NBER WORKING PAPER SERIES

FIRM-TO-FIRM TRADE:
IMPORTS, EXPORTS, AND THE LABOR MARKET

Jonathan Eaton
Samuel S. Kortum
Francis Kramarz

Working Paper 29685
http://www.nber.org/papers/w29685

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2022

The authors have presented versions of this paper at the 2013 IES Summer Workshop at Princeton, the 2016 RIETI Workshop in Tokyo, the 2018 NBER meeting on Trade and Networks, the 2019 Graham Lecture at Princeton, the 2019 China Meetings of the Econometric Society in Guangzhou, the 2019 conference of Baltic economists, and at seminars at EUI, Florence, University of Tokyo, APTS2019, Tokyo, University of Kent, Canterbury, Innovation seminar at College de France, CREST. Gregory Corcos, Gordon Hanson, John McLaren, Xiangliang Li, Yasusada Murata, Cristina Tello Trillo provided helpful comments. Lixing Liang, Etienne Guigue, Lorenzo Kaaks, Gautier Lenfant, Jonathan Libgober, Alexis Maitre, Berengere Patault, Max Perez Leon and Laurence Wicht have provided excellent research assistance. Eaton and Kortum gratefully acknowledge the support of the National Science Foundation under grant numbers SES-0339085 and SES-0820338. Kramarz has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme grant agreement 741467-FIRMNET. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Jonathan Eaton, Samuel S. Kortum, and Francis Kramarz. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Firm-to-Firm Trade: Imports, Exports, and the Labor Market
Jonathan Eaton, Samuel S. Kortum, and Francis Kramarz
NBER Working Paper No. 29685
January 2022
JEL No. F12,F14,F16

ABSTRACT

Customs data reveal heterogeneity and granularity of relationships among buyers and sellers. A key insight is how more exports to a destination break down into more firms selling there and more buyers per exporter. We develop a quantitative general equilibrium model of firm-to-firm matching that builds on this insight to separate the roles of iceberg costs and matching frictions in gravity. In the cross section, we find matching frictions as important as iceberg costs in impeding trade, and more sensitive to distance. Because domestic and imported intermediates compete directly with labor in performing production tasks, our model also fits the heterogeneity of labor shares across French producers. Applying the framework to the 2004 expansion of the European Union, reduced iceberg costs and reduced matching frictions contributed equally to the increase in French exports to the new members. While workers benefitted overall, those competing most directly with imports gained less, even losing in some countries entering the EU.

Jonathan Eaton
Department of Economics
Pennsylvania State University
303 Kern Graduate Building
University Park, PA 16801
and NBER
jxe22@psu.edu

Samuel S. Kortum
Yale University
Department of Economics
P.O. Box 208264
New Haven, CT 06520
samuel.kortum@yale.edu

Francis Kramarz
CREST, ENSAE
5 Avenue Henry Le Chatelier
91120 Palaiseau
FRANCE
kramarz@ensae.fr

1 Introduction

Two long-standing issues in international trade are the nature of frictions hampering exchange over greater distance and the extent to which goods procured abroad compete with local workers. Limitations on data restricted early work on these questions to the sectoral and then to the firm level. Recent access to data on firm-to-firm transactions allows an examination of the assumptions behind this earlier work, providing fresh evidence about international exchange and employment, and guiding new modeling to understand them.[1]

We exploit French customs records to decompose trade not only into the number of exporters to a market and their sales there, but also into their number of buyers. We can then focus on the individual buyer-seller relationship as the most fundamental unit of observation. One finding is that a country's larger sales to a destination involve not only more exporters, but more buyers per exporter.

While a few exporters have more than 100 buyers, most firms engage in only a small number of bilateral relationships. Confronting this granularity leads us to abandon anonymous market interaction in favor of firm-to-firm matching. We extend the Diamond-Mortensen-Pissarides job-creation framework to capture the granular yet polygamous nature of firm-to-firm encounters in international markets.[2]

Trade models have explained gravity, the decline of trade with distance, with the iceberg assumption of Samuelson (1954): Transporting goods over greater distances is costly. The matching framework suggests an alternative explanation proposed by Chaney (2014): Meeting trade partners farther away is hard.

Our model implies that, compared with lower iceberg costs, lower matching frictions raise buyers-per-exporter more than the number of exporters. We can thus disentangle the contributions of iceberg costs and matching frictions to gravity. Matching frictions are as powerful as icebergs in explaining the data, and are even more sensitive to distance. We find firms connect more readily in large markets, suggesting increasing returns to scale in matching.

To produce, the firms in our model need to execute a finite number of tasks, which can be performed either by their employees or with an input produced by another firm, as in

[1] Bernard and Moxnes (2018) survey the first wave of this literature. Blum et al. (2018), using customs data from South America, is an early contribution. Recent studies have explored data on domestic firm-to-firm transactions across locations, within Belgium (Dhyne et al. (2020)), Japan (Fujii et al. (2017), Furusawa et al. (2017), Bernard et al. (2019), Miyauchi (2021)), Turkey (Demir et al. (2021)), India (Panigrahi (2021)), Chile (Arkolakis et al. (2021)).

[2] See Diamond (1982) and Mortensen and Pissarides (1994).

Garetto (2013).[3] A producer's luck in finding inputs thus generates heterogeneity not only in its number of suppliers but in its use of labor. Our matching framework can replicate the vast heterogeneity in labor shares across French manufacturers.

In this setting firms are heterogeneous not only in their underlying efficiency, but in their relationships with suppliers and customers. Better luck in finding cheap suppliers upstream lowers cost downstream, enabling a firm to attract more buyers. In this sense our model resembles Oberfield (2018)'s theory of endogenous buyer-seller networks. We differ in that producers in our model have multiple tasks to perform, and can use either labor or inputs from another firm to perform them. While capturing this rich diversity in firm-level outcomes, our model, like Oberfield's, delivers a solution for the fixed point of firm-to-firm interactions and the consequent distribution of costs. The result is a tractable general equilibrium framework that we can connect with aggregate data on production, trade, and employment.

Our model explains how some workers are more vulnerable to foreign competition than others. We allow different types of labor to specialize in different types of tasks. When a type of task is more easily outsourced, the corresponding type of labor is more susceptible to displacement by imports. We apply our framework to analyze the distribution of income across different types of labor, distinguished in the data by educational attainment. Across a wide range of countries, the most educated workers, specializing in managerial tasks, benefit the most from lower trade barriers while the least educated, specializing in unskilled tasks most readily replaced by intermediates, may actually suffer.

Previous work has confronted a number of facts about the export participation of individual firms. Models developed to explain these facts, including our own, Eaton et al. (2011) (henceforth EKK), rely on firm heterogeneity in efficiency, monopolistic competition, and fixed entry costs, as pioneered by Melitz (2003). Our matching framework explains all of these earlier facts along with new ones revealed by the firm-to-firm data. Profits and fixed costs play no role.[4]

Subsequent work explains firm-to-firm transactions with a fixed cost to individual relation-

[3] A task-based approach to modeling production is not new to the trade literature. See Feenstra and Hanson (1996) and Grossman and Rossi-Hansberg (2008). It also plays a prominent role in the labor literature as described by Acemoglu and Autor (2011). To reflect the heterogeneity in our data we assume, in contrast to these papers, that production involves only an integer number of tasks.

[4] EKK relied on a set of destination-specific entry and demand shocks to explain heterogeneity in a firm's participation across different markets and on the Arkolakis (2010) specification of entry costs to explain why a firm might sell very little in an individual market. Our analysis here explains these facts via a firm's luck in attracting buyers in a destination.

ships rather than to market entry overall. Contributions are Bernard et al. (2018), Furusawa et al. (2017), Lim (2021), Bernard et al. (2021), and Arkolakis et al. (2021).[5] In these models production requires a continuum of intermediate inputs. They thus shut down (i) direct competition between outside suppliers and in-house production and (ii) random sourcing outcomes as a driver of firm heterogeneity, both of which play a central role in our analysis.

Several recent papers follow the same granular approach as ours. Lenoir et al. (2020), exploiting the product dimension of the French Customs data, use a Ricardian variant of the model here to uncover differences in matching frictions across sectors. Miyauchi (2021), in a dynamic framework, models turnover of suppliers across regions of Japan. On the basis of how fast a producer replaces a lost supplier, he finds increasing returns to scale in matching similar in magnitude to what we find here. Panigrahi (2021) models how the network of firm-to-firm links within and across Indian states shapes the size distribution of firms.[6]

We proceed as follows. Section 2 discusses the evidence. Section 3 presents our model of firm-to-firm trade, with Section 4 analyzing its implications for the observations discussed in Section 2. Section 5 describes our procedure for estimating the model, reporting the results and their implications. Section 6 examines two applications of our analysis. Section 7 concludes.

2 Evidence

We focus on production and trade in goods. Critical to our analysis are firm-level data from three sources.

2.1 EU Firm-to-Destination Data

The OECD's "Trade by Enterprise Characteristics" (TEC) reports the number of firms engaged in bilateral trade within the EU, either as exporters or as importers. We narrow our analysis to exporters in the industry and wholesale sectors.

For each of the 27 EU destinations we calculate, using data for 2012, the "exporter count share" as the number of exporters from each other EU source as a share of all EU firms

[5]The last paper introduces firm-to-firm matching through advertising similar to the quality-compatibility framework of Demir et al. (2021). This framework delivers an extensive margin of buyer-seller relationships but not of firm entry or selection.

[6]Eaton et al. (2021a,b) develop dynamic partial-equilibrium frameworks of granular firm-to-firm search and matching, exploiting customs records for U.S. imports of manufactures from Colombia and apparel from China.

exporting there.[7] We then calculate the corresponding "exporter value share" as the value of exports from each other EU source as a share of total EU exports to that destination.[8] Figure 1 plots the exporter count share against the exporter value share. (Figures 1-4 are on log-log scales.) The two shares are highly correlated, with a slope of 0.86 and a standard error of 0.01.

Figure 1: Exporter count shares vs. exporter value shares

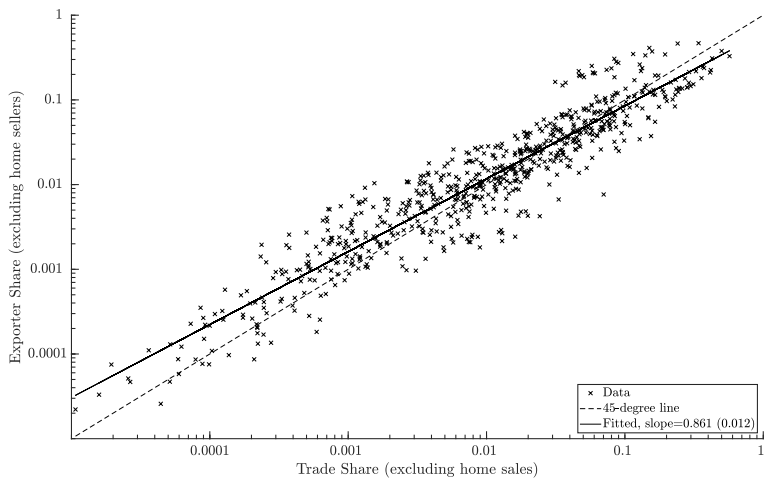

If exporter value share was totally determined by the number of firms exporting, the slope would be 1. That the slope is significantly less reflects the fact that a country with a larger value share in a destination not only has more firms selling there, but also that each firm, on average, sells more there. In our model, bilateral matching frictions account for both the slope of this relationship and its imperfect fit.

2.2 French Firm-to-Firm Data

French Customs report, for 2005, the sales of manufactures by each French exporter to individual buyers in each of 24 EU destinations. We limit ourselves to exporters we regard as "producers," i.e., in either manufacturing or wholesale.[9] Some definitions and identities help in organizing these data.

[7] Appendix A.1 describes how we construct these data. Because the quality and detail of the TEC data have improved over time, we use data for 2012 rather than for 2005, the year of our French firm-to-firm data described next.

[8] Appendix A.4 explains the construction of the value share measure from the World Input-Output Database (WIOD). See Timmer et al. (2015).

[9] In this second dataset we lose France as a destination since the firm-to-firm data from French Customs don't record domestic buyers. We lose Bulgaria and Romania as destinations since they didn't join the EU until 2007. Appendix A.2 describes how we construct these data.

At the aggregate level we observe the value of destination n's total absorption of manufactured goods (purchases from producers), denoted X_n^P, which we call "market size." We also observe X_{nF}^P, spending by n on manufactures from France.[10] The ratio gives us what we call "French market share," $\pi_{nF} = X_{nF}^P / X_n^P$.

At the firm level we observe the number N_{nF} of French exporters to destination n and the average number $\bar{b}_{nF}$ of buyers per French exporter there. We also observe the number F_{nF} of buyers from French exporters in destination n and the average number $\bar{s}_{nF}$ of French sellers per buyer there. Multiplying either of these two pairs together gives us the number R_{nF} of what we call *relationships* between French exporters and their buyers in n:

$$R_{nF} = N_{nF}\bar{b}_{nF} = F_{nF}\bar{s}_{nF}.$$

We also observe average sales $\bar{x}_{nF}$ per relationship with a French exporter. We can thus decompose French exports to n into sales per relationship and number of relationships $X_{nF}^P = R_{nF}\bar{x}_{nF}$.

Table 1 reports the results of regressing R_{nF}, $\bar{x}_{nF}$, N_{nF}, $\bar{b}_{nF}$, F_{nF}, and $\bar{s}_{nF}$ on market size and French market share (all in logs). The first regression shows that the number of relationships R_{nF} in a market varies with an elasticity of 0.81 with respect to market size and nearly in proportion to market share. The R^2 is 0.92. Contrasting column 1 with column 2, relationships, rather than sales per relationship, account for the bulk of the variation in total French exports.

The remaining four columns report the two ways of decomposing relationships. On the sellers' side, the number of French exporters (in column 3) accounts for over half of the variation with respect to market size and nearly two thirds of the variation with respect to market share.[11] Buyers per exporter in column 4 account for the remaining variation. On the buyers' side, the number of buyers (column 5), rather than French exporters per buyer (column 6), account for most of the variation in relationships.

[10] We take X_{nF}^P as the summation over producers in the French firm-to-firm data and X_n^P from WIOD, for 2005, as described in Appendix A.4.

[11] EKK perform the same regression as in column 3 using French data from 1986 with 112 foreign destinations. Coefficients on both market size and market share are somewhat larger. Bernard, Moxnes, and Ulltveit-Moe (2018) perform decompositions similar to ours using data on Norwegian exporters, with similar results. See their Figures 1 and 2 in particular.

