Firm-to-Firm Access in Production Networks*

Jing Cai
University of Maryland, NBER, and BREAD

Wei Lin
Chinese University of Hong Kong

Adam Szeidl
Central European University and CEPR

July 13, 2026

Abstract

To study firm-to-firm access in production networks, we randomized supplier-client referrals among 700 firms in an industry cluster producing the Chinese brush pen. We have two results. (1) A key barrier to access was that firms undervalued potential partners, leading to a low-access trap. The referrals not only created lasting partnerships, but also improved beliefs about potential partner quality and stimulated partner search, leading firms to form additional partnerships beyond those we referred. These effects arose only when firms were subsidized to conduct an initial transaction, allowing them to learn match quality through direct experience. (2) Improving access expanded—rather than merely reallocated—the production network and economic activity. Treated firms increased their total number of partners, upgraded their products, and improved business performance. Indirect effects on prior partners were positive, indicating that gains were not offset by crowding out. Together, these results suggest that low-access traps can constrain industrial development and that improving access can unlock large gains by expanding networks and creating new economic activity.

*Emails: cai516@umd.edu, linwei@cuhk.edu.cn, szeidla@ceu.edu. We thank Jie Bai, David Card, Oleg Itskhoki, Imran Rasul, Yulu Tang, Eric Verhoogen, Chris Woodruff, Daniel Xu, seminar and conference audiences for helpful comments, and Oliver Kiss and Borui Sun for excellent research assistance. We thank the Private Enterprise Development in Low-Income Countries (PEDL), the Structural Transformation and Economic Growth (STEG), the University of Maryland, the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme grant agreement number 724501, and the Austrian Science Fund (FWF) under grant F100600 for funding. This study was registered in the American Economic Association Registry for randomized control trials under trial number AEARCTR-0014601.

1 Introduction

Firms in developing countries often stay small and fail to upgrade their products and practices. A possible reason is barriers to access in the production network. Search-and-matching frictions may constrain the set of a firm's suppliers and clients, limiting business opportunities and slowing growth. Consistent with this view, improving access to distant or foreign partners has been shown to improve firm performance (Jensen and Miller 2018, Bernard, Moxnes and Saito 2019, Atkin, Khandelwal and Osman 2017, Alfaro-Ureña, Manelici and Vasquez 2022).

Given the positive effects of access, an important question is why firms do not improve access on their own. Prior work emphasized the cost of obtaining information about potential partners (e.g., Bernard et al. 2019). This view suggests that lack of access is driven by tangible—spatial or cultural—barriers that limit the flow of information. An alternative, thus far unstudied, view is that firms may persist in a *low-access trap* because they underestimate the gains from new partnerships. Since they undervalue potential partners, they do not search for them, and therefore never obtain the evidence needed to update their beliefs. This view suggests that lack of access may constrain industries even absent tangible barriers. It also identifies a new objective for industrial policy: to correct beliefs about potential partners, which can have amplified effects by stimulating private search. Evidence on this second view is currently lacking.

Another important question—essential for understanding aggregate effects—concerns the way in which improving firm-to-firm access shapes industry equilibrium. Prior work documented that newly formed partnerships can crowd out existing ones, a *reallocation* of the production network in which the negative indirect effects dampen industry-level impacts (Jensen and Miller 2018). However, improved access can also lead to an *expansion* of the production network: an increase in the total number of partnerships, such that the new partnerships create new activity rather than merely displace existing activity. This expansion can generate large direct effects and, through spillovers, positive indirect effects that amplify aggregate gains. Direct evidence on these expansionary forces remains limited.

We investigate both questions using a field experiment that provided referrals between suppliers and clients in a spatially concentrated industrial cluster in China. Combining the experiment with a multi-year survey of firms and networks allows us to trace how improved access reshapes production networks and industry performance. We find that the referrals created lasting new partnerships, but only when firms were subsidized to make a first transaction, allowing them to learn match value through direct experience. Such subsidized referrals improved firms' beliefs about potential partners, stimulated private search, and led to additional partnerships beyond those we referred. These results suggest an undervaluation-driven low-access trap. Turning to industry equilibrium effects, we find that access increased firms' number of partners, improved business performance, and promoted upgrading in product quality and variety. Indirect effects were positive. These results suggest that improving firm-to-firm access can unlock large gains by expanding production networks and generating new economic activity.

In Section 2 we present our context and experimental design. We work with the industry producing the Chinese brush pen, a culturally important and popular product, in a cluster that is the world's leading production center. The cluster is located within a single county, and firms have operated there for a long period, making it unlikely that spatial or cultural barriers are the primary constraint on new partnerships. Most firms in the cluster can be classified into two layers in the supply chain. Input producers (*suppliers*) produce the brush head or the handle; final good producers (*clients*) assemble these inputs and sell the final good. Our sample of about 700 firms covers the vast majority of the cluster. We surveyed these firms five times between 2018 and 2023, collecting detailed data on business performance and the production network.

In 2021 we conducted our intervention. First, we constructed candidate referrals between supplier and client firms. We created potential partnerships that were similar to existing links in the network: we paired firms when one was already transacting with a firm similar in product and price to the other. These pairs were intended to be good matches. Once we constructed the candidate referrals, we created experimental variation in how we delivered them. (i) At the firm level, we created variation among both suppliers and clients in whether the firm received any referrals (*treated* versus *untreated*). This comparison sheds light on the impact of referrals. (ii) At the link level, the main variation concerned how we introduced the firms. For *information-only* referrals, we shared contact details and told firms that the candidate partner would likely be a good match.

For *subsidized* referrals, we in addition offered the client firm a discount for a first transaction. This comparison helps isolate the role of information versus direct experience.

In Section 3 we outline a conceptual framework that guides our empirical analysis. In this framework, supplier-client links are valuable because they create new business opportunities. Firms face a search cost and may undervalue the opportunities offered by potential partners. This framework makes the following main predictions. (i) With sufficient undervaluation, firms may refrain from search and never update their beliefs, creating a low-access trap. Referrals, by providing direct experience with potential partners, correct the undervaluation, leading to increased private search and additional non-referred partners. (ii) Referrals reshape industry equilibrium. Firms may *reallocate* their links towards more productive partnerships, leading to positive direct effects but negative indirect effects. Firms may also—when they have "too few" partners initially—*expand* the number of partners (the firm's degree). The additional partnerships create new business activity, leading to larger direct effects. They also avoid crowding out, and, through spillovers, can have positive indirect effects. (iii) The gains from referrals are larger for firms with lower baseline degree, because each additional partnership represents a larger proportional expansion of the firm's network.

We begin our empirical analysis of the referrals intervention in Section 4 by exploring the barriers to access. We start by looking at the impact of the referrals on the creation of new partnerships, measured with transactions after the end of the subsidy period. We find that subsidized referrals significantly increased the probability of a partnership—by about 41 percentage points—as well as the number and value of transactions. In contrast, information-only referrals had essentially zero effects. Thus, access frictions remain important even absent major tangible barriers, and lack of information about potential partners does not appear to be the main friction.

We then explore the nature of the friction. We first look at the impact of the referrals on firms' satisfaction (expected or actual) with the candidate partner. We find that the subsidized referrals increased satisfaction, suggesting that firms initially undervalued the referred partners. We then show that the referrals also increased firms' beliefs about the quality of potential partners in general, their time spent searching for partners, and their number of non-referred new partners. Together, these results provide direct evidence that the industry was initially in a low-access trap.

A natural question is how undervaluation could persist in a dense cluster where firms likely have chance interactions. We show that learning match quality requires direct trading experience with the partner, implying that chance interactions are unlikely to correct beliefs.

In Section 5 we examine how the referrals reshaped industry equilibrium. We estimate both direct and indirect effects, and organize our analysis around the reallocation and expansion effects introduced above. We begin with the impacts on the production network, using the firm's degree as the main outcome. We measure the direct effect as the impact of the treatment, and the indirect effect as the impact of *exposure*, defined as the share of a firm's prior partners that are treated. We find a positive direct effect on treated firms' degree, inconsistent with pure reallocation but consistent with expansion. We also find that the referrals *did not* have a negative indirect effect on exposed firms' degree. Although, consistent with reallocation, exposed firms did lose some links with treated partners, they offset these losses, in part by reconnecting with previous partners, preserving their overall degree. Combining the direct and indirect effects, we find that the treatment increased the total number of supplier-client links by about 24%. Thus, the referrals did not merely reshuffle existing relationships, but also expanded the equilibrium production network.

We then explore the impact of the referrals on firm performance. We begin with the direct effect. Among suppliers, we find large gains in several measures of firm performance, including a revenue effect of 24 log points. Among clients, we find mostly insignificant effects. Motivated by our framework, which predicts larger effects for low-degree firms, we then split the client sample at the median degree. Consistent with our prediction, low-degree clients exhibit large and significant improvements, including a revenue effect of 32 log points, whereas high-degree clients exhibit no detectable effects. The same heterogeneity emerges when we analogously split the supplier sample by baseline degree. In sum, the evidence suggests that access improved business performance, particularly for low-degree firms.

To examine the mechanism through which access improved performance, we look at treatment effects on intermediate outcomes. Among suppliers, referrals increased product quality, as assessed by independent industry experts. Among clients, referrals increased product variety: treated firms were more likely to introduce a second product and derived a larger share of revenue from it.

Since in this industry, second products typically have higher price and higher quality, our findings suggest complementary upgrading along the production chain. Improved access enabled clients to expand into higher-quality products, while suppliers upgraded quality to support this expansion. The expansion mechanism therefore generated new, higher-quality economic activity.

We next examine indirect effects on firm performance. A key prediction of the expansion channel is that access, when accompanied by upgrading, can generate *positive* indirect effects: improvements in treated firms can spill over to their prior partners. Consistent with this prediction, and in contrast to the reallocation logic, we find that exposure did improve several measures of firm performance, including profit, employment, and total hours worked. The positive indirect effects imply that the gains from improving access were broadly shared, rather than offset by crowding out.

Our conceptual framework suggests that the returns to improving access can be high. The low-access trap implies that firms leave gains on the table; and the expansion mechanism implies that these gains are large and not offset by negative indirect effects. Consistent with this logic, we estimate that the annual increase in reported profits generated by subsidized referrals was 9 to 25 times the cost of the subsidy. Firms receiving the information referrals could have earned these returns, but chose not to. We also estimate, under stronger assumptions, the social return to our intervention, and find comparably large returns. These results suggest that policies improving firm-to-firm access can substantially improve industry performance.

In Section 6, we explore how access can be improved more cost-effectively. Because the referrals in our main experiment were effective only when accompanied by a subsidy, in 2023 we implemented a small follow-up intervention designed to generate interest at lower cost. Specifically, we prepared a marketing offer consisting of a business flyer and, when available, a free sample, and delivered it directly to the client on behalf of the supplier. This intervention was about 40% as effective as the subsidy in creating new partnerships, but cost only 0.5% as much. Thus, careful design of how firms are introduced can substantially improve the cost-effectiveness of policies that expand access.

Our paper builds on a growing literature in development economics studying firm-to-firm access, reviewed in JPAL (2024) and Szeidl (2025). This literature examines technological and institutional ways of improving access, especially between non-local firms: mobile phones (Jensen 2007, Jensen and Miller 2018), flights (Campante and Yanagizawa-Drott 2017, Startz 2025), high-speed trains (Bernard et al. 2019), digital platforms (Bergquist, McIntosh and Startz 2026), the internet (Demir, Javorcik and Panigrahi 2026), firm meetings (Cai and Szeidl 2018, Asiedu, Lambon-Quayefio, Truffa and Wong 2026, Wiles and Houeix 2026), a firm directory (Dillon, Aker and Blumenstock 2025), and training firms to sell (Hjort, de Rochambeau, Iyer and Ao 2025).[1] This work primarily emphasizes informational frictions that limit access, such as lack of information about partner identity, product quality, trustworthiness, or market requirements. A subset of studies (Jensen and Miller 2018, Bergquist et al. 2026) also examine the equilibrium impact of access, and document reallocation towards better partnerships. Our contribution to this research is to document the role of undervaluation in limiting access, and to show that in equilibrium access can also expand the production network, creating new economic activity and broadly shared gains.[2]

Our work also relates to research studying frictions limiting firms' adoption of profitable practices (Verhoogen 2023). An influential literature in management and organizational theory highlighted that organizational learning processes can lead firms to undervalue exploration (Levitt and March 1988, March 1991, Levinthal and March 1993, Denrell and March 2001). More recent work in economics has examined undervaluation in several firm contexts, including consumers' undervaluation of product quality (Bai 2025), retailers' undervaluation of the profitability of products (Bai, Park and Shenoy 2025), and firms' undervaluation of a payment service contract (Gertler, Higgins, Malmendier and Ojeda 2025). Our contribution is to provide randomized evidence showing that undervaluation can limit firm-to-firm access.

Furthermore, we add to the growing literature documenting multiple equilibria and poverty traps in the development context (Balboni, Bandiera, Burgess, Ghatak and Heil 2022, McKelway 2025, Banerjee, Breza, Duflo and Kinnan 2025). Our contribution is to identify an undervaluation-driven low-access trap operating at the firm and industry level, in which suppressed search limits

[1] A related set of studies evaluated the impact of having foreign—presumably better—partners, including foreign buyers (Verhoogen 2008, Atkin et al. 2017, Alfaro-Ureña et al. 2022) and foreign suppliers (Amiti and Konings 2007, Halpern, Koren and Szeidl 2015).

[2] We also build on the macroeconomics literature studying production networks, especially the welfare impacts of policies and shocks, such as supplier churn, in networked economies (Liu 2019, Baqaee, Burstein, Duprez and Farhi 2025, Egger, Faber, Li and Lin 2025). Our contribution is to provide microeconomic evidence on the barriers to access and the impacts of access on the production network.

network formation and prevents firms from realizing profitable opportunities.

Finally, we contribute to the literature on equilibrium effects in development economics. Existing work documents both negative (Rotemberg 2019, Cai and Szeidl 2024) and positive (Breza and Kinnan 2021, Egger, Haushofer, Miguel, Niehaus and Walker 2022) equilibrium effects of credit and subsidy interventions. We show that improving firm-to-firm access can also generate positive equilibrium effects, because it expands rather than merely reallocates the production network.

2 Context, Data, Design

2.1 Context

Our study area is located in a prefectural city in southeastern China, an economic hub with a population of about 6 million and a GDP per capita (PPP) of approximately USD 30,000 in 2023. The city contains a county renowned for its production of the Chinese brush pen, an important product in Chinese culture (YEE 1973).[3] Brush pen production in the county dates back around 1,600 years ago and has earned it the title "Hometown of China's Writing Brush."[4] Demand for the brush pen comes from calligraphy enthusiasts, artists, students, educational institutions, tourists, and the high-end gift market since brush pens are often gifts at significant events. Within the county, production is concentrated in two main urban areas and eight rural areas, encompassing 72 streets and natural villages.[5] The county also hosts a major brush pen trading market that opens nine times per month and attracts about 50,000 traders from all over China annually.

The product. The writing brush consists of two main components: the brush head and the handle. (See Appendix D for details and pictures.) The brush head is generally made of animal hairs or nylon, while the handle is made of bamboo or wood. Product quality varies substantially, and the brush head is the primary determinant of overall quality. Quality depends partly on the type of hair used, with goat and weasel hair being considered of higher quality than mixed hair

[3] In China, prefectural cities generally contain several counties.
[4] In 2021 the county's brush pen was officially recognized as part of China's national cultural heritage (announcement at `https://www.ihchina.cn/project.html`, retrieved October 2024).
[5] Streets and natural villages are administrative units. Streets are sub-jurisdictions of urban communities, natural villages are sub-jurisdictions of rural communities.

composed of animal hair and nylon. It also depends on how well the brush head is crafted, measured by four criteria that—as local experts explained—combine into two dimensions: craftsmanship and durability.[6] Evaluating the former requires visual examination, while evaluating the latter requires writing with the pen.

Supply chain. Firms can be organized into three layers in the supply chain.

1. Raw materials: firms supplying materials such as animal hair and nylon for brush heads, and bamboo, wood, and jade for handles.
2. **Intermediate inputs:** firms producing the brush heads, firms producing the handles, and processing firms that carry out production but do not own their inputs.
3. **Final goods:** firms purchasing intermediate inputs and assembling the final products.

Many firms engage in multiple activities, for example, producing both the brush head and the final good. We classified each firm into one of the layers based on their main activity: in the firm survey we asked them to list their three main products and the associated sales shares, and defined the main product as the one accounting for the highest sales share. Our main focus is on layers 2 and 3, which contain most of the firms. We refer to the firms in these layers as suppliers and clients.

Suppliers and clients can be further subdivided by the type of their main product, which we refer to as the firm's *business type*. Supplier firms have three business types: brush head producers, handle producers, and processing firms. Client firms have two business types: brush pen producers (about 80%) and make-up pen producers. The latter manufacture make-up and painting pens, which are related to but different from the brush pen. These firms generally use mixed hair and place less emphasis on product quality. Among brush head and brush pen producers, we further define the firm's *product type* based on the hair used in its main product: goat, weasel, or mixed hair. Figure 1 illustrates this classification.

In our location, brush head production is generally manual, while handle production may involve machines. The industry largely operates under a household workshop model, with each household

[6] The "four virtues" of a writing brush are that the tip should be sharp (Jian) for precise lines and detailed strokes; the hairs should be even (Qi) for consistent and smooth writing; the brush head should be round (Yuan) to ensure smooth and fluid movement; and the brush should be resilient (Jian) to quickly recover its original shape after use. These virtues are specified in the Chinese National Standard GB/T 34854-2017: Chinese Traditional Stationery - Writing Brush (SAMR 2017).

Figure 1: Structure of data

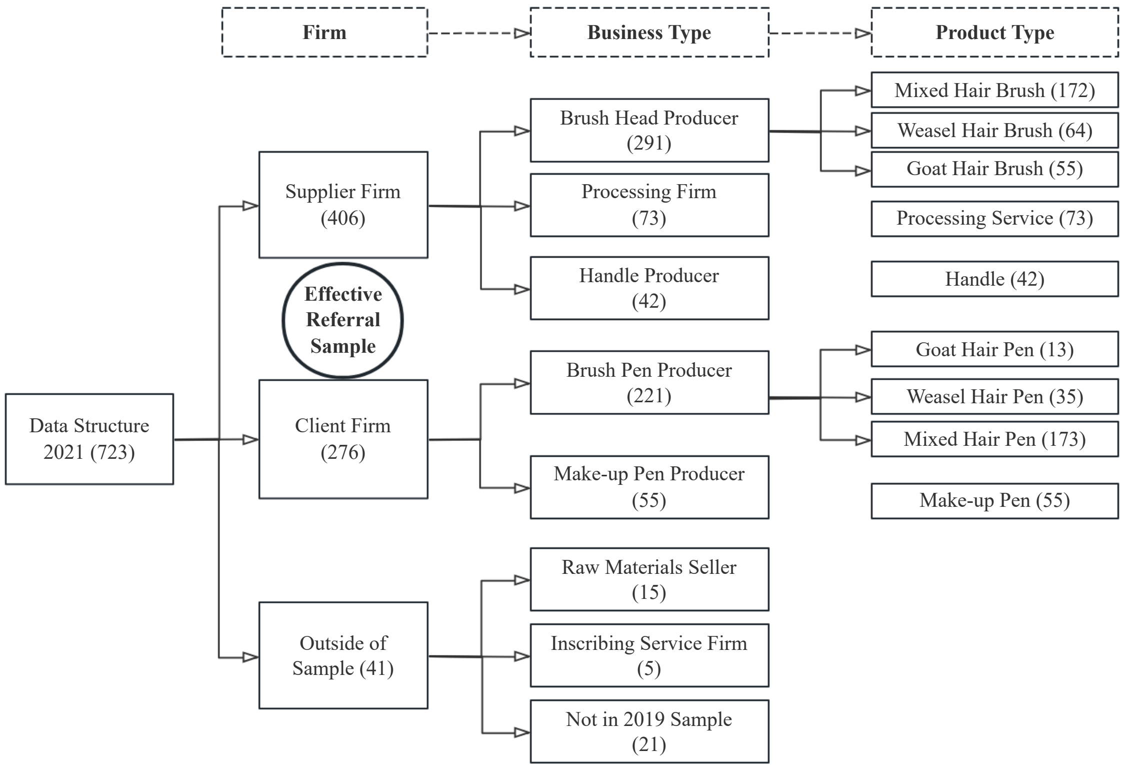

independently producing and selling (see Appendix D for pictures). This decentralization provides flexibility to meet client needs, but frictions may prevent firms from finding the best partners.

2.2 Data

We began working with the industry in the summer of 2018. We started by creating a firm census, building on two datasets we acquired from the government: (i) a list of registered brush pen producers, which primarily contained relatively larger firms; (ii) a population census, which we used to find smaller and unregistered firms. Specifically, we asked the heads of the 72 administrative units (natural villages and streets) to review their local census records and identify all households operating firms in the industry. We then merged the two datasets and asked respected elders in each administrative unit to verify the roster. From the resulting set of firms, we found 805 in our 2018 baseline survey. We believe that these firms cover the vast majority of active firms in the cluster and approximate an industry census.

Our referral intervention took place in the summer of 2021. Partly for the current project, and partly for a different project related to e-commerce, we conducted five in-person surveys: three surveys before our intervention, in 2018, 2019 and 2021 (baseline 1-3); and two surveys after, in 2022 and 2023 (follow-up 1-2). We conducted all surveys in the summer, collecting both firm and network data, except for baseline 1 where we collected firm data in 2018 summer and network data in 2019 spring. We skipped 2020 because of COVID-19, and we conducted a short phone survey in 2024. In each main survey, we collected data on nearly 90% of the firms in our census.

The firm surveys collected five categories of variables. (1) Firm performance. Sales, profit, employment, cost, and other measures. (2) Products. The firm's main products (up to three), their prices and sales shares, and measures of new product introduction. (3) Supply chain. The number of suppliers and clients, satisfaction with existing partners, and the firm's valuation of and search for new partners. (4) Business practices. Time allocation across tasks (e.g., quality control), basic management practices, and borrowing. (5) Managerial characteristics, including demographics and industrial experiences. Not all variables were collected in all survey waves.

We complemented the firm surveys with independent expert assessments of product quality. Since quality is largely determined by the brush head, we purchased one unit of the main product from most brush head and brush pen producers and asked a group of local experts to evaluate their quality. Each product was evaluated by two or three experts, all national award-winning master producers. The experts first grouped items by product type (i.e., goat, weasel, or mixed hair) and then assigned scores within each group along two dimensions: craftsmanship and durability.

An important component of the survey activity was collecting data on the production network. In the network surveys we collected information about firms' *regular* partners, i.e., other firms with which they traded on a regular basis.[7] This was initially challenging, because firms found it tedious to give detailed information about each partner. We therefore compiled lists of potential partners based on the 2018 baseline firm census, separately for different firm types, and typically containing 100 firms.[8] In the first network survey, we asked firms to identify their regular partners from that

[7] Firms also had occasional partners, but they often did not have information about the identity of those partners.

[8] For brush head and handle producer suppliers, the list contained 100 brush pen producer client firms; for brush

list, and then to report any additional regular partners. In the second network survey we followed the same approach but also asked about all previously listed partners. From the third survey on, we no longer used the lists: we asked about all previously identified partners, partners generated through our referrals (after the intervention), and any additional partners. For each partner, in each year, we collected data on the volume and number of transactions, as well as the firm's satisfaction with the relationship. We conducted the first network survey in the spring of 2019, after the first firm survey (summer of 2018), because creating the census and designing the surveys limited our capacity in 2018. In later waves we conducted the two surveys simultaneously.

Figure 1 reports the number of firms of different types in our 2021 survey, which was conducted right before the intervention and is our effective baseline. We surveyed 723 firms in total, from which we selected our main intervention sample of 406 suppliers and 276 clients.[9]

2.3 Design

We conducted our intervention in the summer of 2021. We created (i) firm-level variation in whether a firm receives referrals, and (ii) two forms of link-level variation in the type of referrals: screened versus unscreened, and information versus subsidy. For ease of presentation, we first describe the link-level variations, then the firm-level randomization, and then the implementation of the intervention. We outline the main ideas here and provide additional details in Appendix E.

Link-level randomization 1: screened versus unscreened referrals. This variation was designed to measure the impact of providing a referral "tailored" to be a good match for the firm. We used the network data to screen potential matches. Our idea was that a supplier-client pair may be a good match if a similar connection already exists: if there is a firm similar to the supplier selling to the client, or there is a firm similar to the client buying from the supplier. To implement this, for each firm we defined a set of *similar firms*, as firms that have (i) the same role (supplier or client), and (ii) the same product type for their main product. We measured the closeness of similarity

pen producer client firms, 70 brush head producers, 15 handle producers, and 15 processing firms; for processing firms, 50 brush head producers and 50 brush pen producers; for make-up firms, all 48 make-up firms (all from the 2018 baseline). In all lists, we included the largest firms by revenue.

