

Firm-to-Firm Access and Private Sector Development^{*}

Adam Szeidl
Central European University and CEPR

October 31, 2025

Abstract

Firms in developing countries often stay small and fail to upgrade. A possible explanation is barriers to accessing suppliers and clients in the production network. I present an industry equilibrium model of improving firm-to-firm access, and use this model to review and organize existing empirical work. The model makes four predictions, all of which are consistent with evidence. (1) Improving access improves business performance. (2) Improving access can both expand and reallocate the production network. The former is associated with positive indirect effects, the latter with negative indirect effects. (3) Accounting for these indirect effects, access can generate large aggregate gains. (4) A range of frictions, both external and internal to the firm, imply that private markets often under-provide firm-to-firm access. I conclude by discussing open questions.

JEL codes: O12, O14, O33, L14, L23

Keywords: market access, firm performance, indirect effects, business stealing, spillovers, upgrading

^{*}Email: szeidla@ceu.edu. I thank Vittorio Bassi, Elhanan Helpman, Miklos Koren, Marc Muendler, Chhavi Rastogi, and Meredith Startz for helpful comments. I thank 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.

1 Introduction

Firms in developing countries often stay small and fail to upgrade their products and technologies (Verhoogen 2023). Much research has focused on “supply-side” explanations for these facts that are based on frictions internal to the firm, such as the lack of managerial capabilities (McKenzie, Woodruff, Bjorvatn, Bruhn, Cai, Gonzalez-Uribe, Quinn, Sonobe and Valdivia 2023). There is less work on “demand-side” explanations, which are based on frictions in firms’ access to markets. Access frictions may be especially important in firm-to-firm markets, both because the majority of business transactions are firm-to-firm, and because these markets are decentralized and opaque. Search and matching frictions may prevent “good” firms from reaching their potential by finding enough or appropriate partners (suppliers and clients). They may also enable “bad” firms to survive by limiting their partners’ ability to replace them. Through this logic, lack of access can explain the lack of growth and upgrading. In recent years, a small but growing literature has used well-identified empirical designs to study the role of firm-to-firm access in private sector development.

In this paper, I have two goals. First, I present a simple model that may be helpful for thinking about firm-to-firm access in industry equilibrium. Second, I use this model to organize and review the existing evidence, and to identify knowledge gaps. Both in the model and in reviewing the evidence, I focus on the following main questions. (1) What is the effect of firm-to-firm access on business performance? (2) What is its effect on industry equilibrium? (3) What are the welfare implications? (4) What are the barriers to access? In discussing these questions, I will mainly focus on domestic production networks—though occasionally I discuss trade networks—because production is a necessary ingredient to trade and because most firm-to-firm transactions are between domestic firms. My review will also be shaped by my own interests and knowledge of the literature.

In section 2, I present a model of access in a supplier-client network. In this model, a supplier-client relationship is valuable because it creates a new business opportunity, mod-


elled as a blueprint for a new product variety of the client firm. This is a departure from standard models of supply chains, in which suppliers produce a pre-specified input for the client. The quality of the business opportunity is match specific and drawn from a commonly known distribution. Firms have limited access: they are exogenously provided by a random set of potential partners and can only collaborate with those partners. Moreover, firms have capacity constraints limiting the number of partners they can work with, which reflect either the limits to product differentiation or decreasing returns to scale. It follows that a firm whose access exceeds its capacity has to choose among its potential partners. A partnership is formed if both firms prefer to form it.

The model generates four main implications. First, a firm that experiences an improvement in access improves its business performance. This improvement reflects that greater access comes with more business opportunities, and can materialize through two distinct channels that parallel the expanding variety and quality ladder models of endogenous growth (Grossman and Helpman 1993). As long as the firm is below capacity, a new partner leads to the creation of a new variety and an associated expansion of activity. Once the firm is at capacity, the new partner, with a sufficiently good idea, leads to the replacement of an existing variety, and a reallocation of activity.

Second, these channels shape the impact of access on the structure of the production network. When an intervention increases access between a set of clients and a set of suppliers, the network can change in two qualitatively different ways. When firms are below capacity, the intervention expands the production network: new links do not crowd out existing links. But when firms are at capacity, the intervention reallocates the production network: new and better links crowd out existing links. These different changes in network structure are illustrated by Figure 1. Client firm C_1 's access to a new partner S_2 leads to an expansion of the network in case 1 but a reallocation of the network in case 2.

Third, because access changes the production network, it generates indirect effects on untreated firms that need to be accounted for when evaluating aggregate impact. I show that

Figure 1: Possible effects of access on the production network



when the intervention leads to expanding, indirect effects are either zero, or positive when I allow for upgrading. Intuitively, improved access can induce a firm to upgrade, which then benefits its prior partners. In contrast, when the intervention leads to reallocation, indirect effects are negative. Intuitively, treated firms stop doing business with some prior partners, which harms these prior partners. The direct and indirect effects are illustrated by the vertical arrows in Figure 1. In case 1, access has a positive, in case 2 a negative indirect effect on the prior partner S_1 . Importantly, indirect effects complicate treatment-control comparisons, because they violate the stable unit treatment value assumption (SUTVA). Thus, estimation requires either a setting with a “pure control” arm where firms are not subject to direct or indirect effects, or a model-based empirical approach that exploits variation in exposure to the indirect effect to infer its magnitude.

Fourth, the model captures two frictions, one external, the other internal to the firm, that limit access. The external friction is a search externality: firms may not search because they do not internalize the gain from a new partnership that accumulates to the other party.

The internal friction is pessimistic beliefs: firms may not search because they underestimate the quality of potential partnerships. Because they do not search, they never learn partner quality, allowing their misbeliefs to persist. These frictions imply that access is often undersupplied relative to the social optimum. In addition, the pessimism friction predicts that improving access can have amplified effects, because it can help correct beliefs which then encourage private search.

In section 3, I review evidence that relates to these four implications of the model. There is now a body of work showing that access improves both accounting measures of firm performance, such as revenue and profit, and measures of firm capabilities, such as productivity, quality, and labor specialization. Much of this evidence comes from settings where distance and borders limit access, but recent work shows access frictions even among nearby domestic firms. In my reading, this work convincingly demonstrates that access has positive effects on individual firms across a broad range of settings.

The evidence on industry-level effects is less developed. But there are now a handful of papers that, studying different contexts, have identified industry-level changes consistent with the expanding and the reallocation channels, respectively. These papers also show positive and negative indirect effects as predicted by the different channels. Thus, it appears that these are genuinely different ways in which industries can develop in response to improvements in access. An open question is how to identify circumstances in which access leads to expanding versus reallocation.

Some of the papers studying industry-level outcomes also evaluate aggregate effects. Although the evidence is thin, the results suggest that improving access can generate surprisingly large gains. Thus, improving firm-to-firm access may be an important policy objective. A key question is what policies are effective at improving access.

Finally, a body of evidence shows that access is in short supply, both for reasons external and internal to the firm. On the external side, search costs are important, including lack of information about the identity or the quality of potential partners. Also important are

enforcement costs that reduce the incentive to search. On the internal side, there is evidence that firms are pessimistic about the value of partners, which constrains both search and matching; and that firms do not have the marketing capabilities to demonstrate that they are an appropriate match. But we lack an understanding of why private-sector third parties do not overcome these barriers, e.g., by creating platforms to match firms.

In my reading, this evidence strengthens the case that firm-to-firm access may be an important driver of industry performance. However, as I already suggested, many questions remain. In Section 4, I conclude by discussing some knowledge gaps, including, beyond the ones touched upon above, the impact of specific policies ranging from targeting firms to quality certification; and the relative lack of macro-style evaluations of how access shapes development.

Like any review paper, I had to make choices about what material to cover. My focus on domestic production networks, my model-based approach, and the fact that I cover several non-experimental papers, implies that the present paper complements the recent review on market access by J-PAL (2024). Important work I do not cover in full detail includes the impact of access in cross-border networks, a literature carefully reviewed by Verhoogen (2023). I also discuss only lightly the importance of enforcement frictions and relational contracting in firm-to-firm trade, a literature reviewed by Macchiavello (2022). And I completely neglect the important line of research studying the impact of transportation infrastructure that reduces trade costs (Faber 2014, Donaldson and Hornbeck 2016, Donaldson 2018, Hornbeck and Rotemberg 2024).

2 Model

I present a stylized industry equilibrium model of firm-to-firm access. This model is related to equilibrium models of production networks, including Oberfield (2018), Demir, Fieler, Xu and Yang (2024a), Eaton, Kortum and Kramarz (2022), and Arkolakis, Huneus and

Miyauchi (2025). The goals of the model are to examine the impact of firm-to-firm access on private sector development; to explore the barriers to access; and to help interpret reduced-form evidence.

2.1 Setup

I study an industry producing a differentiated product. The industry 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. Here λ is a continuous variable, but we can think of it as a generalization of having a finite discrete number of potential suppliers.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. On the client side, κ can represent the limit to product differentiation. On the supplier side, it can represent decreasing returns to scale. I assume that the capacity constraints of suppliers and clients are identical to keep the analysis simple. The potential surplus (or value) created from a match between i and j is h_{ij} , which is a random draw from distribution G . I assume that G is 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. I focus on this notion of equilibrium, which I interpret as the long-run equilibrium that emerges after

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 $\lambda \cdot \mu$.

firms have maximally taken advantage of their opportunities. I define the industry surplus as the total value of all active links.

