

PRODUCTION NETWORKS, GEOGRAPHY AND FIRM PERFORMANCE*

Andrew B. Bernard[†]

Tuck School of Business at Dartmouth, CEPR & NBER

Andreas Moxnes[‡]

Dartmouth College, CEPR & NBER

Yukiko U. Saito[§]

RIETI

This Version: July 2014

Abstract

This paper examines the importance of buyer-supplier relationships, geography and the structure of the production network in firm performance. We develop a simple model where firms can outsource tasks and search for suppliers in different locations. Firms located in close proximity to other markets, and firms that face low search costs, will search more and find better suppliers. This in turn drives down the firm's marginal production costs. We test the theory by exploiting the opening of a high-speed (Shinkansen) train line in Japan which lowered the cost of passenger travel but left shipping costs unchanged. Using an exhaustive dataset on firms' buyer-seller linkages, we find significant improvements in firm performance as well as creation of new buyer-seller links, consistent with the model.

Keywords: production networks, trade, productivity, infrastructure.

JEL codes: F14, D22, D85, L10, L14, R12.

*Thanks go to Adam Kleinbaum, Yaniv Dover, and Jonathan Dingel as well as seminar participants at DINR, NOITS, RIETI, World Bank and EITI-Phuket for helpful comments. We gratefully acknowledge the financial support from the Japan Society for the Promotion of Science (No. 25780181). We thank the Center for Spatial Information Science, University of Tokyo (CSIS) for their address matching service. This study is conducted as a part of the Project "Inter-organizational and Inter-inventors Geographical Proximity and Networks" undertaken at the Research Institute of Economy, Trade and Industry (RIETI).

[†]100 Tuck Hall, Hanover, NH 03755, USA; *tel:* +1 603 646 0302; *email:* andrew.b.bernard@tuck.dartmouth.edu

[‡]328 Rockefeller Hall, Hanover, NH 03755, USA; *tel:* +1 603 646 9345; *email:* andreas.moxnes@dartmouth.edu

[§]1-3-1 Kasumigaseki Chiyoda-ku, Tokyo 100-8901 Japan; *tel:* +81 3 3501-8312; *email:* saito-yukiko@rieti.go.jp

1 Introduction

In spite of the widespread perception that firms' success in part depends on their connections with suppliers and customers, relatively little work has been done on the structure, performance and importance of production networks. Even less is known about how geography and trade costs affect links in production networks. Finally, in spite of a large literature on the role of infrastructure on economic outcomes, there is almost no evidence on how infrastructure affects firm-level productivity. This paper examines the importance of buyer-supplier relationships and the structure of the production network in firm performance and yields a novel explanation for variation in economic outcomes across regions.

While there has been an explosion of research on social and economic networks and their formation, to date little of that work has considered the supplier-customer relations between firms. In addition, existing studies are often limited to a particular industry or geography within a country.¹ In this paper we use a comprehensive, unique data set on the production network in Japan. Our data provide supplier-customer links between firms for over 950,000 firms in Japan. This set of firms accounts for the large majority of private sector economic activity in the country. For the large majority of firms in Japan, we can determine their location, suppliers, customers and measures of performance. The fundamental question motivating the research is how firm performance is related to the characteristics of the supply network with a special focus on its geography.

We develop a set of stylized facts about the Japanese production network to guide our model. Many of these facts are also present in cross-border trade networks. Large, more productive firms have more suppliers. Geographic proximity plays an important role in the matching of suppliers and customers. Most connections are local; the median distance to a supplier or customer is 30 kilometers. Larger firms not only have more suppliers, but, on average, have suppliers that are farther away. The production network displays negative degree assortivity; the trading partners of well-connected firms on average are less-well connected themselves. Consider two firms, one with many suppliers (A), the other with few (B). The suppliers to the well-connected firm A will on average have relatively few customers. The few suppliers to firm B will on average have more customers. The negative degree assortivity in the Japanese production network is also found in exporter-importer networks in international trade, e.g. Bernard et al. (2013).

We build a parsimonious model of a domestic economy motivated by the stylized facts. Downstream firms require a continuum of tasks as inputs into the production process, e.g. materials processing, accounting, printing and distribution services. They can produce the tasks themselves or outsource them. Finding suppliers is costly, however, and therefore it may not be profitable for all firms to outsource a given task, even if the market price of a task is lower than the firm's

¹See, for example, the seminal work of Uzzi (1996)

marginal cost of supplying the same task. Our model is closely related to the international sourcing framework in Antràs et al. (2014), but we modify it to focus on a continuum of possible domestic sourcing locations and to allow for the possibility that firms can supply a given task within the boundary of the firm. Downstream firms can observe broad characteristics of potential upstream locations, i.e. average productivity and trade costs, but need to expend resources to observe the prices of individual tasks in a location. In equilibrium, a higher efficiency firm will search across more locations, source more inputs, and have better performance. If the fixed costs of search falls, firms will search more, source more inputs from more distant locations, and firm sales will rise. These effects will be larger in input-intensive industries where the marginal benefit of finding better suppliers is greater. For the aggregate economy, locations with low trade and search costs will have higher performing firms, even if productivity is ex ante identical across all locations. Our framework therefore offers a supply-side microfoundation for why productivity varies widely across locations, as documented in e.g. Sveikauskas (1975), Glaeser and Maré (2001) and Combes et al. (2012).

To examine the predictions of the model we use the 2004 opening of the southern portion of the high-speed rail lines in Japan (Kyushu Shinkansen) as a quasi-natural experiment. The route of this particular extension had been planned at least since 1973 but the actual construction was subject to substantial timing uncertainty due to numerous budgetary and administrative delays, thus limiting the scope for anticipation effects. We examine whether the firms near new Shinkansen stations improved performance after the opening. Estimating a triple difference specification, we find that performance was better for firms near the new stations after the opening and that firms in industries with greater purchased input shares performed better compared to firms in industries with lower purchased input shares.

The model suggests that the firm-level performance improvement after a reduction in the search cost is due to the increased number of suppliers and source locations. We draw on a second cross-section of the Japanese production network in 2010 to examine whether firms in localities near the new stations increase their number of suppliers and the number of source locations more than firms in localities that did not become better connected with the high-speed rail extension. The results show support for the mechanisms emphasized in the model; the number of connections and the number of source locations both increase for firms near the new stations.

This paper is naturally related to a growing literature on the determinants of domestic and foreign sourcing and the impact on firms. Amiti and Konings (2007), Goldberg et al. (2010), Halpern et al. (2011) and Bøler et al. (2014) examine the role of imported inputs in firm productivity where foreign and domestic inputs are imperfect substitutes. Our work is closer to Antràs et al. (2014) who develop and structurally estimate a model where firm performance is positively related to the intensive and extensive margins of purchased imported inputs. In the domestic production network, we find systematic relationships between distance to domestic suppliers and firm performance that

are analogous to those in the international trade context. In this regard our work is related to Fort (2014) who finds an important role for firm heterogeneity and location in the decision to use domestic contract manufacturing services.

The paper is also related to a large literature on the effects of infrastructure on economic development. Governments typically allocate a large fraction of their budgets to infrastructure projects and multilateral institutions similarly emphasize infrastructure in the expenditure allocation. Most research on the effects of infrastructure focuses on the location of economic activity, income and aggregate welfare effects. For example, Donaldson (forthcoming) examines the effects of railroads on income and welfare in India, while Duranton et al. (2013) consider the effects of interstate highways on rural US counties. Redding and Turner (forthcoming) survey the literature on the effects of infrastructure on economic activity. This area of research focuses on the role of infrastructure in reducing transport time and costs for goods between cities and in reducing the travel time for individuals within a city, i.e. commuting time. Our research points to another role for infrastructure in reducing travel time for individuals (as opposed to goods) between regions and the resulting firm-level improvements in performance due to the increased use of low-cost (high-quality) purchased inputs.

Another related strand of recent work focuses on the geography of knowledge transmission across locations. Davis and Dingel (2012) focus on costly idea exchange as the agglomeration force in a model of a system of cities. Our framework focuses on the cost of connecting to others (firms) and the resulting improvements in performance. Cristea (2011) considers the importance of face-to-face meetings in international trade and finds that increased exports raises the demand for business class air travel. Comin et al. (2012) study technology diffusion over time and find that technology diffuses slower to locations that are farther away from technology leaders. Keller and Yeaple (2013) measure the cross-country spatial barriers to the transmission of embodied or disembodied knowledge. They find that person-to-person communication costs increase in distance. Our results suggest that reductions in travel time (costs) increase the likelihood of finding new suppliers, thus reducing the cost of the purchased input bundle and raising firm performance. Hillberry and Hummels (2008) examine trade in intermediate goods as an explanation for highly localized shipments in the U.S.

Our work is also related to Giroud (2013) who examines the effect of new airline connections on within-firm performance of and investment in manufacturing plants. Related work in finance argues that proximity matters for monitoring and relationships.² In contrast to his study which examines reductions in travel costs between headquarters and plants for multi-plant firms, we broaden the scope by exploring all buyer-supplier connections among all firms in the economy. Moreover, our model and empirical strategy emphasize the creation and destruction of linkages in response to infrastructure shocks.

²See Lerner (1995), Coval and Moskowitz (1999, 20010 and Peterson and Rajan (2002).

In the literature on firm-to-firm connections, Oberfield (2013) develops a network theory of search and production where producers potentially sell to many customers but have only one upstream supplier. Downstream firms consider match-specific productivity and price when choosing among available techniques. As in our model, the share of purchased inputs matters for the propagation of shocks in the economy although our focus is on the supplier side rather than the downstream links. Acemoglu et al. (2012) relate these types of microeconomic shocks to aggregate fluctuations in a model of sectoral input-output linkages while Carvalho et al. (2014) use the Japanese production network to study the supply chain disruptions occurring in the aftermath of the 2011 earthquake in Japan.

While our focus is on buyer-supplier matches in the domestic supply network, it is closely related to the nascent literature using matched importer-exporter data. Bernard et al. (2013) consider exporter-importer connections using Norwegian transaction trade data. They find, as we do, negative assortivity in buyer-seller matches and in-degree and out-degree distributions that largely follow power laws. Blum et al. (2012) examine characteristics of trade transactions for the exporter-importer pairs of Chile-Colombia and Argentina-Chile and also find that small suppliers (exporters) typically sell to large (importers) and small importers (buyers) source from large suppliers (exporters).

The rest of the paper is structured as follow. We describe the data in Section 2 and develop a set of stylized facts about buyer-supplier relationships in Section 3. In Section 4, we develop our multi-location model of domestic sourcing. We describe and estimate our natural experiment along with various robustness checks in Section 5 and provide concluding remarks in Section 6.

2 Data

The data employed in this paper comes from two main sources. First, production network data for two moments in time, 2005 and 2010, are assembled by Tokyo Shoko Research, LTD. (TSR). TSR is a credit reporting agency and firms provide information to TSR in the course of obtaining credit reports on potential suppliers and customers or when attempting to qualify as a supplier. The resulting database contains information on more than 950,000 firms in each cross-section, represents more than half of all the firms in Japan and covers all sectors of the economy. The TSR sample is close to the full population of firms with more than 4 employees.³

Each firm provides rank-ordered lists of the most important suppliers (up to 24) and customers (24). TSR also collects information on employment, the number of establishments, the number of factories, up to three (4-digit) industries, three years of sales and profits and a physical address. In addition, the database records TSR's credit score for the firm. Using an address matching service

³Firms with 1 to 4 employees are underrepresented in TSR compared to Census data, while for firms with 5 or more employees, the firm size distribution in TSR is very similar to the distribution in Census data.

provided by the Center for Spatial Information Science at the University of Tokyo, we are able to match a firm’s address to longitude and latitude data.⁴ We use the geo-coded data to create a measure of great circle distance between firms. The top 3 prefectures by counts of firms are Tokyo, Osaka and Aichi (Nagoya) while the top three 2-digit industries by counts are General Construction Work, Specialist Construction Work and Equipment Installation.

Second, firm-level balance sheet data comes from Kigyō Katsudō Kihon Chōsa Houkokusho (The Results of the Basic Survey of Japanese Business Structure and Activities), henceforth Kikatsu, for the period 1998 to 2008. Kikatsu is an annual survey that gives detailed information about firm activities such as sales, employment, capital stock, intermediate purchases and industry affiliation. It covers the full population of manufacturing and non-manufacturing firms with more than 50 employees and with capital of more than 30 million yen.

2.1 Supplier and Customer Connections

The TSR data has both advantages and disadvantages relative to other production network data sets. Among the advantages is the inclusion of firms of all sizes and industries including both publicly listed and unlisted firms. In addition, the TSR firms self-report their most important suppliers and customers; there is no cutoff in terms of sales or purchases.⁵ However, the 24-firm limit for suppliers, customers and owners potentially causes a truncation in the number of relationships in the self-reported data relative to the actual number of such connections.