[9] The 41 firms not included in our intervention were raw material producers (15); firms providing inscribing services, a different type of intermediate input (5); and supplier and client firms for which we lacked the 2019 data that was necessary to construct the referrals (21).

based on whether firms also had a similar price for their main product (price gap below the median across all pairs), with price interpreted as a proxy for quality.[10] Then, for each firm, we identified the set of potential partners: firms that are not yet partners but are either similar to an existing partner or existing partners of a similar firm. We ranked potential partners by the closeness of similarity and the number of supporting similar firms: e.g., a supplier selling to a closely similar firm, or to more similar firms, of a client, received a higher ranking. Finally, we generated referrals using an algorithm that ensured all firms received some referrals and no firms received too many.[11]

For unscreened referrals, we paired firms randomly, subject to the minimal requirement that their main product had to use the same type of hair (if applicable). Thus, for each brush pen producer we randomly chose a supplier which was either a handle producer, a processing firm, or a brush head producer with the same product type (goat, weasel, or mixed hair). For make-up pen producers, since they generally use a mix of hairs, we did not restrict the referred supplier.

Our surveyors, who presented the referrals to firms, were unaware of whether referrals were screened or unscreened and therefore did not communicate them differently.

Link-level randomization 2: information versus subsidized referrals. This variation was designed to shed light on the role of information frictions. In the information referrals, we told the two parties that the researchers are conducting a study in collaboration with the local government, and the researchers think that the referred firm would be a good potential business partner. In the subsidized referrals, in addition to providing the same information, we offered a subsidy to the client firm for the first transaction. The subsidy covered 50% of the transaction value, up to a cap of 1,500 RMB (approximately USD 200). The maximum subsidy per referral amounted to 0.4% of the average annual profit of client firms.

In both information and subsidized referrals, both firms received a coupon which contained the contact details of the referred partner. In the subsidized referrals, the coupon also contained information about the subsidy. The subsidy had to be claimed by the client firm, and required submitting: (1) a video of the face-to-face transaction, (2) both coupons signed by the supplier and

[10] More similar firms should also be closer competitors. However, being a competitor plausibly has another component, spatial closeness, which is not included in our similarity definition.

[11] To do this, we reformulated the problem as one of finding a minimum-cost flow in a weighted network.

Figure 2: Structure of randomization

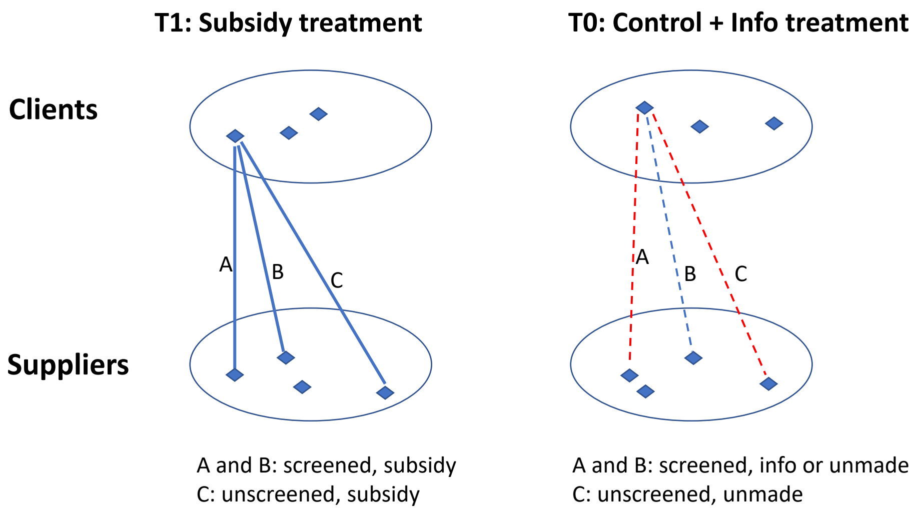

the client, (3) an invoice for the transaction, and (4) an itemized list of all goods purchased.

Firm-level randomization. To generate cross-firm variation in referrals, we randomized all suppliers and all clients into two approximately equal-sized groups, T1 and T0, which we refer to as treated and untreated firms. We then constructed candidate referrals only within groups: one set linking T1 suppliers to T1 clients and another linking T0 suppliers to T0 clients (Figure 2). No cross-group referrals were created. As we describe in more detail below, only a subset of these candidate referrals were converted into actual referrals. For each client firm—whether in T1 or T0—we constructed two screened candidate referrals, labeled A and B in the figure, and one unscreened candidate referral, labeled C.[12] Our referral creation algorithm resulted in all suppliers getting between 1 and 4 screened candidate referrals.

We converted the candidate referrals into actual referrals as follows. We implemented all the T1-to-T1 candidate referrals, all of them as subsidized referrals. Among the T0-to-T0 candidate referrals, we implemented a random half of the B candidate referrals, all of them as information referrals. We did not implement the remaining T0-to-T0 candidate referrals.

If we ignore the information referrals for a moment, then the T0-to-T0 candidate referrals serve

[12] To increase comparability between the screened and unscreened referrals, we made sure that the supplier in the unscreened referral C had the same business type as that in the screened referral A.

as a natural control group for the T1-to-T1 ones, because they were constructed in the same way, but only the T1-to-T1 ones were made as actual referrals. The information referrals seemingly confound this simple structure; but in practice they do not, because it turns out that the information referrals had no effect on subsequent transactions. Indeed, this is why we implemented the information referrals in T0, our intended control group: based on a pilot, we expected them to have no effect. Given this, including them in T0 allowed us to avoid having a third treatment arm—which could have reduced statistical power—and still make firm-level treatment-control comparisons between T1 and T0. Thus, it is valid to think about T0 as a counterfactual for T1.[13] In summary, we have four link-level treatment arms: screened subsidy, unscreened subsidy, screened information, and control; and (effectively) two firm-level treatment arms: treated and untreated.

Implementation. At the time of the 2021 intervention, the cleaned 2021 survey data were not yet available.[14] For this reason, we based our experimental implementation, including the sample definition, the randomization, and the construction of referrals, on the 2019 data. Using firms' reported main products in 2019, we classified firms as suppliers or clients. We then randomized firms into T1 and T0, stratifying by business type and baseline sales. Specifically, we divided client firms into brush pen producers and make-up pen producers; and supplier firms into brush head producers and other input suppliers (handle producers and processing firms). Then, within each of these categories, we divided firms into ten deciles based on their 2019 sales. In each of the resulting 40 strata, half of the firms were randomly assigned to T1. Given the treatment assignment, we created the referrals as described above.

We conducted the intervention in July 2021. Each treated firm received the referrals only after the survey with that firm was completed. Three issues emerged during the implementation. First, the heads of the local administrative units informed us at the beginning of the intervention that 61 firms went out of business. We constructed new referrals to replace the lost candidate partners, in both T1 and T0. Second, some client firms complained that too many of their referrals were processing firms. There were 24 clients for whom all three referrals were processing firms,

[13] Despite its null effect on transactions, we wanted to include the information arm in the intervention, because it helps exclude lack of information about the identity of partners as the barrier to access.

[14] The intervention was implemented immediately after the 2021 survey using the same enumerator team, and we could not delay implementation until the data were cleaned.

and we created three additional non-processing referrals for each of them (two screened and one unscreened), both in T1 and T0.[15] Third, when making the referrals, we learned that 7 firms in T1 had switched roles, from supplier to client or vice versa. We created new referrals for these firms. We were unable to make the analogous changes in T0, because we did not administer subsidized referrals to those firms, and the survey data was not yet ready. Nevertheless, we chose to make these changes in T1—even though they affected only 2% of the T1 sample—because otherwise these firms would have received referrals inconsistent with their role, and in this close-knit community, their dissatisfaction could have spread to other participants. All of our analysis is performed using the new status of these 7 firms. Importantly, balance tests below (using the new status of these firms) show full balance, confirming that these changes are unlikely to affect our results.[16] Our final sample consisted of 406 suppliers and 276 clients.

Finally, we note that in 2019 we introduced another intervention in the same population of firms, in which we encouraged some firms to join an e-commerce platform designed for this industry. We provided firms treated in that intervention both smartphones and training on how to use the platform. The referral treatments are balanced with respect to the e-commerce intervention.

2.4 Summary statistics and balance

Table 1 presents summary statistics in the 2021 baseline survey about both suppliers and clients. Treatment stands for T1. Columns 1-2 are about suppliers, with column 1 reporting the mean of the variable and column 2 the mean difference between treated and untreated firms. Columns 3-4 report summary statistics analogously about clients. Standard errors are in parentheses.

Panel A focuses on firm characteristics. On average, both suppliers and clients have been in operation for over 20 years. Supplier firms have 2.5, client firms 5.6 employees on average. These values include both formal and informal (e.g., family) employment. Supplier firms' average revenue is about RMB 160,000 (USD 25,000) and their average profit is about RMB 80,000 (USD 12,000) so

[15] Among the 61 firms involved in the first issue, 36 were in T1 and 25 in T0; among the 24 firms involved in the second, 12 were in T1 and 12 in T0.

[16] We also re-estimated the main specifications dropping these 7 firms, and found that the results are almost exactly identical to those reported in the paper.

Table 1: Summary Statistics

	Suppliers		Clients	
	Mean	T1-T0	Mean	T1-T0
Panel A: Firm Characteristics				
Firm Age	26.525	0.084	21.333	-1.744
	(13.326)	(1.324)	(11.863)	(1.427)
Number of Employees	2.517	0.296	5.576	-0.117
	(2.201)	(0.218)	(5.571)	(0.672)
Profit (10,000 RMB)	7.744	1.898	37.311	-0.136
	(18.768)	(1.863)	(64.603)	(7.792)
Sales (10,000 RMB)	15.806	1.966	85.509	0.996
	(33.212)	(3.299)	(140.755)	(16.976)
Panel B: Managerial Characteristics				
Gender (1=Female, 0=Male)	0.315	0.016	0.112	-0.009
	(0.465)	(0.046)	(0.316)	(0.038)
Age	53.571	-0.058	49.196	-0.423
	(10.115)	(1.005)	(10.463)	(1.262)
Education- Middle School	0.461	0.034	0.714	0.055
	(0.499)	(0.050)	(0.453)	(0.055)
Panel C: Number of Partners (buyers for supplier, sellers for client)				
In Firm Data	6.116	1.052	5.859	-0.483
	(11.783)	(1.170)	(6.789)	(0.818)
In Network Data, All Partners	2.307	0.261	3.590	-0.006
	(3.015)	(0.299)	(3.793)	(0.458)
In Network Data, Supplier-client Links	1.525	0.163	2.243	-0.098
	(2.654)	(0.264)	(3.310)	(0.399)
Length of Relationship	9.020	0.232	9.621	-0.234
	(4.993)	(0.652)	(4.888)	(0.759)
Panel D: Attrition and Shutdown				
Attrition (2021 to 2022)	0.039	-0.009	0.033	-0.008
	(0.195)	(0.019)	(0.178)	(0.021)
Attrition (2021 to 2023)	0.044	-0.039*	0.072	0.028
	(0.206)	(0.020)	(0.260)	(0.031)
Shutdown (2021 to 2023)	0.022	-0.014	0.011	0.007
	(0.147)	(0.015)	(0.104)	(0.013)
P-val of Joint Significance of Vars in Panels A-C		0.863		0.975
Observations		406		276

Notes: Column 1 reports the mean for suppliers, column 2 reports the difference between treated (T1) and untreated (T0) suppliers, and columns 3 and 4 are analogous for clients. Panels A, B, C report on 2021 baseline data. *** p<0.01, ** p<0.05, * p<0.1.

that profits account for nearly half of revenue.[17] This high share may reflect that reported profits

[17] We convert all baseline RMB values to USD using the 2021 average rate of 0.155 USD.

partly include compensation to the owner, or entry barriers: entry into this artisanal industry requires many years of training, allowing firms to earn high rents. Client firms are larger in terms of financial measures: their revenue is about RMB 860,000 (USD 133,000) and their profit is about RMB 370,000 (USD 57,000). The large standard deviations indicate much cross-firm heterogeneity. Panel B reports managerial characteristics. Essentially all firms are family firms managed by their owner, who is typically a man aged about 50.

Panel C provides summary statistics on the network. We report the mean number of (regular) partners, where partner is defined as a buyer for supplier firms and a seller for client firms. In the firm survey, both suppliers and clients report about 6 partners. In the network survey, suppliers report about 2.3, and clients about 3.6 partners. Averaging across all firms, our network data covers about half of the partnerships reported in the firm survey. Much of the gap comes from firms with many partners, where we do not expect our network data to be comprehensive.[18] For example, among firms with degree ≤ 7 (about 80% of the sample), the network data covers 87% of the number of partnerships reported in the firm survey, indicating that coverage is quite comprehensive for most firms. For further cross-validation, we compared link-level transaction records from the network survey with firm-level revenues and costs from the firm survey. Suppliers' sales in the former account for 81% of their revenue in the latter, and clients' purchases in the former account for 77% of their input cost in the latter. Thus, the network survey covers the most important links well. The third row in Panel C reports on links that are confirmed to be between a supplier and a client firm. Suppliers sell to about 1.5 firms in the client sample, and clients buy from about 2.2 firms in the supplier sample. The difference between this row and the previous reflects that some links in the network data involve a partner whose identity we could not determine, and others capture a non-supplier-to-client transaction, such as a supplier selling to another supplier. Supplier-client links are stable: the average partnership is more than 9 years old.

Consistent with the randomization, there is no significant difference between the treatment arms in any of the variables in Panels A-C. Moreover, as the p-values in the bottom show, regressing the treatment indicator on all the variables in Panels A-C yields no evidence of joint significance.

[18] In the network survey firms had to identify partners individually, and each firm could name up to 10 partners.

Panel D reports attrition and shutdown. Attrition is defined as one in a survey wave if we do not have information about the firm in that wave. This is typically due to the manager not being available at the time of our visit. Shutdown is defined as one in a survey wave if we have information that the firm was not in business at the time of the survey. With these definitions, attrition and shutdown are mutually exclusive. Shutdown is low and balanced across the treatment arms. Attrition is always below 10 percent, is balanced in 2022 for both suppliers and clients, and is balanced in 2023 for clients. However, attrition shows a significant imbalance of 4 percentage points in 2023 for suppliers, raising the concern of selective attrition. To address this concern, we conducted balance tests in the "survival subsample" of firms that are in our data in 2023, that is, excluding all firms affected by attrition or shutdown in 2023. As shown in Appendix Table A1, the treatment is balanced in this subsample, and the summary statistics are very similar to those in Table 1. We conclude that selective attrition is unlikely to bias our results.[19]

3 Conceptual framework

To guide the empirical analysis, we present a conceptual framework—based on the model developed in Szeidl (2025)—about the effects of improving access in the production network. We present the framework verbally here, but include the formal model in Appendix B. We use the framework to organize our empirical analysis of two questions: what barriers limit access, and how improved access affects industry equilibrium.

Setup. Consider an industry consisting of a set of supplier and client firms and a pre-existing supplier-client network. A supplier-client link in this network is valuable because it creates a business opportunity, such as a new product variety the client can produce using the input made by the supplier. The quality of this opportunity is match-specific and drawn from a commonly known distribution. Firms face a capacity constraint on the number of links they can maintain—for example, a client can have at most κ suppliers—which reflects either managerial or technological limits.[20] Firms also face two barriers to accessing new partners. (i) They must incur a search cost,

[19] For completeness, we include some Lee bound estimates in the supplier sample in Appendix Table A10.

[20] The formal model is worked out under the assumption that all firms have the same capacity constraint κ.

which includes both the cost of finding potential partners and that of experimenting to learn the value of a new partnership. (ii) Firms may also hold pessimistic beliefs about the value of new partnerships, i.e., undervalue the associated opportunities.

Implications about barriers to access. In this framework, both of the above barriers can generate under-search relative to the social optimum. The search cost does so through an externality: firms do not internalize the gains from search that accrue to the other party. Undervaluation does so simply by creating a wedge between the (perceived) value and the cost of finding a partner. In both cases, an intervention that reduces the search cost, e.g., by providing referrals, can lead to new partnerships.

Undervaluation can also lead to a network allocation that is privately suboptimal, and make that allocation persistent, a situation we call a *low-access trap.* The logic is straightforward. Under-search driven by undervaluation is privately suboptimal and results in a network where firms have insufficiently many or insufficiently well-matched partners. Furthermore, if undervaluation fully eliminates search, then firms never change their beliefs, maintaining the trap.

A distinctive implication of undervaluation concerns the impact of firms' experience with new partners, as in our subsidized referrals. This experience can lead firms to update their beliefs about potential partners in general. This, in turn, can increase their private (unsubsidized) search, and result in additional partners that were not referred to them. Through this channel, referrals can break the low-access trap. A theory based solely on search costs would not predict these responses.

We summarize the key predictions about barriers to access as follows.

H1 The referrals can create new partnerships.

H2 In a low-access trap, the referrals can increase beliefs about potential partners, private search, and new partnerships.

Implications about industry equilibrium. Assume now that the aforementioned barriers are strong enough that firms do not search for partners on their own. Consider a referrals intervention similar to that we conducted, in which half of the suppliers and half of the clients are treated. The treatment creates new potential partnerships between treated suppliers and treated clients, fully

Figure 3: Possible effects of access on the production network

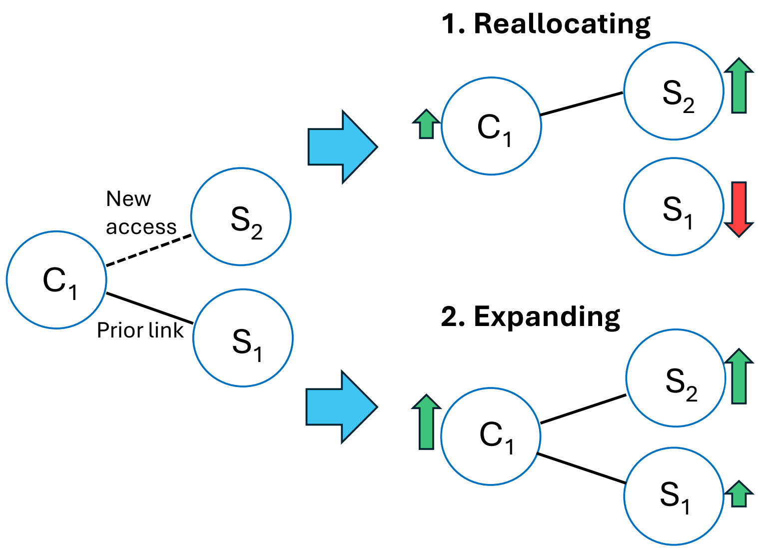

eliminating the barriers between referred firms.[21]

This intervention can shape equilibrium through two qualitatively different channels. The first is a mechanism familiar from many trade models (e.g., Eaton and Kortum 2002): *reallocation.* For an example, consider Case 1 of Figure 3. When a client firm C_1 gets access to a new supplier S_2 that offers a sufficiently attractive opportunity, C_1 replaces an existing supplier S_1 with the new supplier S_2. This change increases the performance of both C_1 and S_2, since otherwise they would not have formed the partnership; but decreases the performance of S_1, who lost a valuable partner. That is, access generates a reallocation of the network which has a positive direct effect on treated firms and a negative indirect effect on their prior partners. Our framework predicts this sort of reallocation when firms are at capacity in terms of their number of partners, because then each new partnership must replace a pre-existing one.

Our framework also highlights a second mechanism, active when firms are below capacity: *expansion.* As depicted in Case 2 of the Figure, when C_1 has free capacity in terms of its number of links, it can take on the new supplier S_2 without crowding out S_1. Then, the referral leads to an expansion of economic activity. This naturally improves the performance of both C_1 and

[21] Thus, these referrals are similar to our subsidized, rather than to our information referrals.

S_2. Moreover, if the expansion in activity leads C_1 to upgrade, e.g., by producing a new variety, then that can positively affect the prior partner S_1 too, e.g., through additional demand for inputs. Thus, when firms are below capacity, access expands the production network, with a positive direct effect and potentially a positive indirect effect. Distinguishing between reallocation and expansion is important because expansion both generates larger direct effects through new economic activity, and avoids the negative indirect effects that dampen aggregate impacts under reallocation. Note that expansion is especially likely with undervaluation, because then firms search too little on their own.

We also have a comparative static prediction about the direct effect of referrals on firm growth: the effect is larger when firms have fewer links to start with. This is for two reasons. First, firms with fewer initial links are more likely to be below capacity, making the expansion mechanism more likely and increasing the direct effect of the referrals. Second, a referral has a proportionally larger effect when the firm has fewer partners to begin with.

Finally, our framework implies that, when the industry is in a low-access trap, the referrals can have high returns. They can have a high private return, because partners are more profitable than believed, and leaving the trap generates further search. If firms are below capacity—likely when they are in the trap—the private return is both larger and not crowded out by negative indirect effects, leading to a high social return.

We summarize the key predictions about industry equilibrium as follows.

H3 When firms are below capacity, the referrals expand the production network.

H4 The referrals have a positive direct effect on firm performance, especially for low-degree firms.

H5 When firms are below capacity, the referrals have a positive indirect effect on prior partners.

H6 When firms are in a low-access trap, the referrals can have large private and social returns.

4 Evidence on barriers to access

In this section, we study the barriers to access empirically, guided by hypotheses H1 and H2. We begin by examining firms' take-up of the referrals. We then ask whether the referrals lead to the creation of new production links. Finally, we examine their effects on beliefs and search behavior to assess whether the industry is in a low-access trap.

4.1 Take-up

For subsidized referrals, we define take-up by whether the referred firms used the subsidy for an initial transaction.[22] Take-up provides an initial indication of firms' interest in forming a new partnership. To examine the treatment effect on take-up, we estimate firm-level regressions of the form:

$$Takeup_i = \mu \cdot \mathit{Treatment}_i + Controls + \nu_i. \tag{1}$$

Here each observation is a firm i, the outcome variable is an indicator for whether the firm used the subsidy for any of its referred links, and the key independent variable is an indicator for assignment to the treatment group (T1). Recall that all firms in the treatment group received subsidized referrals. The controls include strata fixed effects; we cluster standard errors by firm.

Table 2 reports the results. Column 1 shows that among clients, 61 percent of treated firms used the subsidy on at least one of their referred links. Column 2 shows that among suppliers, 63 percent of treated firms used the subsidy on at least one of their referred links. The coefficients are highly significant. Thus, a majority of firms chose to experiment with the subsidized referrals.

We next examine take-up at the link level, where we can separately examine screened and unscreened referrals. We estimate regressions analogous to (1) in the sample of all *candidate referrals*, that is, all T1-to-T1 and T0-to-T0 referrals we constructed. The main independent variable is the subsidy treatment, with T0-to-T0 referrals serving as the control group. We control for the business type of both the supplier and the client and cluster standard errors by supplier and by client. Column 3 shows that firms used the subsidy on 48 percent of the links for which it was offered. Column 4 shows that take-up was essentially the same for screened and unscreened

[22] For information referrals there is no notion of take-up.

Table 2: Take-up of subsidy

Dep. var.: Take-up	Clients (1)	Suppliers (2)	Links (3)	Links (4)
Subsidy	0.609***	0.631***	0.484***	
	(0.041)	(0.034)	(0.039)	
Screened Subsidy				0.474***
				(0.039)
Unscreened Subsidy				0.504***
				(0.047)
Strata FE	Yes	Yes		
Business type FE			Yes	Yes
Observations	276	406	856	856

Notes: In column 1 sample is client firms, in column 2 supplier firms, in columns 3-4 all candidate referrals (made and not made). Standard errors are clustered by firm in columns 1-2, and by supplier and client in columns 3-4. *** p<0.01, ** p<0.05, * p<0.1.

referrals.[23] Screening did not improve take-up, suggesting either that it did not improve match quality or that firms could not assess match quality before gaining experience with the partner.

4.2 Link creation

We next examine whether the referrals created new partnerships. We measure partnerships using transactions that occurred after the end of the subsidy period and were therefore not subsidized. We estimate:

$$y_{ijt} = \beta_S \cdot \mathit{Screened\ subsidy}_{ij} + \beta_U \cdot \mathit{Unscreened\ subsidy}_{ij} + \beta_I \cdot \mathit{Information}_{ij} + \mathit{Controls} + \varepsilon_{ijt}. \quad (2)$$

Each observation is a supplier-client pair ij in a follow-up survey wave $t \in \{2022, 2023\}$. As before, the sample includes all candidate referrals we constructed, both made and not made: the subsidized referrals (T1-to-T1), the information referrals (one quarter of T0-to-T0), and the remaining unmade T0-to-T0 referrals. The main independent variables are indicators for the three types of made

[23] All subsidized transactions, with one exception, had a value of at least 3,000 RMB, and consequently received the maximal subsidy of 1,500.