2.2 Discussion of Model Assumptions

Access generates ideas. A central assumption is that new partnerships generate new business ideas, such as blueprints for new product varieties. This approach is different from the traditional way of thinking about suppliers as producing a particular input that the client needs (Grossman, Helpman and Redding 2024). In the traditional perspective, a new supplier which is better at producing the particular input would replace an existing supplier. The approach proposed here incorporates this logic as a special case: when $\lambda > \kappa$, taking on a new supplier means replacing an existing supplier with one that has a higher-match specific value. But it adds by allowing new suppliers to expand the product scope of the client firm, similarly to expanding variety models of trade and growth (Grossman and Helpman 1993).

Capacity. A natural interpretation of the capacity constraint for clients is that it measures the extent of potential product differentiation. Thus, in settings in which the final good is fairly homogeneous, such as trade networks in agriculture, κ may be low; while in settings in which there is scope for differentiation, such as production networks in manufacturing, κ may be high. This interpretation is somewhat clouded by my assumption that the same κ governs the capacity of supplier firms. It is possible to work out the model with different values of κ for suppliers and clients, but it does not seem to yield major new insights. Thus, for simplicity I focus on the symmetric case, and interpret κ as a measure of potential product differentiation.

Production, pricing, demand. The above model does not incorporate production, quantity and pricing decisions, demand, and consumers. This is to focus the discussion on access in production networks. However, it is straightforward to combine the model with CES demand, and with constant-returns-to-scale production that uses labor. The main novelty introduced

by such a framework is business stealing in the client-to-consumer market, similarly to that modeled in Cai and Szeidl (2024). However, even in that framework, aggregate welfare will be shaped by the industry surplus introduced above, mediated by the substitution elasticities. The same elasticities determine the extent to which the welfare gain accrues to producers or consumers. Because these forces are relatively well-understood, here I focus on the determination of the surplus in the presence of access frictions.

Firm-level differences in quality. In the above framework all suppliers and all clients are identical, and all differences in quality are match specific. In reality, firms differ in quality and sort by quality in the production network (Demir et al. 2024a). I have worked out a version of the model with quality differences. The key novelty from that version is that the extent of quality sorting depends on access: higher access implies more sorting. However, making quality match-specific, as I have done here, improves tractability, and delivers the main insights I discuss below. Thus, I leave the development of a tractable model of access with firm heterogeneity for future research.

2.3 Results

Equilibrium. I begin by characterizing the equilibrium.

Proposition 1 (Equilibrium). *1. For $\lambda < \kappa$, firms take all of their potential partners.*

Each firm has degree λ . An increase in λ by a small factor γ yields an increase in the number of links and in industry surplus by the same factor γ .

2. For $\lambda > \kappa$, firms take the top κ/λ share of their potential partners. Each firm has degree κ . An increase in λ by a small factor γ yields no increase in the number of links and an increase in the industry surplus by a factor $\frac{\kappa}{2\lambda - \kappa}\gamma$.

All proofs are in the Appendix. The Proposition shows that for low values of λ/κ all potential links are accepted, while for high values of λ/κ only links of sufficiently high value

are accepted. This is intuitive: once the number of potential links exceeds the firm’s capacity constraint, the firm needs to select among potential partners.

The Proposition also characterizes the impact of an increase in access by a small factor γ (that is, in a first-order approximation). In the first case, access increases the number of links and industry surplus by the same factor γ . Intuitively, in this range, each new potential link generates a new business idea, leading to the expanding of the production network and an associated increase in surplus. In the second case, the impacts are more muted. New potential links are either rejected or used to replace an existing business idea with a better version of it, leading to the reallocation of the production network. The increase in surplus is now proportional to $\frac{\kappa}{2\lambda-\kappa}$. This fraction equals 1 when $\lambda = \kappa$, because at this point the links that get replaced are of zero value, so it is almost as if new ideas are being created. But it is lower than 1 when $\lambda > \kappa$ and converges to zero as λ grows. Intuitively, once there is sufficient access to obtain quite good ideas, additional improvements in access lead to small additional gains in idea quality.

Experiment in the model. In empirical work studying firm-to-firm access, researchers often evaluate a treatment that improves access between a subset of suppliers and clients. Motivated by this type of empirical setting, I conduct an *experiment* within the model. I take an equilibrium of the model, which I refer to as “baseline,” and then introduce new potential links. Specifically, I 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. I then characterize the emerging *short-run* equilibrium: the maximal symmetric threshold equilibrium that can be reached by adopting some of the new links and eliminating some of the links firms already formed at baseline. I do not allow firms to go back to prior potential links that they had rejected in the process of arriving at the baseline equilibrium, reflecting my focus on short-run effects. Symmetry now requires that firms with the same treatment status use the same cutoff rule for link decisions.

In this experiment, given the randomness and the continuity of the framework, each firm will have exactly half of its baseline partners treated. This symmetry helps characterize the equilibrium. But I am 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, I allow a zero mass of firms to have exposure S different from 0.5. I allow these firms to 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.

I characterize equilibrium degree and surplus in a first-order approximation for small δ . I denote $\gamma = \delta/\lambda$, so that we can think of γ as the proportional increase in access for treated firms. The subscript 0 denotes outcomes at baseline. I characterize outcomes as a function of T and S , where T is indicator for the firm being treated, and S , as noted, 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.*

1. For $\lambda < \kappa$, log degree is

$$\log d(T, S) \approx \log d_0 + T \cdot \gamma, \quad (1)$$

log firm surplus is

$$\log V(T, S) \approx \log V_0 + T \cdot \gamma, \quad (2)$$

and log total surplus is

$$\log W \approx \log W_0 + \frac{\gamma}{2}. \quad (3)$$

2. For $\lambda > \kappa$, log degree is

$$\log d(T, S) \approx \log d_0 - S(1 - T) \cdot \gamma, \quad (4)$$

log firm surplus is

$$\log V(T, S) \approx \log V_0 + T \cdot \frac{\kappa}{2\lambda - \kappa} \gamma - S(1 - T) \cdot \frac{2\lambda - 2\kappa}{2\lambda - \kappa} \gamma \quad (5)$$

and log total surplus is

$$\log W \approx \log W_0 + \frac{1}{2} \frac{2\kappa - \lambda}{2\lambda - \kappa} \gamma. \quad (6)$$

Begin with Case 1, i.e., low access relative to capacity $\lambda < \kappa$. Then, the results are qualitatively similar to the long-run effect of treating all firms as characterized by Proposition 1. The experiment increases degree and surplus for treated firms by a factor γ , and total surplus by a factor $\gamma/2$ because only half of the firms are treated. There are no exposure effects. In this range, improving access only creates new ideas, but has no indirect effects.

Consider next Case 2, i.e., high access relative to capacity $\lambda > \kappa$. By equation (4), an experiment of further improving access has no effect on the degree of treated firms, because these firms purely reallocate. However, it 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. Thus, perhaps surprisingly, the total effect of improving access is to reduce network density.

Turning to the surplus, equation (5) shows that the experiment increases the surplus of treated firms. The increase is by a factor (given γ low) of $\gamma\kappa/(2\lambda - \kappa)$. This equals γ only when $\lambda = \kappa$, in which case the partners of a treated firm that are crowded out generated essentially zero value at baseline. Otherwise, the treatment effect is smaller: replacing is not as effective as creating. As λ goes to infinity, the treatment effect converges to zero.

In addition to these gains to treated firms, (5) shows that the experiment also generates losses to 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. The magnitude of the exposure effect is $\gamma(2\lambda - 2\kappa)/(2\lambda - \kappa)$. Thus, when

$\lambda = \kappa$ there is no exposure effect: intuitively, the links lost by untreated firms created essentially zero value. But when $\lambda > \kappa$ the exposure-induced losses are positive, and as λ goes to infinity, they converge to γ . At a very high level of access, the experiment creates little additional value through replacement, but meaningful losses through link destruction.

Combining these effects, (6) shows that the impact of the experiment on aggregate welfare may be positive or negative. When $\lambda = \kappa$, the impact is $\gamma/2$, as in the low-access case. But for $\lambda > 2\kappa$, the effect is actually negative, because the gains from replacement are smaller than the losses from link destruction. The reason why improving access can reduce welfare is that there are missing markets: I do not allow the firms who lose links to switch to new partners.²

Impact evaluation in the experiment. A key question is how a researcher who conducts the above experiment should evaluate impacts. In general, the answer is not obvious due to the presence of indirect effects. To think about this issue systematically, consider the following three measures.

1. The impact of the experiment on treated firms.
2. The impact of the experiment on untreated firms.
3. The difference between the performance of treated and untreated firms.

The first two measures are obtained by comparing outcomes in the presence versus the absence of the intervention, i.e., relative to a “pure control” group which corresponds to the baseline case in the model. The third measure is a “naive” treatment effect estimation that ignores potential indirect effects.