To mitigate this issue, we combine both self-reported and other-reported information for each firm in the data and use the union of own-reported and other-reported information. For firms A and B, we consider A to be a supplier of B if both firms are in the TSR data *and* either (i) A reports B as customer *or* (ii) B reports A as supplier. Note that some firms that are reported as suppliers and customers are outside the TSR set of firms (NTSR), i.e. they are domestic Japanese firms but are not customers or clients of TSR.

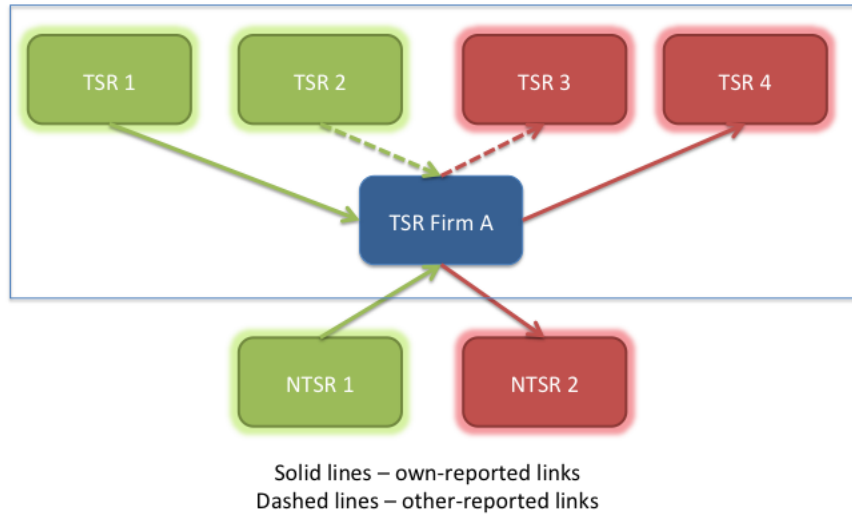
In Figure 1 we show possible suppliers and customers for a firm (Firm A) in the TSR database. Firm A’s reports that it has two customers, TSR4 and NTSR2, and two suppliers, TSR1 and NTSR1. Other firms also report connections to Firm A: TSR2 reports Firm A as a customer while TSR3 reports Firm A as a supplier. In determining Firm A’s in-degree, the number of suppliers, and its out-degree, the number of customers, we ignore the NTSR links and include both own-reported and other-reported connections. Thus, Firm A has an in-degree of 2 (TSR1 and TSR2) and an out-degree of 2 (TSR3 and TSR4).

Many firms report either no suppliers and/or no customers among the TSR firms. This does

⁴As each firm only reports one address, the geographic information for multi-establishment firms is likely to reflect the location of the headquarters.

⁵In their analysis of US production networks, Atalay et al. (2011) use Compustat data on publicly listed firms and their major customers defined as firms that purchase more than 10 percent of the seller’s revenue.

Figure 1: Supplier and Customer Connections: An Example



not mean they recorded no suppliers or customers on their forms but instead all their reported connections are outside the TSR set of firms. A report of no TSR suppliers or no TSR customers might occur for several reasons. A firm might appear to have no TSR customers because all the domestic firms that are customers are outside the TSR database, all its customers are foreign firms or all its customers are non-firms, e.g. the public or government. A firm might appear to have no TSR suppliers because all the domestic firms that are suppliers are outside the TSR database or all its suppliers are foreign.⁶ We choose to work only with the set of TSR firms with a positive in-degree or positive out-degree (links to other TSR firms) and find no evidence of systematic bias in the the sample of firms with positive TSR degree. Using NTSR+TSR data, the distribution of firms with TSR degree equal zero is virtually identical to the overall sample of firms, i.e. the mean and variance of NTSR+TSR out-degree and in-degree distributions are the same.

3 The Production Network

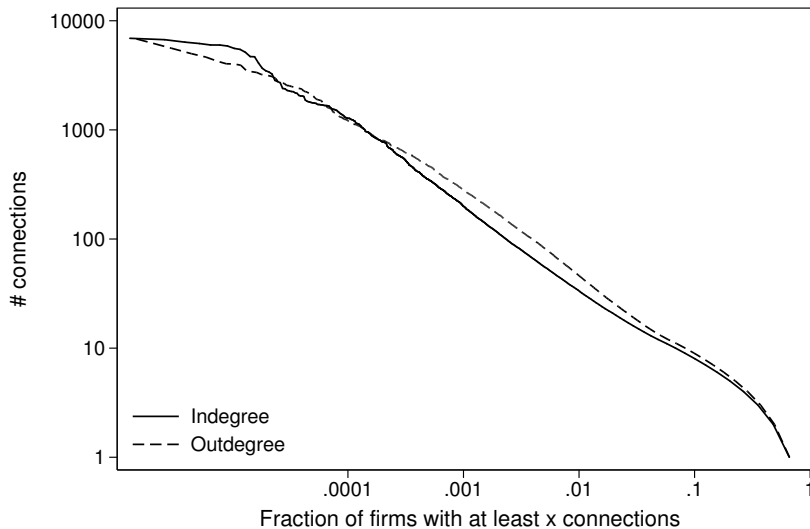
In this section we begin to explore the domestic production network in Japan.⁷ There are 961,318 firms (nodes) in the TSR production network with 3,783,711 supplier-customer connections (directed edges). Of those nodes, 771,107 (676,320) nodes have positive in-degree (out-degree) among TSR firms. For firms with positive in-degree, the mean number of suppliers is 4.9 and the median is 2. For firms with positive out-degree, the mean number of customers is 5.6 and the median is one.

The cdf of the in-degree and out-degree distributions is given in Figure 2. The distributions are

⁶It seems implausible to imagine that an operating firm has no actual domestic suppliers. This is supported by the fact that more TSR firms report no customers than report no suppliers.

⁷All descriptive statistics refer to the 2005 cross-section. Some of these network characteristics are also presented in Saito et al. (2007) and Ohnishi et al. (2010).

Figure 2: In-degree and Out-degree CDFs



Note: 2005 data, log scales.

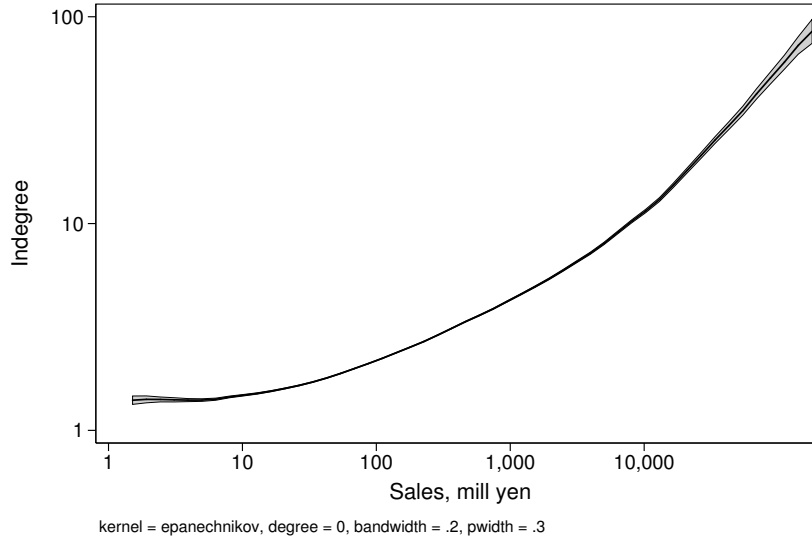
well-approximated by a Pareto (power law) distribution. The estimated Pareto shape parameter is -1.32 for the in-degree distribution and -1.50 for the out-degree distribution. Deviations from the Pareto are found in the extreme tails of the distribution. Firms with very large numbers of connections are somewhat under-represented while firms with small numbers of connections appear in greater numbers. These deviations from a power law distribution are comparable to those found in exporter-importer degree distributions by Bernard et al. (2013) but are much smaller in magnitude compared to those found by Atalay et al. (2011) for supplier-customer connections derived from data on large US firms and their large customers.

As with many socio-economic networks, the Japanese production network has short path lengths between nodes: the mean, median and maximum path length are, 4.7, 4, and 17 respectively. In addition, 99 percent of the nodes are connected to each other in that they belong to the same component. However, unlike networks that exhibit “small worlds” features, clustering coefficients of the undirected network are quite small, total clustering is 0.002 while average clustering is 0.064. This means that if firms a and b are connected to firm c , either as suppliers or customers, firms a and b are not likely to be connected themselves. Instead, firm connections are more likely to be hierarchical.

3.1 Stylized Facts

In this section, we document four facts from the data that will guide the development of the model in Section 4. We explore the relationship between firm characteristics, connections in the production

Figure 3: Size and In-degree.



Note: 2005 data. The figure shows the kernel-weighted local polynomial regression of firm-level log in-degree (vertical axis) on log sales (horizontal axis). Gray area denotes the 95 percent confidence bands. Linear regression slope is 0.24. Sample is first trimmed by excluding the 0.1 percent lowest and highest observations of sales.

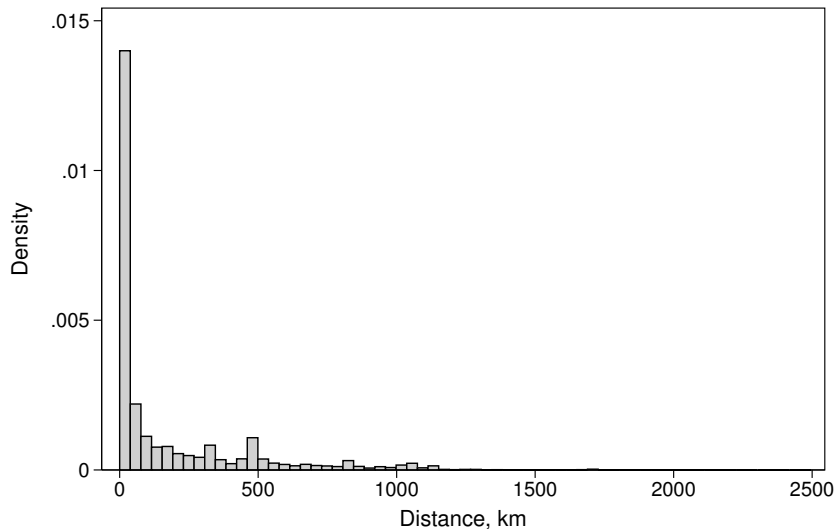
network and geography.

Fact 1: Larger firms have more suppliers. Higher sales are associated with a larger number of supplier connections. Figure 3 plots the kernel-weighted local polynomial regression of a firm’s in-degree (vertical axis) on sales (horizontal axis), both in logs. The linear regression slope is 0.24, meaning that a 10 percent increase in sales is associated with a 2.4 percent increase in the number of suppliers. A similar positive relationship exists between a firm’s sales and out-degree, mirroring the findings in Bernard et al. (2013). Also, the positive relationship is not a result of industry composition; controlling for industry (3 digit JSIC) fixed effects generates the same positive relationship.

Fact 2: The majority of connections is formed locally. Distance is important in the formation of links. We start by calculating the distance between any supplier-customer pair ij and show the density of distance in Figure 4. Geolocation is based on a firm’s headquarters, so for multi-plant firms the interpretation is distance between headquarters. The median (mean) distance is 30 (172) km. Hence, the majority of connections is formed locally. Even so, a few connections span very long distances, so that the average distance is much greater than the median.

Fact 3. Larger firms have suppliers in more locations and their distance to suppliers is higher. Figure 5 shows that larger firms tend to have suppliers in more municipalities. A firm in the 1st decile of the sales distribution has suppliers in 1.5 locations while a firm in the 9th decile has suppliers in

Figure 4: Density of distance across buyer-seller pairs.