Table 1
French Firm Entry into EU Destinations

	$\ln R_{nF}$	$\ln \bar{x}_{nF}$	$\ln N_{nF}$	$\ln \bar{b}_{nF}$	$\ln F_{nF}$	$\ln \bar{s}_{nF}$
constant	-2.80	2.80	-1.39	-1.41	-4.38	1.58
	(0.99)	(0.99)	(0.59)	(0.55)	(0.87)	(1.24)
market size	0.81	0.19	0.47	0.34	0.83	-0.02
	(0.06)	(0.06)	(0.04)	(0.03)	(0.05)	(0.01)
French market share	1.02	-0.02	0.64	0.38	0.85	0.17
	(0.19)	(0.19)	(0.11)	(0.11)	(0.17)	(0.05)
Number of Observations	24	24	24	24	24	24
R^2	0.92	0.33	0.91	0.86	0.93	0.40

Figure 2: French entry and market size

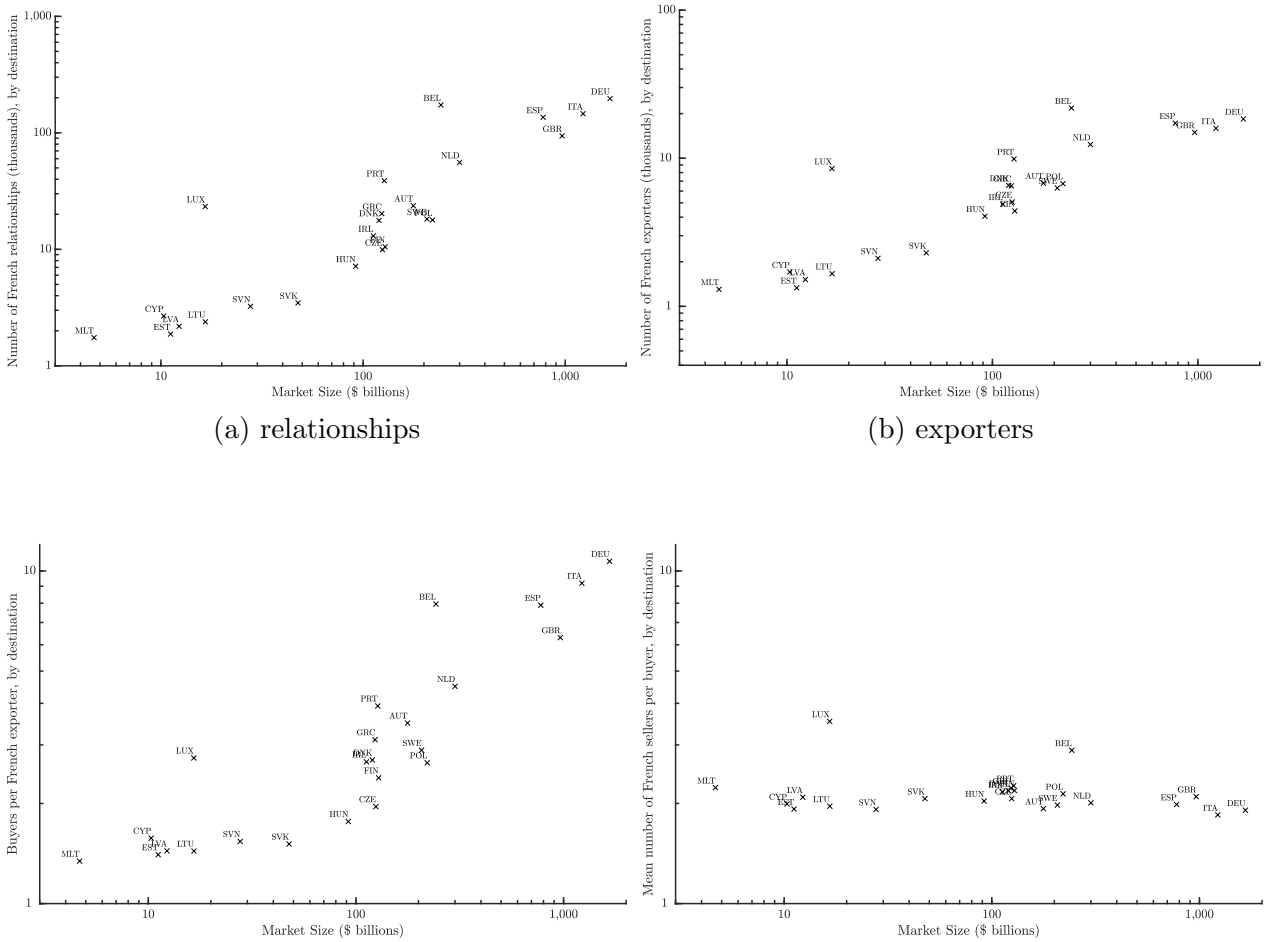

(a) relationships

(b) exporters

(c) buyers per French exporter

(d) French exporters per buyer

Figure 3: Number of partners (50th and 99th percentiles) and market size

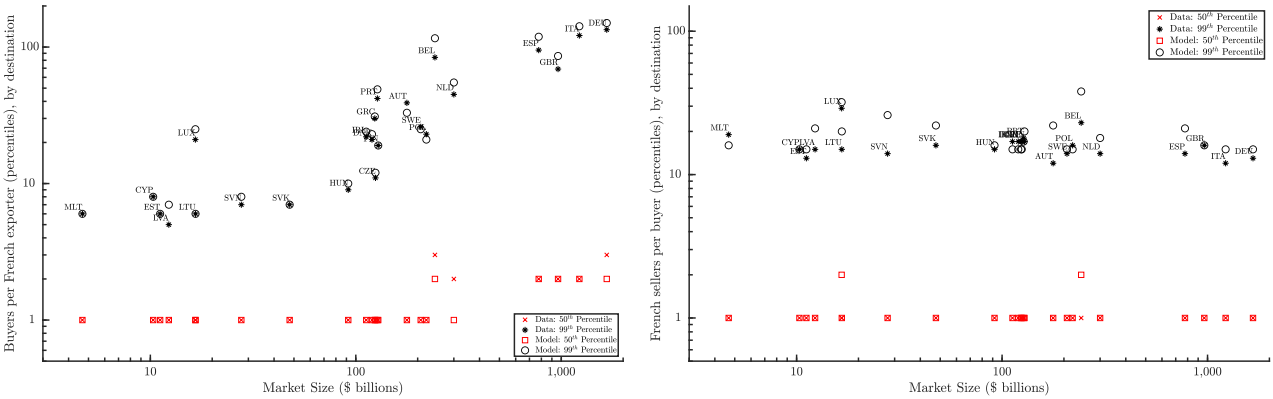

(a) buyers per French exporter

(b) French exporters per buyer

Illustrating the second row of Table 1, Figure 2 plots relationships R_{nF}, number of exporters N_{nF}, buyers per French exporter $\bar{b}_{nF}$, and French exporters per buyer $\bar{s}_{nF}$ against market size. The slopes are in line with the regression coefficients. Luxembourg and Belgium, with their large French market share, are notable positive outliers. (Table 10 reports our country codes.)

The mean number of buyers per French exporter or French exporters per buyer masks vast heterogeneity across firms within a destination. Figure 3 shows that the median ($\mathbf{x}$) is just 1 in most destinations, while the 99th percentile ($*$) slopes similarly to the mean.

German buyers. To examine further the distribution of buyers within a market we focus on Germany, the largest market. The $\mathbf{x}$ in Figure 4a plots the fraction of French exporters to Germany (on the y-axis) against their number of German buyers (on the x-axis). The frequency distribution has a mode of 1, generally declines, and has a long upper tail. Switching perspectives, Figure 4b shows the distribution of French exporters per German buyer. The shape is similar but steeper.

A French exporter's number of buyers in a destination correlates with its export activity elsewhere. The $\mathbf{x}$ in Figure 4c plots on the y-axis the average number of buyers in Germany of a French firm that also exports to the market indicated by the three-letter abbreviation. The x-axis reports the number of such French exporters to both Germany and that other market. Where the destination is DEU (Germany itself) the figure simply reports the average number of buyers per seller from Figure 2c (about 11) against the total number of French exporters to Germany (around 20,000). But for the roughly 1,300 that also export to Estonia (the second least popular alternative destination), the average number of buyers in Germany is nearly 40.

Figure 4: German buyers and French sellers

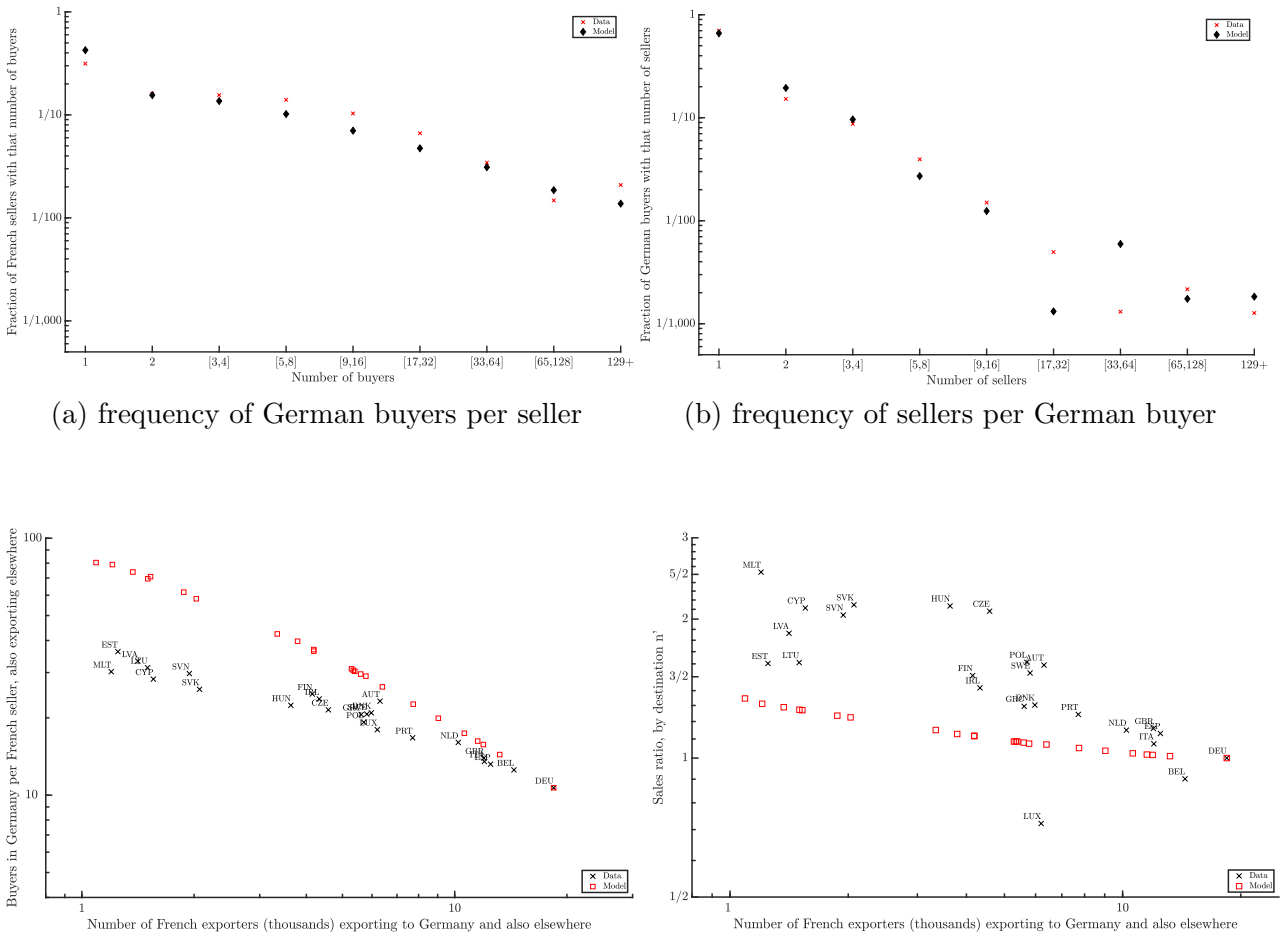

(a) frequency of German buyers per seller

(b) frequency of sellers per German buyer

(c) buyers per seller also exporting elsewhere

(d) sales ratio ($\bar{x}_{nF|n'}/\bar{x}_{nF}$) in n = Germany

As the number selling to the third market declines, the average number of buyers per exporter in Germany rises: French firms that succeed in penetrating a less popular market typically succeed in finding more buyers in Germany.

Just as a French firm that also penetrates a less popular market finds more customers in Germany, it also sells more to each of those German customers. We denote the average sales per relationship in n of a French firm also selling in market n' as $\bar{x}_{nF|n'}$. The **x** in Figure 4d plots the ratio $\bar{x}_{nF|n'}/\bar{x}_{nF}$ on the y-axis, where n is Germany and n' the indicated country, against the number of firms selling both to Germany and to destination n', on the x-axis. As in Figure 4c, the relationship is downward sloping: As fewer firms sell to the firm's other destination, the French firm sells more per customer in Germany.

Figure 5: Distribution of labor shares in production costs

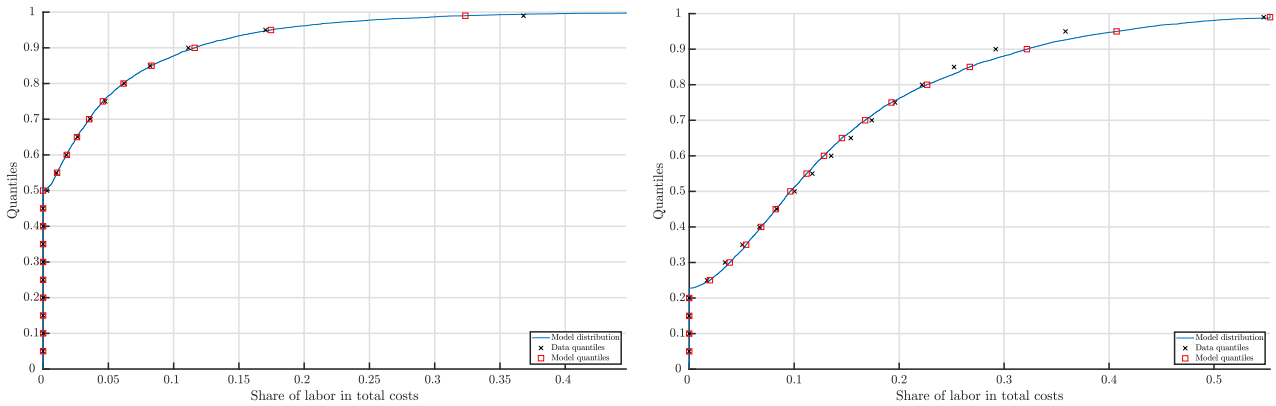

(a) unskilled labor

(b) skilled labor

2.3 French Firm Labor-Share Data

Our model pertains not only to the connections between firms and their customers in different destinations, but also to how firms hire labor or procure inputs from other firms to perform individual tasks. Standard general equilibrium models treat the production function as common across categories of firms, with the prediction that firms in the same category facing common factor prices in competitive input markets employ inputs in the same proportion.

We support our alternative approach with evidence from the Declaration Annuelle des Données Sociales (DADS) for 2005. We measure payments to production labor by French manufacturing firms as a fraction of their total variable costs, defined as the sum of intermediate purchases and payments to production labor. The distribution of the share of unskilled production workers across manufacturing firms appears in Figure 5a and of skilled production workers in Figure 5b. Both shares show enormous heterogeneity. For unskilled workers the share varies from 0 at the 45th percentile to 37 percent at the 99th. For skilled workers the share varies from 0 at the 20th percentile to 55 percent at the 99th.