In Case 1, equation (2) shows that measure 1 is equal to γ , measure 2 is equal to zero, and measure 3 is equal to γ . Intuitively, there are no indirect effects, thus untreated firms do

² If they could go back to some of the partners in their initial pool of λ potential partners, then these welfare losses may be smaller or completely vanish.

not change in response to the intervention, SUTVA holds, and the naive evaluation correctly estimates the treatment effect. Given the absence of indirect effects, it is not surprising that the naive treatment effect also correctly measures the effect of scaling the intervention to the industry, which, as characterized in Proposition 1, is also γ . Thus, in the low access range, the “naive approach” can be used to estimate the treatment effect and evaluate aggregate impact.

In Case 2, however, the naive approach fails. As we have seen above, untreated firms are negatively affected by the experiment, so that SUTVA fails. It can be verified that the naive comparison of treated and untreated firms leads to a difference of $\gamma\lambda/(2\lambda - \kappa)$, which overestimates the effect of the experiment on the treated by a factor of λ/κ , because it ignores the negative effect on the untreated. This bias is especially severe for λ high, in which case, as we saw above, most of the effect of the intervention lies in reducing the surplus of untreated firms. Neither does this naive comparison correctly measure the welfare effect of the experiment (which may be negative) or the welfare effect of scaling up the experiment (which is measured by $\gamma\kappa/(2\lambda - \kappa)$ by Proposition 1). Thus, in this case the naive econometric analysis overestimates the impact of improving access.

These results suggest two ways to empirically evaluate the impact of the experiment. First, the researcher could use a pure control arm, which would correspond to the baseline case in this model. If such an arm is available, the researcher can simply compare both treated and untreated firms with their analogues in the pure control. If feasible, this approach is preferable because it requires few assumptions about the economic model. However, it is often infeasible, for example when the intervention is at scale, so that most firms in the industry are directly or indirectly impacted.

The second approach is to estimate a fully-specified model, which may be a structural model, or reduced-form equations derived from such a model. In our case, this amounts to estimating a regression like (5) that includes both treatment and exposure. This regression allows us to correctly measure the treatment effect as the coefficient of T . Moreover, in

the current model, that treatment effect also correctly measures the impact of scaling up the intervention, which is given by $\gamma\kappa/(2\lambda - \kappa)$ by Proposition 1. Intuitively, when scaled up, the exposure effect on untreated is no longer operating because there are no untreated firms. More broadly, by matching the fully specified model to the data, we can recover deep parameters that allow for counterfactuals and the evaluation of aggregate impacts.

2.4 Search Effort

I turn to explore the incentives to search. This is relevant from a policy perspective: if firms are willing to search on their own, policy interventions to improve access may not be necessary.

I explore search in the baseline model without the experiment. I assume that firms already found their initial potential partners (for example via kinship networks), and that initial partnerships have been formed. I explore the incentives for additional search. For simplicity, I only allow clients to search, not suppliers; and I assume that each client firm, at a cost, can find δ new potential partners. I write the cost of search $c \cdot g_H/2$, so that it is measured in proportion to the average quality of a potential link. Purely for convenience, I 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, I introduce a within-firm search friction: I assume that firms are pessimistic about the distribution of link values. Each firm believes that the distribution is uniform on the interval $[0, \beta g_H]$ where $\beta \leq 1$. Pessimism does not affect firms' views of their existing relationships, only those they obtain by search. As I discuss below, there is evidence on this sort of pessimism in firm-to-firm search.

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

1. 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

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

2. 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

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

Consider Case 1. Note that all the bounds are on c/δ , which is a measure of the cost-to-benefit ratio of search. Comparing the first two bounds shows that when firms underestimate the value of partners by a sufficient margin, they avoid search even when it is privately

optimal. This is intuitive: pessimism reduces search. 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. This is because of a search externality: searching by clients also benefits suppliers, but clients do not internalize these gains.

Consider Case 2. Comparing the first two bounds shows that pessimism again reduces search incentives. However, β now reduces the bound by more, so that pessimism has a larger negative effect on search. This is because in this range the benefit of search is due to replacing partners, and if firms falsely think that some of their current partners (with value above βg_H) can never be replaced, then their incentive to search is further reduced. For comparison, in case 1, the benefit of search is expanding partnerships, so the value of current partners has no bearing on search incentives.

Comparing the second and the third bounds shows that the comparison between private and social search incentives depends on λ . It is easy to see that for $\lambda \in [\kappa, (3/2)\kappa]$, private incentives are weaker than social incentives, but for $\lambda > (3/2)\kappa$ the converse is true. Intuitively, there are two forces. On the one hand, the search externality reduces search incentives relative to the social optimum. On the other hand, new links crowd out the links of some existing partners, a negative effect that the searching firm does not take into account. When λ is high, the second effect dominates, because then the firm already has quite good partnerships, so that the private gains from further improving them are dominated by the losses from link destruction.

The key takeaway from the Proposition is that both because of search externalities and potential within-firm frictions, there may be a case for policies to improve firm-to-firm access. The Proposition also suggests that the case for such policies is stronger in the expanding range (case 1) than in the reallocating range (case 2), since in the latter case, due to negative indirect effects firms may have an incentive to over-search.

An interesting implication of the Proposition is that with pessimistic beliefs, improving access can induce private search. Suppose that we are in case 1, $c/\delta < 1/2$ so that it would

be privately optimal to search, but β is sufficiently low that firms do not search on their own. Now introduce the experiment. As firms experience the true distribution of link values, they should update to $\beta = 1$. But then firms should be willing to search on their own as well. Thus, the experiment can lead to private-sector improvements in access via changing beliefs about the value of search.

2.5 Upgrading

A natural intuition is that improving access can induce firm upgrading. For example, having a higher-quality partner may induce quality upgrading, and having more partners may induce specialization within the firm. In Appendix A2 I introduce a variant of the model that captures peer-induced upgrading in a simple way, by making it dependent on the number of the firm's partners. The value of each link is then shifted up by the upgrading performed by both the supplier and the client.

Proposition 4 in the Appendix characterizes the equilibrium in this setting. There are two key novelties. First, I obtain spillover effects even in the low access (expanding) environment, and these spillover effects are *positive*. This follows because new links induce treated firms to upgrade, which in turn benefits their preexisting partners. Through this logic, exposure actually improves firm performance, which is the opposite of what the model predicts in the reallocation environment. As a result, in the expanding environment, the gains from improved access can be broadly shared. This is the scenario depicted in the first case of Figure 1.

Second, in the high access environment, the experiment has a negative effect even on firms which are untreated and unexposed. That is, controlling the firm's treatment status and exposure status is not sufficient to pick up all direct and indirect effects. This is because an untreated firm with untreated partners is indirectly affected through its untreated partners. These partners are themselves exposed, leading them to lose some links, which in

turn reduces their upgrading. Intuitively, we have a second-degree exposure effect, coming from the (treated) partners of (untreated) partners of the firm of interest. Regression (5) does not control for this second-degree effect; capturing it would require variation in, and controlling for, the firm’s second-order partners’ treatment status. The failure to do so leads to an “intercept problem.” A takeaway here is that models with access can generate rich equilibrium feedback effects. Because these effects can be subtle, accounting for them may be easiest by estimating a fully specified structural model.

3 Evidence on Improving Firm-to-Firm Access

3.1 Firm Level Impacts

A central implication of the model is that improvements in firm-to-firm access lead to more productive partnerships that improve business performance and may drive upgrading. I begin by reviewing evidence on this implication. I focus on papers that feature clean identification, exploit rich data ideally involving networks, and are recent. The papers I discuss all have the same basic research design: they compare “treated” firms that experience improved access with “control” firms that do not, and assume that “control” firms are not affected by the indirect effects of access. This assumption is defensible because the control firms are sufficiently removed from the treated firms in all contexts. Thus, in the language of the model, the papers compare treated firms under the experiment versus under the baseline.

Although my main focus is domestic production networks, I start with a pioneering article on international trade networks. Atkin, Khandelwal and Osman (2017) conducted a randomized controlled trial to evaluate the causal effect of exporting on firm performance. They worked with a sample of microenterprises: rug producers in Egypt. Their intervention was to create experimental variation in access to buyers from a high-income country for a random subset of these firms. Finding buyers was challenging. The authors relied on the

help of a U.S. NGO and an Egyptian intermediary, and recruited buyers via international gift fairs and by directly connecting the Egyptian intermediary to foreign importers and retailers. It took 2 years to create sustained export orders.

The article shows that treated firms experienced 16-26% gains in profit relative to control firms. In a nice methodological contribution, the authors also directly collected data on product quality by asking for the evaluation of independent experts. The treatment created large gains in quality, but a reduction in output quantity per hours worked. These results are plausibly explained by foreign buyers demanding higher-quality products, which take a longer time to produce. The article further shows that the quality improvements were due, at least in part, to learning-by-exporting. That is, the firms genuinely acquired new skills as a result of the export orders.