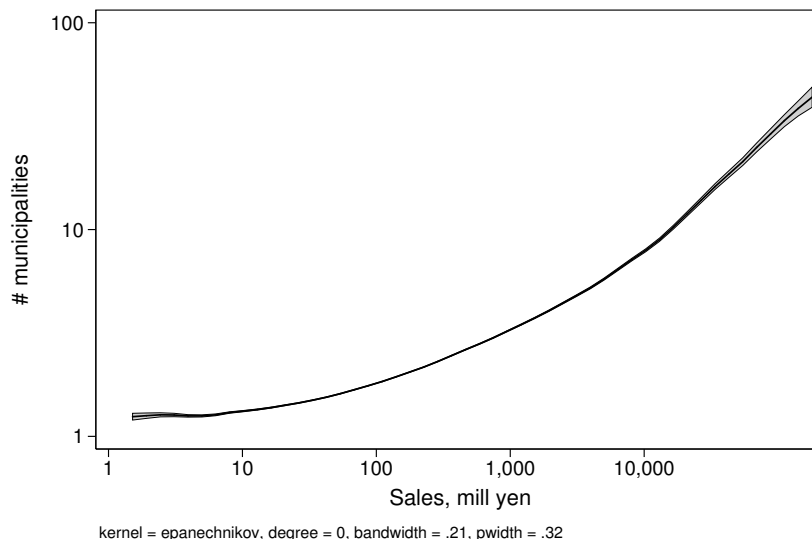
Note: 2005 data. The figure shows the density of distance in km for all buyer-seller pairs.

roughly 4 locations.⁸ At the same time, larger firms have more remote connections; Figure 6 plots the fitted values from a kernel-weighted local polynomial regression of a firm’s maximum distance to its suppliers on its sales (both in logs). The maximum distance to suppliers is around 50km for firms in the 1st decile of the sales distribution, while the max distance is more than three times as high (160 km) for firms in the 9th decile of the sales distribution. Plotting median or average distance against sales produces a similar relationship, although the slope coefficients are different. A similar positive relationship also exists between a firm’s sales and median distance to its customers. The positive relationship is not a result of industry composition; controlling for industry (3 digit JSIC) fixed effects generates the same positive relationship. A potential concern is that larger firms have more plants, so that distance from the relevant plant to a supplier may not be higher for larger firms. We investigate this by removing all multi-plant firms. The positive correlation between firm size and distance to suppliers is equally strong for this subset of firms.

Fact 4: *There is negative degree assortativity among sellers and buyers.* One distinguishing feature of networks is the extent to which a well-connected node is linked to other well-connected nodes, known as degree assortivity. While there is an extensive body of research on degree assortivity in technical and social networks, these relationships are less well documented in economics networks. We find that the better connected a firm, the less well-connected is its average connection. Figure 7 provides an overview of degree assortivity in the Japanese production network. The figure shows

⁸There are in total 1410 municipalities in our dataset, see also Section 5.1.

Figure 5: Size and number of supplier locations.



Note: 2005 data. The figure shows the kernel-weighted local polynomial regression of firm-level log number of municipalities with suppliers (vertical axis) on log sales (horizontal axis). Gray area denotes the 95 percent confidence bands. Linear regression slope is 0.27. Sample is first trimmed by excluding the 0.1 percent lowest and highest observations of sales.

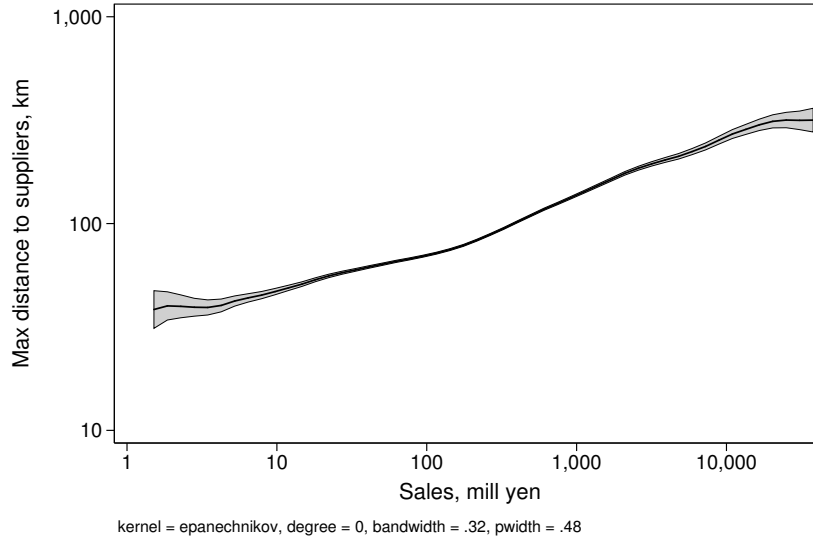
all possible values of the number of suppliers per Japanese firm, a , on the x-axis, and the average number of (customer) connections of these suppliers, $b(a)$, on the y-axis. The interpretation of a point with the coordinates (10,1) is as follows: For a Japanese firm sourcing from 10 suppliers, the average supplier has one customer. The fitted regression line has a slope of -0.19, so a 10 percent increase in number of suppliers is associated with a 2 percent decline in the average supplier’s number of customers.⁹

This results suggests that the best firms, those with the most customers, are selling to firms who on average have fewer customers themselves. Interestingly, social networks typically feature positive assortative matching, that is, highly connected nodes tend to attach to other highly connected nodes, while negative correlations are usually found in technical networks such as servers on the Internet (Jackson and Rogers, 2007).¹⁰ In a recent paper, Bernard et al. (2013) also find negative assortativity between trading firms using Norwegian exporter data matched to foreign importers and Colombian importer trade data matched to foreign exporters.

⁹The correlation between degree and mean degree of connections is a standard measure of assortativity in networks (Jackson and Rogers, 2007).

¹⁰In the friendship network among prison inmates considered by Jackson and Rogers (2007), the correlation between a node’s in-degree and the average in-degree of its neighbors is 0.58. The correlation in our data is -0.31. Serrano and Boguna (2003) find evidence of negative sorting in the network of trading *countries*; i.e. highly connected countries, in terms of trading partners, tend to attach to less connected countries.

Figure 6: Size and max distance to suppliers.



Note: 2005 data. The figure shows the kernel-weighted local polynomial regression of firm-level log max distance to suppliers (vertical axis) on log sales (horizontal axis). Gray area denotes the 95 percent confidence bands. Linear regression slope is 0.32. Sample is first trimmed by excluding the 0.1 percent lowest and highest observations of sales.

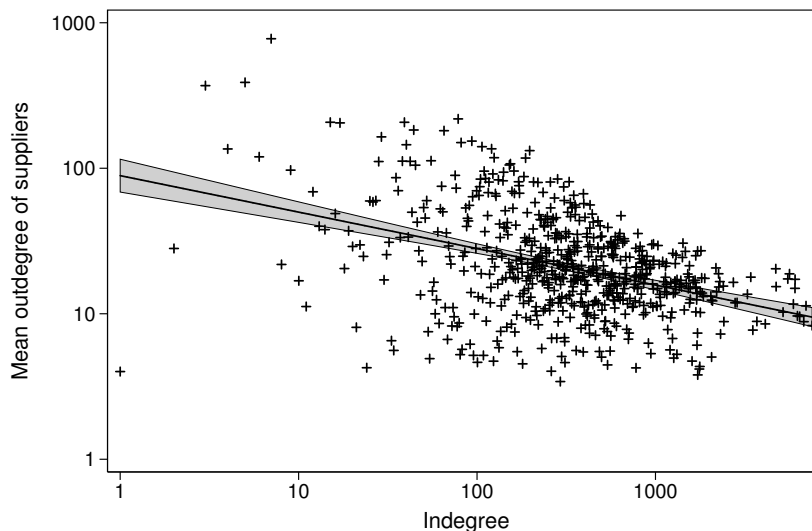
One concern is that the finding of negative degree assortivity may be influenced by the location of the firm. To check this we control for the location of both the firm and its connections. Specifically, we ask whether firms in Tokyo (i) with many suppliers in Osaka (j) are trading with less well-connected firms in Osaka, in terms of their number of customers in Tokyo. We include fixed effects for both the location of firm as well as its suppliers,

$$\overline{Supplier\ Outdegree}_{ij} = \alpha_i^1 + \alpha_j^2 + \beta Indegree_{ij} + \epsilon_{ij} \quad \beta = -0.120 \quad \text{s.e. (0.003)}$$

Again we find negative assortivity: firms with more suppliers in a destination market are sourcing from firms with fewer customers. Controlling for destination countries, Bernard et al. (2013) estimate a comparable coefficient of -0.13 when considering buyer-seller matches in Norwegian exporter data and -0.20 in matches from Colombian importer data. Our findings of negativity assortivity is not limited to this specific configuration of in- and out-degree. We find similar relationships when using customer out-degree or an undirected measure of the total number of connections.¹¹

¹¹In our data, the correlation is negative and significant for (in-degree, average in-degree of suppliers), (in-degree, average in-degree of customers), (in-degree, average out-degree of customers), (out-degree, average out-degree of suppliers), (out-degree, average in-degree of suppliers), (out-degree, average in-degree of customers), (out-degree, average out-degree of customers).

Figure 7: Degree Assortivity - Suppliers and Customers of Suppliers



Note: 2005 data. The figure shows all possible values of the number of suppliers per firm, a , on the x-axis, and the average number of customer connections of these suppliers, $b(a)$, on the y-axis. Axes scales are in logs. The fitted regression line and 95% confidence intervals are denoted by the solid line and gray area. The slope coefficient is -0.19.

4 The Model

We develop a parsimonious model of outsourcing in a domestic economy motivated by the facts in the previous section. The basic structure is as follows. Firms require a continuum of tasks as inputs into the production process, e.g. materials processing, accounting, printing and mailing services. They can produce the tasks themselves or outsource them. Finding suppliers is costly, however, and therefore it may not be profitable for all firms to outsource a given task, even though the market price of a task is lower than the firm's marginal cost of supplying the same task. This setup will produce theoretical predictions that are consistent with the empirical regularities documented in Section 3 and will guide the development of the empirical methodology in Section 5. Our model is closely related to the framework in Antràs et al. (2014), but we modify it in several directions. First, we allow for the possibility of in-house production, i.e. that firms can supply a given task within the boundary of the firm. Second, since we have thousands of unique locations in our data, geography in the model will be continuous. We combine this with distributional assumptions which allow us to obtain sharp analytical results. Third, our model is a framework for understanding domestic, and not international, sourcing. Since productivity differences are typically much smaller within a country than across countries, and since labor is typically much more mobile within a country than across countries, we assume that wages and productivity are common across locations.

4.1 Setup

Geography, sectors and market structure. The economy consists of a unit continuum of locations $i \in \mathcal{S}$. Each location consists of an upstream and a downstream sector. Downstream firms combine labor and a unit continuum of tasks and sell their output to final consumers. Upstream firms produce a single task using labor only. Within a location i and for a given task ω , there are many identical firms producing ω at the same marginal cost. Hence, the upstream sector is characterized by perfect competition. Downstream firms are monopolistically competitive and produce a differentiated good with efficiency z which varies across firms.

Production technology. The production function of a downstream firm is

$$y = z l^\alpha v^{1-\alpha},$$

where l is labor, α is the labor share and v is a CES composite of the unit continuum of tasks. The CES price index is

$$P^{1-\rho} = \int_0^1 p(\omega)^{1-\rho} d\omega,$$

where $p(\omega)$ is the price of an individual task ω and ρ is the elasticity of substitution between tasks. The firm can potentially produce all tasks in-house. If so, the firm's efficiency in producing a task ω is a realization of a random variable ϕ from the Frechet distribution $F(\phi) = e^{-T_0\phi^{-\theta}}$, where T_0 determines the average efficiency in producing a task and $\theta > \rho - 1$ is inversely related to dispersion in task productivity. $F(\phi)$ is identical across all downstream firms, hence, total factor productivity z is the only source of heterogeneity. In equilibrium, the price $p(\omega)$ will depend on whether the firm outsources a task or not and, if outsourced, which location it will outsource from. We turn to this below.

The production function of an upstream firm in location i is $y_u(\omega, i) = \phi(\omega, i)l$. The efficiency of producing a task ω is a realization of a random variable ϕ from the Frechet distribution $F_u(\phi) = e^{-T\phi^{-\theta}}$. The parameter T governs the average productivity. In order to keep the model tractable we assume that average productivity T is identical across locations i . Upstream firms in i selling to j are subject to iceberg trade costs $\tau(i, j) \geq 1$. The cost of supplying ω from i to j is therefore $w\tau(i, j)/\phi(\omega, i)$. We assume that final goods are costlessly traded.

Labor. The economy is inhabited by a measure \bar{L} of workers and each location has population $L(i)$. Consumers derive utility from consumption of the downstream good. They have identical CES preferences with an elasticity of substitution σ . There is perfect labor mobility across regions. Since final goods are costlessly supplied to consumers across locations, nominal wage equalization is sufficient to leave workers indifferent between locations.

Entry. There is a fixed measure of downstream firms in each location, $m(i)$. As there is no free entry, the production of final goods leaves rents. We follow Chaney (2008) and assume that

consumers derive income not only from labor but also from the dividends of a global mutual fund. Each consumer owns w shares of the fund and profits are redistributed to them in units of labor. Total worker income in location i is then $w(1 + \psi)L(i)$, where ψ is the dividend per share of the global mutual fund.