3 A Model of Firm-to-Firm Trade

To understand these facts, which concern production and trade in goods, our model focuses on the goods sector. We incorporate services later to provide a general equilibrium closure. Our world has $i = 1, ..., \mathcal{N}$ countries, each with L_i^l workers of type $l \in \Omega^L$.

3.1 Producers

Production requires performing K types of tasks, with each type $k = 0, 1, 2, ..., K$ having country-specific Cobb-Douglas share $\beta_{k,i}$. (Because of constant returns to scale, the $\beta_{k,i}$ sum to one across k for each i.) Our assumptions allow each type k to differ according to both the type of labor $l(k)$ that can perform such a task and how readily an intermediate can replace labor in performing that task. Within each type, individual tasks combine with an elasticity of substitution σ.

Producer j, located in country i, has output:

$$Q_i(j) = z(j) \prod_{k=0}^{K} \left(\frac{1}{\beta_{k,i}} \left(\sum_{\omega=1}^{m(j)} q_{k,i}(j,\omega)^{(\sigma-1)/\sigma} \right)^{\sigma/(\sigma-1)} \right)^{\beta_{k,i}}, \tag{1}$$

where $z(j)$ is the producer's efficiency, $m(j)$ its number of tasks per type (the same for each type), and $q_{k,i}(j,\omega)$ the quantity of input chosen for task ω of type k.[12]

Any task ω of type $k \geq 1$ can be performed either in house by the appropriate type of labor, $l(k)$, or with an input made by another producer. The labor required to perform the task in house is $a_k(j,\omega)$. From producer j's perspective, the appropriate type of labor and the available inputs from other producers are perfect substitutes for performing a given task. It chooses whichever is cheapest.

Producers hire labor in a standard Walrasian market in which labor of type l in country i has a wage w_i^l. Given the mapping $l(k)$ between types of tasks and types of workers, we define the wage for a task of type k in country i as $w_{k,i} = w_i^{l(k)}$.

In finding intermediates, however, a buyer meets only an integer number of potential suppliers. In purchasing an intermediate, the buyer pays the seller's unit cost.[13]

We treat a task ω of type $k = 0$ differently. Such a task uses a combination of the

[12] In the special case in which $m = 1$, equation (1) reduces to a Cobb-Douglas production function. The point of introducing more than a single task of each type is to match more flexibly our data on firm-to-firm trade. Allowing heterogeneity across producers along this dimension captures the observation that some firms have a very large number of suppliers and others very few. An elasticity of substitution greater than one among tasks explains why a given buyer tends to spend more when buying from a low-cost seller.

[13] A justification is that the buyer and seller engage in Nash bargaining, with the buyer having all the bargaining power. An implication is that there are no variable profits. Our model thus can't accommodate fixed costs, either of market entry as in Melitz (2003) or in accessing markets for inputs, as in Bernard, Moxnes, and Ulltveit-Moe (2018) or Antràs et al. (2017). An alternative, which would allow for variable profits and hence fixed costs, is Bertrand pricing. While we found this alternative analytically tractable, we deemed the added complexity not worth the benefit.

appropriate type of labor $l(0)$ and a services input, with a combined unit cost $w_{0,i}$. For producer j, services-equipped labor required is $a_0(j,\omega)$.[14]

To summarize, the $\beta_{k,i}$'s are common across producers within a country while the number of types of tasks $K+1$ and σ are universal. Producers differ exogenously in their overall efficiency $z(j)$, the number of tasks $m(j)$ they require of each type, and labor required to perform each task $a_k(j,\omega)$. Differences across producers in unit cost derive not only from these exogenous sources of variation, but also from their fortune in finding suppliers.

Let $\tilde{c}_{k,i}(j,\omega)$ denote the lowest price available to producer j in country i for an intermediate to perform task ω of type $k \geq 1$. Its cost to perform this task is thus:

$$c_{k,i}(j,\omega) = \min\left\{a_k(j,\omega)w_{k,i}, \tilde{c}_{k,i}(j,\omega)\right\}.$$

From above, for a task ω of type $k=0$: $c_{0,i}(j,\omega) = a_0(j,\omega)w_{0,i}$. The producer's cost of delivering a unit of its own output to destination n is:

$$c_{ni}(j) = \frac{d_{ni}}{z(j)} \prod_{k=0}^{K} \left(\left(\sum_{\omega=1}^{m(j)} c_{k,i}(j,\omega)^{-(\sigma-1)} \right)^{-1/(\sigma-1)} \right)^{\beta_{k,i}}, \tag{2}$$

where $d_{ni} \geq 1$ is the iceberg cost of shipping a unit of output to destination n from source i, with $d_{ii} = 1$ for all i.

To derive a closed form solution for the distribution of costs, we impose specific distributions on the exogenous sources of producer heterogeneity:

First, following EKK, each country i has a Pareto measure of potential producers. Specifically, the measure with efficiency $z(j) \geq z$ and $m(j) = m$ tasks of each type is:

$$\mu_i^Z(z;m) = \frac{p(m)}{g_i(m)} T_i z^{-\theta}, \tag{3}$$

where $T_i \geq 0$ reflects the magnitude of country i's endowment of technologies and $\theta \geq 0$ their similarities. Here $p(m)$ is a probability distribution on the positive integers, with mean $\bar{m}$. (We choose the terms $g_i(m)$ below to neutralize the effect of heterogeneity in m on the distribution of costs.)

Second, we assume that labor requirements $a_k(j,\omega)$ are drawn, independently over ω, k,

[14] Task 0 guarantees that production always requires some labor, preventing unit cost from collapsing to zero.

and j, from the probability distribution:

$$F(a) = 1 - e^{-a^{\phi}}, \tag{4}$$

where $\phi \geq 0$ governs the similarity of labor requirements across tasks and producers.

Our specifications of the heterogeneity in producer efficiency given in (3), the distribution of labor requirements given in (4), and the distribution of numbers of tasks of each type $p(m)$ are primitives of the model, with T_i, θ, and ϕ exogenous parameters.

Our assumptions about technology, along with the specification of firm-to-firm matching in the next section, deliver the third distributional outcome. The measure of potential producers from i that can produce at unit cost below c is:

$$\mu_{ii}(c) = T_i \Xi_i c^{\theta}, \tag{5}$$

where $\Xi_i \geq 0$ is endogenous, as we show in Section 3.4.

3.2 Retailers

Production would have no point if all output simply served as input into further production. To give producers purpose, we introduce another type of firm, a retailer, which buys from producers (both domestic and foreign) but sells only locally.

Retailers have the same input structure as producers so also have the unit cost function (2), but valid only for $n = i$. They have the same distribution of labor productivity across tasks as (4) and the same distribution of tasks per type $p(m)$. Since their relative size plays no role in our model, we assign retailers a common efficiency $z = 1$.

We treat the measure F_i^R of retailers in country i as exogenous. Sales of individual retailers combine into a CES retail aggregate, with elasticity of substitution σ', that's purchased as final consumption by local households and as an intermediate by the service sector.

3.3 Buyer-Seller Matching

Unlike the measure of retailers, the measure F_i^P of producers is the endogenous outcome of random matching between a producer as a potential seller and either another producer or a retailer as a potential buyer.

Even though there are a continuum of possible sellers and buyers, an individual seller matches with only an integer number of potential buyers and, for any task, an individual buyer matches with only an integer number of potential sellers.

A producer is a buyer only if it is active in having, itself, found at least one buyer. We derive the measure F_n^P of active producers in Section 3.5. Our assumptions imply that whether a producer is active or not doesn't depend on its m. Hence, with the total measure of firms $F_n = F_n^P + F_n^R$ and with the average number of tasks of each type $\bar{m}$, the measure of potential purchases in market n for tasks of type k is $\bar{m}F_n$.

We now turn to the specification of firm-to-firm matching that underlies our analysis. For a task of type k, the intensity with which a given buyer (either a producer or retailer) in n meets a given producer from i that can deliver at unit cost c is:

$$\lambda_{k,ni}(c) = \lambda_k \lambda_{ni} B_n^{-\varphi} S_n(c)^{-\gamma}. \tag{6}$$

The parameter λ_{ni}, the bilateral matching intensity, reflects the ease with which buyers in n match with sellers from i. The parameter λ_k reflects the ease with which a buyer can find a supplier for a task of type k. We normalize the λ_k's to sum to K, restricting $\lambda_0 = 0$. The final two terms in (6) require a more extensive explanation.

The matching literature (e.g., Mortensen and Pissarides, 1994; Petrongolo and Pissarides, 2001) typically posits that, as the measure of buyers and sellers in a market increases, the likelihood of a match between any given buyer and seller falls. To capture such a "congestion effect" on the buyers' side we define the presence of buyers in n as:

$$B_n = \sum_k \lambda_k \bar{m} F_n = K\bar{m}F_n. \tag{7}$$

The parameter $\varphi \geq 0$ in (6) governs the extent to which buyers crowd each other out in meeting sellers.

From (5), the measure of potential producers from i who can deliver to n at a cost below c is:

$$\mu_{ni}(c) = \mu_{ii}(c/d_{ni}) = d_{ni}^{-\theta} T_i \Xi_i c^{\theta}. \tag{8}$$

To capture a congestion effect on the sellers' side we define the presence of sellers with unit

cost below c in market n as:

$$S_n(c) = \sum_{i'} \lambda_{ni'} \mu_{ni'}(c) = \Upsilon_n c^{\theta}, \tag{9}$$

where:

$$\Upsilon_n = \sum_i \lambda_{ni} d_{ni}^{-\theta} T_i \Xi_i.$$

The parameter $\gamma \in [0, 1)$ in (6) governs the extent to which low-cost sellers in n crowd out those with higher costs.

Our specification (6) of matching intensity embodies two asymmetries in the treatment of buyers and sellers. First, only active producers in need of inputs contribute as buyers to congestion in matching, while all potential producers, whether they make a sale or not, contribute as sellers to congestion in matching. Second, all buyers contribute to congestion symmetrically, while sellers congest only matches involving higher cost sellers, giving lower-cost sellers an advantage in matching.

An implication of (6) is that, for a task of type k, the number of encounters between buyers in destination n and a given seller from source i with unit cost *exactly* c is distributed Poisson with parameter:

$$e_{k,ni}(c) = \lambda_{k,ni}(c) \frac{B_n}{K} = \lambda_k \lambda_{ni} \frac{B_n}{K} B_n^{-\varphi} S_n(c)^{-\gamma}. \tag{10}$$

This equation delivers the following expressions for the measure of matches between:
(i) buyers in n and sellers from i with unit cost *below* c for a task of type k:

$$M_{k,ni}(c) = \int_0^c e_{k,ni}(c') d\mu_{ni}(c') = \frac{\lambda_k}{1-\gamma} \lambda_{ni} \mu_{ni}(c) \frac{B_n}{K} B_n^{-\varphi} S_n(c)^{-\gamma},$$

(ii) buyers in n and sellers from anywhere with unit cost below c for a task of type k:

$$M_{k,n}(c) = \sum_i M_{k,ni}(c) = \frac{\lambda_k}{1-\gamma} \frac{B_n}{K} B_n^{-\varphi} S_n(c)^{1-\gamma},$$

and (iii) buyers in n and sellers from anywhere with unit cost below c for *any* task:

$$M_n(c) = \sum_k M_{k,n}(c) = \frac{1}{1-\gamma} B_n^{1-\varphi} S_n(c)^{1-\gamma}. \tag{11}$$

In this last expression the measure of matches is a Cobb-Douglas combination of seller and

buyer presence.[15] The sum $\varphi + \gamma$ governs (negatively) returns to scale in matching, with a value of 1 implying constant returns.

Consider now a buyer in n seeking the cheapest input for a task of type k. From the matching intensity (6), the number of *quotes* below price c that it receives from sellers in i is distributed Poisson with parameter:

$$\rho_{k,ni}(c) = \frac{M_{k,ni}(c)}{B_n/K} = \frac{\lambda_k}{1-\gamma}\lambda_{ni}\mu_{ni}(c)B_n^{-\varphi}S_n(c)^{-\gamma}. \tag{12}$$

Aggregating across potential suppliers from each source i, the number of such quotes from anywhere is distributed Poisson with parameter:

$$\rho_{k,n}(c) = \frac{M_{k,n}(c)}{B_n/K} = \nu_{k,n}c^{\theta(1-\gamma)}, \tag{13}$$

where, using (9):

$$\nu_{k,n} = \frac{\lambda_k}{1-\gamma}B_n^{-\varphi}\Upsilon_n^{1-\gamma}.$$

Evaluating the Poisson distribution at zero, the probability that the buyer encounters no supplier with unit cost below c is $e^{-\rho_{k,n}(c)}$. Buyer j can also perform task ω with labor at unit cost $a_k(j,\omega)w_{k,n}$ which, from equation (4), will exceed c with probability $1 - F(c/w_{k,n})$. Since the two events are independent, the distribution of the lowest cost to fulfill such a task is:

$$G_{k,n}(c) = 1 - \exp\left[-\left(w_{k,n}^{-\phi}c^{\phi} + \nu_{k,n}c^{\theta(1-\gamma)}\right)\right].$$

In order to derive a closed-form expression for the distribution of production costs, we restrict $\phi = \theta(1-\gamma)$ so that the parameter governing heterogeneity of costs of intermediates becomes the same as the parameter governing heterogeneity of labor requirements (4). The distribution of the cost of fulfilling a task of type k simplifies to:

$$G_{k,n}(c) = 1 - \exp\left(-\Phi_{k,n}c^{\theta(1-\gamma)}\right), \tag{14}$$

where:

$$\Phi_{k,n} = \nu_{k,n} + w_{k,n}^{-\theta(1-\gamma)}.$$

[15] The matching function (11) resembles standard formulations in the labor literature, as reviewed by Petrongolo and Pissarides (2001).

(We drop ϕ from the notation in all that follows.)

3.4 Deriving the Cost Distribution

From the distribution of the cost of performing a task (14) we can derive the distribution of the cost of supplying a good. Using (3) and (2), the measure of potential producers from i with exactly m tasks of each type that can produce at a unit cost below c is:

$$\begin{aligned}\mu_{ii}(c;m) &= \frac{p(m)}{g_i(m)} T_i c^\theta \prod_k \left(\int_0^\infty ... \int_0^\infty \left(\sum_{\omega=1}^m c_\omega^{-(\sigma-1)} \right)^{\theta\beta_{k,i}/(\sigma-1)} dG_{k,i}(c_1)...dG_{k,i}(c_m) \right) \\ &= p(m) T_i c^\theta \prod_k \Phi_{k,i}^{\beta_{k,i}/(1-\gamma)}, \end{aligned} \tag{15}$$

where we've rigged $g_i(m)$, which first appeared in (3), to disappear (so that costs are independent of m).[16]

Aggregating over m gives:

$$\mu_{ii}(c) = \sum_m \mu_{ii}(c;m) = T_i c^\theta \prod_k \Phi_{k,i}^{\beta_{k,i}/(1-\gamma)},$$

confirming our conjecture about the cost distribution in equation (5), that $\mu_{ii}(c) = T_i \Xi_i c^\theta$, with:

$$\Xi_i = \prod_k \Phi_{k,i}^{\beta_{k,i}/(1-\gamma)} = \prod_k \left(\frac{\lambda_k}{1-\gamma} B_i^{-\varphi} \Upsilon_i^{1-\gamma} + w_{k,i}^{-\theta(1-\gamma)} \right)^{\beta_{k,i}/(1-\gamma)}. \tag{16}$$

The second equality employs (14) and (13).

