE (2014) Measuring the impacts of teachers II: Teacher value-added and student outcomes in adulthood. *American Economic Review* 104(9):2633–2679.

Elfenbein DW, Sterling AD (2018) (When) is hiring strategic? Human capital acquisition in the age of algorithms. *Strategy Science* 3(4):668–682.

Fang L, Hayunga DK (2024) The impact of real estate agents' expertise on house prices and TOM. *Real Estate Economics* 52(1):45–72.

Gilbukh S, Goldsmith-Pinkham P (2024) Heterogeneous real estate agents and the housing cycle. *The Review of Financial Studies* 37(11):3431–3489.

Goodman-Bacon A (2021) Difference-in-differences with variation in treatment timing. *Journal of Econometrics* 225(2):254–277.

Groysberg B, Lee LE (2009) Hiring stars and their colleagues: Exploration and exploitation in professional service firms. *Organization Science* 20(4):740–758.

Groysberg B, Lee LE, Nanda A (2008) Can they take it with them? The portability of star knowledge workers' performance. *Management Science* 54(7):1213–1230.

Gubler T (2019) Connected, but qualified? Social affiliations, human capital, and service professional performance. *Organization Science* 30(5):912–936.

Han L, Strange WC (2015) The microstructure of housing markets: Search, bargaining, and brokerage. *Handbook of Regional and Urban Economics* 5:813–886.

Han S, Ravid SA (2020) Star turnover and the value of human capital—evidence from broadway shows. *Management Science* 66(2):958–978.

Hendricks JL, Call ML, Campbell EM (2023) High performer peer effects: A review, synthesis, and agenda for future research. *Journal of Management* 49(6):1997–2029.

Huckman RS, Pisano GP (2006) The firm specificity of individual performance: Evidence from cardiac surgery. *Management Science* 52(4):473–488.

Jackson CK, Bruegmann E (2009) Teaching students and teaching each other: The importance of peer learning for teachers. *American Economic Journal: Applied Economics* 1(4):85–108.

KC DS, Terwiesch C (2009) Impact of workload on service time and patient safety: An econometric analysis of hospital operations. *Management Science* 55(9):1486–1498.

Kehoe RR, Lepak DP, Bentley FS (2018) Let's call a star a star: Task performance, external status, and exceptional contributors in organizations. *Journal of Management* 44(5):1848–1872.

Lazear EP (2009) Firm-specific human capital: A skill-weights approach. *Journal of Political Economy* 117(5):914–940.

Lazear EP, Shaw KL, Stanton CT (2015) The value of bosses. *Journal of Labor Economics* 33(4):823–861.

Lee DS (2009) Training, wages, and sample selection: Estimating sharp bounds on treatment effects. *Review of Economic Studies* 76(3):1071–1102.

Levinthal DA, Wu B (2025) Resource redeployment and the pursuit of the new best use: Economic logic and organizational challenges. *Strategy Science* 10(1):32–47.

Malhotra N, Morris T (2009) Heterogeneity in professional service firms. *Journal of Management Studies* 46(6):895–922.

Mas A, Moretti E (2009) Peers at work. *American Economic Review* 99(1):112–145.

Mawdsley JK, Meyer-Doyle P, Chatain O (2022) Client-related factors and collaboration between human assets. *Organization Science* 33(2):518–540.

Mawdsley JK, Somaya D (2016) Employee mobility and organizational outcomes: An integrative conceptual framework and research agenda. *Journal of Management* 42(1):85–113.

Mindruta D, Bercovitz J, Mares V, Feldman M (2025) Stars in their constellations: Great person or great team? *Management Science* 71(3):2170–2191.

Minni V (2026) Making the invisible hand visible: Managers and the allocation of workers to jobs. *The Quarterly Journal of Economics* qjag017.

Morris SS, Alvarez SA, Barney JB (2021) Dancing with the stars: The practical value of theory in managing star employees. *Academy of Management Perspectives* 35(2):248–264.

Oettl A (2012) Reconceptualizing stars: Scientist helpfulness and peer performance. *Management Science* 58(6):1122–1140.

Prato M, Ferraro F (2018) Starstruck: How hiring high-status employees affects incumbents' performance. *Organization Science* 29(5):755–774.

Raffiee J (2017) Employee mobility and interfirm relationship transfer: Evidence from the mobility and client attachments of United States federal lobbyists, 1998–2014. *Strategic Management Journal* 38(10):2019–2040.

Ray C, Nyberg AJ, Maltarich MA, Reilly G (2026) On the move: The impact of internal mobility and internal comobility on unit- and organization-level outcomes. *Academy of Management Journal* 69(1):76–98.

Rosen S (1981) The economics of superstars. *The American Economic Review* 71(5):845–858.

Schmidt J, Priem R, Zanella P (2024) Customers, markets, and five archetypical value creation logics: A review of demand-side research in strategic management. *Journal of Management* 50(6):2309–2342.

Somaya D, Williamson IO, Lorinkova N (2008) Gone but not lost: The different performance impacts of employee mobility between cooperators versus competitors. *Academy of Management Journal* 51(5):936–953.

Tan TF, Netessine S (2014) When does the devil make work? An empirical study of the impact of workload on worker productivity. *Management Science* 60(6):1574–1593.

Von Nordenflycht A (2010) What is a professional service firm? Toward a theory and taxonomy of knowledge-intensive firms. *Academy of Management Review* 35(1):155–174.

Waldinger F (2012) Peer effects in science: Evidence from the dismissal of scientists in Nazi Germany. *The Review of Economic Studies* 79(2):838–861.

Weller I, Hymer CB, Nyberg AJ, Ebert J (2019) How matching creates value: Cogs and wheels for human capital resources research. *Academy of Management Annals* 13(1):188–214.