Outsourcing. The downstream firm located in j can choose to produce a task ω itself or outsource it. The firm can observe average productivity T and trade costs $\tau(i, j)$ from source i . Observing individual prices for all ω , however, requires effort. We therefore assume that the firm must incur a fixed cost $f(j)$ paid in terms of labor to observe individual prices in a location i .¹² As we will see, more productive firms find it optimal to search a wider range of locations because the marginal profits from search are higher for high z firms, while the marginal cost $f(j)$ is constant. Given that $f(j)$ does not vary by source, each location i can be ranked according to its attractiveness as a supplier location, where attractiveness is defined by $\tau(i, j)^{-\theta}$ (see Antràs (2014)). A firm in j will therefore search all locations i where $\tau(i, j)$ is lower than some threshold value (to be defined below). As in Eaton and Kortum (2002), conditional on a set of search locations, firm z 's share of purchases from location i is

$$\chi(z, i, j) = \frac{T\tau(i, j)^{-\theta}}{\Phi(z, j)}.$$

$\Phi(z, j)$ is a measure of market access,

$$\Phi(z, j) = T_0 + \int_1^{\bar{\tau}(z, j)} T\tau^{-\theta} g(\tau, j) d\tau, \quad (1)$$

where $\bar{\tau}(z, j)$ is the highest cost location that z located in j is willing to search. g is the density of trade costs to location j .

The share of tasks outsourced is

$$o(z, j) = 1 - \frac{T_0}{\Phi(z, j)}.$$

Adding more locations to search will raise $\bar{\tau}$ and Φ . More search therefore gives more outsourcing o . As in Eaton and Kortum (2002), the task price index is $P(z, j) = \lambda w \Phi(z, j)^{-1/\theta}$ where λ is a constant.¹³ Hence, more outsourcing leads to lower input costs P with an elasticity $1/\theta$. Searching an additional location means that the firm can observe a new set of prices for all tasks ω . The probability of finding at least one task with a lower price than the existing one is strictly positive, and therefore the aggregate price index must go down.

4.2 Optimal Search

The maximization problem of the firm is then

$$\max_{\bar{\tau}} \{ \pi(z, j) - wf(j)n(z, j) \},$$

¹²In order to keep the problem tractable, we do not allow a source-specific f .

¹³ $\lambda^{1-\rho} = \Gamma\left(\frac{\theta-(\rho-1)}{\theta}\right)$ where Γ is the Gamma function.

where $\pi(z, j)$ is gross profits of firm z located in j and $n(z, j)$ is the measure of locations to search. Total sales of the downstream firm can be written $r = Ap^{1-\sigma}$ where A is a demand shifter and profits are proportional to sales, $\pi = r/\sigma$. Appendix A derives the solution to the problem of the firm as well as the second order condition. The solution to $\bar{\tau}$ is

$$\bar{\tau}(z, j) = \kappa_1 \left(\frac{T}{w^\sigma} \frac{A}{f(j)} \right)^{1/\theta} \Phi(z, j)^{-k/\theta} z^{(\sigma-1)/\theta} \quad (2)$$

where $k = 1 - (\sigma - 1)(1 - \alpha)/\theta$ and κ is a constant.¹⁴ For an arbitrary geography $g(\tau, j)$, one can jointly solve equations (1) and (2).

The expression for the hurdle $\bar{\tau}$ has a number of interesting features. First, better market access Φ leads to more search when $k < 0$ and less search when $k > 0$. The model of Antràs (2014) has the same property and describe this as the complements and substitutes case respectively. Keeping Φ constant, lower search costs f and trade costs τ lead to more search (higher $\bar{\tau}$). Higher efficiency z and more demand A also lead to more search (higher $\bar{\tau}$).

4.3 Model and Data

Let us now go back to the stylized facts presented in Section 3 and relate them to the model. The proofs are found in Appendix B.

First, more productive firms outsource more tasks and therefore have more suppliers:

$$\frac{\partial o(z, j)}{\partial z} > 0,$$

because $\partial \Phi(z, j)/\partial z > 0$. Given that more productive firms search more, they are more likely to find a sourcing option for a given task ω at a lower cost than the cost of producing in-house. This is consistent with the evidence in Figure 3, that larger firms tend to have more suppliers. Note that, according to the model, higher efficiency z leads to both increased sales and in-degree, while higher in-degree itself leads to greater sales. Hence, the level of sales for a given firm is determined by both the direct effect of core efficiency z and the indirect effect of in-degree. The positive correlation shown Figure 3 is a result of both the direct and indirect effect.

Second, more productive firms search more and costlier locations:

$$\frac{\partial \bar{\tau}}{\partial z} > 0.$$

High z firms have a greater incentive to search more locations because the potential cost savings are larger for more productive firms. As a consequence, more productive firms have higher maximum and average trade costs to suppliers. This is consistent with the evidence in Figures 5 and 6, that larger firms tend to have suppliers in more locations and higher maximum distance to their suppliers.

¹⁴ $\kappa_1 = \left(\frac{(\bar{m}\lambda^{1-\alpha})^{1-\sigma}}{\sigma} \frac{(\sigma-1)(1-\alpha)}{\theta} \right)^{1/\theta}$

Third, more productive firms reach suppliers in markets that are on average less well connected. Specifically, consider a firm with efficiency z in location j , sourcing from the marginal location $\bar{\tau}(z, j)$. The expected measure of customers from j among upstream firms in z 's marginal location is decreasing in firm efficiency z ,

$$\frac{\partial C(z, j)}{\partial z} < 0.$$

This reflects the fact that higher z firms reach costlier locations and the suppliers there are on average not very competitive in z 's home market. This is consistent with the evidence in Figure 7 on degree assortativity.¹⁵

4.4 Testable Predictions

As we discuss later in Section 5, we will exploit a natural experiment where a large shock to infrastructure lowered passenger travel time (but not goods travel time) between many location-pairs in Japan. This empirical exercise allows us to quantify the impact on large-scale infrastructure projects on firm performance and to evaluate the importance of the theoretical mechanism emphasized in this paper. In order to guide the subsequent empirical work, this section details the consequences of such a shock according to the model.

First consider the impact on firm sales of lower search costs $f(j)$. Lower $f(j)$ leads to sales growth of a downstream firm in j . Holding final goods demand A constant,

$$\frac{\partial \ln r(z, j)}{\partial \ln f(j)} = \frac{(\sigma - 1)(1 - \alpha)}{\theta} \frac{\partial \ln \Phi(z, j)}{\partial \ln f(j)} < 0.$$

The elasticity $\partial \ln \Phi(z, j) / \partial \ln f(j)$ measures the fall in marginal costs from an increase in $f(j)$ (Appendix B.2). Second, consider how $\partial \ln r(z, j) / \partial \ln f(j)$ varies across industries with different labor intensities α :

$$\frac{\partial^2 \ln r(z, j)}{\partial \ln f(j) \partial \alpha} = \frac{\sigma - 1}{\theta} \left(-\frac{\partial \ln \Phi(z, j)}{\partial \ln f(j)} + (1 - \alpha) \frac{\partial^2 \ln \Phi(z, j)}{\partial \ln f(j) \partial \alpha} \right).$$

The cross elasticity is the sum of a direct and indirect effect. The direct effect is that a percent reduction in input costs $P(z)$ will have a stronger positive effect on sales in industries where inputs constitute a large share of total costs. The indirect effect is that input-intensive firms may search more or less intensively relative to labor-intensive firms when $f(j)$ falls (the cross elasticity $\partial^2 \ln \Phi(z, j) / \partial \ln f(j) \partial \alpha$). Appendix C shows that the direct effect will always dominate when $g(\tau, j)$ is inverse Pareto with support $[1, \tau_H]$. An inverse Pareto captures the notion that a location j has few nearby markets and many remote markets; Appendix D provides empirical evidence that

¹⁵Note that the dependent variable in Figure 7 is *average* out-degree whereas $C(z, j)$ is *marginal* out-degree, i.e. out-degree in the least profitable location i . In practice, this makes little difference because in the model the average is pinned down by the marginal out-degree.

the inverse Pareto is a good approximation of the empirical density in our dataset. Hence, the total effect is

$$\frac{\partial^2 \ln r(z, j)}{\partial \ln f(j) \partial \alpha} > 0,$$

so that sales growth is stronger for input-intensive firms relative to labor-intensive firms. We summarize this in the following proposition.

Proposition 1. *(i) Lower search costs $f(j)$ lead to growth in sales among downstream firms in j . (ii) Sales growth is stronger in input-intensive (low α) relative to labor intensive (high α) industries.*

Part (ii) of Proposition 1 forms the basis of our identification strategy in Section 5.1.

Second consider the impact on supplier connections among firms in j of lower search costs $f(j)$. Lower search costs $f(j)$ lead to more outsourcing and suppliers from new locations among downstream firms in j , see Appendix B.2,

$$\frac{\partial o(z, j)}{\partial f(j)} < 0 \quad \text{and} \quad \frac{\partial \bar{\tau}}{\partial f(j)} < 0.$$

Lower $f(j)$ means that the cost of obtaining information about prices is lower. Firms therefore search additional locations ($\bar{\tau}$ increases). There is positive probability of finding a task at a lower price compared to the price of in-house production. Hence, outsourcing must also increase. We summarize this in the following proposition.

Proposition 2. *Lower search costs $f(j)$ lead to more outsourcing and suppliers from new locations (higher $\bar{\tau}$) among downstream firms in j .*

5 Production Networks and Productivity: A Natural Experiment

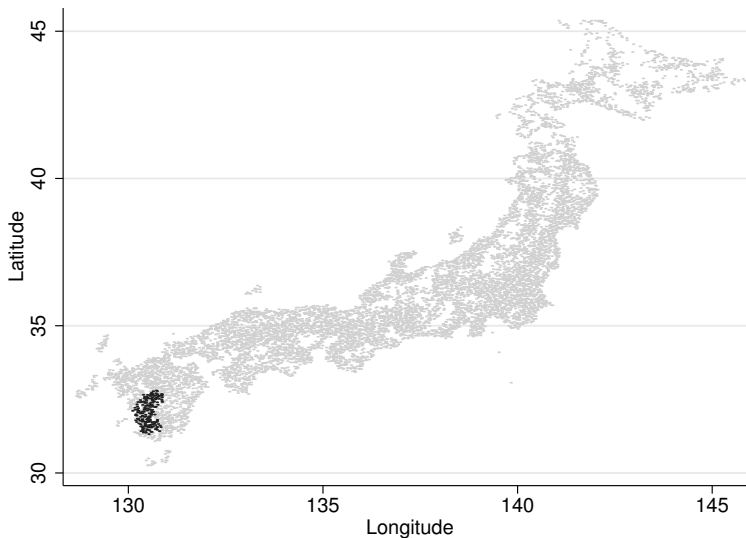
This section details our identification strategy for estimating the impact of lower search costs on firm performance and linkages in the production network. We start by providing some background on the natural experiment.

The southern portion of the high-speed (bullet) train network in Japan was expanded in March 2004 (Kyushu Shinkansen). This resulted in a dramatic reduction in travel time between major cities in the area. For example, travel time between Kagoshima and Shin-Yatsushiro declined from 130 minutes to 35 minutes, and travel time between Hakata and Kagoshima declined from 4 hours to just 2 hours. Figure 8 gives an overview of the geography. The black dots are locations within 30km of a new Shinkansen station, whereas the gray dots are all other localities in the dataset. Although the geographical scope is somewhat limited, the new rail line extended Shinkansen service to two prefectures with a total population of 3.5 million, roughly that of Connecticut.

The Shinkansen expansion offers several advantages for assessing the impact on infrastructure on linkages and firm performance. First, the plan of the expansion started already in 1973, making

it relatively unlikely that firms in our sample could influence the timing and location of stations. Moreover, the timing of completion was subject to substantial uncertainty starting in 1991, limiting the scope for anticipation effects. Nevertheless, our empirical methodology will address endogeneity concerns in a variety of ways (see below). Second, goods in Japan do not travel on the Shinkansen and there was no contemporaneous reduction in travel time for goods along this southern route. Hence, the shock is well suited to study the impact of lower search costs.

Figure 8: Kyushu Shinkansen treated cells.



5.1 Economic Integration and Firm Performance

In this section we ask whether and to what extent the Shinkansen expansion improved performance among firms in affected regions. As shown in Proposition 1, the model suggests a simple identification strategy. Lower search costs $f(j)$ improves firm sales because it allows firms to find lower cost or higher quality suppliers. Moreover, the impact is greater for input-intensive firms (low α firms) relative to labor-intensive firms. Intuitively, improved travel time has no impact on marginal costs in an industry that does not rely on inputs, i.e. when $\alpha = 1$. Of course, sales may improve in $\alpha = 1$ industries as well, because improved travel time allows firms to find new customers. Our empirical strategy will difference out this mechanism, i.e. the methodology only identifies marginal cost effects and not demand side effects (for more on this, see below).