These results support the model's implication that access can lead to higher-value partnerships and learning-based quality upgrading. Moreover, the challenges experienced in finding foreign buyers are consistent with large search and matching frictions. However, the article is silent on the aggregate effects of improving access: as shown in Proposition 2, the new partnerships created may have crowded out some existing partnerships of the client firms.

Thus, access frictions seem important in international markets. Are they similarly important in domestic markets? After all, in domestic markets, both trade and cultural barriers are lower. We explored this question in Cai and Szeidl (2018) using an experiment in which we randomized SMEs in Nanchang, a city in China, into small groups whose managers held monthly meetings for one year. Thus, this was an intervention that improved access in *managerial networks*, which are broader than production networks. We find that the meetings increased revenue, profit and other performance measures, as well as a management score. We also document that firms shared business-relevant information with each other. These results are consistent with the importance of access frictions in the domestic context, and with positive spillovers through information sharing. In addition, we find that the treatment

increased firms' number of partners, and directly generated new partnerships in the meetings. These positive effects are consistent with the expanding mechanism identified in the model. However, the paper does not provide evidence on aggregate effects because it does not have information on the prior partners of treated firms.

A remaining question is whether access frictions are important in domestic *production networks*. Three studies explore this using natural experiments. First, Bernard, Moxnes and Saito (2019) study the 2004 opening of a new 127 km portion of a high-speed train line in Japan (Kyushu Shinkansen). The authors use a triple difference-in-differences strategy, in which they compare firms near versus far from the new stations, before versus after the opening of the new train line, in industries with a higher versus a lower share of purchased inputs. They find that firm revenue and revenue total factor productivity (TFPR) increased. They also find that firms reallocated their business partners closer to the newly connected stations, but they do not find impacts on the number of business partners. Second, Alfaro-Ureña, Manelici and Vasquez (2022) evaluate the effect of becoming a supplier to multinationals using VAT transactions data from Costa Rica. This study uses event study regressions to document that domestic suppliers, in response to selling to multinationals, increased employment by 26% and TFP by 4-9%. These changes were accompanied by a drop in suppliers' sales to other buyers, followed by gradual recovery. Four years after the initial event, sales to other suppliers were higher by 20%, and much of this growth was due to the acquisition of better (e.g., larger) buyers. Third, Demir, Javorcik and Panigrahi (2024b) exploit the expansion of fiber optic cables to study the impact of fast internet in Turkey. They find that when two provinces become connected by fast internet, firms in the origin province increase the share of input spending allocated to the destination province, as well as the number of suppliers they have in the destination province.

These studies are consistent with the importance of access frictions even in domestic production networks. Moreover, the null effect on degree in the Bernard et al. (2019) study, and the crowding out effect in the Alfaro-Ureña et al. (2022) study, are consistent with the

presence of capacity constraints and the reallocation mechanism of the model. These results again raise the question of aggregate effects: the newly established links may have crowded out prior partnerships.

If access is a driver of upgrading, then the lack of access should be associated with the lack of upgrading. Bassi, Lee, Peter, Porzio, Sen and Tugume (2023) is an interesting study documenting precisely this point. They show that even relatively large manufacturing firms in Uganda do not feature labor specialization: they resemble a collection of individual entrepreneurs sharing a workspace. They attribute the absence of specialization to the firms receiving too few orders of any given type. This result is broadly consistent with Proposition 4 where upgrading is limited by the number of partners. Finally, I note that technology can reduce the cost of access. Using online meetings to improve access in managerial networks, as shown by Asiedu, Lambon-Quayefio, Truffa and Wong (2023), or in trade networks, as shown by Wiles and Houeix (2025), improves business performance. A key open question is how to exploit technology to improve access at scale at the industry level.

In summary, the prior work provides evidence that spatial or cultural barriers limit firm-to-firm access, and that improving access improves business performance and upgrading, particularly for supplier firms. However, this evidence is silent about the nature of industry equilibrium effects, because it does not examine impacts on untreated firms.

3.2 Industry Level Impacts: Expanding vs Reallocation

The model predicts that improvements in access can shape industry structure in two ways, by either expanding or reallocating the production network. The evidence on this topic is still thin; I now review the existing work, starting with expanding to follow the order of the model, though I note that reallocation has received more attention.

Expanding. In Cai, Lin and Szeidl (2024), we document evidence on access expanding production networks using a field experiment in China. We study the industry producing

the Chinese writing brush. The supply chain of this industry has two main layers. Supplier firms produce the intermediate inputs, which are the brush head and the handle, and client firms assemble the final good and sell on to retailers and consumers. We study an experiment that attempts to create new firm-to-firm links by making referrals between treated suppliers and treated clients. This experiment closely parallels the experiment of the model.

We find that the referrals created new supplier-client links. Importantly, although the referrals did crowd out some existing links, new link creation far exceeded crowding out. Furthermore, untreated firms who lost links made up for them by reconnecting with prior business partners. As a result of these effects, the treatment increased the number of links by about 21%. This evidence is consistent with the expanding channel of access. The referrals also generated 20-35% increases in revenue and profit. These results confirm that improvements in access improve business performance.

We also study the mechanism underlying these gains. We find that treated suppliers increased product quality, while treated clients expanded product variety by introducing a second product or expanding their sales of that second product. Since in this industry second products tend to be of higher quality, these results suggest complementary upgrading between suppliers and clients. Our findings align with the logic of expanding in the model, which predicts that the new links create new business ideas that expand product variety.

We then look at indirect effects. In a result we found initially surprising, we estimate *positive* indirect effects: having a higher share of prior partners treated increased profit and employment. These results are puzzling from the perspective of the reallocation channel, but are consistent with expanding: Proposition 4 predicts positive exposure effects under expanding, because the upgrading created by the treatment should spill over to firms' prior partners. In summary, this paper provides support for the expanding channel, and suggests that through this channel access can generate broadly shared gains.

The reason for broadly shared gains under expanding is the presence of upgrading spillovers in the production network. There is a growing body of evidence on such spillovers.

In early work, Kee (2015) shows that a trade-policy-driven expansion of FDI in the Bangladesh garment sector led to the improved performance of domestic firms that shared the suppliers with foreign producers. More recently, Boehm, Dhingra and Morrow (2022) and Rastogi (2025) study the removal of entry barriers in India, which improved access to high-quality intermediate inputs. Improved access led to increased product variety, and to shifting the product mix towards products with greater scope for quality upgrading. The mechanism underlying such upgrading spillovers could be complementarities, a plausible force in some of these studies, as well as sharing information, as in Chaurey, Nayyar, Sharma and Verhoogen (2025), or sharing machines and tools, as in Bassi, Muoio, Porzio, Sen and Tugume (2022).

In summary, there is evidence consistent with the expanding channel of the model, and evidence suggesting that access-driven expanding can lead to broadly shared gains.

Reallocation. The second channel for industry-level impacts identified by the model is reallocation. This is a relatively more studied channel. In an important article, Jensen and Miller (2018) document evidence for reallocation using the natural experiment of expanding mobile phone service in the boat building industry in Kerala, India. In their setting, boat builders are the suppliers and fishermen are the clients in the fishing industry. At baseline, the boat builders consisted of many small firms serving their local markets: roughly, each landing spot had one associated boat builder. Importantly, there was much quality variation across builders, as measured with the life span of the boat. For example, the best builder made boats that lasted twice as long as those made by the worst builder.

The article studies how the spread of mobile phones improves access between fishermen and boat builders. Mobile phones, as shown by Jensen (2007), enabled fishermen to travel to different markets to sell their fish. This allowed fishermen to meet other fishermen with boats made by different builders, and thus learn about the quality of non-local boat builders. Importantly, because of the staggered introduction of mobile phones across places, this is a setting that has pure control areas where mobile phones were not introduced, to which the outcomes in the treated areas can be compared.

The article finds that improved access between fishermen and boat builders resulted in substantial reallocation in the production network. High-quality boat builders gained market share and size, while low-quality boat builders lost market share and many exited. By the end of the sample period, the number of boat builders shrank by about 60%. Since this is a setting where each client firm requires just one supplier, it corresponds to the low- κ environment of the model. Thus, consistent with the empirical results, the model predicts reallocation. A difference is that the model features match-specific quality, rather than firm-specific quality. Nevertheless, the ideas that improvements lead to reallocating links, and that reallocation creates both winners and losers, are captured by Proposition 2.

The article also shows that the growth of high-quality firms was accompanied by improvements in productivity and increased labor specialization. For example, at baseline the average worker performed 7-8 major tasks, while at endline half as many. These results are consistent improved access generating upgrading, as characterized by Proposition 4.

This article features a setting with a technological constraint on the number of suppliers—a fisherman needs only one boat—which naturally invites reallocation. A similar constraint applies in trade networks over homogenous goods, where the buyer mainly cares about obtaining the good at the lowest price, suggesting that reallocation may emerge in such settings too. Bergquist, McIntosh and Startz (2024) study one such setting: the trade of agricultural goods in Uganda. They use an RCT to evaluate the impact of a mobile platform, Kudu, that links potential buyers and sellers. The platform was randomized across 110 subcounties with a population of roughly 3 million people. Given the at-scale nature of the intervention, this is a setting which lacks a pure control arm.