Expressions (16) and (9) deliver the system of equations:

$$\Upsilon_n = \sum_i \lambda_{ni} d_{ni}^{-\theta} T_i \prod_k \left(\frac{\lambda_k}{1-\gamma} B_i^{-\varphi} \Upsilon_i^{1-\gamma} + w_{k,i}^{-\theta(1-\gamma)} \right)^{\beta_{k,i}/(1-\gamma)}, \tag{17}$$

for $i, n = 1, 2, ..., \mathcal{N}$. The solution, given the vectors B of buyers and w of wages, gives us

[16] Changing the variables of integration in (15) from c_ω to $x_\omega = \Phi_{k,i} c_\omega^{\theta(1-\gamma)}$ delivers:

$$g_i(m) = \prod_k \left(\int_0^\infty e^{-x_1} ... \int_0^\infty e^{-x_{m-1}} \int_0^\infty e^{-x_m} \left(\sum_{\omega=1}^m x_\omega^{-\frac{\sigma-1}{\theta(1-\gamma)}} \right)^{\frac{\theta\beta_{k,i}}{\sigma-1}} dx_m dx_{m-1} ... dx_1 \right),$$

which depends only on parameters. Appendix B.1 shows that $g_i(m)$ is finite as long as $\beta_{k,i} < 1-\gamma$ for all k.

the vector of Υ's.[17] Substituting the Υ's into equation (16) gives us Ξ_i, the term in the cost distribution (5).

Both B_i and $w_{k,i}$ are endogenous. We turn now to the measure of producers B_i, deferring $w_{k,i}$ until we characterize the aggregate equilibrium in Section 3.8.

3.5 From Potential to Active Producers

As mentioned earlier, a producer becomes an active firm only if it can find at least one buyer, either another active producer or a retailer. Consider a potential producer from i with unit cost c in market n. Its number of encounters with potential buyers needing to perform a task of type k is distributed Poisson with parameter $e_{k,ni}(c)$ given by (10).

But it's not enough for our producer just to encounter a buyer. To make a sale it has to beat out the competition (whether another supplier or labor). The probability that our producer with unit cost c in market n is the lowest cost among both the suppliers the buyer encountered for this task and labor is simply $1 - G_{k,n}(c)$, with $G_{k,n}(c)$ given by (14). This producer's number of buyers in n for tasks of type k is therefore distributed Poisson with parameter:[18]

$$\eta_{k,ni}(c) = e_{k,ni}(c)(1 - G_{k,n}(c)) = e_{k,ni}(c) \exp\left(-\Phi_{k,n} c^{\theta(1-\gamma)}\right).$$

Summing across k, this producer's number of buyers in market n is distributed Poisson with parameter:

$$\eta_{ni}(c) = \sum_k \eta_{k,ni}(c) = \lambda_{ni} B_n^{1-\varphi} \Upsilon_n^{-\gamma} c^{-\theta\gamma} \sum_k \frac{\lambda_k}{K} \exp\left(-\Phi_{k,n} c^{\theta(1-\gamma)}\right) \tag{18}$$

The expected number of buyers falls with c for two reasons: because of congestion a higher cost supplier encounters fewer buyers, and it's less likely to make the sale among those it does encounter.

[17] Restricting $\lambda_0 = 0$ (with $\beta_{0,i} > 0$) guarantees a unique solution for the Υ_n's, as shown in Appendix B.2, allowing us to use an iterative procedure to compute them. (Without this restriction it could be so easy to find input suppliers for all tasks that the cost of production would collapse to zero.) The uniqueness proof also gives us comparative statics. Each element of Υ increases in technology T_i anywhere, and with matching intensities λ_{ni} between any two countries. Each element of Υ decreases with trade costs d_{ni} between any two countries, with any task-specific wage $w_{k,i}$ in any country, and with the measure of buyers B_i in any country. Recall that these comparative statics take wages $w_{k,i}$ and the measures of buyers B_i as given.

[18] While a buyer in our model can make multiple purchases, the probability that more than one is from the same seller is zero. Hence a producer's number of sales is the same as its number of buyers.

A producer in i with unit cost c at home has unit cost cd_{ni} in market n. Summing across all markets, the number of buyers of such a producer is distributed Poisson with parameter:

$$\eta_i(c) = \sum_{n=1}^{\mathcal{N}} \eta_{ni}(cd_{ni}).$$

The probability that the producer has at least 1 buyer is $1 - e^{-\eta_i(c)}$, allowing us to express the measure of active producers in i as:

$$F_i^P = \int_0^\infty \left(1 - e^{-\eta_i(c)}\right) d\mu_{ii}(c). \tag{19}$$

Adding in the exogenous measure of retailers gives us $F_i = F_i^P + F_i^R$, delivering, from expression (7), the measure of buyers B_i.

Together the systems of equations (19) and (17) allow us, for given wages w_i^l around the world, to solve for the Υ_i's and B_i's.

3.6 Labor Shares

Consider a firm in n seeking to fulfill a task of type $k \geq 1$. From (14), the probability that the firm outsources the task using an intermediate is $\varpi_{k,n} = \nu_{k,n}/\Phi_{k,n}$. Hence with probability

$$1 - \varpi_{k,n} = \frac{w_{k,n}^{-\theta(1-\gamma)}}{\Phi_{k,n}}, \tag{20}$$

it hires workers of type $l(k)$ to perform this task. Note that this probability doesn't depend on the unit cost c of the input suppliers. Since there are a continuum of firms, $1 - \varpi_{k,n}$ is also the share of labor in performing this type of task in country n. The elasticity $\theta(1-\gamma)$ of the labor share with respect to the wage reflects heterogeneity in labor requirements.

3.7 Trade Shares

Say that the firm instead outsources the task, having found a supplier able to deliver at a cost below c. The probability that this supplier is from country i is the ratio of (12) to (13):

$$\pi_{ni} = \frac{\rho_{k,ni}(c)}{\rho_{k,n}(c)} = \frac{\lambda_{ni} d_{ni}^{-\theta} T_i \Xi_i}{\Upsilon_n} = \frac{\lambda_{ni} d_{ni}^{-\theta} T_i \Xi_i}{\sum_{i'} \lambda_{ni'} d_{ni'}^{-\theta} T_{i'} \Xi_{i'}}. \tag{21}$$

This probability is the same for any type of task $k \geq 1$ and any cost c. While the amount purchased by the firm in n depends on the cost, the probability that the source is i doesn't. Since there are a continuum of buyers, π_{ni} is also the bilateral trade share of source i in destination n's total absorption of goods.

As in many other trade models, the parameter θ is the trade elasticity with respect to iceberg trade costs. But, unlike in most other trade models, bilateral trade depends both on iceberg trade frictions, through $d_{ni}^{-\theta}$, and on matching frictions, through λ_{ni}, with the two interacting multiplicatively. A goal of our quantitative analysis is to isolate the contribution of each.

We now turn to the determination of wages w_i^l.

3.8 Aggregate Equilibrium

Our focus is firm-to-firm trade in goods. But our modelling is general equilibrium and services occupy a large share of the economy: (i) supplying final output to households; (ii) providing intermediates for making goods; (iii) employing labor; and (iv) using goods as inputs. To capture these relationships succinctly we exploit Cobb-Douglas. We introduce the shares of services α_n^S (and goods $\alpha_n^G = 1 - \alpha_n^S$) in final spending, the share of services in making goods β_n^{GS}, the share of goods in providing services β_n^{SG}, and the shares of different types of labor in services $\beta_n^{S,l}$ (which sum across l to the overall labor share in services $\beta_n^{S,L}$).[19] We set productivity in services in all countries to one. Recall that the service sector, like a final consumer, buys goods from retailers.

We assume that the goods sector uses services only to perform tasks of type 0, together with labor of type $l(0)$, in Cobb-Douglas combination. The fixed fraction of task 0 outsourced to services is thus:

$$\varpi_{0,n} = \beta_n^{GS} / \beta_{0,n}.$$

The share of labor of type l in goods production is the sum across its share in each type of task $k \in \Omega_l$ for which it's appropriate:

$$\beta_n^{G,l} = \sum_{k \in \Omega_l} \beta_{k,n} \left(1 - \varpi_{k,n}\right).$$

[19] We assume constant returns to scale in services so that $\beta_n^{S,L}$ and β_n^{SG} sum to one. Defining service sector output as net of services used in services lets us set $\beta_n^{SS} = 0$.

The overall labor share in the goods sector $\beta_n^{G,L}$ is the sum of the $\beta_n^{G,l}$ across l. The share of goods intermediates in the goods sector is thus $\beta_n^{GG} = 1 - \beta_n^{G,L} - \beta_n^{GS}$. Even though our basic technology in (1) is Cobb Douglas across types of tasks, the labor and intermediates shares in the goods sector depend on wages and deeper parameters.

As modelled in Section 3, goods are internationally traded. We treat retail and services as nontraded. To accommodate the data, we introduce exogenous services trade deficits D_n^S as well as exogenous goods trade deficits D_n^G for each country n. Total labor income (which corresponds to GDP) is:

$$Y_n = \sum_l w_n^l L_n^l.$$

Final spending X_n^F is GDP plus the overall deficit $D_n = D_n^G + D_n^S$. Final spending on goods is $\alpha_n^G X_n^F$ and on services $\alpha_n^S X_n^F$.

We denote the output of producers in country i as Y_i^P and the absorption of goods in country n, excluding the nontraded output of the retail sector, as X_n^P. Equilibrium in the world production of goods thus solves:

$$Y_i^P = \sum_n \pi_{ni} X_n^P. \tag{22}$$

Spending on labor of type l in country i is:

$$w_i^l L_i^l = \beta_i^{G,l} Y_i^G + \beta_i^{S,l} Y_i^S \tag{23}$$

where Y_i^G is output of the goods sector, including retail, and Y_i^S is output of services.

Equations (22) and (23) determine wages of each type of labor l, trade and labor shares, outputs of each sector, and final spending in each country. Appendix B.3 consolidates these equations into two sets of equilibrium conditions amenable to numerical solution.

3.9 Welfare and the Gains from Trade

Consider a household in country i with a budget W_i facing a retail price index P_i^R and service price index P_i^S. Its utility is $U_i = W_i \left(P_i^R\right)^{-\alpha_i^G} \left(P_i^S\right)^{-\alpha_i^S}$ where the price indices are as follows.

Indexing retailers by j, each with unit cost $c^R(j)$, the CES price index for retail, derived

in Appendix B.4, is:

$$P_i^R = \left[\int_0^{F_i^R} c_i^R(j)^{1-\sigma'} dj \right]^{1/(1-\sigma')} = g_i^R \Xi_i^{-1/\theta},$$

where g_i^R is a constant. Services combine inputs from retail and labor. The services price index is:

$$P_i^S = \left(P_i^R\right)^{\beta_i^{SG}} \prod_l \left(w_i^l\right)^{\beta_i^{S,l}} = \left(g_i^R\right)^{\beta_i^{SG}} \Xi_i^{-\beta_i^{SG}/\theta} \prod_l \left(w_i^l\right)^{\beta_i^{S,l}}.$$

To connect our analysis to the literature on the gains from trade consider the special case in which there is only one type of labor that can perform any type of task, so that each country i has only one wage w_i. Imposing balanced trade, w_i is also the budget. Incorporating the single wage into the price indices, the representative household has utility:

$$U_i = \left(w_i^\theta \Xi_i\right)^{(\alpha_i^G + \alpha_i^S \beta_i^{SG})/\theta},$$

which solves:

$$U_i = \prod_{k \geq 1} \left(\frac{\lambda_k}{1-\gamma} O_i U_i^{\theta(1-\gamma)/(\alpha_i^G + \alpha_i^S \beta_i^{SG})} + 1 \right)^{\beta_{k,i}(\alpha_i^G + \alpha_i^S \beta_i^{SG})/[\theta(1-\gamma)(1-\beta_i^{SG}\beta_i^{GS})]}, \tag{24}$$

with:

$$O_i = B_i^{-\varphi} \left(\frac{\lambda_{ii} T_i}{\pi_{ii}} \right)^{1-\gamma}.$$

Hence welfare U_i is increasing in O_i.[20] As in Eaton and Kortum (2002), welfare rises with T_i/π_{ii}. In the model here domestic matching intensity λ_{ii} enhances the contribution of domestic technology T_i while buyer congestion diminishes it. As in the class of models considered by Arkolakis et al. (2012), more openness in the form of a lower domestic trade share π_{ii} enhances welfare. Given λ_{ii} and T_i, whether greater openness (a lower π_{ii}) comes through lower trade barriers or lower bilateral matching frictions is irrelevant.

4 Implications for Firm-to-Firm Trade

How does our model relate to the observations in Section 2?

[20] The two sides of (24), as functions of U_i, have a single crossing point with the right crossing the left from above. The right-hand side is increasing in O_i.

4.1 Relationships

We start with the measure of relationships R_{ni} between buyers in destination n and sellers from source i. Consider a producer from i that can deliver to n at a unit cost c. Its expected number of customers there is $\eta_{ni}(c)$ given by (18). Integrating over the distribution of costs in n for firms from i:

$$R_{ni} = \int_0^\infty \eta_{ni}(c) d\mu_{ni}(c) = \pi_{ni} \overline{\varpi}_n B_n. \tag{25}$$

Here $\overline{\varpi}_n$ is the average of the $\varpi_{k,n}$ over $k \geq 1$, so that $\overline{\varpi}_n B_n$ is the measure of intermediate purchases undertaken by buyers in n.[21]

Equation (25) connects the model to column 1 of Table 1. The measure of relationships between buyers from n and sellers from i is proportional to i's trade share π_{ni}, given the overall size of destination n (captured by X_n^P in Table 1 and by B_n in expression (25)).

4.2 Sellers' Side

We follow Table 1 in next examining how relationships translate into sellers and their buyers.

4.2.1 Sellers

Our model delivers an expression for the measure N_{ni} of sellers to destination n from source i. A producer will sell in n if it has at least one customer there. Consider again a producer from i that can deliver to n at a unit cost c. The probability that it has at least one customer in n is $1 - \exp(-\eta_{ni}(c))$.