Consider the following regression,

$$\ln y_{fjrt} = \alpha_f^1 + \alpha_{rt}^2 + \beta Station_f \times H_j \times I[t \geq 2004] + \beta X_{fjrt} + \epsilon_{fjrt}, \quad (3)$$

where y_{fjrt} is a measure of firm performance for firm f in industry j located in region r at time t . We focus on the 8 year period 2000 to 2008, i.e. 4 years before and after the infrastructure shock. Our measures of firm performance are: sales, sales per employee and revenue TFP.¹⁶ According to our model, input prices are firm-specific, and the infrastructure shock is expected to lower input prices by more for input-intensive firms than labor-intensive firms (see Section 4.4). Hence, as long as unobserved firm-specific input prices are not properly controlled for, revenue TFP is expected to increase by more for input-intensive firms than labor-intensive firms. We discuss this more formally in Appendix E.

The main independent variable is the interaction between $Station_f$, which is one if firm f is within 30 km of a new station, H_j which is the input intensity of the industry in 2003 and $I[t \geq 2004]$ which is an indicator variable taking the value 1 from 2004 and onwards. The 30 km threshold for station is chosen so that total travel time is significantly affected and that Shinkansen dominates alternative modes of transport. For example, for a firm 60 km from a station, car travel time to the station would amount to 40 to 60 minutes, and hence the percentage drop in total travel time would be significantly less compared to a firm located near the station. We also check the results with other thresholds in Section 5.1.2. Input intensity is defined as 1 minus the labor share of the 3-digit industry in 2003. The labor share is the industry's wage costs relative to total costs, see Appendix F. α_f^1 and α_{rt}^2 are firm and prefecture-year fixed effects. There are 47 prefectures in Japan.

The covariates in X_{fjrt} are the remaining interactions $Station_f \times I[t \geq 2004]$ and $H_j \times I[t \geq 2004]$. In addition, since prefectures are relatively large geographic areas, we introduce a second geographic control. Each prefecture is further divided into local administrative units called municipalities. We have in total 1410 municipalities in the dataset, making the average population of a municipality roughly 90,000 (Japan's total population was 127.8 million in 2005) The high number of municipality-year pairs means that municipality-year fixed effects are computationally infeasible. We can, however, include a variable for average performance in a municipality-year.¹⁷ For example, if $\ln y_{fjrt}$ is sales, we include average log sales excluding firm f in the municipality-year as a control variable.

The regression in equation 3 is a triple differences model and the intuition for identification is as follows. The Shinkansen expansion is expected to bring higher performance gains for an input-intensive firm located close to a new station compared to a labor-intensive firm located close to a new station. The triple differences empirical strategy is to compare the growth of input-intensive firms before and after 2004 (1st difference) to the growth of labor-intensive firms (2nd difference), and compare this differential effect in locations with a new station relative locations without a new

¹⁶Firm performance is measured relative to industry (3 digit)-year means. There are in total 315 3 digit industries in the data. Hence, industry-year fixed effects are unnecessary in the estimating equation. TFP is estimated using the Olley and Pakes (1996) methodology, see Appendix E.

¹⁷This approach is similar to Giroud (2013).

station (3rd difference).

The triple differences approach resolves a number of potential concerns. First, performance growth due to demand-side effects (i.e. growth among labor-intensive firms due to new customers) is differenced out because demand side effects are expected to affect labor-intensive and input-intensive firms similarly.¹⁸ Second, a potential concern is that input-intensive firms may grow faster than labor-intensive firms even in the absence of the Shinkansen expansion. The methodology controls for this because the triple interaction coefficient β will only capture the differential impact (input intensive relative to labor intensive) for firms close to a station to the differential impact for firms remote from a station. Hence, if input-intensive firms grow faster in every location, then β will be zero.¹⁹ Third, a potential concern is that the new Shinkansen line was introduced in high growth regions. Again, as we only compare the differential growth for input relative to labor-intensive firms, endogeneity is not a concern as long as the Shinkansen line was not targeted particularly for input-intensive firms, which in our view is highly unlikely.

5.1.1 Results

Table 1 shows regression results from estimating equation (3). Column (1) uses log sales relative to the industry-year as the dependent variable. The triple interaction term β is positive and significant at the 5 percent level, indicating that the Shinkansen expansion boosts firm sales for input-intensive firms relative to labor-intensive firms. The magnitudes are economically significant: a coefficient of 0.47 means that a Shinkansen stop increased sales by 0.47 log points more for a firm with $H_j = 1$ relative to a firm with $H_j = 0$. A firm in the 9th decile of the H_j distribution ($H_j = 0.92$, e.g., industrial plastic products, JSIC 183) increased sales by roughly 0.10 log points more than a firm in the 1st decile of the H_j distribution ($H_j = 0.70$, e.g., general goods rental and leasing, JSIC 701). Columns (2) and (3) uses labor productivity and TFP as the dependent variable. Again, the triple interaction term is positive and significant, suggesting that the infrastructure shock improved firm's productivity. The magnitudes are slightly smaller compared to sales, a firm in the 9th decile of the H_j distribution improves labor productivity by 0.09 log points faster than a firm in the 1st decile.

The fact that an infrastructure project *unrelated to* transportation of goods can improve firm performance by this magnitude is indeed remarkable. More broadly, our findings suggest that domestic trade costs dampen economic activity by limiting buyer-supplier linkages and that reducing these barriers will help development and growth. From a policy perspective, the results may also point to important positive externalities of large-scale infrastructure projects - that infrastructure projects can bring efficiency gains from freer flow of information across firms - benefits that are typically neglected from cost-benefit analyses.

¹⁸The common demand side effect is captured in $Station_f \times I[t \geq 2004]$.

¹⁹The common industry effect is captured in $H_f \times I[t \geq 2004]$.

5.1.2 Robustness

In this section, we explore a number of robustness checks. First, a potential concern is that input-intensive firms near a new station tend to grow faster than labor-intensive firms near a new station (but not in other locations), i.e. that there are pre-trends in the treatment relative to the control group. A simple way to check for this is to conduct a falsification test. We estimate equation (3) on the five year period 1998 to 2002 and incorrectly set the Shinkansen expansion to 2000, i.e. we replace $I[t \geq 2004]$ with $I[t \geq 2000]$. The results are shown in Table 2. For all three dependent variables, the triple interaction term is not significantly different from zero, hence there are no pre-trends in the data.

A second concern is that the results are sensitive to the chosen 30 km threshold for whether a firm belongs to a new station or not. We therefore estimate the model with $Station_f$ taking the value one if firm f is within 10 km of a new station. The results in Table 3 show that the results are relatively close to the baseline - the impact on sales is slightly stronger and the impact on sales per employee and TFP is roughly similar. A third potential concern is that the results are primarily driven by the construction sector, e.g. that construction firms grow because of increased demand related to building the new infrastructure. Note, however, that the triple differences approach should control for this, since identification is based on comparing industries with different input intensities.²⁰ Nevertheless, we re-run the regression excluding firms belonging to construction industries (Table 4). Overall the results are very similar to the baseline results in Table 1.

5.2 Economic Integration and Firm Linkages

The empirical results from Section 5.1 show that the infrastructure shock improves firm performance and that the performance effects are stronger among input-intensive firms, consistent with the model. Moreover, the empirical strategy differences out demand side shocks, so that the estimated effects isolate the impact on marginal costs. In this section, we explore in more detail the economic mechanism behind this decline in marginal costs. Our model suggests that the input price index of the firm, $P(z)$, falls because treated firms search new locations, outsource more tasks and find better suppliers for existing tasks (Sections 4.3 and 4.4). However, one could think of alternative explanations, e.g. that treated firms work more efficiently with their existing suppliers, or that treated firms are better able to hire and attract talent after the Shinkansen expansion.

According to our model, lower search costs $f(j)$ lead to more supplier connections of downstream firms in j , see Section 4.4. The aim of this section is therefore to test whether the Shinkansen expansion affected the growth in supplier connections between the 2005 and 2010 cross-sections of the TSR data.

²⁰Moreover, any potential bias would be negative because the construction demand shock occurred before 2004, not after.

A potential concern with the TSR data is that there is no guarantee that a firm surveyed in 2005 will also be surveyed in 2010. Hence, at the firm-level, the number of supplier connections could change simply because a supplier may be a TSR firm in 2005 and an NTSR firm in 2010. In order to mitigate this issue, we aggregate the TSR data as follows. We divide Japan into a grid consisting of 500×500 cells; each cell is a square roughly 5.6 kilometers on a side. Next, we define C_{ijt} as the number of suppliers in i serving customers in j at time t , where $t = \{2005, 2010\}$. As many cells are covering water and other non-populated areas, the dataset is reduced to roughly 8,000 cells, or localities, after removing these regions. We assign a dummy variable $Treat_i = 1$ to a locality if one or more firms in i are within 30 kilometers of a new Shinkansen stop. We also calculate great circle distances between the center of cells i and j , $Dist_{ij}$. Note that the sample selection issue described above is greatly reduced because C_{ijt} is the sum of supplier connections among all firms in j , so any sample selection noise in the in-degree of a given firm is likely to be averaged out.²¹

In the data, a location j may get new connections for a variety of reasons which may be correlated with the infrastructure shock. The empirical strategy is therefore to test whether location-pairs ij where either i or j or both gets a new station increase their number of connections relative to location-pairs where neither i nor j gets a new station. We therefore estimate the following model,

$$\Delta \ln C_{ij} = \xi_i^1 + \xi_j^2 + \beta_1 Both_{ij} + \beta_2 One_{ij} + \gamma X_{ij} + \epsilon_{ij}, \quad (4)$$

where the dependent variable is the change in the log number of connections between suppliers in i and buyers in j , $\Delta \ln C_{ij} = \ln C_{ij2010} - \ln C_{ij2005}$. The main independent variables are $Both_{ij}$ which equals one if both i and j get a new station and One_{ij} which equals one if either i or j (but not both) gets a new station.²² Location-pairs where neither i nor j gets a new station are the omitted group. ξ_i^1 and ξ_j^2 are source and destination fixed effects respectively and X_{ij} is a vector of covariates (see below).

According to the model, a reduction in $f(j)$ should boost the number of connections from all sources i . Empirically, this would be captured in the buyer fixed effect ξ_j^2 as there are no source-specific components to the fixed cost. The empirical model is more flexible than the theory in allowing for bilateral changes in the search costs. The inclusion of $Both_{ij}$ allows for the possibility that the impact on location-pairs with both ends being newly connected to the Shinkansen train network may be stronger than if only one of them is connected. The inclusion of One_{ij} allows for a new Shinkansen station close to i (but not j) to impact both suppliers in i (connections from i to j) and customers in i (connections from j to i). According to the model, search only occurs among downstream firms, and hence a new station in location i will not impact the number of connections from i suppliers to j customers.

²¹The average number of firms in a cell is 104, see Table 5.

²²Formally, $Both_{ij} = I[Treat_i = 1 \cap Treat_j = 1]$ and $One_{ij} = I[(Treat_i = 1 \cap Treat_j = 0) \cup (Treat_i = 0 \cap Treat_j = 1)]$.

The empirical framework has the flavor of a gravity model of the extensive margins of trade. Consider the non-differenced version of equation (4), i.e. a specification with $\ln C_{ij}$ as the dependent variable and with source-year, destination-year and location-pair fixed effects. That model would identify the impact of the Shinkansen expansion on the number of connections from the change in travel time for a location-pair relative to unaffected location pairs. The source-year and destination-year fixed effects would capture trends in economic activity which may differ across locations. The location-pair fixed effects would control for time-invariant determinants of bilateral trade between locations. Due to the large number of fixed effects, we use the log change in C_{ij} and obtain the estimating equation (4). Because identification is through changes in bilateral passenger travel time, the effect is coming entirely through the movement of individuals, and not through alternative mechanisms such as shipping and transportation costs.

The production network is observed at two moments in time, 2005 and 2010. The timing is not ideal for our purposes, because the Shinkansen extension occurred in March 2004. The underlying assumption is therefore that the impact of the expansion had not fully materialized when firms were surveyed in 2005. Our view is that finding new suppliers is a slow and costly process, so that it is unlikely that firms had adjusted fully after one year. Note that the bias in the coefficients is negative if firms partially adjusted before 2005, i.e. we would estimate a smaller effect compared to the true impact.

5.2.1 Results

We start by documenting a few basic facts about the locality connections dataset. Table 5 show descriptive statistics for the number of connections in 2005 and 2010. The average (median) number of connections between a locality-pair was 7.05 (2) in 2005 and increased slightly to 7.51 (2) in 2010. There are roughly 8,000 localities and almost 400,000 locality-pairs with positive flows in both 2005 and 2010. This implies that many locality-pairs have zero connections, i.e. the number of locality-pairs with positive transactions is much smaller than the theoretical maximum $(L^2 - L)$.²³

Table 6 presents the results from estimating equation (4). Column (1) estimates the model without any fixed effects, while columns (2) includes source and destination fixed effects. The inclusion of these fixed effects controls for changes in the number of firms in the localities. Column (3) and (4) includes log distance and log distance interacted with $Both_{ij}$ and One_{ij} as additional independent variables. Our preferred specifications are (3) and (4), as there is evidence of agglomeration over time, i.e. that local connections grow faster than remote connections (the distance coefficient being negative). It is important to control for this because $Both_{ij}$ is negatively correlated with distance.