Table 3: Link creation

Dep. Var.:	Link			Num Transactions	Value (RMB)	E[Satisfaction] with Price	E[Satisfaction] with Transaction
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Subsidy	0.410***						
	(0.030)						
Screened Subsidy			0.474***	2.656***	7,311.812***	0.573***	0.687***
			(0.036)	(0.265)	(924.774)	(0.100)	(0.102)
Unscreened Subsidy			0.288***	1.258***	3,358.945***	0.311***	0.333***
			(0.033)	(0.191)	(651.112)	(0.096)	(0.097)
Information		0.004	0.010	0.074	268.187	-0.064	-0.023
		(0.004)	(0.006)	(0.064)	(249.739)	(0.065)	(0.064)
Business type, year, unscreened FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,707	544	1,707	1,707	1,707	1,707	1,707

Notes: The sample in column 2 consists of T0 screened candidate referrals; in all other columns, the sample consists of all candidate referrals (both made and not made). In columns 6 and 7, we standardized the outcome by dividing by the standard deviation of the untreated sample. Standard errors clustered by supplier and client. *** p<0.01, ** p<0.05, * p<0.1.

referrals: screened subsidized, unscreened subsidized, and information referrals (all of which were screened). The controls include indicators for whether the referral was screened or unscreened, indicators for the firm's business type, and year fixed effects. Standard errors are clustered by supplier and by client.

Table 3 reports the results. In columns 1-3 the outcome is a link indicator equal to one if i and j transacted at least once in t after the subsidy period. In column 1 we focus on the overall impact of subsidized referrals, pooling the screened and unscreened arms. We find that a subsidized referral increased the probability of a link by a significant 41 percentage points. Here, the information referrals are included in the control group. In column 2 we evaluate the effect of the information referrals, restricting attention to the subsample of T0-to-T0 candidate screened referrals, the natural comparison group. We find a precisely estimated zero effect, validating our approach of treating these referrals as effectively unmade. In column 3, we separate the subsidized referrals into screened versus unscreened, while also including the information referrals. This is our main specification. Screened referrals created links at a significantly higher rate (47pp) than did unscreened referrals (29pp), while information referrals, unsurprisingly, again have a zero effect. Columns 4 and 5 repeat this regression using as outcomes the number of yearly transactions, and the total value of yearly transactions. The same qualitative pattern emerges.

These results have several implications. The significant effects of subsidized referrals are consistent with hypothesis H1 that referrals should create new partnerships. Moreover, their large magnitude suggests that partnering frictions are important. The null effect of the information referrals suggests that lack of information about the identity of potential partners is not the main partnering friction. Instead, there seems to be an additional friction that prevents link formation even conditional on knowing about a potential partner. Undervaluation is one possible explanation for the null effect of information, but this evidence is only indirect.

The larger effect of screened than unscreened referrals suggests the importance of match-specific value. To see why, recall that screened and unscreened referrals were constructed so that, on average, they provided access to similar firms. Unscreened referrals were essentially random, while screened referrals were allocated to make sure all firms got some, but none got too many referrals, implying that both types of referrals were approximately representative of the pool of potential partners.[24] Consistent with this logic, Appendix Table A3 column 1 shows that screened and unscreened referrals were on average indistinguishable in most firm characteristics, with only a slight difference in the log employee count for supplier firms.[25] Because the firms were similar on average, the higher rate of link formation in screened referrals must reflect higher match quality. Table A3 column 2 provides more direct evidence on match-specific value by showing that screened referrals paired firms that were more closely aligned in price, a proxy for quality alignment.[26] Moreover, this alignment predicts link formation: Appendix Table A4 shows that the subsidy treatment was more likely to generate links among price-aligned firms. We conclude that partnerships had match-specific value,

[24] Specifically, each client received two screened and one unscreened referral, ensuring that the average client in screened and unscreened referrals was essentially identical. Suppliers got 1-4 screened referrals, and were randomly chosen for unscreened referrals, so that the average supplier in screened versus unscreened referrals was also similar.

[25] The difference is 0.08, which is small compared to the standard deviation of log employment of 0.57. The difference may be due to screened referrals oversampling large suppliers, which were better matches for more clients.

[26] Column 3 shows that screened referrals were less closely aligned in revenue. This is likely driven by our matching algorithm. Since large firms tend to be good matches for more partners, we likely assigned them systematically to small partners who otherwise would not be matched.

partly driven by quality alignment, and that screening identified higher-quality matches.

Interestingly, although screened referrals were more likely to generate sustained partnerships (Table 3), firms took up the subsidy at similar rates for screened and unscreened referrals (Table 2). This contrast suggests that learning match-specific value required direct experience with the partner. That is, even after learning the identity of the partner, firms faced a residual cost of experimentation to learn about match quality. This experimentation cost was low, since a subsidy amounting to 0.4% of profit induced half of the firms to experiment with the partner, but apparently positive.

4.3 Nature of the friction and the low-access trap

As we noted, the fact that subsidized referrals created many partnerships whereas information referrals created none is consistent with the undervaluation friction. However, this result alone does not identify the mechanism. Similar behavior could arise if firms correctly valued potential partners but faced some other barrier to experimentation. We therefore turn to evidence designed to distinguish undervaluation from alternative explanations.

In the follow-up survey, we asked firms in the information treatment about their reasons for not pursuing the referrals.[27] The most common answer was uncertainty about whether the referred partner would be a good match, reported by 49% of clients and 33% of suppliers. Although descriptive, this evidence is consistent with the undervaluation hypothesis.[28]

To build causal evidence, we follow hypothesis H2 and examine the treatment effect on beliefs and search. We begin with beliefs about satisfaction with the partner. In the follow-up surveys, for each candidate referral, we asked both firms to rate on a 10-point scale how satisfied they were, or expected to be, with the referred partner. Importantly, we asked the question not only for referrals that were made, but also for the T0-to-T0 referrals that we constructed but did not make, explicitly naming the counterfactual partner. Thus, the score measures beliefs conditional on treatment for

[27] 94% of these firms said that they did not know their information referral, indicating that the reason for not transacting was not that they already knew the other party.

[28] The second most salient answer, given by 23% of clients and 46% of suppliers, was that firms did not know how to approach the partner and were afraid of rejection, suggesting an additional friction related to the lack of marketing skills (Hjort et al. 2025).

firms in T1, and conditional on no treatment for firms in T0, allowing us to estimate the treatment effect on beliefs. We focused on two dimensions of satisfaction: price, and the transaction, including partner reliability as well as satisfaction with the product (for clients) or demand (for suppliers).[29]

Columns 6 and 7 in Table 3 report treatment effects on expected satisfaction. To construct the outcome, for each dimension (price and transaction) we averaged the expected satisfaction of the supplier and the client over the link, and then standardized the resulting measure. The screened subsidy improved expected satisfaction by a significant 0.6 standard deviations, the unscreened subsidy by a smaller but still significant 0.3 standard deviations. The information referral had no effect, which is unsurprising since it did not generate transactions. The positive effects on expected satisfaction are consistent with H2 and provide direct evidence that firms initially undervalued potential partners.[30]

The evidence so far is consistent with undervaluation, but does not yet show that it limits partnering through reduced search. To explore whether it does, we estimate the effect of the referrals on search and partnering. The logic of H2 implies that if the improved beliefs resulting from the referrals stimulate search, then firms must have been searching too little absent the intervention. If this additional search also generates new partnerships, then undervaluation limited partnering, creating a low-access trap.

We estimate the firm-level regression

$$y_{it} = const + \beta_1 \cdot Treated_i + Controls + \varepsilon_{it} \tag{3}$$

in the follow-up data. Each observation represents a firm i in a year t. The outcomes measure beliefs, search, and partnering; most were only collected in 2023. The key independent variable is an indicator for the firm's treatment status (T1). The regression includes a set of controls: the baseline value of the outcome when available, year effects when we observe the outcome in both follow-up years, indicators for business type, and strata effects.

[29] Some firms that did not recognize the named partner preferred not to evaluate the named partner. In those cases, we instead asked how satisfied they expected to be with an unknown partner from the county, to capture their unconditional beliefs about partner quality.

[30] The fact that some firms who did not know the referral evaluated an unknown partner does not affect this interpretation. We estimate a positive treatment effect even for unscreened referrals, which were essentially random and therefore closely resemble an unknown partner.

Table 4: Search

Dep. Var.:	Belief: Profit gain from new partner (%)	Search (hrs/month)	Placebo belief: partner on other side (%)	Placebo search: on other side (hrs/month)	Num non-referred new partners
	(1)	(2)	(3)	(4)	(5)
Treatment	7.540***	6.179***	0.252	0.985	1.230*
	(0.943)	(0.981)	(0.679)	(0.884)	(0.745)
Baseline Control, Year FE					Yes
Indicator for client, Strata FE	Yes	Yes	Yes	Yes	Yes
Control Mean	3.941	3.868	3.422	4.018	5.857
Observations	632	632	632	632	1283

Notes: Sample is all firms. Data in columns 1-4 are from 2023, as the outcome was collected only in that year; data in column 5 are from 2022 and 2023. The outcomes in columns 3 and 4 are sellers (for supplier firms) and buyers (for client firms). The outcome in column 5 is the number of partners in the firm data minus the number of referred partners in the network data. The control mean is reported at follow-up because baseline data are unavailable. Standard errors clustered by firm. *** p<0.01, ** p<0.05, * p<0.1.

Table 4 reports the results. In column 1 the outcome is beliefs about new partners in general, measured as the firm's expected percentage increase in profit generated by a new partner. Here partner means seller for a client firm and buyer for a supplier firm. The treatment significantly increased firms' beliefs about new partners in general, not only about the referred partners. Column 2 shows that the treatment also significantly increased the amount of search, suggesting that firms acted on these improved beliefs. Columns 3 and 4 report placebo outcomes for partners on the "other side" of the production chain: buyers (e.g., wholesalers) for client firms and sellers (e.g., raw material producers) for supplier firms. Because we did not refer such partners, we expect, and find, no effects, providing an internal validity check. Finally, column 5 shows a significant increase in the number of firms' non-referred partners, measured as the total number of partners they reported in the firm survey minus the number of referred partners with whom they reported a transaction in the network survey. Thus, the referrals stimulated additional partnering beyond the partnerships we created.

In summary, the referrals increased beliefs, search, and partnering. These results are consistent with hypothesis H2, but not with a theory based purely on search costs, which would not predict a change in beliefs or search. Taken together, our findings provide direct evidence that firms were initially in a low-access trap.

It may seem puzzling that firms persist in a low-access trap given the low tangible search costs in our setting. Firms in the cluster are likely to have chance interactions with potential partners, making it easy to identify and contact them. However, identifying a potential partner is not the same as learning whether the match is valuable. As discussed in Section 4.2, firms take up the subsidy for screened and unscreened referrals at similar rates, yet screened referrals are substantially more likely to generate partnerships. This pattern suggests that learning match quality requires direct trading experience and thus involves a positive experimentation cost. It follows that chance interactions are unlikely to eliminate the trap, because they rarely provide the experience needed to assess match quality.[31]

Why did undervaluation arise in the first place? We cannot answer this question conclusively, but the organizational learning literature suggests one possible explanation. This literature identifies a *competence trap*, in which organizations concentrate on their established activities, generating learning and experience that reinforce reliance on those activities but discourage the exploration of alternatives (Levitt and March 1988, March 1991, Levinthal and March 1993). Similarly, firms in our setting may have concentrated their business within their existing networks, neglecting outside opportunities. Over time, this may have contributed to firms' undervaluation of unfamiliar partners. Consistent with this hypothesis, firms in our sample maintained long-term relationships with a recurring set of partners, but had limited information about firms outside their existing networks. The average relationship age in our sample was 9 years (Table 1), and although 46% of the links present in 2018 were not present in 2019, about 40% of these (temporarily inactive) links re-appeared in 2021.[32] At the same time, firms did not know 94% of the referred partners in the information referrals. These patterns are consistent with the view that reliance on existing partners contributed to the undervaluation of outside partnerships.

[31] A natural question is whether the experimentation cost itself—rather than undervaluation—can be the main barrier to access. As noted above, a pure experimentation-cost explanation does not predict that firms increase search in response to the treatment. Nor can it explain why firms did not take up the information referrals, because, as we show in Section 5.6, the profit gain from take-up is several times the cost of the subsidy that induced take-up.

[32] We also show in Table 5 below that when firms lost partnerships due to crowding out, they made up at least some of those losses by turning to prior partnerships in their network.

Given the evidence on undervaluation, a natural question is which characteristics of potential partners firms underestimated. Our results on satisfaction suggest undervaluation in multiple dimensions. Additional evidence comes from survey data on firms' overall satisfaction with all their partners—not just the referrals—across a range of dimensions. In Appendix Table A7 we report firm-level treatment effects on these outcomes, and find, for both suppliers and clients, significant improvements in satisfaction with partners' price, the business transaction (product quality or demand), and partners' reliability, but not with their flexibility or trade credit. These findings suggest that correcting undervaluation requires updating beliefs along multiple dimensions, highlighting the importance of direct experience with the partner.

Importantly, the undervaluation friction has different policy implications from the information frictions emphasized in prior work. When firms undervalue unfamiliar partners, simply providing information is unlikely to improve access. Instead, effective policies may need to create opportunities for firms to learn through direct experience or other forms of demonstration—for example, through subsidized trial transactions or product samples.

5 Evidence on industry performance

We next examine the effects of the referrals on industry performance, focusing on the reallocation and expansion mechanisms identified by our framework. We begin by studying how the intervention reshaped the production network. We then analyze the effects on business performance, first estimating the direct effects on treated firms and examining their underlying mechanisms, and then estimating the indirect effects on treated firms' prior partners. Finally, we estimate the private and social returns implied by our findings.

5.1 Effects on the production network

We study the impact of the referrals on the production network using firms' degree as the main outcome. Expansion is captured by the direct effect of treatment on treated firms' degree, whereas reallocation is captured by the indirect effect on the degree of treated firms' prior partners. To identify the indirect effect, we introduce a new variable, *exposure*, defined as the share of the firm's baseline partners in the network data—sellers for client firms, buyers for supplier firms—that are treated. Exposure, although not itself a treatment arm, is as good as random, because each partner firm is treated with the same 50% probability.[33] It is important to note that exposure is only defined for firms that at baseline have at least one (regular) partner in the network data, about 55% of the sample. For the remaining firms, which had occasional partners but no regular supplier-client relationships at baseline, we set exposure equal to its sample mean. In all regressions, we appropriately control for the firm not having regular baseline partners, to ensure that the effects of exposure are not identified from those firms.[34]

We estimate the following firm-level regression:

$$\begin{aligned} y_{it} = const + \beta_1 \cdot Treated_i + \beta_2 \cdot Exposure_i * Treated_i \\ + \beta_3 \cdot Exposure_i * Untreated_i + Controls + \varepsilon_{it}. \end{aligned} \tag{4}$$

Each observation represents a firm i in a year $t \in \{2022, 2023\}$. The outcome variables measure the firm's degree. The key independent variables are an indicator for the firm's treatment status (T1), and exposure interacted with indicators for the firm being treated or untreated. Exposure is demeaned so that the treatment effect is identified for a firm with average exposure. Thus, β_1 measures the direct effect of referrals, while β_2 and β_3 measure the indirect effect of referrals for treated and untreated firms, respectively. The regression also includes a set of controls, especially an indicator for the firm not having (regular) partners at baseline, and the interaction of this variable with the treatment. These controls ensure that both the treatment and the exposure effect are identified from variation among firms with baseline partners.

Table 5 presents the results. We begin with the direct effect, measured by the treatment coefficient. In column 1, the outcome variable is degree in the network data: the number of buyers for a supplier firm and the number of sellers for a client firm. Treated firms experienced a significant increase in degree of approximately 1 link, representing a 30% increase relative to control firms.

[33] Balance tests in Appendix Table A2 confirm that exposure is quite balanced with respect to baseline variables. A Borusyak and Hull (2023) type re-centering is not necessary because there is no non-random variation in exposure: each partner firm has a 0.5 probability of being treated, implying that the exposure of each firm has a mean (across randomization draws) of 0.5. We confirmed that mean exposure is approximately 0.5 by drawing counterfactual treatment assignments and computing mean exposure.

[34] The results are similar when restricting the sample to firms that reported partners at baseline.

Table 5: Effects on the network

Dep. Var.:	Degree	Baseline links active	Non-baseline links	Reactivated prior links
	(1)	(2)	(3)	(4)
Treated	0.945***	-0.336***	1.281***	-0.045
	(0.168)	(0.090)	(0.125)	(0.099)
Exposure * Treated	-0.067	0.078	-0.145	-0.023
	(0.305)	(0.133)	(0.257)	(0.139)
Exposure * Untreated	-0.107	-0.349***	0.241*	0.268**
	(0.194)	(0.120)	(0.133)	(0.116)
Controls, strata FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	
Control Mean	1.784	1.784	0.344	
Observations	1283	1283	1283	651

Notes: Sample is all firms. Data in columns 1-3 are from both follow-up waves; column 4 uses a cross-section. Exposure is demeaned share of baseline partners treated. Controls include an indicator for client firm, an indicator for the firm having regular partners at baseline, an interaction of the latter with the treatment, and the number of baseline partners. In column 4, outcome is number of partners in either follow-up (2022/23) that were not partners at baseline (2021) but were partners before (2018 or 19). Control mean is at baseline. Standard errors clustered by firm. *** p<0.01, ** p<0.05, * p<0.1.

This result provides direct evidence of an expansion of the production network. The next two columns decompose total degree into continuing baseline links and newly formed links. Consistent with the reallocation logic, treated firms discontinued some prior links (column 2); but they more than offset these losses through the formation of new links (column 3). These new links include the referral partners but also additional connections, since we know from Section 4.3 that treated firms increased search and established more non-referred new partners. Overall, the treatment induced some reallocation but also a net expansion of treated firms' production networks.

We next turn to the indirect effects. Column 1 shows that exposure has no statistically significant effect on total degree, for either treated or untreated firms. At first glance, this absence of indirect effects may seem surprising in light of the reallocation effect just discussed: if treated firms discontinued some prior links, their former partners should be negatively affected. Indeed, column 2 confirms that exposure reduced untreated firms' links with prior partners, consistent with this crowd-out. However, untreated firms appear to have reacted to these lost links. As column 3 shows, exposure simultaneously increased new link formation among these firms. Exposed firms therefore appear to have rewired their networks in response to the losses, explaining the insignificant effect on their total degree. It may seem puzzling that firms were able to rewire their networks despite the partnering frictions identified in the previous section. Column 4 sheds light on what happened, by showing that some of the new connections were reactivated dormant links, i.e., partners the firms had before the baseline but not in the baseline year.[35] In summary, although the treatment did crowd out some prior links, firms made up for these losses, partly by reconnecting with prior partners, resulting in no detectable indirect effect on overall degree.

The combination of positive direct effects and negligible indirect effects implies that the intervention expanded the equilibrium production network, consistent with hypothesis H3 and the expansion mechanism of our conceptual framework. Using the coefficients from column 1, we estimate the effect of the intervention on the total number of supplier-client links. Under the assumption that all effects of the intervention on the network are captured by the regression, we find that the intervention increased the total number of supplier–client links by approximately 24%.[36] Thus, the referrals meaningfully expanded the equilibrium production network.[37]

[35] Using the network data, we define reactivated links as those present in at least one year during 2018–2020 and in at least one year during 2022 or 2023, but absent in the baseline year 2021. Note that we pool 2022 and 2023 for power, so that the coefficient in column 4 is not directly comparable to those in columns 1-3.

[36] Since 50% of firms are treated, the direct effect yields a change of $0.5 \cdot 0.945 = 0.473$ links per firm. Since the intervention increases average exposure by 0.5, the indirect effect—averaged over treated and untreated firms—yields a change of $0.5 \cdot (-0.067 - 0.107)/2 = -0.044$ links per firm. For firms without regular partners, the direct effect is not significantly different (not reported in the table) and there are no exposure effects, so we assume the total effect is the same. The overall effect is then 0.43 links relative to the control mean of 1.784, or about 24%.

[37] We explore some heterogeneous effects in the Appendix. In Appendix Table A5 we ask how baseline link survival depends on whether the supplier, the client, or both are treated. In all three cases, survival declines by roughly 20 percentage points relative to untreated pairs. In Appendix Table A6 we estimate heterogeneity by firm and link characteristics. Although power is limited, we find some evidence that survival is higher when either firm is more satisfied with the partner; when the supplier is larger; when the firms are kin; and when the relationship is older. These results are consistent with the importance of match quality.

5.2 Direct effect on firm performance

We next investigate the effect of the referrals on firm performance. We begin with the direct effect on treated firms, which we measure using the regression

$$y_{it} = const + \beta_1 \cdot Treated_i + \beta_2 \cdot y_{i0} + Controls + \varepsilon_{it}. \tag{5}$$

Each observation is a firm i in a time period $t \in \{2022, 2023\}$. The outcome variable is a measure of firm behavior or performance, such as log revenue. The key independent variable is an indicator for the firm's treatment status (T1). To maximize power and reduce noise, following McKenzie (2012) we estimate a single treatment effect β_1 across the two follow-up waves, and include the baseline value of the outcome variable as a control. We always include strata and time fixed effects, and depending on the specification we may include other controls as well.

This regression identifies the direct effect of treatment on firm performance under the assumption that indirect effects do not confound the comparison between treated and untreated firms. We explicitly incorporate indirect effects in Section 5.4 below, and show that accounting for them does not affect the estimated direct effect.[38] For inference, because the treatment is randomized across firms in a non-clustered fashion, we follow the recommendation of Abadie, Athey, Imbens and Wooldridge (2022) and cluster standard errors by firm.[39]

Table 6 reports impacts on key firm performance measures, separately for supplier and client firms. Beginning with suppliers, we find a significant increase in revenue of 24 log points, and a significant increase in profit of 15,000 RMB, equivalent to about 19% of the baseline profit of control firms.[40] We find no detectable impact on operating cost, but columns 4 and 5 show significant increases in log employment and the log number of hours worked by all individuals in

[38] This is consistent with the result in Section 5.1 that indirect effects on degree were not different for treated and for untreated firms, so that they do not confound the comparison across treatment arms.

[39] Indirect effects may induce a correlation between the outcomes of different firms, but such a correlation is not a reason to cluster (Abadie et al. 2022). For robustness, in Appendix Table A8 we report clustered standard errors where clusters are defined within product type, based on similarity in price and geography, using hierarchical clustering, and find that our results are robust.

[40] Because profit can be negative, we did not take the logarithm. Instead, to address outliers, we winsorized it, across all firms and all five waves, at the 99th percentile, which corresponds to a profit level of 3 million RMB, or 50 times the median profit of 60,000 RMB. Appendix Table A9 shows that the profit effects are robust to alternative adjustments, including winsorizing by year ($p < 0.10$); by year separately for suppliers and clients ($p < 0.10$); taking the inverse hyperbolic sine ($p < 0.01$); and using the raw profit ($p = 0.13$).

Table 6: Firms: main outcomes

Dep. Var.:	Log Sales	Profit (10,000 RMB)	Log Operating Cost	Log Employ-ment	Log Total Hrs Worked / Day	Num of Partners	Satisfaction with Partners
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Suppliers							
Treatment	0.244***	1.501*	0.051	0.068*	0.128***	1.876**	0.219***
	(0.074)	(0.824)	(0.098)	(0.040)	(0.047)	(0.761)	(0.064)
Baseline Control	Yes	Yes	Yes	Yes		Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	1.890	6.804	1.296	0.711	2.740	5.595	5.059
Observations	761	763	652	763	763	763	761
Panel B: Clients							
Treatment	0.042	0.152	0.023	0.050	0.062	1.608**	0.241***
	(0.113)	(3.916)	(0.123)	(0.067)	(0.082)	(0.672)	(0.071)
Baseline Control	Yes	Yes	Yes	Yes		Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	3.340	37.38	2.740	1.349	3.290	6.102	5.308
Observations	520	520	502	520	519	520	496

Notes: Sample in Panel A is suppliers, in Panel B is clients, in both follow-up waves. In columns 2 and 6 outcome is winsorized at the 99th percentile in the full sample (all firms in all waves). In column 7, outcome is standardized by its standard deviation in the baseline control sample, separately for suppliers and clients. In columns 6 and 7, partner refers to clients in Panel A and suppliers in Panel B. Baseline value of the outcome is not included in column 5 because we did not collect it. Control mean is at baseline. Standard errors clustered by firm. *** p<0.01, ** p<0.05, * p<0.1.

the firm. Thus, treated suppliers substantially increased their activity. Consistent with the network results, treated suppliers also report an increase in the total number of partners (measured in the firm survey) and higher overall satisfaction with their partners.[41]

Panel B turns to client firms. Here, the estimates are generally insignificant: we find no detectable impact on revenue, profit, cost, or employment. One possible explanation is that much of the effect is concentrated in a subset of firms. Our conceptual framework predicts larger gains for

[41] Because attrition among suppliers is somewhat lower in the treatment group than in the control group (Table 1), we report Lee (2009) bounds in Appendix Table A10. The bounded estimates remain significant for revenue, the inverse hyperbolic sine (IHS) of profit, total hours worked, number of partners, and partner satisfaction, though not for level profit or log employment.

low-degree firms, i.e., firms with initially limited network access. Motivated by this prediction, we split clients into two groups based on whether their number of baseline partners is ≤ 1 or > 1, where 1 is the sample median, and estimate treatment effects separately for the two groups. Although this subgroup definition was not pre-registered, it was motivated by our conceptual framework.