To calculate N_{ni}, we integrate this probability over the distribution of costs in n for firms from i:

$$N_{ni} = \int_0^\infty (1 - e^{-\eta_{ni}(c)}) d\mu_{ni}(c) = d_{ni}^{-\theta} \int_0^\infty (1 - e^{-\eta_{ni}(c)}) d\mu_{ii}(c). \tag{26}$$

[21] The derivation is as follows:

$$\begin{aligned}
\int_0^\infty \eta_{ni}(c) d\mu_{ni}(c) &= \lambda_{ni} d_{ni}^{-\theta} T_i \Xi_i B_n^{1-\varphi} \Upsilon_n^{-\gamma} \int_0^\infty \sum_k \frac{\lambda_k}{K} \exp\left(-\Phi_{k,n} c^{\theta(1-\gamma)}\right) \theta c^{\theta(1-\gamma)-1} dc \\
&= \pi_{ni} B_n \frac{1}{K} \sum_{k=1}^K \nu_{k,n} \int_0^\infty \exp\left(-\Phi_{k,n} c^{\theta(1-\gamma)}\right) \theta (1-\gamma) c^{\theta(1-\gamma)-1} dc \\
&= \pi_{ni} B_n \frac{1}{K} \sum_{k=1}^K \frac{\nu_{k,n}}{\Phi_{k,n}},
\end{aligned}$$

where the first equality uses (8) and (18), the second uses (21) and (13), and the third applies integration.

The iceberg cost d_{ni} has a proportionate effect on N_{ni} while, from (18), the effect of λ_{ni} diminishes. The measure of sellers with $b \geq 1$ buyers is:

$$N_{ni}(b) = \int_0^\infty \frac{e^{-\eta_{ni}(c)} \left[\eta_{ni}(c)\right]^b}{b!} d\mu_{ni}(c), \tag{27}$$

connecting the model to Figure 4a.

4.2.2 Buyers per Seller

Dividing the measure of relationships (25) by the measure of sellers (26) gives us buyers per seller:[22]

$$\bar{b}_{ni} = \frac{R_{ni}}{N_{ni}} = \frac{\overline{\varpi}_n \tilde{\lambda}_{ni}}{\int_0^\infty (1 - e^{-\tilde{\lambda}_{ni}\tilde{\eta}_n(x)})dx}, \tag{28}$$

where:

$$\tilde{\lambda}_{ni} = B_n^{1-\varphi/(1-\gamma)} \lambda_{ni}, \tag{29}$$

and:

$$\tilde{\eta}_n(x) = x^{-\gamma} \sum_k \frac{\lambda_k}{K} \exp\left(\frac{\lambda_k}{(1-\gamma)\,\varpi_{k,n}} x^{1-\gamma}\right). \tag{30}$$

Icebergs vanish in (28), leaving matching intensity as the sole bilateral determinant of $\bar{b}_{ni}$.

The elasticity of $\bar{b}_{ni}$ with respect to λ_{ni} is the fraction of sellers with more than one buyer:[23]

$$\frac{\partial \ln \bar{b}_{ni}}{\partial \ln \lambda_{ni}} = 1 - \frac{N_{ni}(1)}{N_{ni}} > 0.$$

[22] To derive this result, change the variable of integration in (26) to $x = B_n^{-\varphi/(1-\gamma)} \Upsilon_n c^\theta$ to obtain:

$$\begin{aligned} d\mu_{ni}(c) &= \frac{\pi_{ni} B_n}{\tilde{\lambda}_{ni}} dx, \\ \eta_{ni}(c) &= \tilde{\lambda}_{ni} \tilde{\eta}_n(x), \\ N_{ni} &= \frac{\pi_{ni} B_n}{\tilde{\lambda}_{ni}} \int_0^\infty (1 - e^{-\tilde{\lambda}_{ni}\tilde{\eta}_n(x)})dx. \end{aligned}$$

[23] Taking logs of (28) and differentiating:

$$\frac{\partial \ln \bar{b}_{ni}}{\partial \ln \lambda_{ni}} = 1 - \frac{\int_0^\infty \tilde{\lambda}_{ni}\tilde{\eta}_n(x) e^{-\tilde{\lambda}_{ni}\tilde{\eta}_n(x)} dx}{\int_0^\infty (1 - e^{-\tilde{\lambda}_{ni}\tilde{\eta}_n(x)})dx}.$$

The result then follows by setting $b = 1$ in (27) and rewriting as in footnote 22:

$$N_{ni}(1) = \frac{\pi_{ni} B_n}{\tilde{\lambda}_{ni}} \int_0^\infty \tilde{\lambda}_{ni}\tilde{\eta}_n(x) e^{-\tilde{\lambda}_{ni}\tilde{\eta}_n(x)} dx.$$

If most sellers have only one buyer the elasticity is close to zero: lowering matching frictions just expands the set of sellers.[24] But, as more sellers have more than one buyer, more new contacts will go to existing sellers, raising buyers per seller with an elasticity approaching 1.

Together, expressions (28) and (25) help in interpreting Figure 1. Relationships vary in proportion to market share π_{ni}, while buyers per seller increase with λ_{ni}. To the extent that higher bilateral matching intensities contribute to larger trade shares π_{ni}, they also contribute to more buyers per seller $\bar{b}_{ni}$. Hence, as in Figure 1, a source country's share of exporters to a market rises less than in proportion to its share of sales there.

4.2.3 Buyers per Seller, Conditional on Selling Elsewhere

Figure 4c shows that French exporters sell to more German buyers when they also sell to a third market that's harder for French firms to penetrate. We now show how our model explains such a relationship.

We calculate the measure of firms from i selling in n that also sell in n':

$$N_{ni(n')} = \int_0^\infty (1 - e^{-\eta_{ni}(c)})(1 - e^{-\eta_{n'i}(cd_{n'i}/d_{ni})}) d\mu_{ni}(c)$$

and the measure of relationships in n for firms from i that also sell in n':

$$R_{ni(n')} = \int_0^\infty \eta_{ni}(c) \left(1 - e^{-\eta_{n'i}(cd_{n'i}/d_{ni})}\right) d\mu_{ni}(c).$$

Buyers per seller conditional on selling in n' is simply:

$$\bar{b}_{ni|n'} = \frac{R_{ni(n')}}{N_{ni(n')}}. \tag{31}$$

A firm's ability to sell in n' indicates that it's likely to have a lower cost in n than the typical firm from i that sells there. A smaller number of firms from i selling in n' indicates a stronger selection effect, so that their average number of buyers in n is higher.

4.3 Buyers' Side

We now turn to how relationships translate into buyers and their suppliers. The measure F_{ni} of buyers in n purchasing from i is the counterpart to the measure N_{ni} of sellers to n from

[24] Applying L'Hôpital's rule: $\lim_{\lambda_{ni}\to 0} \bar{b}_{ni} = 1$ and $\lim_{\lambda_{ni}\to 0} \frac{N_{ni}(1)}{N_{ni}} = 1$.

i, while the random number s_{ni} of country-i sellers to a buyer in n is the counterpart to the random number b_{ni} of buyers in n from a country-i seller.

4.3.1 Buyers

As with N_{ni}, our model delivers an expression for F_{ni}. For a firm with m tasks of each type, the probability that it buys no task of type k from a supplier in i is $(1-\pi_{ni}\varpi_{k,n})^m$, where, recall, $\varpi_{k,n}$ is the probability a task of type k is outsourced and $\pi_{ni}\varpi_{k,n}$ is the probability it's outsourced to a supplier from i. The probability that a firm in n has no suppliers from i is thus:

$$\Pr\left[s_{ni}=0\right]=\sum_{m}p(m)\prod_{k=1}^{K}\left(1-\pi_{ni}\varpi_{k,n}\right)^{m}. \tag{32}$$

The measure of firms in n that buy from i is $F_{ni}=\left(1-\Pr\left[s_{ni}=0\right]\right)F_n$. The measure of importers in n from any foreign source is:

$$I_n=\left(1-\sum_{m}p(m)\prod_{k=1}^{K}\left(1-(1-\pi_{nn})\varpi_{k,n}\right)^{m}\right)F_n, \tag{33}$$

(which is less than the sum over $i\neq n$ of F_{ni} since the same firm may import from more than one country).

The measure $F_{ni}(s)$ of firms in n with $s\geq 1$ suppliers from i is

$$F_{ni}(s)=\left(\sum_{m=1}^{\infty}p(m)\Pr\left[s_{ni}=s|m\right]\right)F_n. \tag{34}$$

Defining $s_{k,ni}$ as the number of sellers from i fulfilling a task of type k, we can express:

$$\Pr\left[s_{ni}=s|m\right]=\sum_{s_1=0}^{s}\sum_{s_2=0}^{s-s_1}\quad\ldots\quad\sum_{s_K=0}^{s-s_1-s_2\ldots-s_{K-1}}\prod_{k=1}^{K}\Pr\left[s_{k,ni}=s_k|m\right]$$

where each $s_{k,ni}$ has a binomial distribution with parameters $(m,\varpi_{k,n}\pi_{ni})$. Expression (34) shows why we need heterogeneity across firms in their number m of tasks per type to generate the dispersion of sellers per buyer evident in Figures 3b and 4b.

4.3.2 Sellers per Buyer

Figures 2d and 3b in Section 2 report data on French suppliers per buyer across EU markets. Using equation (32), our model implies that the mean number of suppliers from i per buyer in n is:

$$\bar{s}_{ni} = \frac{R_{ni}}{F_{ni}} = \frac{B_n \pi_{ni} \overline{\varpi}_n}{(1 - \Pr[s_{ni} = 0]) F_n} = \frac{\bar{m} K \pi_{ni} \overline{\varpi}_n}{1 - \sum_m p(m) \prod_{k=1}^{K} (1 - \pi_{ni} \varpi_{k,n})^m}. \tag{35}$$

Note that n's market size doesn't appear (consistent with the coefficient of -0.02 on market size in the 6th column on Table 1).

In contrast with expression (28), in which buyers per seller depended on bilateral matching intensity λ_{ni} but not on iceberg costs d_{ni}, sellers per buyer depends on the overall trade share π_{ni}. But similar to buyers per seller, the elasticity of sellers per buyer with respect to market share is simply the fraction of buyers with more than one seller:[25]

$$\frac{\partial \ln \bar{s}_{ni}}{\partial \ln \pi_{ni}} = 1 - \frac{F_{ni}(1)}{F_{ni}}.$$

5 Estimation

Having shown that the model captures key properties of the French firm-to-firm data, we now turn to the quantification of its parameters. To do so we need to get specific about how our types of tasks connect to types of labor. We assume three types of each.

Our labor types follow WIOD-SEA's classification of workers by educational attainment: high (tertiary or t), medium (secondary or s), and low (primary or p). For tasks, we treat type 0 tasks as administrative, managerial, or engineering activities (which we call "managerial"), and assign tertiary labor to them. We assign primary p labor to type 1 tasks (which we call "unskilled production") and secondary s labor to type 2 tasks (which we call "skilled

[25] Taking logs of (35) and differentiating:

$$\frac{\partial \ln \bar{s}_{ni}}{\partial \ln \pi_{ni}} = 1 - \frac{\sum_m p(m) \sum_k \left(\prod_{k' \neq k} (1 - \pi_{ni} \varpi_{k',n})^m \right) m (\pi_{ni} \varpi_{k,n}) (1 - \pi_{ni} \varpi_{k,n})^{m-1}}{\Pr[s_{ni} \geq 1]}.$$

The last three terms in the numerator represent the probability of exactly one task of type k being purchased from a producer in i by a firm in n with m tasks per type. Taking account of all types of tasks and all values of m, the entire numerator turns out to be $\Pr[s_{ni} = 1]$. The result follows after multiplying the numerator and denominator by the measure F_n of firms in n.

production"). In short, $K = 2$, $\Omega^L = \{t, p, s\}$, $t = l(0)$, $p = l(1)$, and $s = l(2)$.

Having defined types of labor and tasks, we group the parameters of our model into five sets:

1. Our data don't identify θ, governing heterogeneity in producer efficiency. We set $\theta = 4$, centered in the range of estimates from other studies.[26]

2. The Cobb-Douglas shares $\{(\alpha_n^G, \alpha_n^S), (\beta_n^{S,t}, \beta_n^{S,p}, \beta_n^{S,s}, \beta_n^{SG})\}$ (where the parameters within parentheses sum to one) govern how services enter final demand, employ labor, and use goods as inputs in country n.

3. The share parameters $\{\beta_n^{GS}, \beta_n^{G,t}, (\beta_{0,n}, \beta_{1,n}, \beta_2)\}$ (where $\beta_{0,n} = \beta_n^{GS} + \beta_n^{G,t}$) govern goods production in country n.

4. The parameter σ governs substitutability between tasks within any type. The probabilities $p(m)$ govern the number of tasks per type, with support $\Omega^M = \{1, 4, 16, 64, 256, 1024, 4096\}$.

5. The parameters $\{\gamma, \varphi, \lambda_1, \lambda_2, \lambda_{ni}\}$ govern congestion in matching and match intensities.

5.1 Procedure for Estimation

We describe how we estimate these parameters according to the data that we use to identify them. The procedure is, to a large extent, modular, making more transparent which features of the data identify which parameters. Appendix C provides additional details.

5.1.1 Values Calibrated from WIOD

The WIOD gives us the service sector shares $\{(\alpha_n^G, \alpha_n^S), (\beta_n^{S,t}, \beta_n^{S,p}, \beta_n^{S,s}, \beta_n^{SG})\}$ and the good sector shares $(\beta_n^{G,t}, \beta_n^{G,p}, \beta_n^{G,s}, \beta_n^{GG}, \beta_n^{GS})$, as explained in Appendix A.4.[27] These shares are parameters for the service sector but are endogenous for the goods sector. We obtain the task shares in the goods sector, which are parameters, by estimating β_2 as described below and calculating:

$$\beta_{0,n} = \beta_n^{G,t} + \beta_n^{GS}; \quad \beta_{1,n} = 1 - \beta_{0,n} - \beta_2.$$

[26] The parameter θ corresponds to the trade elasticity with respect to an *ad valorem* tariff, as shown by (21). Using variation in tariffs, Caliendo and Parro (2015) obtain values between 3.5 and 4.5 when pooling across sectors. See also Head and Mayer (2014) and Imbs and Mejean (2015). Our data also fail to identify σ', the elasticity of substitution between goods, but it doesn't matter for our analysis.

[27] For France, to be consistent with our labor-share data, we use DADS rather than WIOD-SEA to split up goods-sector labor into the three types.

The implied outsourcing shares are:

$$\varpi_{0,n} = 1 - \frac{\beta_n^{G,t}}{\beta_{0,n}}; \quad \varpi_{1,n} = 1 - \frac{\beta_n^{G,p}}{\beta_{1,n}}; \quad \varpi_{2,n} = 1 - \frac{\beta_n^{G,s}}{\beta_2};$$

with $\overline{\varpi}_n$ the average of the last two.

To get the matching intensities λ_1 and λ_2, we compute the odds of outsourcing a task of type $k \in \{1, 2\}$:

$$o_{k,n} = \frac{\varpi_{k,n}}{1 - \varpi_{k,n}} = \frac{\nu_{k,n}}{w_{k,n}^{-\theta(1-\gamma)}} = \lambda_k \frac{B_n^{-\varphi} \Upsilon_n^{1-\gamma}}{(1-\gamma)\, w_{k,n}^{-\theta(1-\gamma)}},$$

where the second step uses (20). Taking the odds ratio for skilled to unskilled:

$$\frac{\lambda_2}{\lambda_1} = \frac{o_{2,n}}{o_{1,n}} \left(\frac{w_{2,n}}{w_{1,n}} \right)^{-\theta(1-\gamma)}.$$

We use this expression to solve for λ_1 and λ_2 (imposing our restriction that $\lambda_1 + \lambda_2 = 2$) by taking the geometric mean of the right-hand side across countries. We take the skill premium $w_{2,n}/w_{1,n}$ from WIOD-SEA, setting $\theta = 4$ and using γ as estimated below. Given the skill premium, if the odds of outsourcing are higher for an unskilled task we infer $\lambda_1 > \lambda_2$.