Overall, the results indicate that localities becoming connected by new stations ($Both_{ij} = 1$)

²³As the dependent variable is $\Delta \ln C_{ij}$, we drop pairs with missing $\Delta \ln C_{ij}$ ($C_{ij2005} = 0$ or $C_{ij2010} = 0$). The 8000 localities i in the data therefore have either $C_{ij} > 0$ or $C_{ji} > 0$ for $j \neq i$. ii pairs are also dropped from the dataset because distance is zero for these pairs and log distance is used as an independent variable in several regressions.

increased their number of connections by roughly 40 percent relative to unconnected localities. Recall that the average number of connections is 7, so for the average locality-pair, the Shinkansen extension caused 3 new connections between newly connected localities. Our preferred specifications in columns (3) and (4) suggest that the impact is roughly half as large when only one of the locations in a pair is connected ($One_{ij} = 1$). Perhaps surprisingly, the interaction terms are close to zero. We find no evidence that the infrastructure shock benefited remote connections more than local connections.

Robustness. In the baseline results, we used the threshold of 30 kilometers from a new station to classify cells as treated or untreated. To check the sensitivity of the results, we instead use a threshold of 10 kilometers. The results are presented in Table 7 and indicate that the results are robust to this change. A potential concern in the baseline results is that location-pairs with either $C_{ij2005} = 0$ or $C_{ij2010} = 0$ are dropped because of the log transformation. We address this by estimating by replacing the dependent variable $\Delta \ln C_{ij}$ with a $\Delta I [C_{ijt} > 0]$, i.e. a dummy variable taking the value one if location-pair ij starts trading between 2005 and 2010, zero if there is no change, and minus one if ij stops trading. Including all the zeros in the dataset results in an extremely high number of observations (71 million), so estimation with joint source and destination fixed effects becomes computationally infeasible. Table 8 therefore reports results with either source or destination fixed effects in columns (2) and (3) respectively. The estimates confirm the findings in the baseline specification; location-pairs with new stations are more likely to start trading (and less likely to stop trading) between 2005 and 2010.

Summing up, the reallocation in the production network caused by the extension of the high-speed train network documented here is consistent with the firm performance gains found in Section 5.1. Of course, we cannot completely rule out that other economic mechanisms are also partly driving the performance gains. Nevertheless, the body of evidence presented here suggests that supplier linkages play a key role.

6 Conclusions

This paper examines how firm performance is related to the characteristics of the supply network with a special focus on geography. Using a comprehensive, unique data set on supplier-customer links among 950,000 Japanese firms, we develop a set of facts about the production network. Geographic proximity plays an important role for the matching of suppliers and customers as most connections cover relatively short distances. Large, more productive firms both have more suppliers and, on average, have suppliers that are farther away. While large firms have more suppliers, the trading partners of those large, well-connected firms on average are less-well connected themselves.

Guided by these facts, we develop a simple model where firms can outsource tasks and search for suppliers in different locations. Firms located in close proximity to other markets, and firms

that face low search costs, will search more and find better suppliers. This in turn drives down the firm's marginal production costs. We test the theory by exploiting the opening of a high-speed (Shinkansen) train line in Japan which lowered the cost of passenger travel but left shipping costs unchanged.

We find compelling evidence that the supply network matters for firm performance. A shock to infrastructure which reduces travel time for passengers (but not shipping costs) generates significant performance gains, especially for firms in industries that have large shares of purchased inputs. We also provide evidence that these gains are related to new buyer-seller linkages as predicted by the model.

While there is a large literature on the link between infrastructure and improvements in regional economic outcomes, this paper provides the first direct evidence on the role of infrastructure on firm performance. We highlight a novel transmission mechanism for the effects of improved infrastructure where reductions in search costs allow firms to match with more and better suppliers, thus lowering the marginal cost of production. The resulting geographic variation in marginal costs for otherwise ex ante identical firms yields systematic patterns of economic activity across space. Firms in more geographically central locations have lower marginal costs and produce more. Although, this paper is focused on a domestic production network, it is closely related to recent work on the effects of international trade and outsourcing on firm outcomes.

The focus in this work is on one particular mechanism, the extensive margin of new suppliers, that links firm performance to the supply network. However, the results suggest that future research might fruitfully focus on how heterogeneous firms sort into locations, how reduced travel time might affect intensive margins at existing suppliers, and how buyer-supplier connections form and evolve.

References

- Acemoglu, Daron, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi**, “The Network Origins of Aggregate Fluctuations,” *Econometrica*, 2012, *80(5) September*, 1977–2016. 1
- Amiti, Mary and Jozef Konings**, “Trade Liberalization, Intermediate Inputs and Productivity: Evidence from Indonesia,” *American Economic Review*, 2007, *97(5)*, 1611–1638. 1
- Antràs, Pol**, *Global Production: A Contracting Perspective* 2014. 4.1, 4.2
- , **Teresa C Fort, and Felix Tintelnot**, “The Margins of Global Sourcing: Theory and Evidence from U.S. Firms,” 2014. 1, 4
- Atalay, Enghin, Ali Hortacsu, James Roberts, and Chad Syverson**, “Network Structure of Production,” *Proceeding of the National Academy of Science*, 2011, *108(13)*, 5199–202. 5, 3
- Bernard, Andrew B., Andreas Moxnes, and Karen Helene Ulltveit-Moe**, “Two-sided Heterogeneity and Trade,” Technical Report, CEPR Discussion Paper No. 9681 2013. 1, 3, 3.1, 3.1
- Blum, Bernardo S., Sebastian Claro, and Ignatius J. Horstmann**, “Import Intermediaries and Trade Costs: Theory and Evidence,” mimeo, University of Toronto 2012. 1
- Bøler, Esther Ann, Andreas Moxnes, and Karen Helene Ulltveit-Moe**, “Technological Change, Trade in Intermediates and the Joint Impact on Productivity,” 2014. 1
- Carvalho, Vasco M., Makoto Nirei, and Yukiko U. Saito**, “Supply Chain Disruptions: Evidence from Great East Japan Earthquake,” Discussion Paper, Research Institute of Economy, Trade and Industry (RIETI) 2014. 1
- Chaney, Thomas**, “Distorted gravity: the Intensive and Extensive Margins of International Trade,” *American Economic Review*, 2008, *98(4)*, 1707–1721. 4.1
- Combes, Pierre-Philippe, Gilles Duranton, Laurent Gobillon, Diego Puga, and Sébastien Roux**, “The productivity advantages of large cities: Distinguishing agglomeration from firm selection,” *Econometrica*, 2012, *80(6)*, 2543–2594. 1
- Comin, Diego, Mikhail A. Dmitriev, and Esteban Rossi-Hansberg**, “The spatial diffusion of technology,” 2012. 1
- Cristea, Anca**, “Buyer-seller relationships in International trade: Evidence from U.S. States’ exports and business-class travel,” *Journal of International Economics*, 2011, *84*, 207–220. 1

- Davis, Donald R. and Jonathan I. Dingel**, “A Spatial Knowledge Economy,” Technical Report, NBER Working Paper 18188 2012. 1
- Donaldson, Dave**, “Railroads of the Raj: Estimating the Impact of Transportation Infrastructure,” *American Economic Review*, forthcoming. 1
- Duranton, Gilles, Peter Morrow, and Matthew A. Turner**, “Roads and Trad: Evidence from the US,” *Review of Economic Studies*, 2013. 1
- Eaton, Jonathan and Samuel Kortum**, “Technology, Geography, and Trade,” *Econometrica*, Sept 2002, 70 (5), 1741–1779. 4.1, 4.1
- Fort, Teresa C**, “Technology and Production Fragmentation: Domestic versus Foreign Sourcing,” Technical Report, Tuck School of Business at Dartmouth 2014. 1
- Giroud, Xavier**, “Proximity and Investment: Evidence from Plant-level Data,” *Quarterly Journal of Economics*, 2013, 128, 861–915. 1, 17
- Glaeser, Edward L. and David C. Maré**, “Cities and Skills,” *Journal of Labor Economics*, 2001, 19 (2), 316–42. 1
- Goldberg, Pinelopi Koujianou, Amit Kumar Khandelwal, Nina Pavcnik, and Petia Topalova**, “Imported Intermediate Inputs and Domestic Product Growth: Evidence from India,” *The Quarterly Journal of Economics*, November 2010, 125 (4), 1727–1767. 1
- Halpern, Lazlo, Miklos Koren, and Adam Szeidl**, “Imported Inputs and Productivity,” CeFIG Working paper, Central European University 2011. 1
- Hillberry, Russell and David Hummels**, “Trade responses to geographic frictions: A decomposition using micro-data,” *European Economic Review*, 2008, 52 (3), 527–550. 1
- Jackson, Matthew O. and Brian W. Rogers**, “Meeting Strangers and Friends of Friends: How Random Are Social Networks?,” *American Economic Review*, 2007, 97 (3), 890–915. 3.1, 9, 10
- Keller, Wolfgang and Stephen R Yeaple**, “The Gravity of Knowledge,” *American Economic Review*, 2013, 103(4), 1414–1444. 1
- Klette, Tor Jakob**, “Market power, scale economies and productivity: estimates from a panel of establishment data,” *The Journal of Industrial Economics*, 1999, 47 (4), 451–476. E
- Oberfield, Ezra**, “Business Networks, Production Chains, and Productivity: A Theory of Input-Output Architecture,” Technical Report, Princeton University 2013. 1

- Ohnishi, T., H. Takayasu, and M. Takayasu**, “Network Motifs in Inter-firm Network,” *Journal of Economic Interactions and Coordination*, 2010, *5(2)*, 171–180. 7
- Olley, G. S. and A. Pakes**, “The Dynamics of Productivity in the Telecommunications Equipment Industry,” *Econometrica*, 1996, *64*, 1263–1297. 16, E
- Redding, Stephen J. and Matthew A. Turner**, “Transportation Costs and the Spatial Organization of Economic Activity,” in “Handbook of Urban and Regional Economics” forthcoming, chapter Transportation Costs and the Spatial Organization of Economic Activity. 1
- Saito, Yukiko U., Tsutomu Watanabe, and Mitsuru Iwamura**, “Do larger firms have more interfirm relationships?,” *Physica A: Statistical Mechanics and its Applications*, 2007, *383* (1), 158–163. 7
- Serrano, Ma Angeles and Marian Boguna**, “Topology of the World Trade Web,” *The Accounting Review*, 2003, *68* (1). 10
- Sveikauskas, Leo**, “The productivity of cities,” *The Quarterly Journal of Economics*, 1975, *89* (3), 393–413. 1
- Uzzi, Brian**, “The Sources and Consequences of Embeddedness for the Economic Performance of Organizations: The Network Effect,” *American Sociological Review*, 1996, *61(4)*, 674–698. 1

Table 1: Firm Performance.

	(1) Sales	(2) Sales/employee	(3) TFP
$Station_f \times H_j \times Post2004_t$	0.47** (2.12)	0.42* (1.76)	0.29** (2.44)
Firm and municipality controls	Yes	Yes	Yes
Prefecture-year FE	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
# obs	148,264	146,466	145,058
# firms	18,068	18,068	18,018
R-sq	0.97	0.92	0.94

Note: Robust t-statistics in parentheses. Dependent variables are in logs and are measured relative to industry-year means. *** significant at the 0.01 level, ** significant at the 0.05 level, * significant at the 0.1 level.

Table 2: Firm Performance: Falsification test.

	(1) Sales	(2) Sales/employee	(3) TFP
$Station_f \times H_j \times Post2000_t$	-0.30 (1.05)	-0.05 (0.22)	-0.02 (0.17)
Firm and municipality controls	Yes	Yes	Yes
Prefecture-year FE	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
# obs	66,756	66,756	66,487
# firms	14,165	14,165	14,158
R-sq	0.99	0.94	0.95

Note: Robust t-statistics in parentheses. Dependent variables are in logs and are measured relative to industry-year means. *** significant at the 0.01 level, ** significant at the 0.05 level, * significant at the 0.1 level.

Table 3: Firm Performance: 10km threshold.

	(1) Sales	(2) Sales/employee	(3) TFP
$Station_f \times H_j \times Post2000_t$	0.60** (2.28)	0.39* (1.80)	0.29** (2.57)
Firm and municipality controls	Yes	Yes	Yes
Prefecture-year FE	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
# obs	148,264	146,466	145,058
# firms	18,068	18,068	18,018
R-sq	0.97	0.92	0.94

Note: Robust t-statistics in parentheses. Dependent variables are in logs and are measured relative to industry-year means. *** significant at the 0.01 level, ** significant at the 0.05 level, * significant at the 0.1 level.