Table 7 presents the results. Panel A focuses on brush pen producers with below-median baseline networks (zero or one regular partner). For this group, the treatment effects are economically large and statistically significant. Treated firms experienced a revenue gain of about 31 log points, and a profit gain of about RMB 93,000, which is about 36% of the control mean. Operating costs also increased significantly. The impacts on employment and total working hours are positive but not statistically significant. Moreover, these firms experienced a significant increase in the number of suppliers and satisfaction with their suppliers. In contrast, firms with above-median baseline degree experienced no detectable gains.

A concern with these results is that we selected the subsample of low-degree client firms after the intervention, raising the possibility that the estimated effects reflect sampling variation rather than a genuine treatment effect among low-degree firms. We address this concern in several ways. First, as noted, the heterogeneity follows directly from our conceptual framework, which predicts larger effects for firms facing tighter access constraints, proxied by low baseline degree. It therefore reflects the guidance of the theory, not data mining. Second, when we apply the same median-degree split to the sample of supplier firms, we likewise find that the performance gains are concentrated among low-degree firms (Appendix Table A11). This parallel pattern provides independent support for baseline degree as a key determinant of treatment effects. Third, supporting the interpretation that the treatment had positive effects throughout the client sample, with larger gains for low-degree clients, we obtain significant treatment effects in the full client sample when we use a less noisy outcome: firms' self-assessment of their performance on a 5-point scale. In the 2024 phone survey, we asked firms to rate various dimensions of their performance since 2020 on this bounded scale, which is likely to yield more precise estimates. Appendix Table A12 shows significant gains in revenue and profit, among other outcomes.[42] Together, these findings suggest that our estimates

[42] These estimates may be mildly attenuated by a marketing-offer intervention introduced among untreated firms in 2023, described below in Section 6. Because that intervention increased link formation, it likely biases the estimated

Table 7: Client firms: main outcomes by baseline network size

Dep. Var.:	Log Sales	Profit (10,000 RMB)	Log Operating Cost	Log Employ- ment	Log Total Hrs Worked / Day	Num Partners	Satisfaction with Partners
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Baseline partners <=1							
Treatment	0.321**	9.297***	0.290*	0.064	0.071	2.500***	0.376***
	(0.148)	(3.377)	(0.165)	(0.084)	(0.099)	(0.702)	(0.086)
Baseline Control	Yes	Yes	Yes	Yes		Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	2.739	25.91	2.128	1.143	3.77	4.506	4.609
Observations	304	304	288	304	304	304	284
Panel B: Baseline partners >1							
Treatment	-0.181	-7.586	-0.174	0.043	0.060	0.218	0.045
	(0.186)	(8.460)	(0.209)	(0.117)	(0.151)	(1.292)	(0.132)
Baseline Control	Yes	Yes	Yes	Yes		Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	4.125	52.10	3.484	1.614	3.556	8.150	6.845
Observations	216	216	213	216	215	216	211

Notes: In Panel A, the sample consists of clients with at most one regular supplier at baseline; in Panel B, of clients with more than one regular supplier at baseline. Data are from both follow-up waves. The outcome in column 2 is winsorized, in column 6 it is winsorized, and in column 7 it is standardized. In columns 6 and 7, partner refers to suppliers. The baseline value of the outcome is not included in column 5 because it was not collected. Control mean is at baseline. Standard errors clustered by firm. *** p<0.01, ** P<0.05, * p<0.1.

reflect a genuine treatment effect for low-degree clients.

In summary, we find large treatment effects on business performance, especially for low-degree firms. These results are consistent with hypothesis H4, which predicts gains in business performance that are largest for firms with limited initial network access.

treatment effects downward. But since our outcomes are for the entire 2020-24 period, these changes in the last few months should have only minor effects.

5.3 Intermediate outcomes and mechanism

Why did the new partnerships improve business performance? The expansion mechanism in our conceptual framework suggests that they did so by expanding business activity and upgrading firm capabilities. We turn to explore these forces by looking at treatment effects on intermediate outcomes. As described in Section 2.2, we measured the quality of the firm's main product—if it was a brush head or a brush pen—using evaluations by independent experts. The experts scored products along two dimensions: craftsmanship and durability. We residualized these scores by product type (i.e., hair type) and standardized them, generating two quality measures that serve as our main quality outcomes. In addition, we collected information on the firm's (up to) three most important products, including their prices and revenue shares, as well as data on time devoted to quality control. Together, these measures allow us to assess whether the intervention shifted firms toward higher-value products and greater investment in quality.

Table 8 reports impacts on these product-related outcomes. Among supplier firms, the referrals generated significant improvements in both dimensions of quality, accompanied by increased time devoted to quality control. Among client firms, the main effects are in product variety. The treatment increased the likelihood that the firm offered a second product and the share of firm revenue coming from the second product. The treatment also increased the average price of the firm's product portfolio. This increase reflects the fact that firms' second products are, on average, more expensive: second products in the industry tend to be of higher quality. Consistent with client firms expanding their sales of a higher-quality second product, they also spent more time on quality control.[43]

These findings suggest the mechanism that treated clients used intermediate inputs from their new suppliers to expand the production of their second product. Because this product is of higher quality, its expansion required improvements in input quality, inducing treated suppliers to upgrade. Thus, suppliers and clients upgraded in a complementary fashion: suppliers increased product quality because clients expanded product variety, and clients expanded product variety because

[43] Consistent with the logic of our conceptual framework, Appendix Table A15 shows that these gains in client firms are more pronounced among those with low baseline degree.

Table 8: Product upgrading

Dep. Var.:	Quality score main product		Quality check	If 2nd	Share 2nd	Avg Log
	Craftsmanship	Durability	hrs/wk	product	product	Price
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Suppliers						
Treatment	0.264***	0.258***	4.329***	0.029	2.889*	0.095
	(0.087)	(0.090)	(0.794)	(0.035)	(1.541)	(0.080)
Baseline Control	Yes	Yes	Yes	Yes	Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	-0.198	-0.191	1.759	0.449	14.05	0.558
Observations	408	408	763	763	763	624
Panel B: Clients						
Treatment	0.217*	-0.023	2.446**	0.091**	5.777***	0.187*
	(0.121)	(0.132)	(0.959)	(0.036)	(1.586)	(0.113)
Baseline Control	Yes	Yes	Yes	Yes	Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	-0.351	-0.382	6.033	0.715	19.33	1.388
Observations	292	292	520	520	518	496

Notes: Sample includes firms in both follow-up waves. Quality score in columns 1-2 is evaluated by experts, residualized by product type and standardized. In column 5, the outcome is the revenue share of the firm's second product, zero for firms that do not have a second product. In column 6, the outcome is the sales-weighted average of the log prices of the firm's main products for which we collect data on, up to three. Control mean is at baseline. Standard errors clustered by firm. *** p<0.01, ** p<0.05, * p<0.1.

suppliers increased product quality. These patterns are consistent with quality complementarities in production networks (Kremer 1993, Demir, Fieler, Xu and Yang 2024). More broadly, they suggest that frictions in firm-to-firm access may be an important barrier to industrial upgrading in developing economies (Verhoogen 2023).

5.4 Indirect effects on firm performance

Estimating indirect effects on firm performance is important for assessing the intervention's overall impacts on the industry. Our conceptual framework predicts negative indirect effects under the reallocation mechanism and positive indirect effects under the expansion mechanism. Thus, the sign of the indirect effects is a further test of these mechanisms. Given that our network-level

Table 9: Indirect effects

Dep. Var.:	Log Sales	Profit (10,000 RMB)	Log Operating Cost	Log Employ-ment	Log Total Hrs Worked / Day	Num of Partners	Satisfaction with Partners
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Exposure	0.083	4.861*	-0.009	0.125**	0.115*	1.036	-0.018
	(0.119)	(2.582)	(0.152)	(0.062)	(0.068)	(0.927)	(0.082)
Treatment	0.169***	1.185	0.063	0.064*	0.106**	1.753***	0.240***
	(0.063)	(1.686)	(0.076)	(0.036)	(0.043)	(0.513)	(0.047)
Baseline Control	Yes	Yes	Yes	Yes		Yes	Yes
Controls, Strata, Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	2.468	19.05	1.921	0.966	2.964	5.798	5.159
Observations	1281	1283	1154	1283	1282	1283	1257

Notes: Sample is all firms in both follow-up waves. Exposure is the demeaned share of baseline partners treated. The outcomes in columns 2 and 6 are winsorized, and in column 7 are standardized. Controls include indicators for the firm having no regular partners at baseline and for being a client firm. The baseline value of the outcome is not included in column 5 because it was not collected. Control mean is at baseline. Standard errors clustered by firm. *** p<0.01, ** p<0.05, * p<0.1.

results in Section 5.1 suggest expansion, we expect indirect effects on performance to be positive.

We estimate indirect effects using a specification analogous to equation (4) in Section 5.1, regressing outcomes on treatment and exposure. To maximize power we (i) pool suppliers and clients, and (ii) omit the interaction between treatment and exposure, since Table 5 shows no differential exposure effects on degree by treatment status.

Table 9 reports the results. Exposure has broadly positive effects on firm performance, with statistically significant impacts on profit, employment, and hours worked.[44] These findings are consistent with the expansion mechanism and reinforce our earlier evidence on network expansion and upgrading. Specifically, as shown in Table 5, exposure did not reduce firms' degree, limiting negative indirect effects through crowd-out. At the same time, exposure created upgrading (Table 8), increasing the potential for positive indirect effects.[45] Overall, these findings provide strong

[44] As shown by Table A2, exposure is not balanced at baseline with respect to partner satisfaction, so we prefer not to interpret the (insignificant) impact on satisfaction.

[45] The logic of positive spillovers further predicts that more exposed firms should have "better" partners in the follow-up waves, as the upgrading of the treated among these partners is the source of the positive spillover. Appendix Table A16 confirms that exposure positively predicts average partner quality in the follow-up data.

support for hypothesis H5.

The results, presented in Appendix Table A17, show that the treatment effect estimates remain similar in magnitude and statistical significance to those reported in Tables 6 and 7, indicating that indirect effects do not confound our estimates of the direct effects.[46]

Indirect effects in final goods market. A second type of indirect effect may arise from business stealing in the final goods market. If treated client firms upgrade, they may attract consumers away from competitors. We expect this effect to be weak in our sample. Brush pens are produced and traded all over China, and both intuition and evidence suggest that business stealing effects on tradable goods are not localized (Rotemberg 2019). Moreover, in standard monopolistic competition models, business stealing that results from treated firms' upgrading affects treated and untreated firms equally to a first-order approximation (Cai and Szeidl 2024).[47] Thus, business stealing in the final goods market is unlikely to bias our treatment-control comparisons.[48]

5.5 Concerns with measurement, identification, and interpretation

We discuss some remaining potential concerns with our results.

Experimenter demand. One concern is that our firm-level treatment effects reflect experimenter demand: firms, in exchange for the subsidy, may have over-reported their performance. We address this concern using two sets of outcomes plausibly robust to experimenter demand. In the endline survey we collected data on the owner's phone use, including their monthly phone bill, the number of calls they made, and share of business-related calls. In the same survey, we also asked our enumerators—who were unaware of firms' treatment status—to record indicators of the firm's activity during their visit, including the number of customers present, the number of calls made during the visit, and the number of employees on site. Enumerators also rated how busy the owner

[46] The exposure effects are less significant than in Table 9 because of the smaller sample sizes.

[47] This follows because upgrading by treated firms affects the market by reducing the (quality-adjusted) price index which governs the level of demand for all firms, including both treated and untreated firms. Consistent with this logic, in Cai and Szeidl (2024) the estimated business stealing effects on treated and untreated firms are not significantly different.

[48] We also test for this effect directly, by including in the regression the share of a firm's competitors that are treated. Competitors are defined by refining the similarity measure used to construct the referrals with spatial closeness. Table A18 shows that the share of competitors treated has an insignificant effect on firm performance. However, we note that given the small sample size for client firms these estimates come with wide confidence intervals.

was and the firm's overall operations on a five-point scale.[49] Appendix Table A19 shows positive and statistically significant treatment effects on several of these outcomes for both suppliers and low-degree clients, consistent with genuine increases in business activity.

A related concern is that experimenter demand could bias our network-level findings if firms overstated transactions with referred partners. However, this explanation cannot account for the larger link-formation effects of screened relative to unscreened referrals. The two types of referrals were presented identically to firms, and the enumerators administering the intervention were themselves unaware of whether a referral had been screened. Thus, neither firms' expectations nor experimenter behavior could have differed systematically across the two referral types. We conclude that experimenter demand is unlikely to explain our results.

Income effect of subsidies. Another concern is that the subsidy for the first transaction generated an income effect or relaxed credit constraints, driving impacts through a different mechanism. However, the subsidy cap per link was only about 0.4% of average client profit—which is the relevant benchmark, since subsidies were received by client firms. Among subsidy-receiving client firms, the average subsidy amounted to 0.89% of average profit, and to 1.09% among low-degree clients. These magnitudes are small and unlikely to generate meaningful income effects. Moreover, our baseline survey evidence suggests that firms were not credit constrained: only 38 firms (5.6%) identified credit constraints as their primary growth barrier, while barriers related to production and market access were much more common. We conclude that the subsidy itself is unlikely to explain the observed performance gains.

Measured vs. economic profits. A concern with our measured profit effect is that it may partly reflect an increase in owner labor income rather than economic profit. Our profit measure follows the standard approach in the literature on firms in development economics. As discussed by de Mel, McKenzie and Woodruff (2009), directly eliciting reported profits performs at least as well as reconstructing profits from more detailed questions on revenues and expenses. We are not able to net out owners' labor income from profits: we do not have data on the wages of owner-managers, and we do not observe owner work hours after the intervention, because we dropped this question after

[49] On this scale, 1 represents "extremely poor, on the verge of closing down," while 5 represents "highly prosperous, business is thriving."

the 2021 survey to shorten the questionnaire. However, the 2021 baseline hours data suggest that increased owner labor is unlikely to be the primary driver of our results. Before the intervention, owner-managers worked an average of 8.4 hours per day, seven days per week, leaving limited scope for a substantial increase in owner labor.

Even if the increase in reported profits partly reflects greater owner effort, it still provides evidence that firms expanded operations. In fact, additional owner effort may represent one mechanism through which firms managed larger production networks. This interpretation is consistent with the broader pattern of evidence, including increases in revenue, the number of partners, product upgrading, and positive indirect effects.

Dynamics. Even if the referrals improved business performance initially, these improvements may have been short-lived: after experimenting with the referrals, firms may have switched back to their old partners. To explore this, in Appendix Table A20 we estimate the impacts separately in the two follow-up surveys of 2022 and 2023. Power is lower in these year-specific regressions, but the estimated effects are similar in magnitude across the two years and generally not significantly different from each other. A second piece of evidence comes from our 2024 phone survey, in which we asked firms to report their performance over 2020-24 on a 5-point scale. As we discussed above, Appendix Table A12 shows significant gains for revenue, profit, and a number of other outcomes. Taken together, these findings suggest that the referrals generated persistent improvements.

Interactions with e-commerce treatment. Our intervention was cross-randomized with an e-commerce experiment that we implemented in the same population of firms in 2019. Since the treatments were cross-randomized, the e-commerce intervention does not confound our results. However, it may affect the interpretation of our results: the positive effects of the referrals could have depended on the increase in demand generated by the e-commerce intervention. To explore this, in Appendix Table A21 we re-estimate our main firm-level regressions, adding the interaction between the referral and the e-commerce treatment. For suppliers, the interaction does not affect the results, which is not surprising because few of these firms engage in e-commerce. For low-degree clients, the referral effects are imprecisely estimated both among firms that did and did not receive the e-commerce treatment. We are underpowered to separately estimate these impacts. Thus, we cannot rule out that sufficient access to final good markets is a prerequisite for firm-to-firm access to improve business performance.

5.6 Private and social return

Hypothesis H6 says that in a low-access trap, the referrals can yield high returns. We turn to estimate the private and social return generated by the intervention.[50]

Private return. We define the private return as the expected increase in reported profit gain from a subsidized referral, relative to the cost of the subsidy. Firms in the information treatment could have earned this expected return by self-financing the subsidy on their information referral. We estimate the private return separately for suppliers and for clients. For suppliers, we compute the profit gain by multiplying the profit effect from column 2 of Table 6 with the number of suppliers, and compare it to the total subsidy cost we paid out. We find that the annual profit gain is about 9.4 times the subsidy cost. For clients, we focus on low-degree firms and compute the profit gain as the product of the profit effect from column 2 of Table 7 and the number of firms in this subgroup. This gain is about 25 times the total cost of the subsidy we paid out (to all firms). Thus, in the sample of suppliers and low-degree clients, we estimate private returns ranging between 940-2,500% per year.

These financial returns substantially exceed existing estimates of the private return to business capital in the development context, which are on the order of 100 percent per year (de Mel, McKenzie and Woodruff 2008, Banerjee and Duflo 2014, Cai and Szeidl 2024). Our large returns further support the interpretation that the barrier to taking up the information referral was not a well-calibrated cost-benefit assessment, and that firms may have been stuck in a low-access trap.

Social return. We next develop an estimate of the social return, but we note that this estimate relies on more assumptions and should be interpreted as suggestive. The social return needs to account for three additional welfare impacts not included in the private return: the consumer surplus; spillover effects on other firms; and the cost of collecting the data used to create the

[50] The producer-surplus component of our return calculations is based on reported profits. To the extent that reported profits include compensation to owner-managers, these calculations do not net out the opportunity cost of additional owner effort.

referrals. We account for each of these effects in turn.

The gain in consumer surplus in the brush pen market comes from client firms upgrading their product portfolio, including improvements in variety and quality. To estimate this gain, we build on Feenstra (1994) and Cai and Szeidl (2024) and in Appendix C derive an expression for the gain in the consumer surplus under a CES demand system, which is valid even in the presence of business stealing in the final good market. This expression depends only on the treatment effect on log revenue and the elasticity of substitution among final goods.[51] We use this formula for low-degree clients, and—following Cai and Szeidl (2024)—employ an elasticity of substitution $\sigma = 5$, which is a conservative choice given recent estimates of the retail elasticity of substitution (Atkin, Faber and Gonzalez-Navarro 2018, Dolfen, Einav, Klenow, Klopack, Levin, Levin and Best 2023). Turning to spillovers on other firms, we estimated positive spillovers in the input market, which we conservatively bound with zero. We do not have a reliable estimate of spillovers in the final good market, but our model in Appendix C implies that any such business stealing effect is dominated by the gain to treated clients. That is, the overall producer surplus of client firms, including spillovers, should not shrink, and can therefore be conservatively bounded by zero. Finally, we compute returns relative to the total cost of conducting the intervention, including both the subsidy and the survey costs.

We find that the gain in consumer surplus was 5.7 times the cost of the intervention. The gain in producer surplus for suppliers, measured with the treatment effect on profits, was 4.2 times the cost of the intervention. Given the above bounds of zero for the other effects, our conservative estimate of the social return is 990%. We emphasize that this calculation relies on multiple assumptions and should be taken as suggestive. Nevertheless, the high estimate is consistent with hypothesis H6 that in a low-access trap, improving access can generate large returns; and suggests that improving access can generate substantial value by expanding production networks and creating new economic activity.

[51] The gain in consumer surplus is increasing in the treatment effect on log revenue which measures the extent to which consumers shift their demand in response to the upgrading of treated clients; and decreasing in the elasticity of substitution which measures the extent to which this demand shift reflects a gain in welfare.

6 Marketing offer intervention

From a policy perspective, an important question is whether there are lower-cost policies than the subsidy to overcome partnering frictions. Recall from Section 4.3 that our survey of firms participating in the information referrals suggested two main frictions: (i) undervaluation of potential partners and (ii) lack of marketing skills. To address both, we conducted a small-scale follow-up experiment in 2023 evaluating an alternative way to make screened referrals. Specifically, we introduced a *marketing offer* intervention. In collaboration with the supplier, we prepared a one-page flyer describing the supplier's product, including its type and price, and a free sample produced by the supplier if one was available. We then delivered these materials to the client on behalf of the supplier. The marketing offer was expected to overcome undervaluation by providing verifiable product information, and to alleviate marketing constraints by initiating contact with the client.

We implemented the new experiment in the sample of T0-to-T0 candidate screened referrals, which consisted of the two candidate screened referrals we constructed for each T0 (untreated) client in the main experiment. This sample included both the "control" referrals that we constructed but did not make, and the information referrals. We included the latter because they did not lead to subsequent transactions. We introduced two treatment arms:

- EI: Enhanced information. We provided the same information as in the information arm of the main experiment, and in addition explained that similar referrals to other firms led to subsequent trades, and encouraged the firm to contact the proposed partner.

- MO: Marketing offer. In addition to EI, we delivered the marketing offer.

We randomized the treatments as follows. For a random third of client firms, we administered the EI treatment to both of their screened referrals. For the remaining two-thirds, we administered the MO treatment to one screened referral (the A referral in Figure 2), and the EI treatment to the other (the B referral). This design allows us to identify both the direct effect of the marketing offer and potential within-firm spillovers from the MO to the EI referral. Both undervaluation and the lack of marketing skills imply possible spillovers: the former through learning the value of partners, the latter through learning marketing skills.

We use the following estimating equation:

$$y_{ij} = \delta_M \cdot \textit{Marketing offer}_{ij} + \delta_C \cdot \textit{Client spillover}_{ij} + \delta_S \cdot \textit{Supplier spillover}_{ij} + \textit{Controls}_{ij} + \varepsilon_{ij}. \quad (6)$$

This is a link-level regression using the full sample of T0-to-T0 screened referrals. The outcome is an indicator of whether a transaction occurred, measured in the 2024 phone survey. On the right-hand side, the main treatment variable is an indicator for the MO treatment. Two other variables measure spillovers. *Client spillover* is an indicator for an MO-untreated referral where the client has another referral assigned to MO. This variable captures the possibility that the client learns about partner quality from exposure to its other (MO-treated) referral. *Supplier spillover* is defined analogously for the supplier: it equals one for an MO-untreated referral where the supplier has another referral treated by MO, capturing potential learning on the supplier side. The treatment creates variation in both spillover variables. This variation is effectively random on the client side: each client has two referrals, and at most one is randomly treated by MO, creating opportunity for the client to learn about the other. But the variation is not purely random on the supplier side, as suppliers connected to more screened referrals are mechanically more likely to have at least one MO-treated referral. To isolate the random component, we include controls for the number of screened referrals associated with each supplier.

Table 10 reports the results. Column 1 shows that the marketing offer increased the probability of a subsequent transaction by a significant 19 percentage points. Columns 2 and 3 show that it also significantly increased the number of transactions and the transaction value. The spillover coefficients are always positive, but only significant in one case, providing suggestive evidence for learning effects.[52] We conclude that the marketing offer was effective in creating connections.

To gauge the magnitude of the effect of the marketing offer, we compare it to that of the subsidized referral in the main experiment. A direct comparison with Table 3 is complicated by the fact that in the main intervention we evaluated impacts over a year, whereas here only about

[52] Since we lack a pure control arm, we are unable to identify the impact of enhanced information. However, take-up in that arm was only about 6%, much lower than the marketing offer effect, suggesting that even intensive information treatments have limited capacity to generate new partnerships.

Table 10: Marketing offer intervention

Dep. Var.:	Link	Num Transactions	Value (RMB)
	(1)	(2)	(3)
Marketing Offer	0.191***	0.457**	791.613**
	(0.063)	(0.181)	(358.221)
Client Spillover	0.077	0.126	130.973
	(0.060)	(0.171)	(338.328)
Supplier Spillover	0.141*	0.287	526.271
	(0.081)	(0.231)	(457.996)
Controls	Yes	Yes	Yes
Business type FE	Yes	Yes	Yes
Observations	218	218	218

Notes: Sample is T0 candidate referrals. Client Spillover is an indicator equal to 1 for EI referrals where the client received a marketing offer on another referral; Supplier Spillover is defined analogously for suppliers. Controls include indicators for the number of times a supplier has been referred. Standard errors clustered by firm. *** p<0.01, ** p<0.05, * p<0.1.