Comparing treated and untreated local markets, Bergquist et al. (2024) find that there is more trade and more price convergence of the goods in treated than in untreated markets. However, the paper emphasizes that—as Proposition 2 also shows—this “naive” comparison overestimates the true impact of access in the presence of reallocation. The paper presents exposure regressions, similar to (5), to document these negative reallocation effects.

As expected, exposure effects reflect the reallocation of trade.³ Moreover, consistent with equation (5) and its underlying logic, the reallocation effects are stronger for untreated markets, plausibly because treated markets have more trade opportunities and consequently are less affected by exposure. The paper does not discuss impacts on upgrading by farmers or traders, which could be an interesting direction of future research. In summary, the results align with the predictions of the model under the reallocation channel.

I conclude that there is evidence on both the expanding and reallocation channel of improved access. A key question is what are the aggregate effects of improved access under these channels.

3.3 Measuring Aggregate Impacts

Accounting for aggregate effects is challenging due to indirect effects. I discuss three different ways of dealing with this challenge.

I begin with a reduced-form approach that gives a welfare bound under the expanding channel. This is the approach we use in our paper on the industry of the Chinese writing brush (Cai et al. 2024). Because in our setting—due to expanding—indirect effects are mildly positive, ignoring them makes the welfare evaluation conservative. This simplifies the analysis, because it shuts down business stealing in the supplier-to-client market. There is still a possibility of business stealing in the client-to-consumer market, but we find no evidence for such an effect and make the strong assumption that it can be ignored. Then, we can read off the producer surplus directly from the treatment effect regressions. We also account for the impact on the consumer surplus using a simple model of consumer demand. Intuitively, treated client firms experience revenue gains if consumers switch to their products. This switch reflects gains to the consumer surplus, but the magnitude of this gain depends on the elasticity of substitution between products (Feenstra 1994). Using this

³ In their setting, the sign of the effect is influenced by whether the partner market is in surplus or deficit for the crop, but the signs are always consistent with the logic of reallocation.

approach, we estimate annual welfare gains, including both the producer and the consumer surplus, that are several times larger than the cost of the intervention.

A second approach is to match a structural model to the data. Using a model-based approach seems essential for studying the reallocation channel. The Bergquist et al. (2024) paper on Ugandan agricultural markets is a beautiful demonstration of how such an approach can work. Their approach is to estimate a structural model, and then use the model to simulate the counterfactual in which no intervention takes place, i.e., the pure control arm. Comparing the simulated outcomes under the scenario capturing the actual intervention and under the scenario of the pure control reveals both the direct and the indirect effect of the intervention.

To implement this approach, the authors write down a fully specified trade model in which traders decide which markets to serve. The platform is modeled as a reduction in the fixed cost of serving a market. This model is matched to the data, in a way that takes advantage of both the treatment and exposure regressions discussed above. The authors evaluate the model's goodness of fit by simulating the actual intervention in the estimated model, and demonstrating that the treatment effect regressions in the actual and the simulated data align closely. They then use the model for aggregate evaluation. They find meaningful gains: for example, overall trade volume increases by about 1%. However, these estimates are smaller than what would be obtained under a naive evaluation: for example, the naive impact on trade volume would be 154% larger than the actual impact. The paper finds that overall welfare improves by about 0.02%. This is small both because maize, the good they study, is a small share of total consumption, and because under the reallocation channel access creates both winners and losers.

Demir et al. (2024b), in their study of the impact of fast internet in Turkey, present another model-based approach to evaluate equilibrium effects. At the level of the firm, their model only allows for the reallocation but not the expanding channel, because firms combine a set of pre-defined tasks to produce a final good. But at the level of a firm buying from a

specific province, their model allows for both channels in that firms may both expand and reallocate the set of suppliers in the province. They estimate the model exploiting variation in supplier choice across provinces driven by the introduction of fast internet. Armed with the estimated model, they then compare outcomes in simulated economies with versus without the introduction of fast internet, and estimate large welfare gains of about 2%.

Finally, I discuss a third “growth accounting” approach that can assess aggregate impacts under both expanding and reallocation, but does not permit counterfactual analysis. Baqaee, Burstein, Duprez and Farhi (2023) investigate the contribution of “supplier churn”—changes in the set of firms’ suppliers—to GDP growth. The idea of their approach is to express GDP growth in an accounting identity as a function of multiple terms corresponding to improvements in the performance of individual firms. One of these terms captures supplier churn, i.e., changes in the set of a firm’s suppliers. These changes affect GDP growth because suppliers are assumed to affect the firm’s marginal cost. Using this decomposition requires an estimate of the impact of supplier additions and separations on a firm’s marginal cost. The paper obtains this estimate using data from Belgium, exploiting the entry and exit of suppliers. This approach accounts for both expanding and replacing suppliers, i.e., both the expanding and the reallocation effects. The paper finds that supplier churn accounts for roughly 50% of productivity growth in Belgium, a large effect suggesting that firm-to-firm access is potentially important for macro outcomes.

In summary, although there is still much we do not know, the evidence suggests that firm-to-firm access can substantially improve aggregate outcomes.

3.4 Nature of the Friction

If access indeed generates large gains, then a key question is why firms do not obtain access themselves. Proposition 3 highlights two qualitatively different answers, with different policy implications. One answer is that firms do not gain access because of a rational cost-benefit

calculation, as a result of high search costs. Then, the case for policies hinges on inefficiencies outside of firm optimization, with the search externality identified in the model being a prime candidate; and policies should aim to reduce search costs. Another answer is that firms do not gain access because of a firm-level friction, such as pessimism in the model. Then, policies may need to target the within-firm friction.

External barriers. Search costs are a key external barrier. These costs may be high because of lack of information about potential partners, including both the identity and the quality of these partners. Concerning the identity of partners, Bergquist et al. (2024) compare the impact of Kudu, the platform that connects potential buyers and sellers, with that of a treatment arm providing only price information. Unlike Kudu, providing only price information does not lead to price convergence. Concerning the quality of partners, the evidence in Jensen and Miller (2018) about switching to better boat builders is strongly suggestive of this barrier. Further evidence comes from Startz (2024) who shows that Nigerian importers can access information about new product variety sooner by visiting their suppliers in China. And randomized evidence comes from Wiles and Houeix (2025), who show that information in WhatsApp groups of Senegalese importers about the identity of potential suppliers in Turkey contribute to lasting partnerships. Lack of information is a plausible barrier in all these contexts because suppliers and clients are separated by distance and borders.

A second external barrier is an enforcement friction: the supplier may not deliver the good that the client has paid for. This friction can reduce the incentive to search. A growing literature, beginning with McMillan and Woodruff (1999) and reviewed in Macchiavello (2022), studies enforcement frictions and relational contracts between firms in the development context. In a recent contribution, Boehm, South, Oberfield and Waseem (2024) document in data from India and Pakistan that contracting frictions, measured with court congestion, increase relationship duration in relationship specific industries, suggesting that these frictions limit search. And Startz (2024), by structurally estimating a model of relational trade,

finds that the enforcement friction is as important as the information friction in explaining why Nigerian importers visit their suppliers in China. Plausibly, like the search friction, the enforcement friction is also relatively more important for longer-distance partnerships.

Within-firm frictions. There are many important contexts where potential suppliers and clients are not spatially far from each other, and are not separated by borders or cultural barriers. A leading example is firms in a city. In such contexts, information and enforcement frictions seem less binding. Are there other, potentially within-firm frictions that limit access?

In Cai et al. (2024) we provide evidence for pessimistic beliefs as the friction. First, we rule out the information friction by showing that providing information about the identity of potential partners does not lead to subsequent transactions, but subsidizing a first transaction does. Second, we show that firms’ beliefs about the value of partners causally increased as a result of the treatment. Such an increase is inconsistent with Bayesian updating under a well-calibrated prior, because then the posterior should on average agree with the prior. But it is consistent with miscalibrated and pessimistic beliefs. Third, we show that in response to the treatment, firms increased their beliefs about the value of search, the amount of time they allocated to search, and the number of partners they obtained outside of the experiment. Thus, firms seem to have held excessively pessimistic beliefs before the experiment. Importantly, such beliefs are consistent with a self-confirming equilibrium: pessimistic beliefs prevent firms from seeking out potential partners, and thus from finding out that their beliefs are miscalibrated.

Our study also suggests a second within-firm friction: the lack of marketing capability. Asking firms in the information treatment why they did not engage with the referred partners led to two main responses. One was that they did not think the partners would be sufficiently good (consistent with pessimism); the other that firms—especially suppliers, which are on average smaller—did not know how to contact the potential partners. The second answer suggests that some firms lacked the marketing capability to contact potential

clients. Consistent with this intuition, Hjort, de Rochambeau, Iyer and Ao (2024) show that teaching Liberian firms how to market their products to large buyers generates meaningful impacts on firm performance. Treated firms win three times as many contracts as control firms; and the quarter of firms for whom impacts are strongest, three years after the intervention, continue to win attractive contracts and have higher employment. These results confirm that the lack of marketing capabilities is another barrier.