5.1.2 Parameters Estimated from French Firm-to-Firm Data

We use data on N_{nF}, $N_{nF}(b)$, F_{nF}, $F_{nF}(s)$, and R_{nF}, as described in Appendix A.2, to estimate the seller congestion parameter γ and the distribution of tasks per firm $\{p(m)\}$.

1. We estimate the parameter of seller congestion γ by minimizing the distance between the observed distribution of buyers per French exporter in different EU destinations n and the model expression:

$$\frac{N_{nF}(b)}{N_{nF}} = \frac{\int_0^\infty (\tilde{\lambda}_{nF} \tilde{\eta}_n(x))^b e^{-\tilde{\lambda}_{nF} \tilde{\eta}_n(x)} dx}{b! \int_0^\infty (1 - e^{-\tilde{\lambda}_{nF} \tilde{\eta}_n(x)}) dx}. \tag{36}$$

Here we've applied the change of variable in footnote 22 to (27). The parameter γ enters these expressions through $\tilde{\eta}_n(x)$, as shown in (30). A higher γ gives low-cost suppliers an advantage in meeting buyers, which complements their advantage in making a sale to the buyers they meet, generating a longer right tail in the distribution of buyers per seller in any destination n. We use the values of $\varpi_{k,n}$ from above to evaluate $\tilde{\eta}_n(x)$. We calculate the $\tilde{\lambda}_{nF}$ needed in (36) by inverting (28) with $\bar{b}_{nF} = R_{nF}/N_{nF}$.

2. We estimate $\left\{p(m) | m \in \Omega^M\right\}$ to minimize the distance between the observed distribu-

tion of French suppliers per buyer in different EU destinations n and the model's implication, from expressions (32) and (34):

$$\frac{F_{nF}(s)}{F_{nF}} = \frac{\sum_{m\in\Omega^M} p(m)\Pr\left[s_{nF}=s|m\right]}{1-\Pr\left[s_{nF}=0\right]}. \tag{37}$$

Explaining why some buyers have many French suppliers requires that a very large m is possible. But most importers in the data have only one French supplier, indicating that they have only a small number of tasks. The estimated $\{p(m)\}$ yield a value for the mean number of tasks of each type $\bar{m}$.[28]

5.1.3 Parameters Estimated from French Firm Labor-Share Data

We estimate the elasticity of substitution across tasks, σ, and the share of skilled tasks, β_2, to minimize the distance between the model's implied distributions of labor shares and the distributions in the data, shown in Figure 5a for unskilled ($k = 1$) and Figure 5b for skilled ($k = 2$). While our model doesn't deliver closed-form expressions for these distributions, they're easy to simulate (as described in Appendix C.1). The parameter β_2 is identified by the upper support of the distributions of labor shares. The identification of σ is more subtle. If $\sigma = 1$, given m, the distributions of labor shares are binomial with probability of "success" $1-\varpi_{k,n}$ for $k = 1, 2$. A higher σ gives a firm more leeway to substitute its spending into tasks with lower cost inputs, generating the smoother distribution we observe.

5.1.4 Parameters Estimated from EU Firm-to-Destination (TEC) Data

We now turn to how we estimate the bilateral matching intensities λ_{ni} and the buyer congestion parameter φ. We use $\varpi_{1,n}$, $\varpi_{2,n}$, γ, and $\{p(m)\}$ from above; measures of bilateral exporters N_{ni}, total importers I_n, and non-retail importers I_n^P as described in Appendix A.1; data on 2012 bilateral trade shares π_{ni} as described in Appendix A.4; and a model-based estimate of the measure of bilateral relationships.

Steps 1-3 show how we estimate the number of buyers per seller $\bar{b}_{ni}$ for $n \neq i$. Step 4 uses the estimated $\bar{b}_{ni}$ to back out the $\tilde{\lambda}_{ni}$. Step 5 addresses $\tilde{\lambda}_{ii}$. Step 6 uses the $\tilde{\lambda}_{ni}$ to estimate the buyer congestion parameter φ and to recover the λ_{ni}

[28] Appendix C.1 explains the calculation of the probabilities $\Pr\left[s_{nF}=s|m\right]$, which makes use of the 2005 French market share data π_{nF} described in Appendix A.4.

1. We invert (33) to back out an estimate of the measure of firms $\hat{F}_n$ in destination n, where I_n is the number of firms (producers and retailers) in n reporting imports from other EU members. Since $\hat{F}_n$ incorporates noise both in our importer data and in our calculation of the probability of importing, we project $\ln \hat{F}_n$ onto $\ln X_n^P$ to obtain our measure of firms F_n. (Figure 8 in the appendix depicts the projection.)

2. Our measure of buyers B_n in destination n is F_n from Step 1 times average tasks per firm $K\bar{m}$, as given in (7).

3. An initial estimate of buyers per seller $\widehat{\bar{b}}_{ni}$ in market n from source i is:

$$\widehat{\bar{b}}_{ni} = \frac{\pi_{ni}\overline{\varpi}_n B_n}{N_{ni}}.$$

The numerator relates to R_{ni} as given in (25). This procedure delivers $\widehat{\bar{b}}_{ni} < 1$ for 21 of our 702 bilateral pairs. Since $\bar{b}_{ni}$ must exceed 1, we estimate:

$$\widehat{\bar{b}}_{ni} - 1 = \exp(\beta' x_{ni}) + \varepsilon_{ni}.$$

for $n \neq i$ using Poisson pseudo maximum likelihood. Here the vector x_{ni} includes effects for source i and destination n, distance between n and i (in logs), and other bilateral indicators, as described in Appendix C.2. From this regression, reported in Table 15, we extract our measure of buyers per seller as:

$$\bar{b}_{ni} = 1 + \exp(\hat{\beta}' x_{ni})$$

4. We back out $\tilde{\lambda}_{ni}$ $n \neq i$ to fit $\bar{b}_{ni}$ by inverting (28).

5. To get $\tilde{\lambda}_{ii}$ we first impute the number of producers in country i, F_i^P, from TEC's count of the producers there that import from any EU source (I_i^P) and the count of all firms (producers and retailers) that import from any EU source (I_i), setting:

$$F_i^P = \frac{I_i^P}{I_i} F_i,$$

with F_i taken from Step 1. Rewriting (19), using the change of variable $y = T_i \Xi_i c^\theta$, we get:

$$F_i^P = \int_0^\infty \left[1 - \exp\left(-\tilde{\lambda}_{ii}\tilde{\eta}_i \left(\frac{\tilde{\lambda}_{ii}}{\pi_{ii} B_i} y \right) - \sum_{n \in EU, n \neq i} \tilde{\lambda}_{ni}\tilde{\eta}_n \left(\frac{\overline{\varpi}_n \tilde{\lambda}_{ni}}{\bar{b}_{ni} N_{ni}} y \right) \right) \right] dy. \tag{38}$$

Given F_i^P and the $\tilde{\lambda}_{ni}$ for $n \neq i$ from Step 4, we find the value of $\tilde{\lambda}_{ii}$ that satisfies this expression for each country i.

6. We recover the bilateral matching intensities λ_{ni} from the $\tilde{\lambda}_{ni}$ (from Steps 4 and 5), our estimates of B_n (from Step 2), our estimate of supplier congestion γ, and an estimate of the buyer congestion parameter φ. To estimate φ we make the identifying assumption that $\ln \lambda_{ni}$ is orthogonal to $\ln B_n$. Based on (29) we regress:

$$\ln \tilde{\lambda}_{ni} = S_i + \frac{1-\gamma-\varphi}{1-\gamma} \ln B_n + \varepsilon_{ni}, \tag{39}$$

pooling across n and i. Here S_i is a fixed effect for source country i and ε_{ni} is the residual. Under our identifying assumption, we can infer φ (and hence returns to scale in the matching function $1-\gamma-\varphi$) from the coefficient on $\ln B_n$ (with a coefficient of 0 implying constant returns to scale and a positive coefficient indicating increasing returns). We obtain the bilateral matching intensities as:

$$\lambda_{ni} = \frac{\tilde{\lambda}_{ni}}{B_n^{(1-\gamma-\varphi)/(1-\gamma)}}. \tag{40}$$

5.2 Parameter Estimates

Table 2 reports our share parameters (averaging, except for β_2, across countries). The task shares substantially exceed the labor shares, reflecting outsourcing. The implied average outsourcing probability for an unskilled task is $\varpi_{1,n} = 0.87$ and for a skilled task is $\varpi_{2,n} = 0.73$. Our estimate $\beta_2 = 0.38$ implies upper bounds of 0.45 and 0.55 on the unskilled and skilled French labor shares shown in Figure 5.

Table 2
Preference and Production Shares

	final	labor			tasks			intermediates	
shares:	demand	t	p	s	0	1	2	G	S
service sector	α_n^S	$\beta_n^{S,t}$	$\beta_n^{S,p}$	$\beta_n^{S,s}$				β_n^{SG}	
	0.69	0.29	0.12	0.30				0.29	
goods sector	α_n^G	$\beta_n^{G,t}$	$\beta_n^{G,p}$	$\beta_n^{G,s}$	$\beta_{0,n}$	$\beta_{1,n}$	β_2	β_n^{GG}	β_n^{GS}
	0.31	0.04	0.04	0.10	0.29	0.33	0.38	0.57	0.24
s.e.							(0.03)		

Table 3 reports our estimates of σ, the elasticity of substitution between tasks of a given type, and the $p(m)$'s, the probability that a producer has m tasks per type. The median

number of tasks per type is below 4, while the mean is nearly 17, reflecting right skewness. This heterogeneity across firms in their numbers of tasks helps us fit the mean, median, and upper tail of French sellers per buyer shown in Figures 2d and 3b. The value of $\sigma > 1$ flattens the distribution of labor shares to match those in Figures 5a and 5b.

Table 3
Task Substitutability and Frequency

	σ	$p(1)$	$p(4)$	$p(16)$	$p(64)$ $\times 10$	$p(256)$ $\times 10^2$	$p(1024)$ $\times 10^2$	$p(4096)$ $\times 10^3$	$\bar{m}$
value	2.64	0.049	0.53	0.37	0.48	0.53	0.20	0.60	16.97
s.e.	(0.10)	(0.008)	(0.04)	(0.05)	(0.27)	(0.27)	(0.08)	(0.31)	(3.07)

Table 4 reports our estimates for the seller and buyer congestion parameters and matching intensities for each type of task. Expression (14) shows the role of γ in shifting the effective cost distribution toward lower costs: More seller congestion crowds out high cost sellers, enabling low cost firms to reach and to sell to more buyers, helping deliver the fat tail in Figure 4a. The higher matching intensity of unskilled tasks aligns with the higher probability of outsourcing these tasks.

Table 4
Matching Parameters

	congestion		intensities	
	seller	buyer	unskilled	skilled
	γ	φ	λ_1	λ_2
value	0.34	0.34	1.65	0.35
s.e.	(0.03)	(0.02)	(0.06)	(0.06)

Figure 6 plots our estimates of the $\tilde{\lambda}_{ni}$'s (at 729 too numerous to report individually) against our measure of buyers in each destination (both on log scales).[29] From equation (39), the theoretical slope is $1 - \varphi/(1-\gamma)$. The positive slope in the figure indicates increasing returns.[30] We estimate $\varphi = 0.34$ from this relationship, using our estimate of $\gamma = 0.34$. Our estimates imply that a ten percent increase in both buyers and sellers leads to a nearly

[29]Our procedure for estimating γ from the French firm-to-firm data, described in Section 5.1.2, yields one set of estimates for $\tilde{\lambda}_{nF}$ while our procedure in Section 5.1.4 based on the EU firm-to-destination data yields another. For compatibility across source countries we use the second. The correlation between the two, in logs, is 0.93.

[30]The slope in the figure is 0.48. The slope in the regression, which includes source-country effects, is also 0.48 (standard error 0.03).

fourteen percent increase in matches. Matching is easier in a larger market, implying larger firms, more buyers per seller, and greater sales per buyer.

Figure 6: Identification of buyer congestion φ

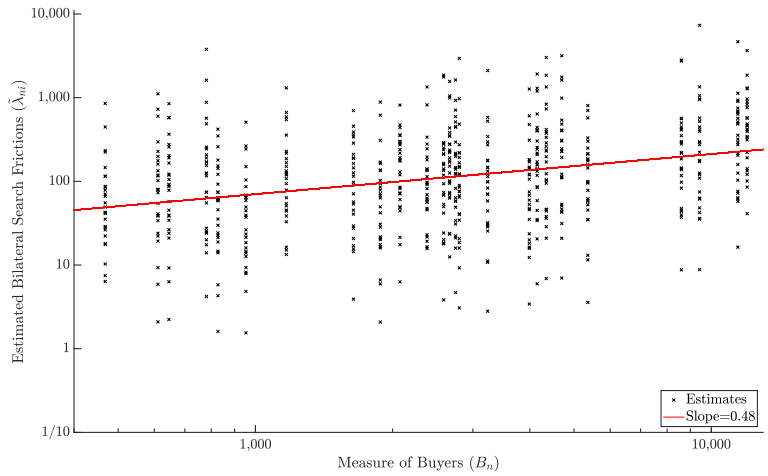

5.3 Implications for Gravity

From equation (21), bilateral matching intensities, along with iceberg costs, govern variation in bilateral trade shares according to:

$$\pi_{ni} = \frac{\lambda_{ni} d_{ni}^{-\theta} T_i \Xi_i}{\Upsilon_n} = \frac{\tau_{ni}^{-\theta} T_i \Xi_i}{\Upsilon_n}. \tag{41}$$

In the metric of iceberg trade costs, $\lambda_{ni}^{-1/\theta}$ is the implied bilateral matching friction. The overall trade friction is then $\tau_{ni} = d_{ni}\lambda_{ni}^{-1/\theta}$. We now examine the relative contributions of the two using two different approaches.

5.3.1 The Head-Ries Index

First, we use the Head-Ries index to assess the overall magnitude and variation of these two trade frictions.[31] We calculate the Head-Ries index:

$$H_{ni} = \sqrt{\frac{\pi_{ni}\pi_{in}}{\pi_{ii}\pi_{nn}}} = \left(\frac{\tau_{ni}\tau_{in}}{\tau_{ii}\tau_{nn}}\right)^{-\theta},$$

[31] See Head and Ries (2001).

using our trade-share data. It measures (inversely) the trade frictions between countries n and i, as can be seen from (41). We can decompose it into bilateral matching frictions and iceberg costs:

$$H_{ni} = \sqrt{\frac{\lambda_{ni}\lambda_{in}}{\lambda_{ii}\lambda_{nn}}} \times (d_{ni}d_{in})^{-\theta} = H^{\lambda}_{ni} \times H^{d}_{ni}.$$

We compute H^{λ}_{ni} directly from the λ_{ni} and then back out $H^{d}_{ni} = H_{ni}/H^{\lambda}_{ni}$.