4 months elapsed between our intervention and the phone survey (from August 2023 to January 2024). Therefore, we re-estimated Table 3 for a four-month horizon (not reported), which was possible because in the follow-up surveys we asked firms about transactions at that horizon as well. We find that the marketing offer was about 51% as effective as the subsidized referral in creating a new transaction, created about 54% as many transactions, and about 40% in transaction value. Thus, the marketing offer appears to be at least 40 percent as effective as the subsidy, while costing only 0.5 percent as much. These findings suggest that improving the design of making connections has the potential to increase the returns from access.

7 Conclusion

In this paper, we studied the impact of firm-to-firm referrals on business networks and performance. We obtained two main results. First, firms' undervaluation of potential partners created a low-access trap. Providing information alone did not create partnerships, whereas combining information with a small subsidy did. Subsidized referrals increased beliefs about partner value, stimulated private search, and led to new non-referred relationships. Second, improving access changed industry equilibrium by expanding, rather than merely reallocating, the production network. Referrals increased firms' number of partners, promoted upgrading in product quality and variety, improved business performance, and generated positive indirect effects. The implied private and social returns were substantial, although the latter rely on strong assumptions.

Both of our main results have broader implications. A key implication of the first is that belief-based traps can meaningfully constrain economic performance even in a reasonably functional manufacturing context in a middle-income country. More broadly, the undervaluation mechanism may help explain a central puzzle in private sector development: firms' under-adoption of profitable technologies and practices (Verhoogen 2023). We document undervaluation for one practice—search for new partners—but see no reason why similar forces could not affect other practices as well. If firms underestimate the returns to experimentation, overcoming undervaluation may require direct experience with a product or practice. This perspective helps explain why demonstrations, such as piloting (Bloom, Eifert, Mahajan, McKenzie and Roberts 2013) or firm visits (Cai and Szeidl 2018), seem to increase adoption.

A key implication of our second result is that barriers to access can create a new form of inefficiency in industrial production: excess sparsity in production networks. This inefficiency may contribute to differences in industry productivity across places and over time. It also has important implications for policy. When production networks are excessively sparse, improving access can generate large gains because the new partnerships create new economic activity rather than merely reallocating existing activity. Moreover, indirect effects can amplify, rather than offset, the direct benefits of intervention.

Our results hold two policy lessons. The first is that firm-to-firm access can be inefficiently under-provided, making a case for industrial policy. Undervaluation causes firms to search too little even from their own perspective. In addition, search externalities and positive spillovers imply that privately optimal search can fall short of the social optimum. Together, these forces imply that policies improving access can create substantial value.

The second policy lesson is that in the presence of undervaluation, industrial policy has a new type of objective: improving firms' beliefs about the value of unfamiliar opportunities. Policies that merely provide information may be insufficient. Instead, effective policies may need to provide firms with direct experience of unfamiliar opportunities. In the context of access, experience can take the form of subsidized trial transactions, as in our main experiment, or product samples and active outreach, as in our marketing intervention. More generally, pilot demonstrations and visits to adopters can improve beliefs and encourage experimentation with unfamiliar opportunities.

Our findings raise three sets of questions for future research. (i) Most closely related to this paper, one question is how to design effective policies to improve access. This includes designing both the form of creating access—e.g., trade fairs, marketing offers—and its targeting, namely which firms and which links to target. Our results indicate that low-degree firms benefit most, and that screened referrals are particularly effective, highlighting the value of network-based targeting. Implementing such targeting at scale may require the use of administrative production network data from VAT filings. (ii) In the context of undervaluation, a key question concerns the causes of undervaluation. One possible mechanism is the competence trap highlighted in organizational learning theory, where firms focus on their main activities, which can lead to the undervaluation of alternatives. More broadly, an important question is whether undervaluation constrains the adoption of management practices, technologies, and other profitable business decisions. (iii) Turning to production networks, understanding the determinants and consequences of network density may help explain persistent differences in industrial performance across places and over time.

References

Abadie, Alberto, Susan Athey, Guido W Imbens, and Jeffrey M Wooldridge, "When Should You Adjust Standard Errors for Clustering?*," *The Quarterly Journal of Economics*, 10 2022, *138* (1), 1–35.

Alfaro-Ureña, Alonso, Isabela Manelici, and Jose P Vasquez, "The Effects of Joining Multinational Supply Chains: New Evidence from Firm-to-Firm Linkages*," *The Quarterly Journal of Economics*, 01 2022, *137* (3), 1495–1552.

Amiti, Mary and Jozef Konings, "Trade Liberalization, Intermediate Inputs, and Productivity: Evidence from Indonesia," *American Economic Review*, December 2007, *97* (5), 1611–1638.

Asiedu, Edward, Monica Lambon-Quayefio, Francesca Truffa, and Ashley Wong, "Female Entrepreneurship and Professional Networks," Working Paper 2026.

Atkin, David, Amit K. Khandelwal, and Adam Osman, "Exporting and Firm Performance: Evidence from a Randomized Experiment," *The Quarterly Journal of Economics*, 02 2017, *132* (2), 551–615.

____, Benjamin Faber, and Marco Gonzalez-Navarro, "Retail Globalization and Household Welfare: Evidence from Mexico," *Journal of Political Economy*, 2018, *126* (1), 1–73.

Bai, Jie, "Melons as lemons: Asymmetric information, consumer learning and seller reputation," *Review of Economic Studies*, 2025, *92* (6), 3574–3610.

____, David Sungho Park, and Ajay Shenoy, "Misperceptions and Product Choice: Evidence from a Randomized Trial in Zambia," Working Paper 2025.

Balboni, Clare, Oriana Bandiera, Robin Burgess, Maitreesh Ghatak, and Anton Heil, "Why do people stay poor?," *The Quarterly Journal of Economics*, 2022, *137* (2), 785–844.

Banerjee, Abhijit and Esther Duflo, "Do Firms Want to Borrow More? Testing Credit Constraints Using a Directed Lending Program," *Review of Economic Studies*, 2014, *81*, 572–607.

____, Emily Breza, Esther Duflo, and Cynthia Kinnan, "Can microfinance unlock a poverty trap for some entrepreneurs?," Working Paper 2025.

Baqaee, David, Ariel Burstein, Cédric Duprez, and Emmanuel Farhi, "Consumer Surplus From Suppliers: How Big Is It and Does It Matter for Growth?," *Econometrica*, 2025, *93* (6), 2043–2081.

Bergquist, Lauren, Craig McIntosh, and Meredith Startz, "Search Costs, Intermediation, and Trade: Experimental Evidence from Ugandan Agricultural Markets," Working Paper 2026.

Bernard, Andrew B., Andreas Moxnes, and Yukiko U. Saito, "Production Networks, Geography, and Firm Performance," *Journal of Political Economy*, 2019, *127* (2), 639–688.

Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts, "Does Management Matter? Evidence from India," *Quarterly Journal of Economics*, 2013, *128*, 1–51.

Borusyak, Kirill and Peter Hull, "Nonrandom Exposure to Exogenous Shocks," *Econometrica*, 2023, *91* (6), 2155–2185.

Breza, Emily and Cynthia Kinnan, "Measuring the Equilibrium Impacts of Credit: Evidence from the Indian Microfinance Crisis," *The Quarterly Journal of Economics*, 05 2021, *136* (3), 1447–1497.

Cai, Jing and Adam Szeidl, "Interfirm Relationships and Business Performance," *The Quarterly Journal of Economics*, 12 2018, *133* (3), 1229–1282.

___ **and** ___ , "Indirect Effects of Access to Finance," *American Economic Review*, August 2024, *114* (8), 2308–51.

Campante, Filipe and David Yanagizawa-Drott, "Long-Range Growth: Economic Development in the Global Network of Air Links*," *The Quarterly Journal of Economics*, 12 2017, *133* (3), 1395–1458.

de Mel, Suresh, David J. McKenzie, and Christopher Woodruff, "Measuring Microenterprise Profits: Must We Ask How the Sausage Is Made?," *Journal of Development Economics*, 2009, *88* (1), 19–31.

de Mel, Suresh, David McKenzie, and Christopher Woodruff, "Returns to Capital in Microenterprises: Evidence from a Field Experiment," *The Quarterly Journal of Economics*, 11 2008, *123* (4), 1329–1372.

Demir, Banu, Ana Cecília Fieler, Daniel Yi Xu, and Kelly Kaili Yang, "O-Ring Production Networks," *Journal of Political Economy*, 2024, *132* (1), 200–247.

___ **, Beata Javorcik, and Piyush Panigrahi**, "Breaking invisible barriers: Does fast internet improve access to input markets?," Working paper 2026.

Denrell, Jerker and James G. March, "Adaptation as Information Restriction: The Hot Stove Effect," *Organization Science*, 2001, *12* (5), 523–538.

Dillon, Brian, Jenny C Aker, and Joshua E Blumenstock, "Yellow Pages: Information, connections and firm performance," *The Economic Journal*, 2025, *135* (670), 2067–2082.

Dolfen, Paul, Liran Einav, Peter J Klenow, Benjamin Klopack, Jonathan D Levin, Larry Levin, and Wayne Best, "Assessing the gains from e-commerce," *American Economic Journal: Macroeconomics*, 2023, *15* (1), 342–370.

Eaton, Jonathan and Samuel Kortum, "Technology, geography, and trade," *Econometrica*, 2002, *70* (5), 1741–1779.

Egger, Dennis, Benjamin Faber, Ming Li, and Wei Lin, "Rural-Urban Migration and Market Integration," Working Paper 34098, National Bureau of Economic Research August 2025.

___ **, Johannes Haushofer, Edward Miguel, Paul Niehaus, and Michael Walker**, "General equilibrium effects of cash transfers: experimental evidence from Kenya," *Econometrica*, 2022, *90* (6), 2603–2643.

Feenstra, Robert C, "New product varieties and the measurement of international prices," *The American Economic Review*, 1994, pp. 157–177.

Gertler, Paul, Sean Higgins, Ulrike Malmendier, and Waldo Ojeda, "Do Behavioral Frictions Prevent Firms from Adopting Profitable Opportunities?," Working Paper, National Bureau of Economic Research 2025.

Halpern, László, Miklós Koren, and Adam Szeidl, "Imported Inputs and Productivity," *American Economic Review*, December 2015, *105* (12), 3660–3703.

Hjort, Jonas, Golvine de Rochambeau, Vinayak Iyer, and Fei Ao, "Informational Barriers to Market Access: Experimental Evidence from Liberian Firms," Working Paper 2025.

James, Gareth, Daniela Witten, Trevor Hastie, and Robert Tibshirani, *An Introduction to Statistical Learning: with Applications in R*, 2 ed., Springer, 2021.

Jensen, Robert, "The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector," *The Quarterly Journal of Economics*, 2007, *122* (3), 879–924.

—— **and Nolan H. Miller**, "Market Integration, Demand, and the Growth of Firms: Evidence from a Natural Experiment in India," *American Economic Review*, December 2018, *108* (12), 3583–3625.

JPAL, Abdul Latif Jameel Poverty Action Lab, "Market Access: Connecting Firms and Entrepreneurs to Markets to Spur Business and Job Growth," Technical Report, J-PAL Policy Insights 2024.

Kremer, Michael, "The O-Ring Theory of Economic Development," *The Quarterly Journal of Economics*, 1993, *108* (3), 551–575.

Lee, David S., "Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects," *The Review of Economic Studies*, 2009, *76* (3), 1071–1102.

Levinthal, Daniel A. and James G. March, "The Myopia of Learning," *Strategic Management Journal*, 1993, *14* (S2), 95–112.

Levitt, Barbara and James G. March, "Organizational Learning," *Annual Review of Sociology*, 1988, *14*, 319–340.

Liu, Ernest, "Industrial policies in production networks," *The Quarterly Journal of Economics*, 2019, *134* (4), 1883–1948.

March, James G., "Exploration and Exploitation in Organizational Learning," *Organization Science*, 1991, *2* (1), 71–87.

McKelway, Madeline, "Women's Self-Efficacy and Economic Outcomes: Evidence from a Two-Stage Experiment in India," Working paper 2025.

McKenzie, David, "Beyond baseline and follow-up: The case for more T in experiments," *Journal of Development Economics*, 2012, *99* (2), 210–221.

Rotemberg, Martin, "Equilibrium Effects of Firm Subsidies," *American Economic Review*, October 2019, *109* (10), 3475–3513.

SAMR, "National Standards of the People's Republic of China GB/T 34854-2017: Chinese Traditional Stationery - Writing Brush," the National Standardization Administration, The State Administration for Market Regulation 2017.

Startz, Meredith, "The value of face-to-face: Search and contracting problems in Nigerian trade," Working Paper, Dartmouth College 2025.

Szeidl, Adam, "Firm-to-Firm Access and Private Sector Development," Working Paper 2025.

Verhoogen, Eric, "Firm-Level Upgrading in Developing Countries," *Journal of Economic Literature*, December 2023, *61* (4), 1410–64.

Verhoogen, Eric A., "Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector*," *The Quarterly Journal of Economics*, 05 2008, *123* (2), 489–530.

Wiles, Edward and Deivy Houeix, "Relational Frictions Along the Supply Chain: Evidence from Senegalese Traders," Working Paper 2026.

YEE, CHIANG, *Chinese Calligraphy: An Introduction to Its Aesthetic and Technique, Third Revised and Enlarged Edition*, Harvard University Press, 1973.

A Supplemental evidence

Table A1: Baseline Balance for Firms Remaining in Sample in 2023

	Suppliers		Clients	
	Mean	T1-T0	Mean	T1-T0
Panel A: Firm Characteristics				
Firm Age	26.483	0.325	21.830	-0.897
	(13.132)	(1.351)	(11.828)	(1.489)
Number of Employees	2.525	0.239	5.312	-0.143
	(2.221)	(0.228)	(5.404)	(0.681)
Profit (10,000 RMB)	7.676	2.271	34.888	3.743
	(19.163)	(1.968)	(62.284)	(7.844)
Sales (10,000 RMB)	15.665	2.363	76.664	7.414
	(33.681)	(3.463)	(129.228)	(16.276)
Panel B: Managerial Characteristics				
Gender (1=Female, 0=Male)	0.322	0.009	0.115	-0.021
	(0.468)	(0.048)	(0.319)	(0.040)
Age	53.565	-0.211	49.625	0.236
	(9.866)	(1.015)	(10.442)	(1.316)
Education- Middle School	0.462	0.009	0.700	0.056
	(0.499)	(0.051)	(0.459)	(0.058)
Panel C: Number of Partners (buyers for supplier, sellers for client)				
In Firm Data	6.211	1.217	5.636	-0.356
	(11.930)	(1.226)	(6.672)	(0.840)
In Network Data, All Partners	2.361	0.189	3.642	0.049
	(3.076)	(0.316)	(3.866)	(0.487)
In Network Data, Supplier-client Links	1.562	0.101	2.273	0.014
	(2.714)	(0.279)	(3.376)	(0.425)
Length of Relationship	8.958	0.373	9.492	-0.059
	(5.005)	(0.673)	(4.912)	(0.800)
P Value of Joint Significance of Vars in Panels A-C		0.884		0.967
Observations	379	379	253	253

Notes: Sample is firms in the second (2023) follow-up wave. All variables are from 2021 baseline data. Column 1 reports the mean for suppliers, column 2 the difference between treated (T1) and untreated (T0) suppliers, and columns 3 and 4 are analogous for clients. *** p<0.01, ** p<0.05, * p<0.1.

Balance. Table A1 presents summary statistics for the subset of firms that were in our sample both in the 2021 baseline and the 2023 endline. That is, firms subject to attrition or shutdown between the baseline and the second follow-up are excluded. We look at this subsample because in the full sample there is a significant difference in attrition in 2023 among supplier firms. The table shows that the treatment remains balanced in this subsample. In particular, the p-value for the joint significance of all variables in Panels A-C among supplier firms exceeds 0.9.

Table A2 presents balance tests for the exposure variable, which is defined as the demeaned share of the firm's baseline partners (in the network data) that are treated. We regress baseline outcomes on treatment, exposure, as well as an indicator for the firm having no regular partners in the network data at baseline. The latter is necessary because exposure is only defined for firms that have regular partners; otherwise, we set its value to zero. The table shows that exposure is well-balanced, with the exception of the satisfaction with partners measure.

Table A2: Baseline balance of exposure

Dep. Var.:	Log Sales	Profit (10,000 RMB)	Log Operating Cost	Log Employ-ment	Num of Partners	Satisfaction with Partners	Degree (network data)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Exposure	-0.122	1.737	-0.001	-0.004	1.679	0.335**	0.163
	(0.144)	(4.318)	(0.171)	(0.081)	(1.606)	(0.138)	(0.292)
Treatment	0.007	0.248	-0.014	0.009	0.443	-0.022	0.184
	(0.086)	(2.775)	(0.100)	(0.045)	(0.738)	(0.075)	(0.177)
Strata FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	2.468	19.05	1.921	0.966	5.798	5.159	1.784
Observations	681	682	615	682	682	681	682

Notes: Sample is all firms at baseline. Exposure is demeaned share of baseline partners treated. The outcome in column 2 is winsorized, in column 5 it is winsorized, and in column 6 it is standardized. In column 7, the outcome is baseline degree in the network survey. Standard errors clustered by firm. *** p<0.01, ** p<0.05, * p<0.1.

Screened versus unscreened referrals. Table A3 compares screened and unscreened referrals by regressing, in the sample of screened and unscreened candidate referrals (both made and unmade) an indicator for screened referral on supplier, client, and joint characteristics. Column 1 shows that

Table A3: Difference between screened and unscreened referrals

Dep. Var:	If Screened Referral				
	(1)	(2)	(3)	(4)	(5)
Supplier Log Sales	0.020				
	(0.020)				
Supplier Log Employee	0.077*				
	(0.040)				
Supplier Firm Age	0.001				
	(0.002)				
Supplier Owner Age	0.002				
	(0.002)				
Supplier Log Num Client	0.002				
	(0.020)				
Client Log Sales	0.000				
	(0.016)				
Client Log Employee	0.004				
	(0.029)				
Client Firm Age	-0.000				
	(0.002)				
Client Owner Age	-0.001				
	(0.002)				
Client Log Num Supplier	-0.001				
	(0.022)				
Diff Price Rank		-0.141*			
		(0.081)			
Diff Sales Rank			0.175**		
			(0.068)		
Log Distance				-0.033*	
				(0.019)	
Kin					-0.002
					(0.075)
Business Type FE	Yes	Yes	Yes	Yes	Yes
Observations	796	646	799	805	858

Note: Sample is all candidate referrals, screened and unscreened, made and not made. Sales rank for suppliers is computed by ranking baseline (2021) sales across suppliers and normalizing the rank to range from 0 to 1. Sales rank for clients, and price rank, are computed analogously. Difference is the absolute value of the difference in rank between the supplier and client for a given link. Standard errors clustered by supplier and client. *** p<0.01, ** p<0.05, * p<0.1.

screened and unscreened referrals are on average indistinguishable in most firm characteristics, with only a slight difference in the employee count for supplier firms (which is probably due to screened referrals oversampling large suppliers). Column 2 shows that the smaller the difference in price rank is associated with a higher likelihood that the referral is screened. Price rank is defined by ranking suppliers and clients separately based on the price of their main product and normalizing each rank to the unit interval. That is, screened referrals are more closely aligned in price, which we interpret as a proxy for alignment in quality. Column 3 looks at sales rank, defined analogously to price rank using baseline sales. We find that a larger difference in sales rank is associated with a higher likelihood that the referral is screened. Since unscreened referrals are effectively random, this result is suggestive of negative assortativity by revenue between suppliers and clients among screened referrals. Such a pattern could be driven by our matching algorithm: since large firms tend to be good matches for more partners, we likely assigned them systematically to small partners who otherwise would not be matched. Column 4 shows that screened referrals were slightly closer in space, which may be driven by a degree of geographic concentration of similar firms.

Heterogeneity in link creation. Table A4 estimates heterogeneous effects in our link creation regression (2). We combine the screened and unscreened subsidy treatments into a single subsidy indicator, which we interact with supplier, client, and pair level variables. The sample is the set of all candidate referrals. Links are more likely to be created when firms are larger; when firms are more similar in terms of price rank, a proxy for similarity in quality; and when firms are kin.

Heterogeneity in link survival. Tables A5 and A6 document various heterogeneities in the crowding out of prior links. In Table A5, the sample is all baseline links, the outcomes are different measures of activity on the link, and on the right-hand-side we include indicators for the supplier only, the client only, or both firms being treated. The omitted category is when both firms are untreated. The results show that all three types of links experience crowding out. In Table A6 we estimate heterogeneous effect versions of the regression in Table A5. The sample is all pre-existing links. The outcomes are different measures of activity on the link with the inverse hyperbolic sine (IHS) of the number and value of transactions also included. On the right-hand-side, we combine the three explanatory variables of the previous table into a single *Either treated* variable. Power is low, but the results suggest that links are more likely to survive when baseline satisfaction is higher, when suppliers are larger, when firms are kin, and when the relationship age is larger.

Table A4: Heterogeneity in link creation

Dep. Var.:	Link	Num Transactions	Value (RMB)	IHS Num Transactions	IHS Value
	(1)	(2)	(3)	(4)	(5)
Panel A					
Subsidy * Supplier log sales	0.086***	0.615***	1,497.083***	0.223***	0.859***
	(0.017)	(0.140)	(539.278)	(0.044)	(0.182)
Subsidy * Client log sales	0.041**	0.310**	969.928*	0.103**	0.400**
	(0.018)	(0.141)	(569.934)	(0.043)	(0.183)
Observations	1,586	1,586	1,586	1,586	1,586
Panel B					
Subsidy * Diff sales rank	0.041	0.888	2,840.432	0.206	0.551
	(0.102)	(0.817)	(3,274.867)	(0.256)	(1.073)
Observations	1,592	1,592	1,592	1,592	1,592
Panel C					
Subsidy * Diff price rank	-0.202**	-1.604**	-7,371.775***	-0.505**	-2.272**
	(0.099)	(0.726)	(2,700.155)	(0.237)	(0.986)
Observations	1,289	1,289	1,289	1,289	1,289
Panel D					
Subsidy * Log distance	-0.033	-0.194	-444.871	-0.085	-0.318
	(0.031)	(0.222)	(830.208)	(0.072)	(0.310)
Observations	1,499	1,499	1,499	1,499	1,499
Panel E					
Subsidy * Kin	0.248**	1.346	4,643.975	0.525**	2.518**
	(0.101)	(0.827)	(3,208.118)	(0.237)	(1.006)
Observations	1,592	1,592	1,592	1,592	1,592

Notes: Each panel–column combination represents a separate regression. The sample consists of all candidate referrals, both made and not made. "Subsidy" refers to subsidized referrals (screened and unscreened). "IHS" denotes the inverse hyperbolic sine transformation. Sales rank for suppliers is computed by ranking baseline (2021) sales across all suppliers and normalizing the rank to range from 0 to 1. Sales rank for clients and price rank are computed analogously. All specifications include the subsidy (uninteracted), the variable defining heterogeneity, and fixed effects for unscreened referral, supplier and client business type, and year. Standard errors clustered by supplier and client. *** p<0.01, ** p<0.05, * p<0.1.

Table A5: Link survival by treatment status of supplier and client

Dep. Var.:	Link	Num Transactions	Value (RMB)
	(1)	(2)	(3)
Client Treated * Supplier Treated	-0.208***	-1.420	-18,457.623**
	(0.050)	(1.083)	(8,256.183)
Client Untreated * Supplier Treated	-0.181***	0.185	-13,424.635*
	(0.044)	(2.042)	(7,065.487)
Client Treated * Supplier Untreated	-0.198***	-1.966**	-21,845.341***
	(0.053)	(0.931)	(6,922.504)
Year 2022	0.103***	2.465***	4,944.023*
	(0.026)	(0.791)	(2,795.800)
Constant	0.706***	5.921***	34,643.898***
	(0.034)	(0.839)	(7,489.296)
Business type, year FE	Yes	Yes	Yes
Observations	1,238	1,238	1,238

Notes: Sample is baseline (2021) links. Standard errors clustered by supplier and client.
*** p<0.01, ** p<0.05, * p<0.1.