I conclude that there is evidence on both search costs and within-firm frictions acting as barriers to access. But I note that in a sense all of these explanations are incomplete. Concerning search costs and marketing capabilities, it is often unclear why market forces do not overcome these frictions. For larger firms, search costs do not appear to be prohibitively high; and for smaller firms, intermediaries could step in and provide matching services. Concerning pessimism, although it can survive in a self-confirming equilibrium, it is unclear how it emerges, and which firms it most likely impacts.

4 Conclusion: Open Questions

I conclude by discussing four research areas I find promising.

First, the evidence is thin on core conceptual issues. I highlight three such issues. (a) We need a better understanding of when access generates expansion versus reallocation, and more generally a better understanding of industry level impacts. (b) The sources of frictions. We lack an understanding of why some barriers, such as pessimism, emerge; and why other barriers, such as lack of information, are not overcome by intermediaries who connect firms for a fee. (c) Targeting. We do not know which firms or industries benefit the most from improving access, and how to identify them.

Second, and closely related to the first point, we need evidence on what policies are effective in improving access at scale. How should we build better firm-to-firm markets? I see promise in three approaches. (a) Use new data and technology. For example, banks

could exploit digital data to screen better and expand lending to SMEs. And governments or private enterprises could create online platforms to reduce search costs. (b) Create quality certification to reduce matching frictions. A trusted certification system could reduce uncertainty and help overcome the pessimism barrier. Such a certification could be combined with an online platform. (c) Provide skills that reduce matching costs.

Third, to assess the aggregate impact of improving access, we need to connect the evidence to quantitative models. An emerging line of research builds equilibrium models of production networks that explicitly incorporate search, such as Demir et al. (2024a) and Arkolakis et al. (2025). We need to connect these sort of models with well-identified evidence to derive aggregate implications. Counterfactual analysis in these models could identify new types of policies to be evaluated using RCTs.

Fourth, we lack evidence on the aggregate importance of firm-to-firm access for development. To what extent does lack of access explain why firms in developing countries stay small? How much does access contribute to differences in firm performance across places and over time? Answering these questions likely requires data on production networks across different levels of development.

References

- 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.
- Arkolakis, Costas, Federico Huneus, and Yuhei Miyauchi**, “Production Network Formation, Trade, and Welfare,” Technical Report, CESifo Working Paper 2025.
- Asiedu, Edward, Monica Lambon-Quayefio, Francesca Truffa, and Ashley Wong**, “Female Entrepreneurship and Professional Networks,” Working Paper 2023.
- 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.

- Baqae, David, Ariel Burstein, Cédric Duprez, and Emmanuel Farhi**, “Supplier churn and growth: a micro-to-macro analysis,” Technical Report, National Bureau of Economic Research 2023.
- Bassi, Vittorio, Jung Hyuk Lee, Alessandra Peter, Tommaso Porzio, Ritwika Sen, and Esau Tugume**, “Self-employment within the firm,” Technical Report, National Bureau of Economic Research 2023.
- , **Raffaella Muoio, Tommaso Porzio, Ritwika Sen, and Esau Tugume**, “Achieving scale collectively,” *Econometrica*, 2022, 90 (6), 2937–2978.
- Bergquist, Lauren, Craig McIntosh, and Meredith Startz**, “Search Costs, Intermediation, and Trade: Experimental Evidence from Ugandan Agricultural Markets,” Working Paper, Yale 2024.
- Bernard, Andrew B., Andreas Moxnes, and Yukiko U. Saito**, “Production Networks, Geography, and Firm Performance,” *Journal of Political Economy*, 2019, 127 (2), 639–688.
- Boehm, Johannes, Ruairidh South, Ezra Oberfield, and Mazhar Waseem**, “The Network Origins of Firm Dynamics: Contracting Frictions and Dynamism with Long-Term Relationships,” Technical Report 2024.
- , **Swati Dhingra, and John Morrow**, “The comparative advantage of firms,” *Journal of Political Economy*, 2022, 130 (12), 3025–3100.
- 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*, 2024.
- , **Wei Lin, and Adam Szeidl**, “Firm-to-firm Referrals,” Technical Report, National Bureau of Economic Research 2024.
- Chaurey, Ritam, Gaurav Nayyar, Siddharth Sharma, and Eric Verhoogen**, “Social Learning among Urban Manufacturing Firms: Energy-Efficient Motors in Bangladesh,” Technical Report, National Bureau of Economic Research 2025.
- 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?,” Technical Report, Working paper 2024.

- Donaldson, Dave**, “Railroads of the Raj: Estimating the Impact of Transportation Infrastructure,” *American Economic Review*, April 2018, 108 (4-5), 899–934.
- **and Richard Hornbeck**, “Railroads and American Economic Growth: A “Market Access” Approach *,” *The Quarterly Journal of Economics*, 02 2016, 131 (2), 799–858.
- Eaton, Jonathan, Samuel S Kortum, and Francis Kramarz**, “Firm-to-firm trade: Imports, exports, and the labor market,” Technical Report, National Bureau of Economic Research 2022.
- Faber, Benjamin**, “Trade integration, market size, and industrialization: evidence from China’s National Trunk Highway System,” *Review of Economic Studies*, 2014, 81 (3), 1046–1070.
- Feenstra, Robert C**, “New product varieties and the measurement of international prices,” *The American Economic Review*, 1994, pp. 157–177.
- Grossman, Gene M and Elhanan Helpman**, *Innovation and growth in the global economy*, MIT press, 1993.
- Grossman, Gene M., Elhanan Helpman, and Stephen J. Redding**, “When Tariffs Disrupt Global Supply Chains,” *American Economic Review*, April 2024, 114 (4), 988–1029.
- Hjort, Jonas, Golvine de Rochambeau, Vinayak Iyer, and Fei Ao**, “Informational Barriers to Market Access: Experimental Evidence from Liberian Firms,” Working Paper, University College, London 2024.
- Hornbeck, Richard and Martin Rotemberg**, “Growth Off the Rails: Aggregate Productivity Growth in Distorted Economies,” Working Paper, NYU 2024.
- J-PAL, Abdul Latif Jameel Poverty Action Lab**, “Market Access: Connecting Firms and Entrepreneurs to Markets to Spur Business and Job Growth,” *J-PAL Policy Insights*, September 2024.
- 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.

- Kee, Hiau Looi**, “Local intermediate inputs and the shared supplier spillovers of foreign direct investment,” *Journal of Development Economics*, 2015, *112*, 56–71.
- Macchiavello, Rocco**, “Relational contracts and development,” *Annual Review of Economics*, 2022, *14* (1), 337–362.
- McKenzie, David, Christopher Woodruff, Kjetil Bjorvatn, Miriam Bruhn, Jing Cai, Juanita Gonzalez-Uribe, Simon Quinn, Tetsushi Sonobe, and Martin Valdivia**, “Training Entrepreneurs,” *VoxDevLit*, September 2023, *1* (3).
- McMillan, John and Christopher Woodruff**, “Interfirm Relationships and Informal Credit in Vietnam,” *Quarterly Journal of Economics*, 1999, *114*, 1285–1320.
- Oberfield, Ezra**, “A theory of input–output architecture,” *Econometrica*, 2018, *86* (2), 559–589.
- Rastogi, Chhavi**, “Investment along the supply chain: removing growth restrictions on firms in India,” Technical Report 2025.
- Startz, Meredith**, “The value of face-to-face: Search and contracting problems in Nigerian trade,” Working Paper, Dartmouth College 2024.
- Verhoogen, Eric**, “Firm-Level Upgrading in Developing Countries,” *Journal of Economic Literature*, December 2023, *61* (4), 1410–64.
- Wiles, Edward and Deivy Houeix**, “Relational Frictions Along the Supply Chain: Evidence from Senegalese Traders,” Technical Report 2025.

Appendix

A1 Proofs

Proof of Proposition 1.

For $\lambda < \kappa$, firms take all partners. The surplus accumulating to each client firm is

$$V_i = (1/2)\lambda \int_0^\infty xg(x)dx = \lambda g_H/4$$

where $1/2$ is the share of the surplus they keep. All firms have degree λ . The total surplus in the industry W is twice the average surplus because the mass of firms is 2, i.e., is $\lambda g_H/2$. The effect of increasing λ by $\delta = \gamma\lambda$ on the log surplus is

$$\log W(\delta + \lambda) - \log W(\lambda) = \log(\delta + \lambda) - \log(\lambda) \approx \frac{\delta}{\lambda} = \gamma.$$

to a first-order approximation when δ (or δ/λ) is small.