Table 5
Head-Ries Decomposition

	$\ln H_{ni}$	$\ln H^{\lambda}_{ni}$	$\ln H^{d}_{ni}$
Mean	-4.96	-2.47	-2.49
Residual variance	0.80	0.31	0.26
Variance decomposition		0.54	0.46

Notes: The number of observations is 351. In calculating the statistics reported in the last two rows we first remove source and destination effects from each series. The last row reports coefficients from regressions of each component on the total.

Table 5 reports various statistics from this decomposition. Implications are that our estimates of matching frictions and iceberg costs are similar in their absolute magnitudes, variation, and contribution to overall trade frictions. The variance in $\ln H_{ni}$ exceeds the sum of the variances of its two components, reflecting positive correlation of $\ln H^{\lambda}_{ni}$ and $\ln H^{d}_{ni}$ across bilateral pairs.

5.3.2 Gravity Regressions

Second, we examine how our measure of matching frictions connects to a standard gravity equation of bilateral trade. To guide this analysis we take logs of (41) to get:

$$\ln \pi_{ni} = \ln T_i \Xi_i - \ln \Upsilon_n + \ln \lambda_{ni} - \theta \ln d_{ni}.$$

We capture $\ln \Upsilon_n$ with destination effects and $\ln T_i \Xi_i$ with source effects.[32] Table 6 presents the results of different specifications. With just fixed effects the R^2 is 0.80 (column 1). Column 2 adds bilateral matching intensity $\ln \lambda_{ni}$ and column 3 ln(distance), each on their own. An explanation for the large absolute values of the coefficients on $\ln \lambda_{ni}$ and ln(distance) in the

[32] In contrast to the Head-Ries indices, observations in a gravity equation include bilateral trade separately in each direction. We continue to eliminate home observations (for which $n = i$).

two regressions is that they're negatively correlated (suggesting a positive correlation between the two trade frictions).[33]

In column 4 we impose the theoretical coefficient of one on $\ln \lambda_{ni}$, which still cuts the residual variance in column 1 by half. Column 5 adds ln(distance) to this regression. The elasticity of trade shares with respect to distance falls from -1.69 (in column 3) to -0.66. An implication is that distance impedes bilateral trade more through matching frictions (with an elasticity of -1.03) than through iceberg costs (with an elasticity of -0.66). (With $\theta = 4$ the distance elasticity of iceberg costs is itself just 0.17.).

Table 6
Gravity Regressions

	dependent variable: $\ln \pi_{ni}$				
	(1)	(2)	(3)	(4)	(5)
$\ln \lambda_{ni}$		1.55		1.00	1.00
		(0.06)			
ln(distance)			-1.69		-0.66
			(0.07)		(0.07)
constant	-5.52	-9.55	8.39	-7.07	-1.00
	(0.27)	(0.21)	(0.55)	(0.16)	(0.55)
R^2	0.80	0.92	0.93	0.90	0.92

Notes: The number of observations is 702. Data on π_{ni} are for 2012 from WIOD. Data on (population-weighted) distance are from CEPII's Gravity database. We include destination and source fixed effects. Standard errors are in parentheses.

6 Applications

We perform two exercises that illustrate the distinct implications of iceberg costs and matching frictions. The first is a counterfactual in which iceberg costs or bilateral matching frictions decline uniformly. The second is a factual in which, for a set of new EU members, bilateral frictions (of both types) change to match the actual changes in French exports to those destinations and in the number of buyers per French exporter in those destinations.

Both exercises exploit "exact hat algebra" described in Appendix D. If x is the baseline value of a variable and x' is its value in the alternative scenario then $\hat{x} = x'/x$ is its change

[33]Since the λ_{ni} are estimated, and hence measured with error, we expect a downward bias in the coefficient in column 2. The fact that the estimate is much larger than the value of 1 implied by theory is thus even stronger evidence that the two types of trade frictions are positively correlated, consistent with the evidence in Table 5.

(with $\hat{x} = 1$ denoting no change). The shocks driving each scenario are some combination of changes in iceberg costs $\hat{d}_{ni}$ and matching frictions $\hat{\lambda}_{ni}^{-1/\theta}$.

6.1 Reducing Trade Frictions

Our counterfactual lowers trade frictions between all countries, first setting $\hat{d}_{ni} = 0.9$ for $n \neq i$ (with $\hat{d}_{nn} = 1$) and then setting $\hat{\lambda}_{ni}^{-1/\theta} = 0.9$ for $n \neq i$ (with $\hat{\lambda}_{nn} = 1$). Table 7 shows the results for three source countries: France, Germany, and Greece. As shown in the first three rows, the two experiments yield nearly the same aggregate outcomes.[34] Each country's domestic market share falls, its average share in foreign markets rises, and real GDP goes up.[35]

Table 7
Reducing Trade Frictions: Aggregate and Producer Outcomes

		10% fall in d_{ni}			10% fall in $\lambda_{ni}^{-1/\theta}$		
		France	Germany	Greece	France	Germany	Greece
Market share at home		0.87	0.86	0.93	0.87	0.86	0.93
Market share in other EU		1.35	1.27	1.54	1.34	1.26	1.55
Real GDP		1.02	1.02	1.01	1.02	1.02	1.01
Active producers:	Total	0.87	0.94	1.05	0.87	0.92	0.97
	Exporters	1.43	1.34	1.36	1.04	0.98	1.12
Relationships:	Total	0.99	1.05	1.00	0.99	1.04	0.97
	Foreign	1.35	1.25	1.44	1.32	1.23	1.43
Buyers per seller:	Total	1.14	1.12	0.95	1.13	1.13	1.00
	Foreign	0.94	0.93	1.05	1.27	1.25	1.27

Notes: The shock applies to all source-destination pairs, with d_{ii} and λ_{ii} unchanged. All values are counterfactual relative to baseline.

Turning to firm-level outcomes, the number of active producers typically declines in parallel with the decline in home market share, as buyers switch from domestic producers to foreign competitors. Similarly, the number of relationships with buyers in foreign markets rises in parallel with the increase in foreign market share. How this increase in foreign relationships

[34] From (41), the changes to the two trade frictions have identical effects on trade shares, given importer and exporter characteristics. Aggregate outcomes differ slightly however. The decline in matching frictions, unlike the decline in iceberg costs, lowers the measure of active producers, reducing buyer congestion.

[35] An explanation for Greece's more robust expansion abroad is that, initially, it exports very little relative to its sales at home. A drop in either type of trade friction raises its wage less (as reflected in its smaller increase in real GDP), so its market share abroad expands more.

comes about differs starkly across the two experiments, however. With lower iceberg costs the additional relationships come with more exporters, leaving buyers per exporter largely unaffected. With lower matching frictions the number of exporters changes only modestly, with foreign buyers per exporter rising by about a quarter. In summary: (i) reducing trade frictions of either sort concentrates economic activity among firms while increasing the number of foreign relationships; (ii) with lower iceberg costs the increase in foreign relationships reflects a jump in the number of exporters; (iii) with lower matching frictions the increase in foreign relationships reflects a jump in foreign buyers per exporter.

How do lower trade frictions affect firms of different efficiency? A message of the Melitz model, explored in EKK, is that globalization, in the form of lower iceberg costs, kills off the least efficient firms as they succumb to foreign competition while allowing the most efficient firms to expand their presence in foreign markets.[36] Our finding here is that globalization in the form of lower bilateral matching frictions magnifies these unequal effects across firms.

Table 8 shows changes in the number of active French producers and exporters by decile of firm efficiency and for the top 1%. (The bottom row reports changes in totals, repeating the numbers in Table 7 for active producers from France.) Note, from the first two columns, that the two experiments have almost identical effects in destroying firms at the low end of the efficiency distribution. But the second two columns show that reduced matching frictions, as opposed to reduced iceberg costs, generates little entry into exporting. Instead, as shown in the last column of Table 9, with lower matching frictions the most efficient firms, which had the most buyers to begin with, gain even more. Total foreign buyers per exporter rises. While the most efficient exporters also gain buyers when iceberg costs fall, as shown in the second to last column of Table 9, this effect is more than offset by the entry of less efficient exporters with few buyers, leaving total foreign buyers per exporter lower in this case.

6.2 Entry into the EU

For our factual, we turn to the ten countries, mostly from eastern Europe, that joined the EU in 2004. As shown in Table 10, French market share rose dramatically in the subsequent year (often by over 50%) in nine of these ten countries, while hardly changing in the other

[36] The results of the corresponding counterfactual from EKK appear in their Tables IV-VI on pages 1493-1495. Despite the fact that the two counterfactuals use a different model, different base years (1986 rather than 2012), and different firm groupings (by sales rather than by efficiency), the implications across firms of lower iceberg costs are strikingly similar.

Table 8

Reducing Trade Frictions: French Firm Entry and Exit by Productivity

Initial productivity percentile:	Number of firms counterfactual to baseline: 10% fall in d_{ni}	Number of firms counterfactual to baseline: 10% fall in $\lambda_{ni}^{-1/\theta}$	Number of exporters: Baseline fraction exporting:	Number of exporters counterfactual to baseline: 10% fall in d_{ni}	Number of exporters counterfactual to baseline: 10% fall in $\lambda_{ni}^{-1/\theta}$
0 to 10	0.58	0.57	0.000	-	-
10 to 20	0.68	0.68	0.000	24.00	1.00
20 to 30	0.76	0.76	0.000	13.60	1.00
30 to 40	0.84	0.83	0.001	11.13	1.07
40 to 50	0.91	0.91	0.003	8.52	1.08
50 to 60	0.97	0.96	0.009	6.48	1.10
60 to 70	1.00	0.99	0.037	4.74	1.12
70 to 80	1.00	1.00	0.166	2.95	1.13
80 to 90	1.00	1.00	0.665	1.39	1.07
90 to 99	1.00	1.00	0.998	1.00	1.00
99 to 100	1.00	1.00	1.000	1.00	1.00
Total	0.87	0.87	0.188	1.43	1.04

Notes: The shock applies to all source-destination pairs, with d_{ii} and λ_{ii} unchanged.

Table 9

Reducing Trade Frictions: French Firm Buyers by Productivity

Initial productivity percentile	Buyers per firm: baseline	Buyers per firm counterfactual to baseline: 10% fall in d_{ni}	Buyers per firm counterfactual to baseline: 10% fall in $\lambda_{ni}^{-1/\theta}$	Foreign buyers per exporter: baseline	Foreign buyers per exporter counterfactual to baseline: 10% fall in d_{ni}	Foreign buyers per exporter counterfactual to baseline: 10% fall in $\lambda_{ni}^{-1/\theta}$
0 to 10	1.1	0.97	0.97	0.0	-	-
10 to 20	1.3	0.92	0.92	1.0	1.00	1.00
20 to 30	1.6	0.87	0.87	1.0	1.00	1.00
30 to 40	2.0	0.82	0.82	1.0	1.00	1.00
40 to 50	2.8	0.78	0.78	1.0	1.01	1.00
50 to 60	4.3	0.77	0.75	1.0	1.03	1.00
60 to 70	7.0	0.78	0.75	1.0	1.09	1.00
70 to 80	12.8	0.82	0.78	1.1	1.30	1.02
80 to 90	29.1	0.89	0.82	2.1	1.92	1.11
90 to 99	185.5	0.96	0.90	33.8	1.64	1.27
99 to 100	3029.3	1.05	1.08	1340.6	1.26	1.34
Total	53.2	1.14	1.13	88.4	0.94	1.27

Notes: The shock applies to all source-destination pairs, with d_{ii} and λ_{ii} unchanged.

fourteen.[37]

[37]Table 10 lists EU members as of 2004 and their year of entry. Numbers for the ten 2004 entrants are in bold. France's market share in Malta, one of the new members, actually tanked. The largest French exporter

As reported in the last row, France's market share across the new members grew by 45 percent while buyers per French exporter grew by 9 percent. From equations (21) and (28) the product of the trade frictions governs bilateral trade shares while matching frictions alone govern buyers per seller. To fit these two facts about French expansion in the new members, our factual exercise sets $\hat{d}_{ni} = 0.92$ and $\hat{\lambda}_{ni}^{-1/4} = 0.92$ whenever n or i is one of the ten entrants. Hence our model interprets EU membership as delivering a 16 percent decrease in trade frictions, with lower matching frictions contributing half the total change.

Table 10 shows what our exercise implies for French market share and buyers per seller among both new and old EU members. Among old members (not targeted) we slightly understate the decline in French market share and miss the slight increase in buyers per French seller. Among the ten entrants, the correlation between data and model for French market share is 0.49 and for buyers per French seller is 0.39.

Using the implied changes in trade frictions from this episode, we ask what our model says about some of its consequences. As shown in Panel A of Table 11, entrants lose home market share but gain market share abroad. We also compute significant gains in real GDP for the entrants. The effects on incumbents are in the same direction but muted.

A feature of our framework is the connection it draws between trade and the labor market. The key distinction among types of labor in our model is the intensity with which their employers in the goods sector match with suppliers of competing intermediates. We assume that tertiary workers are immune from this competition but, as reported in Section 6, we estimate $\lambda_1 = 1.65$ and $\lambda_2 = 0.35$.

As shown in the first three rows of panel B, lower trade frictions imply lower labor shares in the goods sector among our selected countries. Primary workers, whose tasks are most threatened by outsourcing, experience the largest drop. Secondary workers experience a smaller drop while labor overall (which includes tertiary workers with their fixed share) experience a smaller drop still. This ranking of labor types holds across the other outcomes we consider for the goods sector.

The implied real wages of secondary and tertiary workers rise in all our selected countries. Primary workers can actually lose in countries entering the EU. The reason is that lower trade frictions increase competition from intermediates in performing their tasks.

to Malta in 2004 accounts for all of this decline. Appendix A.2 explains how we constructed the data. French firm-to-firm data are available only for EU members, preventing us from observing buyers of French firms in a country before it joined the EU. We thank Gregory Corcos for suggesting that we examine this episode.