Table A6: Heterogeneity in link survival

Dep. Var.:	Link	Num Transactions	Value (RMB)	IHS Num Transactions	IHS Value
	(1)	(2)	(3)	(4)	(5)
Panel A					
Either treated * Supplier satisfactory	0.045	3.773	50,786.415*	0.526*	1.252
	(0.098)	(3.573)	(28,871.926)	(0.314)	(1.174)
Either treated * Client satisfactory	0.315***	-2.717	-15,228.832	0.432	3.065**
	(0.109)	(2.310)	(15,730.362)	(0.348)	(1.180)
Observations	646	646	646	646	646
Panel B					
Either treated * Supplier log sales	0.159***	1.906*	-3,670.175	0.406***	1.567***
	(0.029)	(1.151)	(7,667.440)	(0.089)	(0.310)
Either treated * Client log sales	0.002	-1.940	-5,860.146	-0.098	-0.090
	(0.027)	(1.196)	(3,778.696)	(0.082)	(0.289)
Observations	1,234	1,234	1,234	1,234	1,234
Panel C					
Either treated * Diff sales rank	-0.073	4.694	-4,732.430	-0.204	-1.291
	(0.214)	(6.513)	(28,531.588)	(0.602)	(2.165)
Observations	1,238	1,238	1,238	1,238	1,238
Panel D					
Either treated * Diff price rank	0.235	1.073	10,255.716	0.730	2.541
	(0.200)	(3.397)	(26,794.891)	(0.538)	(2.031)
Observations	1,044	1,044	1,044	1,044	1,044
Panel E					
Either treated * Log distance	-0.027	-0.744	7,894.447	-0.133	-0.182
	(0.043)	(0.793)	(6,360.681)	(0.120)	(0.473)
Observations	1,124	1,124	1,124	1,124	1,124
Panel F					
Either treated * Kin	0.415***	-4.408	-30,317.359	0.664*	3.822***
	(0.112)	(5.434)	(24,597.016)	(0.351)	(1.260)
Observations	1,238	1,238	1,238	1,238	1,238
Panel G					
Either treated * Length of Relationship	0.125*	-0.509	-14,425.165	0.311	1.076
	(0.069)	(1.707)	(16,119.429)	(0.200)	(0.725)
Observations	1,152	1,152	1,152	1,152	1,152

Notes: Each panel in each column is a different regression. Sample is baseline links. "IHS" denotes the inverse hyperbolic sine transformation. "Either treated" is an indicator variable equal to 1 if the supplier, the client, or both are treated. Sales rank for suppliers is computed by ranking baseline (2021) sales across all suppliers and normalizing the rank to range from 0 to 1. Sales rank for clients and price rank are computed analogously. All specifications include the uninteracted "either treated" variable, the uninteracted heterogeneity-defining variable, and fixed effects for supplier and client business type and year. Standard errors clustered by supplier and client.
*** p<0.01, ** p<0.05, * p<0.1.

Satisfaction with partners. In Table A7 we estimate firm-level treatment effect regressions on measures of firms' overall satisfaction with their partners. We find, for both suppliers and clients, significant improvements in satisfaction with partners' price, the business transaction (product or demand), and partners' reliability, but not with their flexibility or trade credit.

Table A7: Satisfaction with partners: different dimensions

Panel A: Suppliers' satisfaction with clients						
Dep. Var.:	Satisfaction with					
	Price	Demand	Reliability	Flexibility	Credit	Overall
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	0.212**	0.338***	0.179**	0.091	0.134	0.268***
	(0.094)	(0.115)	(0.082)	(0.074)	(0.096)	(0.078)
Baseline Control	Yes	Yes	Yes	Yes	Yes	Yes
Strata, year FE	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	5.300	5.074	6.148	6.173	6.429	6.170
Observations	762	763	762	762	740	761
Panel B: Clients' satisfaction with suppliers						
Dep. Var.:	Satisfaction with					
	Price	Quality	Reliability	Flexibility	Credit	Overall
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	0.307***	0.462***	0.211**	-0.016	0.066	0.292***
	(0.095)	(0.094)	(0.085)	(0.086)	(0.115)	(0.085)
Baseline Control	Yes	Yes	Yes	Yes	Yes	Yes
Strata, year FE	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	5.504	6.147	6.547	6.533	6.585	6.453
Observations	496	485	496	496	476	496

Notes: Sample in Panel A is suppliers, in Panel B clients, in both follow-up waves.
Outcomes measure satisfaction with partners in different dimensions on a 10-point scale.
Control mean is at baseline. Standard errors clustered by supplier and client.
*** p<0.01, ** p<0.05, * p<0.1.

Robustness of main firm-level results. In Tables A8 - A10 we present different robustness checks of our main firm-level results.

In Table A8 we show that the significance of the treatment effects for suppliers and low-degree

Table A8: Robustness of firm-level results to clustering

Dep. Var.:	Log Sales	Profit (10,000 RMB)	Log Operating Cost	Log Employ-ment	Log Total Hrs Worked / Day	Num of Partners	Satisfaction with Partners
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Suppliers							
Treatment	0.209**	2.049*	0.059	0.079*	0.159***	2.003*	0.257***
	(0.101)	(1.184)	(0.126)	(0.044)	(0.044)	(1.159)	(0.079)
Baseline Control	Yes	Yes	Yes	Yes		Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	2.118	7.393	1.363	0.755	2.788	6.135	5.047
Observations	627	629	612	629	629	629	627
Panel B: Clients, Baseline partners <=1							
Treatment	0.321**	9.297***	0.290*	0.064	0.071	2.500***	0.376***
	(0.148)	(3.377)	(0.165)	(0.084)	(0.099)	(0.702)	(0.086)
Baseline Control	Yes	Yes	Yes	Yes		Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	2.916	27.92	2.230	1.190	3.124	4.817	4.618
Observations	304	304	288	304	304	304	284

Notes: Sample in Panel A is suppliers, in Panel B is clients, in both follow-up waves. In columns 2 and 6 outcome is winsorized at the 99th percentile in the full sample (all firms in all waves). In column 7, outcome is standardized by its standard deviation in the baseline control sample, separately for suppliers and clients. In columns 6 and 7, partner refers to clients in Panel A and suppliers in Panel B. Baseline value of the outcome is not included in column 5 because we did not collect it. Control mean is at baseline. Standard errors clustered based on the extent to which firms are competitors. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

clients is robust to clustering firms based on the extent to which they are competitors. We define clusters using a measure of "competition distance," which (inversely) measures the strength of competition between two firms. This measure depends on the product type, the price, and the location of firms, and we define it as follows. First, we standardize the latitude and longitude of the village (or street) where the firm is located, as well as the log price of the firm's main product. Then we compute the Euclidean distance between the three-dimensional vectors of these variables for all pairs of firms. To this initial measure of distance, we add a prohibitively high term for firm pairs that have a different product type for their main product. Thus, the resulting distance (or "dissimilarity matrix") would be zero if and only if the two firms sell the same main product, at

Table A9: Robustness of treatment effect on profit to handling outliers

	Profit winsorized at 99 percentile			IHS Profit	Raw profit
	In full sample	By year	By year, supplier, client		
	(1)	(2)	(3)	(4)	(5)
Panel A: Suppliers					
Treatment	1.501*	1.456*	1.222*	0.238***	3.330
	(0.824)	(0.802)	(0.624)	(0.066)	(2.204)
Baseline Control	Yes	Yes	Yes	Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes
Control Mean	6.804	6.804	6.632	2.073	6.804
Observations	763	763	763	763	763
Panel B: Clients, Baseline partners <=1					
Treatment	9.297***	9.103***	11.349***	0.366***	11.488***
	(3.377)	(3.290)	(4.294)	(0.132)	(4.318)
Baseline Control	Yes	Yes	Yes	Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes
Control Mean	25.91	26.82	29.92	2.789	37.34
Observations	304	304	304	304	304

Notes: Sample in panel A is supplier firms, in Panel B is low-degree client firms. In column 1, profit is winsorized at the 99th percentile of the full sample (all firms in all five waves); in column 2, it is winsorized at the 99th percentile by year; and in column 3, at the 99th percentile by year and by supplier/client. IHS denotes the inverse hyperbolic sine. Standard errors clustered by firm. *** p<0.01, ** p<0.05, * p<0.1.

the same price, and are in the same location. Then, using this measure, we implement hierarchical clustering with average linkage to create clusters (e.g., James, Witten, Hastie and Tibshirani 2021). We obtain 60 clusters for suppliers and 36 clusters for low-degree clients. As Table A8 shows, the results are very similar to the version where we cluster by firm.

In Table A9 we show that the significance of the treatment effect on profit is robust to different ways of handling outliers: winsorizing at the 99th percentile in the full sample, by year, and by year separately for suppliers and clients; or using the inverse hyperbolic sine. In column 5 we use raw profit as the outcome; the p-value in the supplier sample is 0.13. In Table A10 we show the robustness of the treatment effect for suppliers to Lee bounds, which is one approach to address

Table A10: Lee bounds for treatment effect on supplier firms

Dep. Var.:	Log Sales	Profit (10,000 RMB)	IHS Profit	Log Operating Cost	Log Employ-ment	Log Total Hrs Worked	Num of Partners	Satisfaction Clients
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Supplier upper bound								
Treatment	0.278***	1.606*	0.268***	0.148	0.079**	0.153***	1.987**	0.255***
	(0.073)	(0.832)	(0.065)	(0.091)	(0.040)	(0.046)	(0.770)	(0.062)
Baseline Control	Yes	Yes	Yes	Yes	Yes		Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	1.890	6.804	2.073	1.296	0.711	2.740	5.595	5.059
Observations	752	755	755	645	755	755	755	754
Panel B: Supplier lower bound								
Treatment	0.217***	0.614	0.210***	0.008	0.051	0.107**	0.993*	0.189***
	(0.073)	(0.615)	(0.065)	(0.096)	(0.039)	(0.046)	(0.527)	(0.063)
Baseline Control	Yes	Yes	Yes	Yes	Yes		Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	1.890	6.804	2.073	1.296	0.711	2.740	5.595	5.059
Observations	752	754	754	643	754	754	754	752

Notes: This table reports Lee (2009) bounds for the main firm effects among suppliers. Given the treatment–control difference in attrition of 3.9 percentage points and the 97.5% response rate in the treatment group, we trim a share of 0.039 / 0.975 = 0.04 of the treated sample, corresponding to 8 firms. Panel A reports results under the upper-bound assumption, where we remove firms with the lowest outcomes; Panel B reports results under the lower-bound assumption, where we remove firms with the highest outcomes. Otherwise, specifications are identical to those in Table 6. Standard errors clustered by firm. *** p<0.01, ** p<0.05, * p<0.1.

the slight imbalance in attrition in 2023. The lower bounds are significant for revenue, IHS profit, hours worked, the number of clients, and satisfaction with clients.

Sample selection by degree. Tables A11 - A15 present some results to address the concern that we selected the low-degree clients subsample after the intervention. The motivation for selecting this subsample was the prediction of our framework that treatment effects should be larger for low-degree firms. Table A11 shows that among suppliers, significant treatment effects are largely driven by low-degree firms. This provides external validity to choosing the low-degree subsample for clients. Table A12 reports treatment effect estimates of firm performance using as outcomes firms' self-reported evaluation of their performance during 2020-24 on a 5-point scale. The main point

Table A11: Supplier firms' main outcomes by baseline degree

Dep. Var.:	Log Sales	Profit (10,000 RMB)	Log Operating Cost	Log Employ-ment	Log Total Hrs Worked / Day	Num Partners	Satisfaction with Partners
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Baseline partners <=1							
Treatment	0.270***	1.201***	-0.000	0.090*	0.110**	1.019**	0.160**
	(0.091)	(0.441)	(0.130)	(0.046)	(0.055)	(0.442)	(0.081)
Baseline Control	Yes	Yes	Yes	Yes		Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	1.576	5.070	0.938	0.598	2.631	4.966	4.907
Observations	519	521	421	521	521	521	520
Panel B: Baseline partners > 1							
Treatment	0.116	1.745	0.140	0.014	0.143	4.493*	0.308***
	(0.138)	(2.449)	(0.163)	(0.076)	(0.100)	(2.318)	(0.109)
Baseline Control	Yes	Yes	Yes	Yes		Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	2.647	10.99	2.050	0.983	2.989	7.117	5.425
Observations	242	242	229	242	242	242	241

Notes: In Panel A, the sample consists of suppliers with at most one regular client at baseline; in Panel B, of suppliers with more than one regular clients at baseline. Data are from both follow-up waves. The outcome in column 2 is winsorized, in column 6 it is winsorized, and in column 7 it is standardized. In columns 6 and 7, partner refers to clients. The baseline value of the outcome is not included in column 5 because it was not collected. Control mean is at baseline. Standard errors clustered by firm. *** p<0.01, ** P<0.05, * p<0.1.

here is that we find significant gains for revenue and profit in the full sample of clients, strengthening the interpretation that client firms gained from the intervention. Table A13 estimates the search regression of Table 4 separately for low-degree and high-degree firms. The treatment increases beliefs and search in both samples, but the effect on the number of new partners is only significant in the low-degree sample. Table A14 estimates Table 5 separately for low- and high-degree firms. Relative to the mean baseline degree in the control group, the growth in degree in response to the treatment is much higher among low-degree firms. Table A15 estimates the upgrading regression of Table 8 for client firms separately by degree. The effects are driven by low-degree firms.

Table A12: Firms' self-reported growth on 5-point scale

Dep. Var.:	Perceived Num Products (1-5)	Perceived Quantity (1-5)	Perceived Avg Quality (1-5)	Perceived Avg Price (1-5)	Perceived Sales (1-5)	Perceived Profits (1-5)
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Suppliers						
Treatment	0.194*	0.202*	0.280***	0.179*	0.240**	0.348***
	(0.108)	(0.108)	(0.086)	(0.092)	(0.110)	(0.120)
Strata FE	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	2.261	2.254	2.944	2.761	2.254	2.190
Observations	288	288	288	288	288	288
Panel B: Clients						
Treatment	0.242	0.203	0.192**	0.154*	0.323*	0.286*
	(0.165)	(0.162)	(0.085)	(0.091)	(0.166)	(0.162)
Strata FE	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	2.375	2.362	3.038	2.813	2.350	2.337
Observations	160	160	160	160	160	160

Notes: Sample in Panel A is suppliers, in Panel B is clients. All outcomes were collected in the 2024 phone survey and measure perceived business performance since 2020 on a 5-point scale. Standard errors clustered by firm. *** p<0.01, ** p<0.05, * p<0.1.

Table A13: Search by baseline degree

Dep. Var.:	Belief: Profit gain from new partner (%)	Search (hrs/month)	Placebo belief: partner on other side (%)	Placebo search: on other side (hrs/month)	Num non-referred new partners
	(1)	(2)	(3)	(4)	(5)
Panel A: Baseline partners <=1					
Treatment	8.362***	7.126***	0.904	1.171	1.508***
	(1.231)	(1.289)	(0.912)	(1.062)	(0.402)
Baseline Control, Year FE					Yes
Indicator for client, Strata FE	Yes	Yes	Yes	Yes	Yes
Control Mean	3.627	3.047	2.677	3.234	4.806
Observations	404	404	404	404	825
Panel B: Baseline partners >1					
Treatment	7.302***	4.666***	-0.775	0.596	0.862
	(1.499)	(1.623)	(1.061)	(1.666)	(2.202)
Baseline Control, Year FE					Yes
Indicator for client, Strata FE	Yes	Yes	Yes	Yes	Yes
Control Mean	4.500	5.327	4.748	5.412	8.698
Observations	216	216	216	216	458

Notes: The sample in Panel A consists of firms with at most one regular partner, and in Panel B of firms with more than one regular partner (where a partner is a supplier for client firms and a client for supplier firms). Data in columns 1-4 are from 2023, as the outcome was collected only in that year; data in column 5 are from 2022 and 2023. The outcomes in columns 3 and 4 are sellers (for supplier firms) and buyers (for client firms). The outcome in column 5 is the number of partners in the firm data minus the number of referred partners in the network data. The control mean is reported at follow-up because baseline data are unavailable. Standard errors clustered by firm. *** p<0.01, ** p<0.05, * p<0.1.

Table A14: Link rewiring and expansion by baseline degree

Dep. Var.:	Degree	Baseline links active	Non-baseline links	Reactivated prior links
	(1)	(2)	(3)	(4)
Panel A: Baseline partners <= 1				
Treated	0.474***	-0.379***	0.852***	-0.100
	(0.166)	(0.066)	(0.147)	(0.065)
Demean Exposure * Treated	-0.392	0.056	-0.448*	-0.124*
	(0.283)	(0.103)	(0.263)	(0.065)
Demean Exposure * Untreated	0.047	-0.151*	0.198	0.154
	(0.194)	(0.088)	(0.146)	(0.107)
Controls, strata FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	
Control Mean	0.347	0.347	0.167	
Observations	825	825	825	415
Panel B: Baseline partners >1				
Treated	1.147***	-0.353**	1.500***	0.005
	(0.256)	(0.145)	(0.180)	(0.160)
Demean Exposure * Treated	0.865	0.270	0.595	0.325
	(0.798)	(0.386)	(0.609)	(0.502)
Demean Exposure * Untreated	0.037	-0.633*	0.670*	0.796**
	(0.555)	(0.353)	(0.386)	(0.392)
Controls, strata FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	
Control Mean	4.442	4.442	0.658	
Observations	458	458	458	226

Notes: The sample in Panel A consists of firms with at most one regular partner, and in Panel B of firms with more than one regular partner (where a partner is a supplier for client firms and a client for supplier firms). Data in columns 1–3 are from both follow-up waves; column 4 uses a cross-section. Exposure is demeaned share of baseline partners treated. Controls include an indicator for client firm, an indicator for the firm having regular partners at baseline, an interaction of the latter with the treatment, and the number of baseline partners. In column 4, outcome is number of partners in either follow-up (2022/23) that were not partners at baseline (2021) but were partners before (2018 or 19). Control mean is at baseline. Standard errors clustered by firms. *** p<0.01, ** p<0.05, * p<0.1.

Table A15: Clients: upgrading by baseline degree

Dep. Var.:	Quality score main product		Quality	If 2nd	Share 2nd	Avg Log
	Craftsmanship	Durability	check hrs/wk	product	product	Price
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Clients, Baseline partners <=1						
Treatment	0.120	-0.276	1.718	0.129**	6.684***	0.382**
	(0.178)	(0.177)	(1.052)	(0.053)	(2.216)	(0.169)
Baseline Control	Yes	Yes	Yes	Yes	Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	-0.410	-0.435	4.922	0.636	18.51	1.028
Observations	138	138	304	304	303	282
Panel B: Clients, Baseline partners >1						
Treatment	0.260	0.145	3.703**	0.044	3.908*	0.015
	(0.176)	(0.192)	(1.798)	(0.048)	(2.319)	(0.151)
Baseline Control	Yes	Yes	Yes	Yes	Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	-0.264	-0.306	7.458	0.817	20.43	1.837
Observations	153	153	216	216	215	213

Notes: The sample in Panel A consists of clients with at most one regular supplier, and in Panel B of clients with more than one regular supplier. Quality score in columns 1-2 is evaluated by experts, residualized by product type and standardized. In column 5, the outcome is the revenue share of the firm's second product, zero for firms that do not have a second product. In column 6, the outcome is the sales-weighted average of the log prices of the firm's main products for which we collect data on, up to three. Control mean is at baseline. Standard errors clustered by firm. *** p<0.01, ** p<0.05, * p<0.1.

Indirect effects. Tables A16 - A18 provide additional evidence related to indirect effects. In Table A16 we regress measures of average partner quality on treatment and exposure. Partners are defined using the network at the time of the outcome (i.e., the follow-up network). Our focus in this regression is on how exposure affects partner characteristics. Column 1 shows that exposure increases partners' revenue, while columns 2 and 3 show that it increases both measures of partner quality. These results are consistent with the hypothesis that the positive effects of exposure on firm performance documented in Table 9 are driven by improvements in the quality of firms' trading partners. The negative effect of the treatment in column 1 is probably driven by the fact that the referrals created many connections with small firms, changing the composition of partners.

Table A16: Indirect effects on partner quality

Dep. Var.:	Partners' Avg Log Sales	Partners' Avg Quality: Craftsmanship	Partners' Avg Quality: Durability
	(1)	(2)	(3)
Exposure	0.229*	2.613***	2.798**
	(0.129)	(0.583)	(1.147)
Treatment	-0.327***	0.447	-0.233
	(0.078)	(0.300)	(0.662)
Baseline Control	Yes	Yes	Yes
Indicator for Client	Yes	Yes	Yes
Strata, Year FE	Yes	Yes	Yes
Control Mean	3.685	28.33	36.59
Observations	620	503	503

Notes: Sample is all firms in follow-up waves. Outcomes are the average characteristics of partners in that wave: log sales in column 1, craftsmanship quality in column 2, and durability quality in column 3. Standard errors clustered by firm. *** p<0.01, ** P<0.05, * p<0.1.

Table A17 estimates our main firm-level treatment effect regression, for suppliers and low-degree clients, but including a control for exposure. The results show that the treatment effects are qualitatively the same when we control for indirect effects. Finally, Table A18 presents estimates of indirect effects in the final good market, coming from close competitors among client firms. Close competitors are defined to be other firms who have the same main product type, are in the same location, and have a similar price (price difference smaller than the median). This is a refinement of the notion of similar firms we used when we constructed the screened referrals. We find no significant business stealing effects, but we note that confidence intervals are wide: this regression is underpowered.

Table A17: Robustness of main firm effects to exposure

Dep. Var.:	Log Sales	Profit (10,000 RMB)	Log Operating Cost	Log Employ-ment	Log Total Hrs Worked / Day	Num of Partners	Satisfaction with Partners
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Suppliers							
Exposure	0.025	0.022	-0.050	0.061	0.037	0.953	0.015
	(0.143)	(1.111)	(0.184)	(0.071)	(0.071)	(1.299)	(0.103)
Treatment	0.249***	1.537*	0.064	0.070*	0.133***	1.913**	0.223***
	(0.074)	(0.811)	(0.097)	(0.039)	(0.046)	(0.749)	(0.064)
Baseline Control	Yes	Yes	Yes	Yes		Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	1.890	6.804	1.296	0.711	2.740	5.595	5.059
Observations	761	763	652	763	763	763	761
Panel B: Clients, Baseline partners <=1							
Exposure	0.236	14.901**	0.175	-0.017	-0.071	-1.297	-0.259*
	(0.262)	(7.440)	(0.326)	(0.140)	(0.156)	(1.398)	(0.150)
Treatment	0.315**	8.809***	0.277*	0.063	0.072	2.503***	0.389***
	(0.149)	(3.185)	(0.165)	(0.084)	(0.099)	(0.705)	(0.086)
Baseline Control	Yes	Yes	Yes	Yes		Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	2.739	25.91	2.128	1.143	3.079	4.506	4.609
Observations	304	304	288	304	304	304	284

Notes: Sample in Panel A is suppliers, in Panel B is low-degree clients, in both follow-up waves. Exposure, defined as demeaned share of baseline partners treated, is included in all regressions. An indicator for the firm having no regular partners at baseline is also included. In columns 2 and 6 outcome is winsorized at the 99th percentile in the full sample (all firms in all waves). In column 7, outcome is standardized by its standard deviation in the baseline control sample, separately for suppliers and clients. In columns 6 and 7, partner refers to clients in Panel A and suppliers in Panel B. Baseline value of the outcome is not included in column 5 because we did not collect it. Standard errors clustered by firm. *** p<0.01, ** p<0.05, * p<0.1.

Table A18: Main firm effects and competition in final goods market

Dep. Var.:	Log Sales	Profit (10,000 RMB)	Log Operating Cost	Log Employment	Log Total Hrs Worked / Day	Num of Partners	Satisfaction Suppliers
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Share of Treated among Close Competitors	-0.000	-7.080	0.002	-0.059	0.087	2.306	-0.011
	(0.278)	(5.793)	(0.330)	(0.157)	(0.180)	(1.514)	(0.200)
Treatment	0.037	-1.316	-0.042	0.028	0.094	2.167***	0.210***
	(0.117)	(3.935)	(0.122)	(0.072)	(0.084)	(0.681)	(0.075)
Baseline Control	Yes	Yes	Yes	Yes		Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	3.340	37.38	2.740	1.349	3.290	6.102	5.308
Observations	452	452	445	452	451	452	443

Notes: Sample is client firms in both follow-up waves. Share of Treated among Close Competitors is the share of close competitors which are treated. Close competitors are defined based on having the same product type for the main product and similarity in price and location, as described in Appendix A. Standard errors clustered by firm. *** p<0.01, ** p<0.05, * p<0.1.

Concerns with measurement, identification, and interpretation. In Tables A19 - A21 we address some additional concerns. Table A19 looks at experimenter demand. We estimate our firm-level specification, for suppliers and low-degree clients, using as outcomes measures that are plausibly robust to experimenter demand. We find significant effects on the owner's (self-reported) phone use, and on some measures of our enumerators' assessment of the business. In Table A20 we look at dynamics by separately estimating impacts for 2022 and 2023. There is less power, but the estimated effects are broadly similar across the two years and differ significantly in only one of the 14 specifications. In Table A21 we look at our firm-level treatment effects interacted with the cross-randomized e-commerce treatment. The interaction is always insignificant, but the referrals effect also becomes less precisely estimated, particularly among low-degree clients. Thus, we cannot rule out that improving market access—or having sufficient market access—in the final good market is necessary for improvements in access in the input market to matter.