For $\lambda > \kappa$, firms take the top κ/λ share of partners. The surplus for each client firm is

$$V_i = (1/2)\kappa \int_{G^{-1}(1-\kappa/\lambda)}^\infty xg(x)dx = \frac{\kappa}{2} \left(1 - \frac{\kappa}{2\lambda}\right) g_H.$$

All firms have degree κ . The total surplus in the industry is

$$W(\lambda) = \kappa g_H \left(1 - \frac{\kappa}{2\lambda}\right).$$

For $\lambda = \kappa$ this collapses to the previous welfare expression. The effect of increasing λ by δ

on the log surplus is

$$\begin{aligned}\log\left(1 - \frac{\kappa}{2(\lambda + \delta)}\right) - \log\left(1 - \frac{\kappa}{2\lambda}\right) &\approx \log\left(1 - \frac{\kappa}{2\lambda} + \frac{\kappa}{2\lambda^2}\delta\right) - \log\left(1 - \frac{\kappa}{2\lambda}\right) \\ &\approx \log\left(1 + \frac{\kappa}{(2\lambda - \kappa)\lambda}\delta\right) \approx \frac{\kappa}{2\lambda - \kappa} \frac{\delta}{\lambda} = \frac{\kappa}{2\lambda - \kappa} \gamma.\end{aligned}$$

Proof of Proposition 2.

Case 1. If $\lambda + \delta < \kappa$, then all new links are accepted. We can write degree compactly as

$$d(T, S) = \lambda + T\delta$$

and log degree as

$$\log d(T, S) = \log(\lambda + T\delta) \approx \log d_0(T, S) + T \cdot \frac{\delta}{\lambda} = \log d_0 + T \cdot \gamma$$

so that

$$\Delta \log d(T, S) \approx T \cdot \gamma.$$

We can write the surplus of a firm as

$$V(T, S) = (1/2)g_H \frac{\lambda}{2} + T(1/2)g_H \frac{\delta}{2} \tag{A1}$$

and log surplus as

$$\log V(T, S) \approx \log V_0 + T \frac{\delta}{\lambda} = \log V_0 + T \cdot \gamma$$

so that

$$\Delta \log V(T, S) \approx T \cdot \gamma.$$

Total surplus is

$$W_1 = g_H \frac{\lambda}{2} + g_H \frac{\delta}{4}$$

and the log surplus is

$$\log W_1 \approx \log W_0 + \frac{\gamma}{2}$$

because only half of the firms experience the treatment.

Case 2: $\kappa < \lambda$. Now there is crowding out. Treated firms take κ links among their $\lambda + \delta$ opportunities. This has two implications.

1. Each potential link of a treated firm is kept with probability $\kappa/(\lambda + \delta)$.
2. The links held by a treated firm have quality above the threshold $g_H (1 - \frac{\kappa}{\lambda + \delta})$.

The first observation implies that among the δ new opportunities of a client, $\delta\kappa/(\lambda + \delta)$ links are kept, which means that among the κ pre-existing links of that client, this many links (which have quality below the new threshold) are dropped. The total volume of links dropped by the set of clients is then $\delta\kappa/2(\lambda + \delta)$ because only half of clients are treated. Half of this is links that connected treated clients to treated suppliers. But treated suppliers will continue to have degree κ because they have the new links; in fact, by the second observation, they also wanted to drop these (lower-quality) links. The other half of these links are to untreated suppliers, and they get dropped. Thus, each untreated supplier loses $\delta\kappa/2(\lambda + \delta)$ links and ends up with degree $\kappa(1 - \delta/2(\lambda + \delta))$. Note, here we are assuming that firms cannot go back to their previous offers in the short run.

As a result of these changes, the shape of the network is the following. Treated clients have κ suppliers, and a share $(\lambda/2 + \delta)/(\lambda + \delta)$ is with treated suppliers. Untreated clients have $\kappa(1 - \delta/2(\lambda + \delta))$ suppliers, and a share $(1/(2(1 - \delta/2(\lambda + \delta))))$ is with untreated suppliers. The patterns for suppliers are completely analogous. Thus, treated firms partner more with treated, and in consequence untreated firms partner more with untreated.

To express degree as a function of exposure, suppose that share S of a firm's partners are treated. As we discussed above, in practice $S = 0.5$ always, but we can entertain the possibility of any S . We can write the number of partners more compactly as

$$d(T, S) = \kappa \left(1 - S(1 - T) \frac{\delta}{\delta + \lambda} \right) \approx \kappa - S(1 - T) \frac{\kappa}{\lambda} \delta.$$

To understand the first expression, note that for a treated firm ($T = 1$) degree is κ independently of exposure. This is because, by observation 2, this firm keeps only links with quality above the new threshold of $g_H \left(1 - \frac{\kappa}{\lambda + \delta} \right)$, and in this decision rule it is in agreement with all treated partners. In contrast, for an untreated firm ($T = 0$), exposure S reduces degree because, by observation 1, a share $\frac{\delta}{\delta + \lambda}$ of preexisting links with each treated partner are dropped. The second expression in the formula follows from a first-order approximation in δ .

We can write log degree as

$$\log d(T, S) \approx \log d_0 - S(1 - T) \frac{\delta}{\lambda} = \log d_0 - S(1 - T) \cdot \gamma.$$

We now turn to the surplus as a function of exposure. Consider the average quality of the links lost. By observation 1, untreated firm loses a share $\delta/(\lambda + \delta)$ of its links to treated firms. The range of qualities of prior links is $[g_H(1 - \kappa/\lambda), g_H]$. Of this, the bottom share is lost, so the remaining range has length $\lambda/(\lambda + \delta)(\kappa/\lambda)g_H$, that is $\kappa g_H/(\lambda + \delta)$. Thus, the average quality of links with treated partners is $g_H(1 - \kappa/2(\lambda + \delta))$. The average quality of links with untreated partners is unchanged at $g_H(1 - \kappa/2\lambda)$.

These considerations imply that the surplus of untreated firms as a function of exposure

is

$$\begin{aligned}
V(0, S) &= (1/2)g_H\kappa(1 - S) \left(1 - \frac{\kappa}{2\lambda}\right) + (1/2)g_H\kappa S \left(1 - \frac{\delta}{(\lambda + \delta)}\right) \left(1 - \frac{\kappa}{2(\lambda + \delta)}\right) \\
&= (1/2)g_H\kappa \left(1 - \frac{\kappa}{2\lambda}\right) - S(1/2)g_H\kappa \left[\left(1 - \frac{\kappa}{2\lambda}\right) - \left(1 - \frac{\delta}{(\lambda + \delta)}\right) \left(1 - \frac{\kappa}{2(\lambda + \delta)}\right) \right] \\
&\approx (1/2)g_H\kappa \left(1 - \frac{\kappa}{2\lambda}\right) - S(1/2)g_H\kappa \frac{\lambda - \kappa}{\lambda^2} \delta.
\end{aligned}$$

The surplus of treated firms does not depend on exposure, because they drop links below the new high cutoff (of observation 2) independently of exposure, i.e., there is agreement on which links to drop. Their surplus is

$$V(1, S) = (1/2)g_H\kappa \left(1 - \frac{\kappa}{2(\lambda + \delta)}\right) \approx (1/2)g_H\kappa \left(1 - \frac{\kappa}{2\lambda}\right) + g_H \left(\frac{\kappa^2}{4\lambda^2}\right) \delta$$

where the last expression is a first-order approximation in δ .

It follows that we can write the surplus of a firm more compactly as

$$V(T, S) \approx (1/2)g_H\kappa \left(1 - \frac{\kappa}{2\lambda}\right) + Tg_H \left(\frac{\kappa^2}{4\lambda^2}\right) \delta - S(1 - T)(1/2)g_H\kappa \frac{\lambda - \kappa}{\lambda^2} \delta. \quad (\text{A2})$$

The log surplus is

$$\begin{aligned}
\log V(T, S) &\approx \log V_0 + T \cdot \frac{\kappa}{2\lambda - \kappa} \frac{\delta}{\lambda} - S(1 - T) \frac{2\lambda - 2\kappa}{2\lambda - \kappa} \frac{\delta}{\lambda} \\
&= \log V_0 + T \cdot \frac{\kappa}{2\lambda - \kappa} \gamma - S(1 - T) \cdot \frac{2\lambda - 2\kappa}{2\lambda - \kappa} \gamma.
\end{aligned}$$

Total surplus is

$$\begin{aligned} W_1 &\approx g_H \kappa \left(1 - \frac{\kappa}{2\lambda}\right) + g_H \left(\frac{\kappa^2}{4\lambda^2}\right) \delta - (1/4)g_H \kappa \frac{\lambda - \kappa}{\lambda^2} \delta \\ &= g_H \kappa \left(1 - \frac{\kappa}{2\lambda}\right) + g_H \frac{\kappa}{4\lambda} (2\kappa - \lambda) \frac{\delta}{\lambda}. \end{aligned}$$

The log surplus is

$$\log W_1 \approx \log W_0 + \frac{1}{2} \frac{2\kappa - \lambda}{2\lambda - \kappa} \gamma.$$

Proof of Proposition 3.

We compute the client-level private and social gains from search as the per client firm values of the analogous gains obtained when administering the intervention.

Case 1. Given equation (A1), search creates private value $g_H \delta / 4$. If firms perceive g_H as βg_H , they will search when $c g_H / 2 > \beta g_H \delta / 4$ or $c / \delta > \beta / 2$, whereas they should search when $c / \delta > 1 / 2$, proving the first two inequalities. Also by (A1), since new links create the same value to suppliers, search creates social value $g_H \delta / 2$, leading to the third inequality.