Table 4: Firm Performance: Construction excluded.

	(1) Sales	(2) Sales/employee	(3) TFP
$Station_f \times H_j \times Post2004_t$	0.52** (2.34)	0.43* (1.81)	0.31*** (2.61)
Firm and municipality controls	Yes	Yes	Yes
Prefecture-year FE	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
# obs	145,641	143,868	142,474
# firms	17,729	17,729	17,681
R-sq	0.97	0.92	0.94

Note: Robust t-statistics in parentheses. Dependent variables are in logs and are measured relative to industry-year means. *** significant at the 0.01 level, ** significant at the 0.05 level, * significant at the 0.1 level.

Table 5: Connections: Descriptive statistics.

	Mean	Median	Std.dev.	min	max
C_{ij2005}	7.05	2	62.35	1	16507
C_{ij2010}	7.51	2	59.93	1	14808
$\Delta \ln C_{ij}$	0.08	0	0.56	-3.47	3.78
$Both_{ij}$	0.01	0	0.08	0	1
One_{ij}	0.02	0	0.14	0	1
Firms per cell	104.21	19	463.30	1	21,207
# sources	7,613				
# destinations	8,054				
# obs	386,294				

Table 6: Shinkansen: Growth in connections.

	(1)	(2)	(3)	(4)
$Both_{ij}$	0.07*** (5.91)	0.12*** (7.91)	0.39*** (20.12)	0.42*** (7.93)
One_{ij}	-0.02*** (3.56)	-0.01 (0.74)	0.19*** (19.87)	0.15*** (6.42)
$\ln Dist_{ij}$			-0.06*** (71.32)	-0.06*** (81.98)
$Both_{ij} \times \ln Dist_{ij}$				-0.01 (0.86)
$One_{ij} \times \ln Dist_{ij}$				0.01* (1.87)
Destination FE	No	Yes	Yes	Yes
Source FE	No	Yes	Yes	Yes
# obs	386,294	386,294	386,294	386,294
# sources		7,613	7,613	7,613
# destinations		8,054	8,054	8,054
R-sq	0.00	0.17	0.18	0.18

Note: Bootstrapped t-statistics in parentheses with 200 replications. Dependent variable is $\Delta \ln C_{ij} = \ln C_{ij2010} - \ln C_{ij2005}$. *** significant at the 0.01 level, ** significant at the 0.05 level, * significant at the 0.1 level.

Table 7: Shinkansen: Growth in connections. 10 km threshold.

	(1)	(2)	(3)	(4)
$Both_{ij}$	0.09*** (3.46)	0.14*** (6.48)	0.44*** (23.87)	0.43*** (5.97)
One_{ij}	-0.05*** (6.02)	-0.01 (1.07)	0.20*** (23.58)	0.15*** (4.76)
$\ln Dist_{ij}$			-0.06*** (76.66)	-0.06*** (104.35)
$Both_{ij} \times \ln Dist_{ij}$				0.00 (0.16)
$One_{ij} \times \ln Dist_{ij}$				0.01* (1.72)
Destination FE	No	Yes	Yes	Yes
Source FE	No	Yes	Yes	Yes
# obs	386,294	386,294	386,294	386,294
# sources		7,613	7,613	7,613
# destinations		8,054	8,054	8,054
R-sq	0.00	0.17	0.18	0.18

Note: Bootstrapped t-statistics in parentheses with 200 replications. Dependent variable is $\Delta \ln C_{ij} = \ln C_{ij2010} - \ln C_{ij2005}$. *** significant at the 0.01 level, ** significant at the 0.05 level, * significant at the 0.1 level.

Table 8: Shinkansen: Extensive margin connections.

$Both_{ij}$	0.71*** (17.30)	0.72*** (17.28)	0.69*** (16.50)
One_{ij}	0.04*** (9.71)	0.05*** (7.87)	0.04*** (6.03)
$\ln Dist_{ij}$	-0.26*** (240.64)	-0.26*** (232.24)	-0.26*** (232.79)
Destination FE	No	No	Yes
Source FE	No	Yes	No
# obs	70,676,571	70,130,526	70,130,526
# sources		8,612	
# destinations			8,612
R-sq	0.00	0.00	

Notes: Coefficient estimates are multiplied by 100. t-statistics in parentheses. Dependent variable is $\Delta I[C_{ijt} > 0]$. *** significant at the 0.01 level, ** significant at the 0.05 level, * significant at the 0.1 level.

Appendix

A Optimal Search

The solution for $\bar{\tau}$ is

$$\bar{\tau}(z, j) = \kappa_1 \left(\frac{T}{w^\sigma} \frac{A}{f(j)} \right)^{1/\theta} \Phi(z, j)^{-k/\theta} z^{(\sigma-1)/\theta}.$$

Proof. The maximization problem of the firm is

$$\max_{\bar{\tau}} \{ \pi(z, j) - wf(j)n(z, j) \},$$

where $\pi(z, j) = Ap(z, j)^{1-\sigma}/\sigma$. A is the demand shifter for final goods, $A = w(1+\psi)\bar{L}Q^{\sigma-1}$, where Q is the CES price index for final goods. Given monopolistic competition and CES preferences, the firm charges a price that is a constant mark-up over marginal costs: $p(z, j) = \bar{m}w^\alpha P(z, j)^{1-\alpha}/z$, where $\bar{m} = \sigma/(\sigma-1)$. Inserting p and P into the profit function then yields

$$\pi(z, j) = \frac{(\bar{m}\lambda^{1-\alpha})^{1-\sigma}}{\sigma} Aw^{1-\sigma} \Phi(z, j)^{(\sigma-1)(1-\alpha)/\theta} z^{\sigma-1}.$$

The expressions for Φ and n are

$$\begin{aligned} \Phi(z, j) &= T_0 + \int_1^{\bar{\tau}(z, j)} T\tau^{-\theta} g(\tau, j) d\tau, \\ n(z, j) &= \int_1^{\bar{\tau}(z, j)} g(\tau, j) d\tau. \end{aligned}$$

Differentiating with respect to $\bar{\tau}$ yields

$$\begin{aligned} \frac{\partial \Phi(z, j)}{\partial \bar{\tau}} &= T\bar{\tau}^{-\theta} g(\bar{\tau}, j), \\ \frac{\partial n(z, j)}{\partial \bar{\tau}} &= g(\bar{\tau}, j). \end{aligned}$$

The first order condition is then

$$\frac{(\bar{m}\lambda^{1-\alpha})^{1-\sigma}}{\sigma} A \frac{(\sigma-1)(1-\alpha)}{\theta} w^{1-\sigma} \Phi(z, j)^{(\sigma-1)(1-\alpha)/\theta-1} z^{\sigma-1} T\bar{\tau}^{-\theta} = wf(j).$$

Rearranging,

$$\bar{\tau}(z, j) = \kappa_1 \left(\frac{T}{w^\sigma} \frac{A}{f(j)} \right)^{1/\theta} \Phi(z, j)^{-k/\theta} z^{(\sigma-1)/\theta},$$

where $\kappa_1 = \left(\frac{(\bar{m}\lambda^{1-\alpha})^{1-\sigma}}{\sigma} \frac{(\sigma-1)(1-\alpha)}{\theta} \right)^{1/\theta}$ and $k = 1 - (\sigma-1)(1-\alpha)/\theta$.

The second order condition is

$$\frac{(\bar{m}\lambda^{1-\alpha})^{1-\sigma}}{\sigma} A \frac{(\sigma-1)(1-\alpha)}{\theta} z^{\sigma-1} w^{1-\sigma} \left(-k\Phi^{-k-1} \left(\frac{\partial \Phi(z, j)}{\partial \bar{\tau}} \right)^2 + \frac{\partial^2 \Phi(z, j)}{\partial \bar{\tau}^2} \Phi(z, j)^{-k} \right) - wf(j) \frac{\partial^2 n(z, j)}{\partial \bar{\tau}^2} < 0.$$

Inserting the expressions for $\partial^2\Phi/\partial\bar{\tau}^2$ and $\partial n/\partial n^2$, this can be rewritten as

$$\frac{(\bar{m}\lambda^{1-\alpha})^{1-\sigma}}{\sigma} A \frac{(\sigma-1)(1-\alpha)}{\theta} z^{\sigma-1} w^{1-\sigma} \Phi^{1-k} \frac{T\bar{\tau}^{-\theta}}{\Phi} \left(-\frac{k}{\Phi} T\bar{\tau}^{-\theta} g^2 - \theta\bar{\tau}^{-1} g + g' \right) - wf(j)g' < 0.$$

Using the first order condition, we know that the following must hold on optimum:

$$\pi(z, j) \underline{\chi}(z, j) = wf \frac{\theta}{(\sigma-1)(1-\alpha)},$$

where $\underline{\chi}$ is the trade share from the marginal location $\bar{\tau}$, $\underline{\chi}(z, j) = T\bar{\tau}(z, j)^{-\theta} / \Phi(z, j)$. This tells us that gross profits from the marginal location $\bar{\tau}$ equals the fixed search cost f multiplied by the factor $\theta/(\sigma-1)(1-\alpha)$. Exploiting this relationship gives us the second order condition

$$k\underline{\chi}(z, j) g(\bar{\tau}, j) + \frac{\theta}{\bar{\tau}} > 0.$$

□

B Predictions of the Model

This section derives implications of the model described in Sections 4.3 and 4.4 of the main text.

B.1 The relationship between $\bar{\tau}$ and z

The cutoff $\bar{\tau}$ is increasing in z , $\partial\bar{\tau}/\partial z > 0$.

Proof. Using equation (1), we have

$$\frac{\partial\Phi}{\partial z} = \frac{\partial\bar{\tau}}{\partial z} T\bar{\tau}^{-\theta} g(\bar{\tau}, j).$$

Using equation (2), we have

$$\frac{\partial\bar{\tau}}{\partial z} = \kappa_1 \left(\frac{T}{w^\sigma} \frac{A}{f(j)} \right)^{1/\theta} \Phi(z, j)^{-k/\theta} z^{(\sigma-1)/\theta} \left(-\frac{k}{\theta} \Phi(z, j)^{-1} \frac{\partial\Phi}{\partial z} + \frac{\sigma-1}{\theta} z^{-1} \right).$$

Substituting in $\partial\Phi/\partial z$ and rearranging yields

$$\frac{\partial\bar{\tau}}{\partial z} = \frac{(\sigma-1)/z}{\theta/\bar{\tau} + k\bar{\chi}(z, j)g(\bar{\tau}, j)},$$

which is positive given that the regularity condition $\theta/\bar{\tau} + k\bar{\chi}(z, j)g(\bar{\tau}, j) > 0$ holds. □

This also implies that $\partial o/\partial z > 0$ because $\partial\Phi/\partial z > 0$.

B.2 The relationship between $\bar{\tau}$ and $f(j)$

The cutoff $\bar{\tau}$ is decreasing in costs $f(j)$, $\partial\bar{\tau}/\partial f(j) < 0$.

Proof. Using equation (1), we have

$$\frac{\partial\Phi}{\partial f} = \frac{\partial\bar{\tau}}{\partial f} T\bar{\tau}^{-\theta} g(\bar{\tau}, j).$$

Using equation (2), we have

$$\frac{\partial\bar{\tau}(z, j)}{\partial f} = -\frac{1}{\theta}\bar{\tau}(z, j) \frac{1}{f} \left(1 + k \frac{\partial\Phi}{\partial f} \frac{f}{\Phi(z, j)} \right).$$

Substituting in $\partial\Phi/\partial f$ and rearranging yields

$$\frac{\partial\bar{\tau}(z, j)}{\partial f} = \frac{-1/f}{\theta/\bar{\tau} + k\bar{\chi}(z, j)g(\bar{\tau}, j)},$$

which is negative given that the regularity condition $\theta/\bar{\tau} + k\bar{\chi}(z, j)g(\bar{\tau}, j) > 0$ holds. \square

Note that we can also express

$$\frac{\partial\Phi}{\partial f} \frac{f}{\Phi} = -\frac{\underline{\chi}(z, j)g(\bar{\tau}, j)}{\theta/\bar{\tau} + k\underline{\chi}(z, j)g(\bar{\tau}, j)} < 0.$$

Furthermore, $\partial o/\partial f < 0$ because $\partial\Phi/\partial f < 0$.

B.3 Assortivity

The expected measure of buyers from j among suppliers in z 's marginal market decreasing in efficiency z .

Proof. The expected measure of buyers from j for a task ω in location with trade costs τ (given the assumption of a unit continuum of tasks) is

$$m(j) \int_{\underline{z}(\tau)} \frac{T\tau^{-\theta}}{\Phi(z, j)} f(z, j),$$

where $f(z, j)$ is the density of productivity in location j and $\underline{z}(\tau, j)$ is the minimum efficiency z required in order to source from a location with trade costs τ .