Table A19: Experimenter demand

Dep. Var.:	Owner's Phone Use			Enumerator Observation in the Field				
	Log Monthly Bill (RMB)	Log Num Calls & Wechat / Month	Share Related to Business Talk (%)	Num Visitors	Num Calls	Busy (1-5)	Num Employees	Assessment of operations (1-5)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Suppliers								
Treatment	0.138**	0.176**	3.513	0.031	0.094**	0.302**	0.351	0.310***
	(0.066)	(0.082)	(2.208)	(0.041)	(0.041)	(0.120)	(0.225)	(0.103)
Strata FE	Yes	Yes	YES	Yes	Yes	Yes	Yes	Yes
Control Mean	4.10	2.35	22.73	0.06	0.10	2.36	0.60	2.38
Observations	372	370	373	363	363	369	363	371
Panel B: Clients Baseline partners <=1								
Treatment	0.363***	0.216	-0.346	0.072	0.028	0.395*	-0.048	0.375**
	(0.099)	(0.159)	(4.410)	(0.064)	(0.071)	(0.201)	(0.188)	(0.174)
Strata FE	Yes	Yes	YES	Yes	Yes	Yes	Yes	Yes
Control Mean	4.19	2.49	29.93	0.06	0.09	2.56	0.73	2.67
Observations	142	141	143	141	141	142	141	143

Notes: Sample in Panel A is suppliers, in Panel B is low-degree clients. All outcomes are collected in the 2023 survey. Columns 4-8 are enumerators' evaluation of the firm. Standard errors clustered by firm. *** p<0.01, ** p<0.05, * p<0.1.

Table A20: Performance over time

Dep. Var.:	Log Sales	Profit (10,000 RMB)	Log Operating Cost	Log Employ-ment	Log Total Hrs Worked / Day	Num of Partners	Satisfaction with Partners
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Suppliers							
Treatment * 2022	0.205**	1.254*	0.104	0.071	0.177***	1.948**	0.256***
	(0.088)	(0.741)	(0.109)	(0.049)	(0.062)	(0.849)	(0.093)
Treatment* 2023	0.283***	1.751	-0.004	0.064	0.079	1.804*	0.181**
	(0.093)	(1.327)	(0.135)	(0.053)	(0.054)	(0.984)	(0.084)
Baseline Control	Yes	Yes	Yes	Yes	No	Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	1.890	6.804	1.296	0.711	2.740	5.595	5.059
P value for 2022 = 2023	0.457	0.718	0.467	0.908	0.145	0.888	0.544
Observations	761	763	652	763	763	763	761
Panel B: Clients, Baseline partners <=1							
Treatment * 2022	0.336*	10.073**	0.213	0.090	0.038	3.697***	0.313**
	(0.192)	(4.553)	(0.202)	(0.109)	(0.130)	(1.141)	(0.125)
Treatment* 2023	0.305*	8.496*	0.368	0.037	0.105	1.265	0.440***
	(0.180)	(4.687)	(0.237)	(0.104)	(0.116)	(0.862)	(0.114)
Baseline Control	Yes	Yes	Yes	Yes	No	Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	2.739	25.91	2.128	1.143	3.079	4.506	4.609
P value for 2022 = 2023	0.890	0.802	0.595	0.693	0.652	0.0971	0.445
Observations	304	304	288	304	304	304	284

Notes: Sample in Panel A is suppliers, in Panel B is low-degree clients. Treatment included separately by wave in all regressions. In columns 2 and 6 outcome is winsorized at the 99th percentile in the full sample (all firms in all waves). In column 7, outcome is standardized by its standard deviation in the baseline control sample, separately for suppliers and clients. In columns 6 and 7, partner refers to clients in Panel A and suppliers in Panel B. Baseline value of the outcome is not included in column 5 because we did not collect it. Standard errors clustered by firm. *** p<0.01, ** p<0.05, * p<0.1.

Table A21: Interactions with the E-commerce treatment

Dep. Var.:	Log Sales	Profit (10,000 RMB)	Log Operating Cost	Log Employ-ment	Log Total Hrs Worked / Day	Num of Partners	Satisfaction with Partners
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Supplier							
Referral * E-commerce	-0.005	1.698	0.139	-0.133	-0.081	-0.140	-0.048
	(0.155)	(2.382)	(0.204)	(0.082)	(0.097)	(1.444)	(0.132)
Referral	0.249**	0.576	-0.021	0.138**	0.171***	1.947**	0.243**
	(0.106)	(0.951)	(0.146)	(0.059)	(0.065)	(0.815)	(0.097)
E-commerce	-0.087	0.844	-0.075	0.025	0.023	0.209	0.059
	(0.107)	(0.918)	(0.131)	(0.055)	(0.066)	(0.621)	(0.092)
Baseline Control	Yes	Yes	Yes	Yes		Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	1.890	6.804	1.296	0.711	2.740	5.595	5.059
Observations	761	763	652	763	763	763	761
Panel B: Clients, Baseline partners <=1							
Referral * E-commerce	0.441	4.303	0.407	0.173	0.207	-0.979	0.261
	(0.284)	(6.468)	(0.324)	(0.167)	(0.192)	(1.414)	(0.188)
Referral	0.083	6.972	0.056	-0.026	-0.046	2.958***	0.248**
	(0.216)	(4.477)	(0.255)	(0.119)	(0.154)	(1.057)	(0.115)
E-commerce	-0.019	0.355	0.067	-0.042	0.083	1.093	-0.177
	(0.224)	(3.808)	(0.252)	(0.128)	(0.139)	(0.872)	(0.139)
Baseline Control	Yes	Yes	Yes	Yes		Yes	Yes
Strata, Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	2.739	25.91	2.128	1.143	3.079	4.506	4.609
Observations	304	304	288	304	304	304	284

Notes: Sample in panel A is supplier firms, in Panel B is low-degree client firms. All regressions include the e-commerce treatment and the interaction between the referrals and the e-commerce treatments. In columns 2 and 6 outcome is winsorized at the 99th percentile in the full sample (all firms in all waves). In column 7, outcome is standardized by its standard deviation in the baseline control sample, separately for suppliers and clients. In columns 6 and 7, partner refers to clients in Panel A and suppliers in Panel B. Baseline value of the outcome is not included in column 5 because we did not collect it. Standard errors clustered by firm. *** p<0.01, ** p<0.05, * p<0.1.

Marketing offer. Table A22 reports summary statistics and balance for the marketing offer experiment. There is no evidence of imbalance.

Table A22: Summary statistics of second-phase experiment

	Clients		
	No Marketing Offer	Marketing Offer	Difference
Panel A: Firm Characteristics			
Firm Age	25.550	24.119	-1.431
	(12.077)	(11.645)	(2.264)
Number of Employees	4.425	4.988	0.563
	(5.560)	(5.628)	(1.077)
Profit (10,000 RMB)	29.313	22.910	-6.403
	(64.566)	(42.580)	(9.731)
Sales (10,000 RMB)	63.967	62.615	-1.353
	(135.737)	(120.040)	(24.066)
Panel B: Managerial Characteristics			
Gender (1=Female, 0=Male)	0.100	0.131	0.031
	(0.304)	(0.339)	(0.063)
Age	52.475	51.298	-1.177
	(10.188)	(9.076)	(1.815)
Education- Middle School	0.700	0.655	-0.045
	(0.464)	(0.478)	(0.091)
Panel C: Number of Partners (buyers for supplier, sellers for client)			
In Firm Data	5.975	5.607	-0.368
	(6.670)	(5.678)	(1.155)
In Network Data, Supplier-client Links	1.750	1.881	0.131
	(2.609)	(2.826)	(0.530)
P-val of Joint Significance of Vars in Panels A-C			0.782
Observations	40	84	124

Notes: Sample is client firms in T0 in 2023 follow-up. Column 1 includes firms that received EI on both their candidate referrals. Column 2 includes firms that received MO on one of their candidate referrals. *** p<0.01, ** p<0.05, * p<0.1.

B Model and hypotheses

We derive our hypotheses within the industry equilibrium model of firm-to-firm access developed in Szeidl (2025). For completeness we include a summary of the model and its main implications—sometimes borrowed almost verbatim from Szeidl (2025)—but we emphasize that this model is not new to the present paper. The novelty here is purely in deriving the hypotheses of Section 3. For ease of presentation we depart from the order in which we described the model's implications in the main text. We first present the setup of the model without search, then introduce referrals, then search, and finally upgrading. We conclude by deriving the hypotheses of Section 3.

B.1 Setup

Consider an industry producing differentiated products, which consists of a unit mass of clients, indexed by i, and a unit mass of suppliers, indexed by j. Each client i is endowed with λ potential suppliers chosen at random, where λ is a continuous variable.[1] Each supplier-client relationship generates a new business idea, such as a blueprint for a new differentiated product variety. Firms have capacity constraint κ: each client firm can work with at most κ suppliers, and each supplier firm can work with at most κ client firms. The potential surplus (or value) created from a match between i and j is h_{ij}, which is a random draw from distribution G, assumed uniform on $[0, g_H]$. The supplier and the client divide the surplus equally.

Once the potential suppliers are drawn and the potential surplus values are realized, clients and suppliers decide which potential partnerships to keep. A potential partnership is kept if and only if both the supplier and the client prefer to keep it. An equilibrium is a symmetric threshold equilibrium if there is a common threshold value such that each firm keeps only links with value above that threshold. A symmetric threshold equilibrium is maximal if there is no other equilibrium in which a strict superset of links form. We use this notion of equilibrium. Industry surplus is the total value of all active links.

Proposition 1 (Baseline equilibrium). *(i) For $\lambda < \kappa$, firms take all of their potential partners. Each firm has degree λ.*

(ii) For $\lambda > \kappa$, firms take the top κ/λ share of their potential partners. Each firm has degree κ.

This result follows from Proposition 1 of Szeidl (2025) and characterizes the "baseline" equilibrium before the referrals intervention. For low values of λ/κ all potential links are accepted, while

[1] Thus, λ is analogous to a density function. A single client j has potential links of measure zero, but a small positive mass of clients $\mu > 0$ have a mass of potential links $\lambda \cdot \mu$.

for high values of λ/κ only links of sufficiently high value are accepted. Intuitively, once the number of potential links exceeds the firm's capacity constraint, the firm needs to select among potential partners.

B.2 Referrals experiment

Consider the referrals experiment within the model. Start with the baseline equilibrium introduced above. Assume that a random half of client firms and a random half of supplier firms are treated. Each treated client firm gets δ new potential suppliers, all of whom are treated suppliers. Thus, by symmetry, each treated supplier firm gets δ new potential clients, all of whom are treated clients. We focus on the emerging *short-run* equilibrium: the maximal symmetric threshold equilibrium that can be reached by taking on some of the new links and eliminating some of the links firms formed at baseline. Given the focus on the short run, firms are not allowed to go back to prior potential links that they had rejected when arriving at the baseline equilibrium. Symmetry now requires that firms with the same treatment status use the same cutoff rule for link decisions.

In this experiment, given randomness and continuity, each firm will have exactly half of its baseline partners treated. This symmetry helps characterize the equilibrium. But we are also interested in the effect of exposure, i.e., of variation in the share of baseline partners of a firm that are treated. To explore exposure effects, suppose that a zero mass of firms have exposure S different from 0.5. These firms make optimal link decisions given the equilibrium behavior of all other firms. Since these firms have mass zero, they do not distort any equilibrium quantity.

The next results characterize degree and surplus in a first-order approximation for small δ. Let $\gamma = \delta/\lambda$, so that we can think of γ as the proportional increase in access for treated firms. The subscript 0 denotes outcomes at baseline. We characterize outcomes as a function of T and S, where T is indicator for the firm being treated, and S is the share of the firm's prior partners that are treated.

Proposition 2 (Experiment)**.** *Degree and surplus in the short-run equilibrium of the experiment depart from their baseline values as follows.*

(i) For $\lambda < \kappa$, log degree is

$$\log d(T,S) \approx \log d_0 + T \cdot \gamma, \tag{A1}$$

log firm surplus is

$$\log V(T,S) \approx \log V_0 + T \cdot \gamma, \tag{A2}$$

and log total surplus is

$$\log W \approx \log W_0 + \frac{\gamma}{2}. \tag{A3}$$

(ii) For $\lambda > \kappa$, log degree is

$$\log d(T,S) \approx \log d_0 - S(1-T) \cdot \gamma, \tag{A4}$$

log firm surplus is

$$\log V(T,S) \approx \log V_0 + T \cdot \frac{\kappa}{2\lambda - \kappa}\gamma - S(1-T) \cdot \frac{2\lambda - 2\kappa}{2\lambda - \kappa}\gamma \tag{A5}$$

and log total surplus is

$$\log W \approx \log W_0 + \frac{1}{2}\frac{2\kappa - \lambda}{2\lambda - \kappa}\gamma. \tag{A6}$$

This is Proposition 2 of Szeidl (2025). In Case (i), with low access relative to capacity $\lambda < \kappa$, which corresponds to the example of expanding in Figure 3, there are no indirect effects, and the size of the treatment effect on both degree and firm surplus is governed by the proportional increase in access γ. The effect on industry surplus is always positive. In Case (ii), with high access relative to capacity $\lambda > \kappa$, there are indirect effects on untreated firms, captured by the interaction between S and $1 - T$. The experiment has no effect on the degree of treated firms, because these firms purely reallocate, but reduces the degree of untreated firms, in proportion to their baseline exposure S, because some of their pre-existing links to treated partners are crowded out. Turning to the surplus, the experiment increases the surplus of treated firms, but by less than in Case (i). At the same time, it has negative indirect effects on untreated firms, in proportion to their exposure. This follows because untreated firms lose some of their connections with treated firms due to these treated firms finding better partners. Combining these effects, (A6) shows that the impact of the experiment on industry surplus can be positive or negative.

B.3 Search Effort

To explore search, consider the baseline model without the experiment. Assume that firms already found their initial potential partners and initial partnerships have been formed, and consider the incentives for additional search. Only clients are allowed to search, not suppliers. Each client firm, at a cost, can find δ new potential partners. Write the cost of search as $c \cdot g_H/2$, so that it is measured in proportion to the average quality of a potential link. Purely for convenience, make this search opportunity analogous to the experiment: only half of the clients (treated clients) can search, and their search is among half of all suppliers (treated suppliers). This formulation could represent a search platform created by the government, onto which half of clients and half of suppliers are boarded. Finally, introduce to this framework the friction that firms undervalue potential links. Each firm believes that the distribution of new matches is uniform on the interval $[0, \beta g_H]$ where $\beta \leq 1$. Undervaluation does not affect firms' views of their existing relationships, only those they obtain by search.

Proposition 3 (Search)**.** *Assume that $\beta > 1 - \kappa/\lambda$.*

(i) When $\lambda < \kappa$ we have that

- *the firm chooses to search if*

$$\frac{c}{\delta} < \beta\frac{1}{2},$$

- *it would be privately optimal to search if*

$$\frac{c}{\delta} < \frac{1}{2},$$

- *it would be socially optimal to search if*

$$\frac{c}{\delta} < 1.$$

(ii) When $\lambda > \kappa$, we have that

- *the firm chooses to search if*

$$\frac{c}{\delta} < \beta\left(1 - \frac{\lambda}{\kappa}(1-\beta)\right)^2 \frac{\kappa^2}{2\lambda^2},$$

- *it would be privately optimal to search if*

$$\frac{c}{\delta} < \frac{\kappa^2}{2\lambda^2},$$

- *it would be socially optimal to search if*

$$\frac{c}{\delta} < \frac{\kappa(2\kappa - \lambda)}{\lambda^2}.$$

This is Proposition 3 in Szeidl (2025). Note that all the bounds are on c/δ, which measures the cost-to-benefit ratio of search. In Case (i), comparing the first two bounds shows that when firms undervalue partners by a sufficient margin, they avoid search even when it is privately optimal. Comparing the second and the third bounds shows that even with correct beliefs, there is a range of costs where search is privately suboptimal but socially optimal. In Case (ii), comparing the first two bounds shows that undervaluation again reduces search incentives. Comparing the second and the third bounds shows that the comparison between private and social search incentives depends on λ. For $\lambda \in [\kappa, (3/2)\kappa]$, private incentives are weaker than social incentives, but for $\lambda > (3/2)\kappa$ the converse is true. This follows because there are two forces. The search externality reduces search incentives relative to the social optimum; but search generates crowding out and negative indirect effects. When λ is high, the second effect dominates, because then the firm already has quite good partnerships, limiting the first effect.

B.4 Upgrading

Upgrading can be added to the model through a learning-by-doing mechanism: having degree d is assumed to improve idea quality by ad. Thus, a firm that has degree d, such that its partners have average degree d', enjoys upgrading gains of $(a/2)(d^2 + dd')$. Here $(a/2)d^2$ is own learning: degree d improves idea quality by ad, which affects all d links, generating surplus ad^2, half of which accumulates to the firm. And $(a/2)dd'$ is peer learning: peers on average gain ad', affecting d links, and half of the resulting add' accumulates to the firm. Assume that $a = \alpha g_H/2$, so that α measures the effect of learning by doing in terms of its share of the surplus from a random match.

Ignore search, and consider the effect of the experiment in the presence of upgrading. Importantly, when firms make their decision on which links to keep, they do not consider the benefits from peer upgrading. This assumption simplifies the analysis because it implies link decision thresholds that do not depend on the peers' degree. It can be justified by assuming that firms, when they make link decisions, do not know or think about the treatment status of their peers.

Let E be an indicator for the presence of the intervention.

Proposition 4 (Upgrading)**.** *Surplus in the experiment departs from its baseline value as follows.*

(i) For $\lambda < \kappa$, log firm surplus is

$$\log V(T,S) \approx \log V_0 + T \cdot \gamma(1+\alpha\lambda) + S \cdot \gamma\alpha\lambda. \tag{A7}$$

(ii) For $\lambda > \kappa$, log firm surplus is

$$\begin{aligned}\log V(T,S) \approx \log V_0 - E \cdot \gamma \frac{\lambda}{2\lambda-\kappa}\frac{\alpha}{2}\kappa + T \cdot \gamma \frac{\kappa}{2\lambda-\kappa}\left(1 - 2\alpha\kappa\frac{\lambda}{2\lambda-\kappa}\right) \\ - S(1-T)\gamma\left(\frac{2\lambda-2\kappa}{2\lambda-\kappa}\left(1-2\alpha\kappa\frac{\lambda}{2\lambda-\kappa}\right) + 3\frac{\lambda}{2\lambda-\kappa}\alpha\kappa\right).\end{aligned}$$

This is Proposition 4 in Szeidl (2025). In Case (i), two things change relative to Proposition 2. The coefficient of the treatment effect is larger, which follows because of a spillover effect: the new links create upgrading, which benefits the relationships over the existing links and thus the firm's surplus. And there is a positive effect of exposure, which follows because of another spillover effect: treated peers, because of their increased degree, upgrade, benefitting the firm.

In Case (ii), several things change. Most importantly, the presence of the intervention has a direct negative effect, even for firms who are neither treated nor exposed. This is the term involving E. This effect emerges because the untreated peers of these firms lose partners, which reduces their upgrading, negatively impacting the firm. This is a second-degree indirect effect. The specific coefficients in the expression for the surplus also change, but their signs remain unchanged, so that the qualitative message of the result is the same as that of Proposition 2.

B.5 Low-access trap and hypotheses

We turn to derive the hypotheses from the model. We begin by formally introducing the low-access trap. Consider the search framework of Proposition 3.

Definition 1. Firms are in the low-access trap if either (i) $\lambda < \kappa$ and $c/\delta \in (\beta/2, 1/2)$, or (ii) $\lambda > \kappa$ and $c/\delta \in (\beta\left(1 - \frac{\lambda}{\kappa}(1-\beta)\right)^2 \frac{\kappa^2}{2\lambda^2}, \frac{\kappa^2}{2\lambda^2})$.

In both of these cases, according to Proposition 3, it would be privately optimal to search, but because of undervaluation ($\beta < 1$) firms choose not to. Because they do not search, they never find out that their beliefs are miscalibrated.

Now suppose that firms are provided with a small amount of referrals. As a result, they update their beliefs. Given the continuity of our framework, we assume that even a small amount of referrals reveal the full distribution of match qualities. Thus, the referrals correct beliefs about match quality, that is, β increases to 1.[2] Then, the actual and the privately optimal incentives to search become identical. Moreover, if the referrals did not provide sufficiently many new partners, then the incentives to search remain in place, and the firm chooses to search privately.

With this preparation, we turn to derive the hypotheses of Section 3.

H1 The referrals can create new partnerships.

This follows from Propositions 3 and 2. Proposition 3 implies that in both cases (i) and (ii), it is possible to reduce the search cost c sufficiently to make search worthwhile. We can conceptualize the referrals as a reduction in c. It then follows that the referrals can lead firms to search, which is equivalent to entertaining the referred partnerships. Proposition 2 shows that when firms do this, new partnerships are formed.

H2 In a low-access trap, referrals can increase beliefs about potential partners and private search.

This follows from our discussion of the low-access trap above.

H3 When firms are below capacity, the referrals expand the production network.

This is Case (i) of Proposition 2, showing that when firms are below capacity, the referrals increase degree.

H4 The referrals have a positive direct effect on firm performance, especially for low-degree firms.

The first part of the statement follows from Proposition 2, which shows that the treatment T has

[2] In practice, a finite number of referrals need not fully correct beliefs, but should move them closer to the truth.

a positive effect on the firm surplus in both cases (i) and (ii). The second part again follows from Proposition 2. In case (i), the treatment effect is proportional to $\gamma = \delta/\lambda = \delta/d_0$ where $d_0 = \lambda$ is baseline degree. Thus, an increase in baseline degree in this range leads to a proportionally smaller effect. When baseline degree reaches κ, we get to case (ii), where the treatment effect is proportional to $\kappa/(2\lambda - \kappa) \cdot \gamma$. This is always weakly smaller than δ/κ, i.e., the lowest possible treatment effect of case (i). This follows because $\kappa/(2\lambda - \kappa) \leq 1$ and $\gamma = \delta/\lambda \leq \delta/\kappa$ for $\lambda \geq \kappa$. Intuitively, the gain is smaller both because fewer referrals are good enough to be taken on, and because they add smaller value since they crowd out existing partners. In summary, the treatment effect is smaller when degree is higher both because the referrals constitute a proportionally smaller gain in potential partnerships, and because in the reallocation case the gains from referrals are smaller.

H5 When firms are below capacity, the referrals have a positive indirect effect on prior partners.

This follows from Proposition 4, case (i), which shows that the coefficient of S is positive. The Proposition also shows that in this case the indirect effect impacts treated and untreated firms equally, i.e., S is not interacted by $1 - T$, justifying our empirical specification in Section 5.4.

H6 When firms are in a low-access trap, the referrals can have large private and social returns.

In the low-access trap defined above, in both case (i) and (ii), the private gain exceeds the cost of search. Thus, the referrals have a "large" private return in the sense that with well-calibrated beliefs, a return of that magnitude would not be left on the table. Furthermore, in case (i), which is more likely in the low-access trap as firms under-search, we know from Propositions 2 and 4 that indirect effects are zero or positive. Thus, the social return exceeds the private return.

C Social return evaluation

An important component of the social return is the consumer surplus. We first describe our model-based methodology for measuring the treatment effect on the consumer surplus, and then our approach to measuring the social return, which implements this methodology and combines it with measuring the other components of the social return.

C.1 Characterizing the consumer surplus

We derive an expression for the gain in the consumer surplus as a result of our intervention. This expression combines a regression coefficient we estimate and an elasticity of substitution we can borrow from prior work, and is thus computable. To do this, we need to model the demand side for the final good, as well as competition between client firms. We do this using a framework that models client firms and their consumers, but not suppliers. Incorporating suppliers into this framework, perhaps along the lines of Szeidl (2025), is beyond the scope of this paper. As a further simplification, we represent the impact of the treatment on client firms in reduced form as an improvement in product quality. We assume there are no indirect effects of the treatment on client firms through interactions in the input market, consistent with the expansion channel characterized in the paper. However, our framework allows for indirect effects due to competition in the final good market. Our framework follows closely that developed in Cai and Szeidl (2024).