Case 2. Given equation (A2), search creates private value $g_H \delta \kappa^2 / (4\lambda^2)$. But if firms underestimate g_H by β , then they have a different view of the private value, because they think that those of their current links that have value above βg_H are unreplaceable. Thus, effectively, they think that access can only lead to replacing a share of their links. And firms believe that the same holds for all other firms in the industry. This (imagined) setting is equivalent to a model in which the distribution of values is bounded by βg_H and the capacity constraint of firms is lower, corresponding to only those links that they consider replaceable. Specifically, a firm's baseline links have value above $(1 - \kappa / \lambda) g_H$ and the firm perceives the distribution of outside values to be bounded from above by βg_H . Thus, only links with value in $[(1 - \kappa / \lambda) g_H, \beta g_H]$ are replaceable. The share of these links among the κ links of the firm is $(\beta - 1 + (\kappa / \lambda)) / (\kappa / \lambda)$, so that the mass of these links is κ times this share, or $\kappa - \lambda(1 - \beta)$.

This is the “effective capacity constraint” of the firms with biased beliefs.

With this capacity constraint, and with the biased belief about g_H , the perceived private value created by search is

$$\beta g_H \delta \frac{(\kappa - \lambda(1 - \beta))^2}{4\lambda^2} = \beta \left(1 - \frac{\lambda}{\kappa}(1 - \beta)\right)^2 g_H \delta \frac{\kappa^2}{4\lambda^2}. \quad (\text{A3})$$

The firm compares the search cost $cg_H/2$ with this, leading to the first inequality of the Proposition. It should compare the search cost to the value of this for $\beta = 1$, leading to the second inequality. Finally, by (A3), the social value of search is given by $g_H \delta \kappa^2 / (4\lambda^2)(2\kappa - \lambda)/\kappa$. We have to normalize this by the mass of client firms treated, which is $1/2$. This gives the third inequality in the Proposition.

A2 Model with upgrading

Setup. Each firm learns by doing: degree d improves idea quality by ad . Thus, a firm that has degree d , such that its partners have average degree d' , gains $(a/2)(d^2 + dd')$. 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. We 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.

We now explore the effect of the experiment in the presence of upgrading. Importantly, we assume that 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 allows us to work with 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.

We 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.*

1. For $\lambda < \kappa$, log firm surplus is

$$\log V(T, S) \approx \log V_0 + T \cdot \gamma(1 + \alpha\lambda) + S \cdot \gamma\alpha\lambda. \quad (\text{A4})$$

2. For $\lambda > \kappa$, log firm surplus is

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

In Case 1, two things change. 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 2, 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 . As noted in the text, this effect emerges because the untreated peers of these firms lose partners, which reduces their upgrading, negatively impacting the firm. 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.

Proof. Case 1. The surplus of a firm is

$$V(T, S) = \frac{g_H \lambda}{4} + T \cdot \frac{g_H \delta}{4} + 2 \frac{a}{2} \lambda^2 + T \cdot \frac{a}{2} ((\lambda + \delta)^2 - \lambda^2 + \delta(\lambda + \delta)) + S \cdot \frac{a}{2} \lambda ((\lambda + \delta) - \lambda).$$

The first two terms are the surplus from the baseline model. The next term captures the effect of learning for a firm that experiences no treatment or exposure. This firm has degree λ with peers having degree λ , yielding $(a/2)\lambda^2$ from own learning and $(a/2)\lambda^2$ from peer learning. The next term represents the gain from the treatment. Treated firms' own learning increases to $(a/2)(\lambda + \delta)^2$, and their peer learning over their δ new links to treated firms increases by $(a/2)\delta(\lambda + \delta)$. The last term represents the effect of baseline exposure. Exposure does not change the number of partners and own learning, but because treated peers invest more, it changes peer learning. The number of treated baseline peers is $S\lambda$, and they increase learning from λ to $\lambda + \delta$.

Simplifying and taking a first-order approximation in δ

$$V(T, S) \approx \frac{g_H \lambda}{4} + a\lambda^2 + T \cdot \left(\frac{g_H \delta}{4} + \frac{a}{2} 3\lambda\delta \right) + S \cdot \frac{a}{2} \lambda\delta.$$

In this model version $V_0 = g_H \lambda / 4 + a\lambda^2$. We can write

$$\log V(T, S) \approx \log V_0 + T \cdot \left(\gamma + \frac{(a/2)\lambda\delta}{\frac{g_H \lambda}{4} + a\lambda^2} \right) + S \cdot \frac{(a/2)\lambda\delta}{\frac{g_H \lambda}{4} + a\lambda^2}.$$

Note that

$$\frac{(a/2)\lambda\delta}{\frac{g_H \lambda}{4} + a\lambda^2} = \frac{(\alpha g_H / 4)\lambda\delta}{\frac{g_H \lambda}{4} + \alpha(g_H / 2)\lambda^2} = \frac{\alpha\lambda}{1 + 2\alpha\lambda} \cdot \frac{\delta}{\lambda} \approx \alpha\lambda\gamma$$

where at the last step we took a first-order approximation in α . With this approximation, we assume that α is small, but of a higher order of magnitude than δ , because we are ignoring terms of order δ^2 but not terms of order $\alpha\delta$. This means that we should not ignore terms of order α^2 .

Using this approximation, we can write

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

Case 2. We compute the effects of own learning and peer learning separately. Under the intervention, the degree of a firm with treatment T and exposure S is $\kappa(1 - S(1 - T)\gamma)$, so that the own learning effect is

$$\frac{a}{2}\kappa^2(1 - S(1 - T)\gamma)^2 \approx \frac{a}{2}\kappa^2 - a\kappa^2 S(1 - T)\gamma.$$

To understand the effect of peer learning, first note that the number of the firm's treated peers is

$$T \frac{\delta\kappa}{\delta + \lambda} + S \frac{\lambda\kappa}{\delta + \lambda}$$

and the number of the firm's untreated peers is

$$(1 - S)\kappa \left(1 - T \frac{\delta}{\delta + \lambda}\right).$$

The first expression follows because, as our earlier calculations show, out of the $T\delta$ newly accesses treated potential partners, a share $\kappa/(\lambda + \delta)$ is kept, while out of the $S\kappa$ baseline partners who are treated, a share $\lambda/(\delta + \lambda)$ is kept. The second expression follows because an untreated firm keeps all its $(1 - S)\kappa$ untreated baseline peers, while a treated firm only keeps a share $\lambda/(\delta + \lambda)$.

Each treated peer has degree κ generating learning of $a\kappa$. Each untreated peer has degree $\kappa(1 - \delta/(2\lambda))$ generating learning of $a\kappa(1 - \delta/(2\lambda))$. Aggregating across all peers, adjusting by the surplus sharing, and using a first-order approximation, the peer learning effect is

$$\frac{a}{2} \left(T \frac{\delta\kappa^2}{\delta + \lambda} + S \frac{\lambda\kappa^2}{\delta + \lambda} \right) + \frac{a}{2}\kappa^2 \left(1 - \frac{\delta}{2\lambda}\right) (1 - S)(1 - T\gamma).$$

Adding up own learning and peer learning, and collecting terms, yields after some calculations the total learning effect of (approximately)

$$a\kappa^2 - \frac{a}{4}\kappa^2\gamma - S(1-T)\frac{3}{2}a\kappa^2\gamma.$$

Here the second term is conceptually new. It is the additional effect of learning on a firm with no treatment or exposure in the presence of the intervention. This effect is negative: the κ untreated peers of such a firm reduce their degree from κ by $\kappa\delta/(2\lambda)$, reducing peer learning. This effect was absent in prior model versions in which a firm without treatment or exposure had the same outcomes as a firm absent the intervention. Also note that there is no direct treatment effect. Treated firms continue to have degree κ so they do not reduce own investment. And while the composition of their peers (in terms of treated or untreated) changes with exposure, this effect is of order δ , and the investment gap between treated and untreated peers is also of order δ , so that the reallocation has second-order effects on treated firms.

To compute the log surplus, note that the baseline surplus absent the intervention now increases by $a\kappa^2$. Using this new baseline surplus to normalize the intervention effect yields

$$\frac{a/4 \cdot \kappa^2\gamma}{(g_H/2)\kappa(1 - \kappa/(2\lambda)) + a\kappa^2} \approx \frac{\alpha}{2} \frac{\kappa\lambda}{2\lambda - \kappa} \gamma$$

where we used an approximation in that we ignored a term of order $\alpha^2\gamma$.

To adjust the formulas from Proposition 2 using the new baseline surplus, we normalize the old baseline surplus with the new one:

$$\frac{(g_H/2)\kappa(1 - \kappa/(2\lambda))}{(g_H/2)\kappa(1 - \kappa/(2\lambda)) + a\kappa^2} \approx 1 - 2\alpha\kappa \frac{\lambda}{2\lambda - \kappa}.$$

Here in the approximation we ignored a term of order α^2 . We should not ignore such terms, but we only use this expression to adjust other terms of order γ and we assume $\gamma\alpha^2$ terms

can be ignored.

We can now write the log surplus as

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