The expected measure of buyers from j among suppliers in z 's marginal market is therefore

$$C(z_0, j) = m(j) T\bar{\tau}(z_0)^{-\theta} \int_{\underline{z}(\bar{\tau}(z_0))} \Phi(z, j) f(z, j).$$

$\partial C/\partial z_0 < 0$ because $\partial\bar{\tau}/\partial z_0 > 0$ (see Section B.1) and

$$\frac{\partial\underline{z}}{\partial\tau} = \frac{1}{\partial\bar{\tau}/\partial z} > 0.$$

\square

C Distributional Assumptions

Assume that τ is inversely Pareto distributed with support $[1, \tau_H]$ and shape $\gamma > \theta$: $g(\tau) = \gamma \frac{\tau_H^{-\gamma}}{1 - \tau_H^{-\gamma}} \tau^{\gamma-1}$. An inverse Pareto captures the empirical fact that a location j has few nearby markets and many remote markets; we show in Appendix D that the inverse Pareto is a reasonable approximation of the empirical distribution of distance in our data. In addition, we assume that a downstream firm's average productivity in task production, T_0 , is related to the average cost of purchasing tasks in the marketplace as follows: $T_0 = \frac{T \tau_H^{-\gamma}}{1 - \tau_H^{-\gamma}} \frac{\gamma}{\gamma - \theta}$.²⁴ Hence, a downstream firm cannot be too efficient in producing tasks itself, otherwise there would be no incentive to outsource. Given these additional assumptions, the hurdle $\bar{\tau}$, equilibrium market access Φ and measure of searched locations n are

$$\bar{\tau}(z, j) = \kappa_2 \left(\frac{A}{w^\sigma f(j)} \right)^{1/\omega} T^{(1-k)/\omega} z^{(\sigma-1)/\omega}, \quad (5)$$

$$\Phi(z, j) = \frac{T \tau_H^{-\gamma}}{1 - \tau_H^{-\gamma}} \frac{\gamma}{\gamma - \theta} \bar{\tau}(z, j)^{\gamma - \theta}, \quad (6)$$

$$n(z, j) = \frac{\tau_H^{-\gamma}}{1 - \tau_H^{-\gamma}} (\bar{\tau}(z, j)^\gamma - 1), \quad (7)$$

where $\omega = \theta + k(\gamma - \theta)$ and κ_2 is a constant.²⁵ The sourcing problem has an interior solution only if the second order condition, $\omega > 0$, is satisfied. Henceforth, we focus exclusively on the interior solution, i.e. $\omega > 0$.

Firm sales. Under the distributional assumption, we get

$$\frac{\partial \ln r(z, j)}{\partial \ln f(j)} = -\frac{(\sigma - 1)(1 - \alpha) \gamma - \theta}{\theta} \frac{\gamma - \theta}{\omega} < 0,$$

and

$$\frac{\partial^2 \ln r(z, j)}{\partial \ln f(j) \partial (1 - \alpha)} = -\frac{\sigma - 1}{\theta} \frac{\gamma - \theta}{\omega} \left(1 + (1 - \alpha) \frac{\sigma - 1}{\theta} \frac{\gamma - \theta}{\omega} \right) < 0,$$

hence the elasticity is more negative when $1 - \alpha$ is high.

D The distribution of trade costs

This section provides empirical support for the assumption that trade costs are inversely Pareto distributed with density

$$g(\tau) = \gamma \frac{\tau_H^{-\gamma}}{1 - \tau_H^{-\gamma}} \tau^{\gamma-1}. \quad (8)$$

²⁴Note that $\int_1^{\tau_H} T \tau^{-\theta} g(\tau, j) d\tau = \frac{T \tau_H^{-\gamma}}{1 - \tau_H^{-\gamma}} \frac{\gamma}{\gamma - \theta} (\tau_H^{\gamma - \theta} - 1)$, so T_0 equals $\int_1^{\tau_H} T \tau^{-\theta} g(\tau, j) d\tau / (\tau_H^{\gamma - \theta} - 1)$.

²⁵ $\kappa_2 = \kappa_1^{\theta/\omega} \left(\frac{\tau_H^{-\gamma}}{1 - \tau_H^{-\gamma}} \frac{\gamma}{\gamma - \theta} \right)^{-k/\omega}$

Let distance d from location i be inversely Pareto distributed with support $[0, d_{H_i}]$ and shape parameter $\kappa > 0$. The cdf is

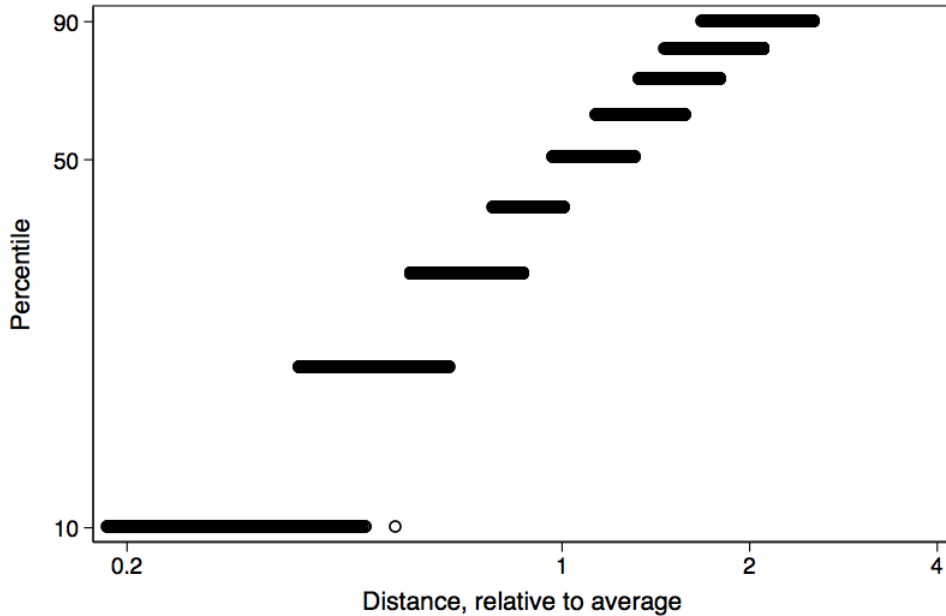
$$H_i(d) = \left(\frac{d}{d_{H_i}} \right)^\kappa.$$

Consider the 500x500 grid dataset described in Section 5.2. We calculate distance for every location pair ij and the empirical distribution of distance for each location i . Due to the large number of location-pairs, we limit the calculations to the 1st, 2nd, ..., 9th deciles of the distance distribution. From this, we obtain the k -th decile in location i , d_{ik} . If the distribution is inverse Pareto, the following must hold:

$$\ln H_{ik} = -\kappa \ln d_{iH} + \kappa \ln d_{ik}, \quad (9)$$

where H_{ik} takes the values 0.1 for $k = 1$, 0.2 for $k = 2$ and so on. The inverse Pareto should fit the data well if the relationship between H_{ik} and the d_{ik} is approximately log linear. Figure 9 plots H_{ik} against $d_{ik} - \bar{d}_i$, where $\bar{d}_i = (1/9) \sum_{k=1}^9 d_{ik}$, on log axes. The normalization removes the constant term d_{iH} which may vary across locations. Overall, the relationship is close to linear, although there is clearly heterogeneity in the distribution across locations. Estimating equation (9) with location fixed effects produces a slope coefficient κ of 1.07.

Figure 9: The cdf of distance across locations.



Note: Axes on log scales.

As is common in the literature, we assume that for large d , $\tau(d)$ is well approximated by the power law $\tau = (\alpha d)^\rho$ with $\alpha > 0$ and $\rho > 0$. Then τ inherits the distribution of d with shape

parameter $\gamma = \kappa/\rho$. Because $\tau \geq 1$, the τ density has support $[1, \tau_H]$, which yields the expression of $g(\tau)$ given in equation 8.

E Firm Performance

This section describes the measures of firm performance used in Section 5.1. As in Klette (1999), all firm level variables are demeaned relative to industry-year means, $\ln \hat{y}_{ijt} = \ln y_{ijt} - \ln \bar{y}_{jt}$, where $\ln \hat{y}_{ijt}$ refers to the demeaned variable for firm i in industry j at time t , $\ln y_{ijt}$ refers to the original log variable and $\ln \bar{y}_{jt}$ refers to the mean of the log variable in industry-year jt . The industry classification is 3-digit JSIC. Demeaning by industry-year has the benefit that it eliminates the need for deflating nominal variables, moreover it allows the technology of an industry to move freely over time.

TFP is estimated by the Olley and Pakes (1996) procedure. We estimate the gross production function

$$\ln Output_{it} = \beta_l \ln Labor_{it} + \beta_m \ln Materials_{it} + \beta_k \ln Capital_{it} + \omega_{it} + \eta_{it}, \quad (10)$$

where ω_{it} is total factor productivity of the firm and η_{it} is either measurement error or a shock to productivity which is not forecastable during the period in which labor can be adjusted. After obtaining the estimates $\hat{\beta}_l$, $\hat{\beta}_m$ and $\hat{\beta}_k$, TFP is calculated by subtracting predicted output from actual output,

$$tfp_{it} = \hat{\omega}_{it} = \ln Output_{it} - \hat{\beta}_l \ln Labor_{it} - \hat{\beta}_m \ln Materials_{it} - \hat{\beta}_k \ln Capital_{it}.$$

According to our model, the input price index $P(z, j)$ is firm-specific. As a consequence, as long as input prices are not properly controlled for in equation (10), tfp_{it} will be a function of input prices. This means that the infrastructure shock is expected to increase tfp_{it} for the treated firms because among these firms $P(z, j)$ falls. Formally, this can be seen as follows. Denote *observed* log materials as $\ln \tilde{m}_{it} = \ln Materials_{it} + \ln P_{it}$, where $Materials_{it}$ is the physical quantity (e.g., the CES aggregate) and P_{it} is the price (e.g., the CES price index). The production function in equation (10) can then be rewritten as

$$\ln Output_{it} = \beta_l \ln Labor_{it} + \beta_m \ln \tilde{m}_{it} + \beta_k \ln Capital_{it} + \tilde{\omega}_{it} + \eta_{it},$$

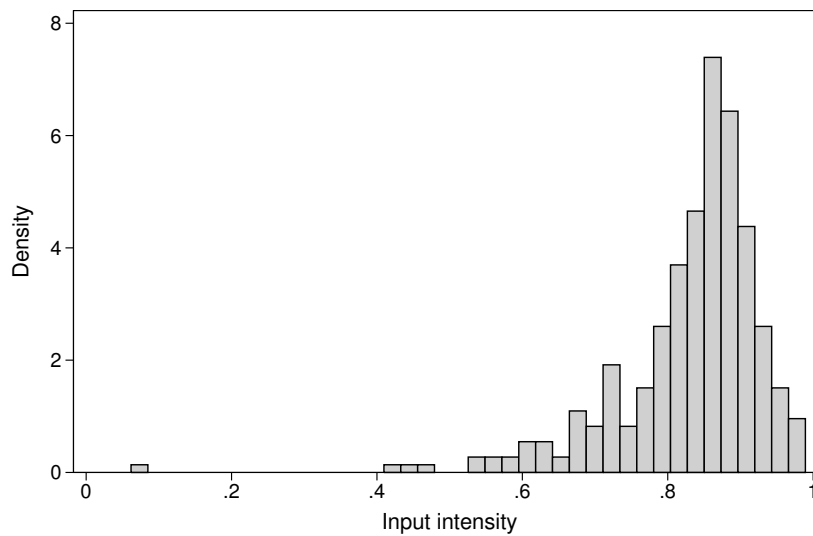
where $\tilde{\omega}_{it} = \omega_{it} - \beta_m \ln P_{it}$. Hence, using \tilde{m}_{it} instead of $Materials_{it}$ in the production function will produce a TFP estimate $\tilde{\omega}_{it}$ that confounds efficiency ω_{it} and the firm-specific input price P_{it} . The Shinkansen natural experiment in Section 5.1 is expected to lower P_{it} for treated firms. Hence, $\tilde{\omega}_{it}$ is expected to increase for treated firms.

F Data Appendix

F.1 Input intensity

Input intensity H_j is calculated as input costs relative to total costs for each JSIC 3-digit industry j in year 2003. Specifically, denote WC_j total wage costs for industry j , $WC_j = \sum_{i \in j} \text{wage costs}_i$ and total costs $TC_j = \sum_{i \in j} \text{total costs}_i$. H_j is then $H_j = 1 - WC_j/TC_j$. Figure 10 shows the density of H_j across all 315 JSIC industries.

Figure 10: Density of input intensity H_j across industries.



Notes: 2003 data.