Setup. We assume that there is a continuum of client firms of a fixed positive mass and index them by $i \in I$. Their output combines into a composite good

$$Q_I = \left[\int_{i \in I} (h_i \cdot Q_i)^{1-1/\sigma} di\right]^{\frac{\sigma}{\sigma-1}} \tag{A8}$$

where h_i is the quality or appeal of the product (or service) of firm i. Consumer utility is

$$\frac{Q_I^\theta}{\theta} + H \tag{A9}$$

where $0 < \theta < 1$ and H is a numeraire good with price normalized to one, produced competitively by labor. Concerning firm production, each firm i has productivity ω_i and uses a constant returns to scale technology to produce using labor. That is, we represent intermediate inputs and their impact on client firm performance in a reduced-form fashion with h_i and ω_i. In the same spirit, we assume that the treatment can be approximated as an increase in h_i for treated firms by a factor $\bar{\gamma}$, so that $\log h_i = \log \tilde{h}_i + \bar{\gamma} T_i$.[3] This is a reduced-form representation of the effect that the referrals improve business performance. From here on, we use the convention that variables absent the introduction of the treatment (i.e., in the counterfactual) are denoted by tilde, e.g., $\tilde{h}_i$

[3] Although $\bar{\gamma}$ does not directly map into the γ of the Szeidl (2025) model described above, it is closely related, hence the choice of notation.

is product quality of firm i absent the treatment, and the treatment effect is denoted by Δ, e.g., $\Delta \log h_i = \bar{\gamma} T_i$.

Proposition 5. *(i) Client firm revenue after the intervention can be written as*

$$\log P_i Q_i = \xi_0 + \xi_1 T_i + \varepsilon_i \tag{A10}$$

where $\xi_1 = \bar{\gamma}(\sigma - 1)$ and ε_i is orthogonal to T_i in the cross-section of client firms.

(ii) The gain in the consumer surplus as a result of the intervention can be written as

$$\Delta CS = \bar{\gamma} \int_{i:T_i=1} \tilde{P}_i \tilde{Q}_i = \frac{\xi_1}{\sigma - 1} \int_{i:T_i=1} \tilde{P}_i \tilde{Q}_i, \tag{A11}$$

that is, the product of $\xi_1/(\sigma - 1)$ and treated firms' revenue absent the treatment.

Proof. Given the CES structure, firms charge a constant markup over marginal cost, so that we can treat prices P_i as effectively exogenous. Define

$$P_I = \left(\int_{i \in I} \left(\frac{P_i}{h_i} \right)^{1-\sigma} di \right)^{\frac{1}{1-\sigma}}$$

as the quality-adjusted price index. Familiar arguments imply that

$$\frac{h_i Q_i}{Q_I} = \left(\frac{P_i / h_i}{P_I} \right)^{-\sigma} \tag{A12}$$

so that the quality-adjusted relative quantity of the product of firm i relates with elasticity $-\sigma$ to its quality-adjusted relative price. Rearranging this implies that $\int_{i \in I} P_i Q_i di = P_I Q_I$. Given a total budget B, maximizing (A9) subject to $H + P_I Q_I = B$ yields

$$Q_I^{\theta - 1} = P_I \tag{A13}$$

which implies that

$$P_I Q_I = Q_I^{\theta} = P_I^{\frac{\theta}{\theta - 1}}. \tag{A14}$$

This effectively solves the model, because it expresses Q_I with (known) prices, and then (A12) expresses Q_i with Q_I and prices.

We can write the revenue of firm i as

$$P_i Q_i = P_I Q_I \left(\frac{P_i}{h_i}\right)^{1-\sigma} P_I^{\sigma-1} = P_I^{\frac{\theta}{\theta-1}+\sigma-1} \left(\frac{P_i}{h_i}\right)^{1-\sigma} \tag{A15}$$

which means that we can write log revenue in the cross section as

$$\log P_i Q_i = \xi_0 + \xi_1 T_i + \varepsilon_i \tag{A16}$$

where $\xi_0 = (\theta/(\theta-1) + \sigma - 1) \log P_I$, $\xi_1 = (\sigma-1)\bar{\gamma}$, and $\varepsilon_i = (1-\sigma)(\log P_i - \log \tilde{h}_i)$. Since P_I is constant in the cross-section, so is ξ_0. And because T_i is randomized, it is orthogonal to ε_i which measures quality-adjusted firm productivity. Thus, we have established the first claim of the proposition, and we can interpret (A16) as a cross-sectional estimator of the treatment effect on log revenue.

This equation may initially suggest that the treatment only has a direct effect on firm performance, but that interpretation is incorrect. This follows because the treatment—though it does not affect firm-level prices and hence ε_i—does affect the constant ξ_0. In particular, the treatment affects the price index: $P_I < \tilde{P}_I$ because qualities improve, implying that $\xi_0 < \tilde{\xi}_0$. This is a business stealing effect that manifests itself as an "intercept problem." Firm revenue for untreated firms falls due to competition from treated peers, captured by a decline in the price index. Importantly, treated firms are affected by this force equally (to a first-order approximation), since they operate in the same more competitive marketplace. This is why business stealing is captured by the constant ξ_0.

We now turn to evaluate the impact on the consumer surplus. Using (A14) we can express the total utility of consumers as

$$\frac{1-\theta}{\theta} P_I^{\frac{\theta}{\theta-1}} + B. \tag{A17}$$

The first term measures the consumer surplus from consuming the products of the industry. The impact of the treatment on the consumer surplus, for $\bar{\gamma}$ small, equals the change in the consumer surplus in response to a change in $\bar{\gamma}$ from zero to positive. The derivative of (A17) with respect to $\bar{\gamma}$ at $\bar{\gamma} = 0$ is

$$-\tilde{P}_I^{\frac{1}{\theta-1}} \frac{1}{1-\sigma} \tilde{P}_I^{\sigma} (1-\sigma) \int_{i:T_i=1} \left(\frac{\tilde{P}_i}{\tilde{h}_i}\right)^{1-\sigma} = -\tilde{P}_I^{\frac{\theta}{\theta-1}} \tilde{P}_I^{\sigma-1} \int_{i:T_i=1} \left(\frac{\tilde{P}_i}{\tilde{h}_i}\right)^{1-\sigma} = \int_{i:T_i=1} \tilde{P}_i \tilde{Q}_i.$$

Thus, to a first-order approximation in $\bar{\gamma}$, the change in the consumer surplus as a result of the treatment is

$$\Delta CS = \bar{\gamma} \int_{i:T_i=1} \tilde{P}_i \tilde{Q}_i = \frac{\xi_1}{\sigma - 1} \int_{i:T_i=1} \tilde{P}_i \tilde{Q}_i. \tag{A18}$$

This establishes the second claim of the Proposition.

C.2 Measuring the social return

The social return consists of the following components. (i) Gain in consumer surplus. (ii) Gain in the producer surplus of treated firms. (iii) Spillovers to the producer surplus of untreated firms. These gains have to be compared to the total cost of the intervention, including the cost of the surveys and the subsidies. We estimate this total cost to be RMB 710,000.

To measure the gain in the consumer surplus, we use Proposition 5, applied to the subsample of low-degree client firms. We take the estimate of ξ_1 from Table 7, column 1, and combine it with the total revenue of treated low-degree client firms at the 2021 baseline, which we use as the proxy for their counterfactual revenue. Following Cai and Szeidl (2024), we use $\sigma = 5$. We find that the gain in the consumer surplus was about 5.7 times the cost of the intervention. To measure the gain in the producer surplus of treated firms, we take the profit treatment effect estimates for suppliers from Table 6, and for low-degree clients from Table 7. We find that the profit gains were about 4.2 times the cost of the intervention for the former, and about 11.3 times the cost of the intervention for the latter group of firms. Finally, we need to assess the role of spillovers. In the input markets, we estimate positive spillovers, which can be bounded from below by zero. In the final goods market, there may be negative spillovers, as we discussed after Proposition 5. However, these negative spillovers never exceed the gains to treated client firms: in our model, improvements in access do not shrink the industry.

Using these estimates, a measure of the social return is to include the effect for suppliers and the consumer surplus, but bound the positive spillover in the input market by zero, and bound the gains for clients plus the possible negative business stealing effect in the final good market—where the former exceeds the latter in our model—by zero. This gives an estimate of 990 percent for the social return.

D The Brush Pen: Production and Quality Evaluation

This appendix provides more details about the brush pen and the quality evaluation procedure.

D.1 The Product and Production

We include some pictures of the product and the production process.

Figure A1: Brush pens of various sizes and types of hairs

Figure A2: Workshop for brush head production

Figure A3: Workshop for handle production

Figure A4: Stores for brush head and handle

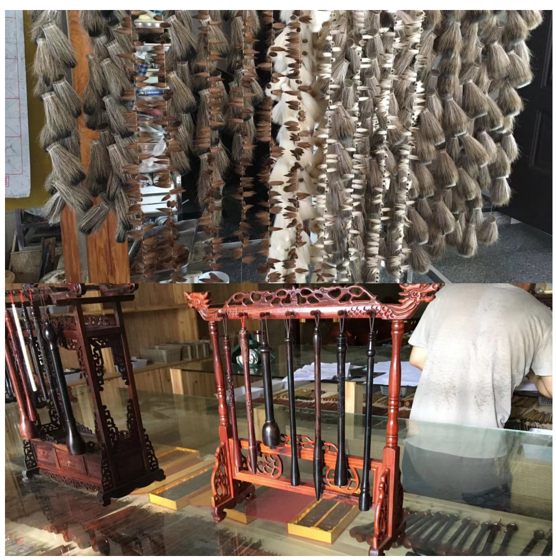

D.2 Brush Pen Structure

The brush pen is divided into five main components, each with specific dimensions:

1. **Brush head top**: The working end of the brush pen that comes into contact with the writing

Figure A5: Assembling and store for brush pens

surface.

2. **Ferrule**: The section that holds the brush head in place and connects it to the handle.

3. **Handle**: The main body of the brush pen, providing grip and balance.

4. **End Cap**: The cap at the end of the handle, often used for aesthetic purposes.

5. **Pen Loop**: A loop for hanging or attaching the brush pen.

Key measurements include:

- l_1: Length of exposed brush head, important for determining the precision and flexibility of the brush.

- l_2: Total length of the brush pen, affecting the overall balance and handling.

- d_1: Diameter of the brush head, influencing the fineness of the strokes.

- d_2: Diameter of the handle, impacting the comfort and grip.

Figure A6: Structure of pen

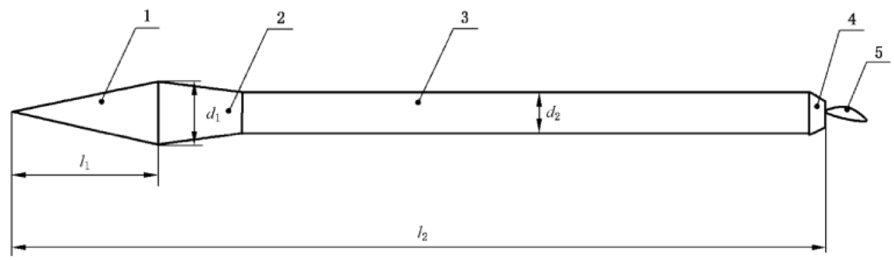

D.3 Evaluation Criteria

D.3.1 Brush Head Quality

1. **Classification of Brush Head Types**

   - Goat Hair Brush: Made entirely from goat hair.
   - Weasel Hair Brush: Made entirely from weasel hair.
   - Mixed Hair Brush: Combines different types of animal hair and nylon.

2. **Grading of Brush Head Quality** We discuss how our experts evaluated the brush head.

   - The criteria are based on:

     (a) Fineness of the hair: Determines the craftsmanship, smoothness, and precision of the brush strokes.

     (b) Straightness and toughness of the hair tip and strands: Affects the durability and consistency of the brush.

   - Each brush head is scored on its craftsmanship (up to 40) and durability (up to 60), with a total score of 100 points. The details are as follows,

     Sharpness (Jian): Whether the tip of the brush head is sharp and well-defined after trimming.

     Evenness (Qi): After the brush is repaired and loosened, whether the bristles form a neat, level surface.

Roundness (Yuan): Whether the entire brush head forms a standard cone shape without concavity or convexity.

Resilience/Elasticity (Jian): Elasticity from brush root to brush head tip, maintaining a stable transition in strength. The root has the greatest elasticity, transitioning smoothly in the middle, and the tip is the softest.

- **Durability/Writing Function** is evaluated on the following criteria, with a total score of 60 points:

  (a) Suitability for writing different scripts.

  (b) Whether the pen brush tip maintains integrity during writing without falling apart.

D.4 Scoring Process

The jury evaluating the pens consisted of six experts. Each product was evaluated independently by 2 or 3 judges. Scoring was conducted anonymously, with products identified only by codes, and no information about the manufacturer or price was provided to the judges. Judges were prohibited from communicating or discussing their evaluations. For each product, we computed the average score of all judges who evaluated that product.

E Detailed Procedure for the Design and Implementation of the Referrals

E.1 Design

E.1.1 Screened referrals

The first step in constructing the screened referrals was the treatment assignment, illustrated in Figure 2. Using the 2019 survey data, and following the stratification procedure described in the text, we divided 295 client firms into two groups: 147 treated, 148 untreated. Similarly, we divided 450 supplier firms into two groups: 223 treated, 227 untreated. Then, for each treated client, we constructed 2 screened candidate referrals that were treated suppliers; and for each untreated client, we constructed 2 screened candidate referrals that were untreated suppliers.

Construction was based on representing the referrals as a network flow problem. At a high level, we did the following. First, we constructed the set of potential referrals that seemed good to us. Then we organized these referrals into an auxiliary directed network. In this network, each network flow represents a particular configuration of supplier-client referrals. Then we chose link capacities, and costs for the links and nodes. These choices encoded our objectives that referrals should be relatively good matches; that all clients should have two referrals; and that all suppliers should have at least one and none should have too many referrals. Finally, we used a standard algorithm to find a minimum-cost maximum flow. We now describe this procedure in detail.

1. *Candidate referrals.* We combined the 2018 and 2019 surveys to create a baseline production network. Then, we classified a supplier-client pair, who were not transacting in this data, to be a potential referral if either the supplier had an existing partner client firm that was similar to the client, or the client had an existing partner supplier that was similar to the supplier. Similar meant that the two firms had the same role (supplier or client) and had the same product type for their main product (the list of all product types is in column 3 of Figure 1). We call this set of links type-2 potential referrals, and call their collection auxiliary network 2. A subset of these links were classified as "closely similar" if the firms also had a similar price, defined to mean that the log price difference was lower than the median across all firm pairs. We call this subset of links type-1 potential referrals, and call their collection auxiliary network 1. We expected that type-2 links would be decent referrals, but that type-1 links would be better. Thus, we aimed to construct referrals that use type-1 links when possible, and type-2 links otherwise. We sometimes refer to the collection of type k-links ($k = 1, 2$) as network k.

2. *Strength of potential referrals.* Besides the difference between type-1 and type-2 potential referrals, we introduced a second notion of the strength of referrals, which was based on the number of similar firms that supported a potential referral. Specifically, a link in network k was classified as strong if the two firms were connected by at least two similar firms (using the

definition associated to network k), and weak if they were connected by only one. Thus, we ended up with two notions of strength, one for type-1, the other for type-2 potential referrals.

3. *Directed network.* We created a directed network as follows. We took network 2 (i.e., the "decent" potential referrals) and made all links directed from suppliers to clients; added an artificial source node s and added links which point from s to each supplier; and added an additional target node t, and added links which point to t from each client. We then assigned all client-to-target links capacity 2, and all supplier-to-client links capacity 1. In this directed network, the maximal flow assigns a value of at most 2 going through each client. If a flow exists that meets this constraint with equality, then that flow creates two referrals per client.

4. *Costs of links and nodes.* We assigned costs to meet our objectives as follows. (i) The first link from s to a supplier was free, every subsequent link expensive. (The exact costs did not matter.) This condition was intended to ensure that many suppliers are reached. (ii) Link costs were ordered as follows: strong in network 1 $>$ weak in network 1 $>$ strong in network 2 $>$ weak in network 2. We did this to include many network 1 links and strong links.

5. *Implementation.* We found the minimum-cost maximum flow using a standard network simplex algorithm. We were able to assign two candidate referrals to all 295 clients and reach all 450 suppliers. In particular, among the 450 suppliers, 326 received one, 112 received two, 8 received three, and 4 received four candidate referrals. For each client, we randomly labelled one of the two screened candidate referrals as the A referral, and the other as the B referral.

E.1.2 Unscreened referrals

To construct the unscreened candidate referrals, for each client we randomly chose a supplier with the same treatment status, subject to the requirements that (i) the supplier had to have the same business type as the type A referral of the client we constructed above; (ii) for brush pen producer clients, the main product of the supplier had to use the same type of hair as that of the client (if applicable), (iii) the supplier was not a partner or a screened referral of the client. That is, we first set business type to be the same as that of the type A referral (as handle producer, processing firm, or brush head producer). Then, for the former two business types, we chose the supplier randomly. For the latter business type, if the client firm was a brush pen producer, we chose the supplier randomly among suppliers using the same hair type (goat, weasel, or mixed). If the client firm was a make-up pen producer—since such firms generally use a mix of hairs—we did not restrict the hair type. And we always ensured that the unscreened referral was distinct from existing partners or screened referrals of the client.

E.1.3 Information treatment

We implemented the information treatment over untreated-to-untreated candidate referrals. We divided the untreated clients into two equal-sized groups, one of which received the information treatment. For each client in this group, we selected the screened candidate referral link B as the information referral. The firms in these referrals then received the information treatment as described in the main text.

E.1.4 Making the referrals

In the information referrals, we told the two parties that the researchers and the local government are conducting a research study, and we think that the referred firm would be a good potential business partner. In the screened and unscreened referrals, we also offered a subsidy for the first transaction, which covered 50% of the transaction value up to a cap of 1,500 RMB per transaction, and which was valid for two months. In both information and subsidized referrals, we gave coupons to both firms which included the contact details of the referred partner. In the subsidized referrals, the coupon also contained information on the subsidy (see Figure A7 for the subsidized coupon, with the identifying information of the firms blurred out). Only the client firm could ask for the 50% subsidy, and to claim it, it was required to submit: (1) a video of the face-to-face transaction, (2) both coupons signed by the supplier and the client, (3) an invoice documenting the transaction, and (4) an itemized list of all goods purchased.

Figure A7: 50% Subsidy Coupon

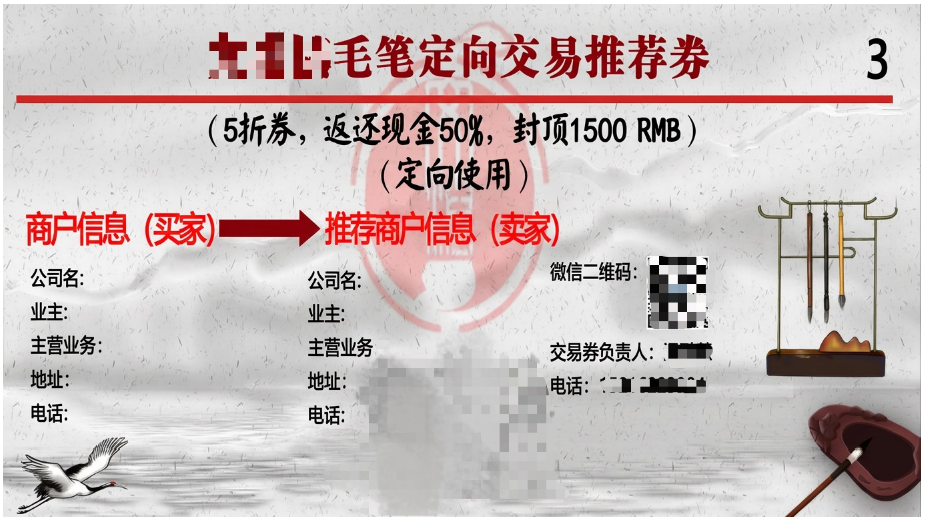

E.2 Implementation

During the implementation of the referrals in 2021, three challenges emerged. 1) 61 firms (36 in T1 and 25 in T0) were reported to us by the leaders of the local administrative units to have gone out of business. These firms were involved in some referrals, which required replacements.[4] 2) For 24 client firms (12 in T1, 12 in T0) all three referrals were processing firms. Some of these clients complained, because they found these matches less desirable. 3) 7 firms, all in T1, switched roles, moving from supplier to client or vice versa, and complained about their referrals. We could not find the analogous switchers in T0, because data entry from the survey was not ready. To address these challenges, we revised our objectives for the screened referrals as follows.

1. All existing clients should have two screened referrals to existing suppliers. All existing suppliers should have at least one screened referral to an existing client.

2. Any replacement referral should have the same business type as the original referral.

3. All clients for whom both screened referrals were processing firms should get two additional non-processing referrals.

[4] We later found 20 of these firms in the surveys, i.e., only 41 of the 61 actually exited. The non-exiting firms received their original referrals.

4. Other than these changes, the original screened referrals should remain unchanged.

E.2.1 Adjusting the algorithm

To meet the revised objectives, we adjusted our algorithm as follows.

1. *Redefining nodes.*
   (a) For all firms that changed sides (from supplier to client or conversely) we dropped prior links and created new links based on their new role. All new supplier-client links were of type 2, because product information for the switchers was not yet available.
   (b) We removed all firms which exited.
   (c) We removed all processing suppliers.
2. *Adjusting links and capacities.*
   (a) We assigned to all new links capacities and costs the same way as in our initial algorithm.
   (b) We removed all initially constructed referrals that had no problems, i.e., that (i) connected firms that had not exited, (ii) involved firms that had not switched business roles, and (iii) did not involve both A and B links being assigned to processing suppliers. We adjusted client firms' capacity from the initial level of 2 down by the number of links of the client we removed. This ensured that we only modified referrals that had challenges.
   (c) If a client had remaining capacity 1, i.e., one problematic referral, we deleted all links from the client to suppliers who did not have the business type of the problematic referral.
   (d) If the client had remaining capacity 2, i.e., two problematic referrals, and at least one of those was the result of exit, then we replaced the client with two client nodes, each with capacity 1, and performed operation (c) for that of the new nodes corresponding to the exiting referral.
   (e) For each supplier, we set its initial counter (which measures the flow going through it, and governs the cost of additional flow) to be the number of referrals we deleted. This operation was to ensure that referrals continue to be spread out across suppliers.

(f) We deleted all unscreened referrals, to make sure an unscreened referral is not used for a screened referral by accident.

We then reran our algorithm of finding a minimum-cost maximum flow to create the new screened referrals. For the unscreened referrals, we replaced any missing referral using the original approach but in the revised pool of firms.

E.2.2 Summary of the modifications

Our revisions resulted in the following changes.

1. 44 out of 295 clients lost one of their screened referrals. All of these were successfully replaced by another screened referral with the same business type.

2. 24 out of 295 clients had both of their screened referrals as processing suppliers. For these firms, the algorithm constructed 2 new non-processing screened referrals. We also constructed new unscreened referrals for these client firms, to maintain that the C referral has the same business type as the A referral. And we made both their initial referrals (if treated) and their new referrals, so that treated clients in this set ended up with 6 referrals in total.

3. 3 clients lost both screened supplier assignments. We constructed replacements using the algorithm. For one (and only one) of these referrals we were unable to match the business type of the original referral, and we therefore switched it from handle to brush head producer.

4. On the supplier side, 2 switchers (in T1), who were initially clients but ended up as processing firms did not receive any referrals, and 39 suppliers (across T1 and T0), all processing firms, were only reached by the original processing referrals to which we added new referrals on the client side. In summary, essentially all suppliers were reached with the full set of referrals.

F Detailed Procedure for the Marketing Offer Experiment

For each untreated client in our main intervention we created two screened referrals, some of which were information referrals, the others unmade referrals. We used this sample of referrals for the marketing offer experiment. We had two treatment arms.

1) Enhanced information (EI). In addition to the same information we provide in the information arm of the main intervention, we tell the firms the following.

- The researchers think that the referred firm would be a good partner for you.
- The referred firm is aware of your business, and understands that the researchers think the two of you would be good partners.
- The researchers have evidence from previous referrals that firms are often more satisfied with their referred partners than they initially expect.
- For a supplier: We encourage you to reach out to the client by sending a free sample along with a price quote. Include information about your main specialty and the number of years you have been in business.
- For a client: We encourage you to reach out to the supplier to request a free sample and a price quote.

2) Marketing offer (MO). In addition to providing the EI treatment, we provided the client (i) a one-page document that included the type and price of the supplier's products, along with the supplier's contact information, and (ii) a free sample produced by the supplier if one was available.

We introduced the treatment in our sample of firms for this experiment as follows. For a random third of the client firms (43), we administered EI to both of their screened referrals. For the remaining two-thirds of client firms (89), we administered EI to exactly one of their screened referrals, and MO to their other screened referral.

Table A22 presents summary statistics and balance using the 2021 baseline survey. Since the randomization is at the client level, the table is about clients. In column 1 we include the firms that received EI only, while in column 2 we include the firms that received both EI and MO. There are no significant differences in any of the variables, providing evidence that the randomization was balanced.