Table 10
EU Expansion: Implications for French Exports

Destination:	EU entry year	French Market Share[1]		Mean Buyers per French Exporter[1]	
		Data[2]	Model[3]	Data[2]	Model[3]
Austria (AUT)	1995	0.95	0.97	1.04	0.98
Belgium (BEL)	1958	0.99	0.99	1.02	1.00
Cyprus (CYP)	**2004**	**1.48**	**1.13**	**1.10**	**1.14**
Czech Republic (CZE)	**2004**	**1.50**	**1.54**	**1.07**	**1.08**
Germany (DEU)	1958	0.96	0.97	1.01	0.99
Denmark (DNK)	1973	0.92	0.97	1.03	0.99
Spain (ESP)	1986	0.93	0.99	1.04	1.00
Estonia (EST)	**2004**	**1.43**	**1.24**	**1.01**	**1.08**
Finland (FIN)	1995	0.90	0.98	1.00	1.00
United Kingdom (GBR)	1973	1.01	0.98	1.01	1.00
Greece (GRE)	1981	0.85	0.98	1.04	1.00
Hungary (HUN)	**2004**	**1.56**	**1.44**	**1.07**	**1.04**
Ireland (IRL)	1973	1.00	0.99	1.08	1.00
Italy (ITA)	1958	0.95	0.99	1.05	1.00
Lithuania (LTU)	**2004**	**1.42**	**1.45**	**1.05**	**1.07**
Luxembourg (LUX)	1958	0.90	0.99	0.99	1.00
Latvia (LVA)	**2004**	**1.65**	**1.29**	**1.06**	**1.11**
Malta (MLT)	**2004**	**0.27**	**1.15**	**1.07**	**1.16**
Netherlands (NLD)	1958	1.02	0.99	1.02	1.00
Poland (POL)	**2004**	**1.38**	**1.61**	**1.16**	**1.11**
Portugal (PRT)	1986	1.02	0.99	1.02	1.00
Slovakia (SVK)	**2004**	**1.86**	**1.55**	**1.07**	**1.08**
Slovenia (SVN)	**2004**	**1.61**	**1.42**	**1.03**	**1.06**
Sweden (SWE)	1995	0.92	0.97	1.01	0.99
Average changes:[4]					
Incumbent EU members		0.96	0.98	1.02	1.00
New EU members		1.455	1.454	1.087	1.088

Notes: Countries entering in 2004 in bold.
[1] Ratio of post-expansion to pre-expansion magnitudes.
[2] From French VAT and WIOD data (ratio of 2005 to 2004). For compatibility between years we use a new version of the VAT data created by French Customs for CASD (Centre d'accès sécurisé aux données).
[3] $\hat{d}_{ni} = 0.92$ and $\hat{\lambda}_{ni}^{-1/4} = 0.92$ whenever n or i is one of the ten entrants.
[4] Weighted by the number of French exporters to that destination in 2005 (VAT data).

As shown in the last six rows of Table 11, in each of our selected countries tertiary workers shift to the goods sector while primary and secondary workers move into services. Our Cobb-Douglas assumptions tie down the relative sizes of the two sectors in terms of the value of final output and of value added. But trade liberalization, by enhancing access to intermediates for

goods producers, changes the composition of inputs in that sector. Primary and secondary workers, replaced by intermediates, move to services. But since relative value added in the two sectors can't change, tertiary workers, with their immunity to competition from imports, move in the other direction. The decline in the wages of primary and secondary workers relative to tertiary workers accommodates the rise of their employment in services.

Table 11
EU Expansion: Aggregate and Labor-Market Outcomes

		Selected Entrants			Selected Incumbents		
		Poland	Czech Rep.	Hungary	Austria	Germany	France
Panel A: Aggregate outcomes							
Home market share		0.85	0.75	0.69	0.96	0.97	0.99
Average market share abroad		1.81	1.68	1.64	1.01	1.03	1.03
Real GDP		1.03	1.04	1.05	1.01	1.00	1.00
Panel B: Labor-market outcomes							
Labor share in	Overall	0.91	0.84	0.79	0.97	0.98	0.99
goods sector:	Primary	0.84	0.77	0.69	0.96	0.97	0.99
	Secondary	0.88	0.81	0.74	0.97	0.97	0.99
Real wage:	Tertiary	1.06	1.10	1.12	1.01	1.01	1.00
	Primary	1.01	0.99	0.99	1.00	1.00	1.00
	Secondary	1.01	1.02	1.01	1.00	1.00	1.00
Goods sector	Tertiary	1.03	1.05	1.04	1.01	1.01	1.00
employment:	Primary	0.90	0.89	0.82	0.98	0.98	0.99
	Secondary	0.95	0.91	0.86	0.99	0.99	0.99
Services sector	Tertiary	0.99	0.99	0.99	1.00	1.00	1.00
employment:	Primary	1.04	1.09	1.12	1.01	1.01	1.00
	Secondary	1.04	1.06	1.10	1.01	1.01	1.00

Notes: Post-expansion to pre-expansion magnitudes, with $\hat{d}_{ni} = \hat{\lambda}_{ni}^{-1/4} = 0.92$ if n or i is one of the ten entrants.

7 Conclusion

Our framework, taking into account the granularity of individual buyer-seller relationships, expands the scope for firm heterogeneity in a number of dimensions. Regardless of their underlying efficiency, firms' fortunes differ in procuring cheap inputs, contributing to differences in their costs and their use of labor. Within each market firms have different success in connecting with buyers. A firm may happen to sell a lot in a small, remote market while striking

out in a large one close by.

We've used the framework to show that matching frictions contribute as much to gravity as iceberg costs, and rise even more with distance. We find matching easier in larger markets, suggesting increasing returns. Trade expansion affects workers differently depending on how easily their employers can replace them with foreign inputs.

As more data become available, the framework can address a vastly wider set of issues. Observing both domestic and international firm-to-firm connections can help identify the factors inhibiting cross-border trade. Expanding on the product dimension can reconcile individual firms' idiosyncratic purchases and sales with aggregate input-output analysis. Keeping track of firm-to-firm connections over time can provide insight into the short versus long run effects of trade policy.

References

Acemoglu, D. and D. Autor (2011): "Skills, Tasks and Technologies: Implications for Employment and Earnings," in *Handbook of Labor Economics*, ed. by O. Ashenfelter and D. Card, Elsevier, vol. 4, chap. 12, 1043–1171.

Antràs, P., T. C. Fort, and F. Tintelnot (2017): "The Margins of Global Sourcing: Theory and Evidence from U.S. Firms," *American Economic Review*, 107, 2514–64.

Arkolakis, C. (2010): "Market Penetration Costs and the New Consumers Margin in International Trade," *Journal of Political Economy*, 118, 1151–1199.

Arkolakis, C., A. Costinot, and A. Rodríguez-Clare (2012): "New Trade Models, Same Old Gains?" *American Economic Review*, 102, 94–130.

Arkolakis, C., F. Huneeus, and Y. Miyauchi (2021): "Spatial Production Networks," Unpublished, Yale University.

Bergounhon, F., C. Lenoir, and I. Mejean (2018): "A Guideline to French Firm-Level Trade Data," Unpublished, CREST-Ecole Polytechnique.

Bernard, A. B., E. Dhyne, G. Magerman, K. Manova, and A. Moxnes (2021): "Origins of Firm Heterogeneity: A Production Approach," Forthcoming, *Journal of Political Economy*.

Bernard, A. B. and A. Moxnes (2018): "Networks and Trade," *Annual Review of Economics*, 10, 65–85.

Bernard, A. B., A. Moxnes, and Y. U. Saito (2019): "Production Networks, Geography, and Firm Performance," *Journal of Political Economy*, 127, 639–688.

Bernard, A. B., A. Moxnes, and K. H. Ulltveit-Moe (2018): "Two-Sided Heterogeneity and Trade," *Review of Economics and Statistics*, 100, 424–439.

Blum, B. S., S. Claro, and I. J. Horstmann (2018): "Trade Costs and the Role of International Trade Intermediaries," in *Handbook of International Trade and Transportation*, ed. by B. A. Blonigen and W. W. Wilson, Edward Elgar Publishing, chap. 11, 337–367.

Caliendo, L. and F. Parro (2015): "Estimates of the Trade and Welfare Effects of NAFTA," *Review of Economic Studies*, 82, 1–44.

Chaney, T. (2014): "The Network Structure of International Trade," *American Economic Review*, 104, 3600–3634.

Demir, B., A. C. Fieler, D. Y. Xu, and K. K. Yang (2021): "O-Ring Production Networks," Unpublished, Bilkent University.

Dhyne, E., A. K. Kikkawa, M. Mogstad, and F. Tintelnot (2020): "Trade and Domestic Production Networks," *Review of Economic Studies*, 88, 643–668.

Diamond, P. A. (1982): "Wage Determination and Efficiency in Search Equilibrium," *Review of Economic Studies*, 49, 217–227.

Eaton, J., M. Eslava, C. Krizan, D. Jinkins, and J. Tybout (2021a): "A Search and Learning Model of Export Dynamics," NBER Working Paper No. 29100.

Eaton, J., D. Jinkins, J. Tybout, and D. Xu (2021b): "Two-Sided Search in International Markets," Pennsylvania State University.

Eaton, J. and S. Kortum (2002): "Technology, Geography, and Trade," *Econometrica*, 70, 1741–1779.

Eaton, J., S. Kortum, and F. Kramarz (2011): "An Anatomy of International Trade: Evidence from French Firms," *Econometrica*, 79, 1453–1498.

Feenstra, R. C. and G. H. Hanson (1996): "Globalization, Outsourcing, and Wage Inequality," *American Economic Review*, 86, 240–245.

Fujii, D., Y. Ono, and Y. U. Saito (2017): "Indirect Exports and Wholesalers: Evidence from Interfirm Transaction Network Data," *Japan and the World Economy*, 44, 35–47.

Furusawa, K., T. Inui, K. Ito, and H. Tang (2017): "Global Sourcing and Domestic Production Networks," CES IFO Working Paper Series No. 6658.

Garetto, S. (2013): "Input Sourcing and Multinational Production," *American Economic*

Journal: Macroeconomics, 5, 118–51.

Grossman, G. M. and E. Rossi-Hansberg (2008): "Trading Tasks: A Simple Theory of Offshoring," *American Economic Review*, 98, 1978–97.

Head, K. and T. Mayer (2014): "Gravity Equations: Workhorse, Toolkit, and Cookbook," in *Handbook of International Economics*, ed. by G. Gopinath, E. Helpman, and K. Rogoff, Elsevier, vol. 4, chap. 3, 131–195.

Head, K. and J. Ries (2001): "Increasing Returns Versus National Product Differentiation as an Explanation for the Pattern of U.S.-Canada Trade," *American Economic Review*, 91, 858–876.

Imbs, J. and I. Mejean (2015): "Elasticity Optimism," *American Economic Journal: Macroeconomics*, 7, 43–83.

Lenoir, C., J. Martin, and I. Mejean (2020): "Search Frictions in International Goods Markets," Unpublished, CREST.

Lim, K. (2021): "Endogenous Production Networks and the Business Cycle," Unpublished, University of Toronto.

Melitz, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71, 1695–1725.

Miyauchi, Y. (2021): "Matching and Agglomeration: Theory and Evidence from Japanese Firm-to-Firm Trade," Unpublished, Boston University.

Mortensen, D. T. and C. A. Pissarides (1994): "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, 61, 397–415.

Oberfield, E. (2018): "A Theory of Input-Output Architecture," *Econometrica*, 86, 559–589.

Panigrahi, P. (2021): "Endogenous Spatial Production Networks: Quantitative Implications for Trade and Productivity," Unpublished, University of California – Berkeley.

Petrongolo, B. and C. A. Pissarides (2001): "Looking into the Black Box: A Survey of the Matching Function," *Journal of Economic Literature*, 39, 390–431.

Samuelson, P. A. (1954): "The Transfer Problem and Transport Costs, II: Analysis of Effects of Trade Impediments," *Economic Journal*, 64, 264–289.

Timmer, M. P., E. Dietzenbacher, B. Los, R. Stehrer, and G. J. de Vries (2015): "An Illustrated User Guide to the World Input–Output Database: the Case of Global Automotive Production," *Review of International Economics*, 23, 575–605.

On-Line Appendices

A Data Sources

We use four distinct datasets. The first three, discussed in Section 2, cover only the goods sector (defined for our purposes as manufacturing, wholesale, and retail). The fourth incorporates the goods sector into general equilibrium.

A.1 EU Firm-to-Destination Data (TEC)

These data, from the OECD's Trade by Enterprise Characteristics data (henceforth TEC), report the total number of firms in each of 27 EU members exporting goods to each of the other EU members in 2012.[38]

We use this dataset to construct our measure of N_{ni} (number of exporters from i selling in n) used in Figure 1 and in step 4 of Section 5.1.4 and our measures of I_n (number of importers in n) and I_n^P (number of importing producers) in Steps 1 and 7 of Section 5.1.4. Constructing these measures requires:

1. the number of firms exporting to at least one EU destination, as reported by the exporting country.

2. the bilateral number of exporters to each EU destination, as reported by the exporting country.

3. the number of firms importing from at least one EU source, as reported by the importing country.

Constructing N_{ni} presents two problems. One is that, for some exporter firms, the data don't report the destination. We denote the number of such exporters as N_i^0. The second is that the bilateral data don't record the sector of the exporting firm at our level of disaggregation. Hence we can only get the counts for a broader class of exporting firms (combining

[38]These data are available at https://stats.oecd.org/Index.aspx?DataSetCode=TEC1_REV4. They may undercount the total number of exporters because of the following loophole:

> The exemption threshold defines the value above which the parties (taxable persons) are obliged to provide Intrastat information. Member States are required to determine this threshold each year. The threshold is expressed in annual values and it is set in order to ensure that the information provided is such that at least 97% of the total dispatches and at least 93% of the total arrivals, expressed in value, of the relevant Member State's taxable persons is covered.

Eurostat (2017): Compilers guide on European statistics on international trade in goods - 2017 edition, p.179, available at: https://ec.europa.eu/eurostat/documents/3859598/8021340/KS-02-17-333-EN-N.pdf.

exporters in industry, wholesale, and retail). We denote this count as N_{ni}^B. The multilateral export data do, however, provide exporter information both at this broader level of aggregation and at our preferred level (industry and wholesale). Denoting the first by E_i^B and the second by E_i we compute:

$$N_{ni} = \left(\frac{E_i}{E_i^B}\right)\left(\frac{N_i^0 + \sum_n N_{ni}^B}{\sum_{n'} N_{n'i}^B}\right) N_{ni}^B.$$

Our measure I_n^P is the number of importers in country n in industry and wholesale while I_n includes retail as well.[39]

A.2 French Firm-to-Firm Data (VAT)

Table 12
2005 VAT Data, Initial vs. Final sample

	Initial dataset	Final dataset[1]		
Sectors of exporters	All	Manufacturing & Wholesale	Manufacturing	Wholesale
Products	All	Manufactured	Manufactured	Manufactured
Observations[2]	3,984,909	3,145,709	1,756,616	1,389,093
Relationships[3]	1,295,446	1,019,987	658,040	361,947
Exporters	46,928	30,787	18,444	12,343
Buyers	571,149	481,832		
Sales ($ millions)	257,302	190,501	164,680	25,821

Source: VAT Data; Year: 2005;
[1] The final dataset excludes sellers belonging to sectors others than Manufacturing and Wholesale (i.e., Retail, Agriculture, Extraction).
[2] An observation corresponds to a seller-buyer-product triad.
[3] A relationship is a seller-buyer dyad with sales aggregated over all products.

These data report the sales of each French firm to each of its buyers in each of the other 24 EU member countries in 2005. French Customs collects the data in administering the EU's value-added tax. Bergounhon et al. (2018) provide a thorough description.

We restrict exporters to French firms in the manufacturing sector and to firms in the wholesale sector that ship manufactures. Sales include only shipments of manufactures.[40] We include all buyers as the data don't report their sector.

[39] Due to changes in reporting thresholds, the counts of importers can occasionally take dramatic jumps up or down between 2012 and an adjacent year. To obtain the broadest measure of importers we take the maximum value that the importing country reports in any year from 2011 to 2013.

[40] The smallest exporters aren't required to report the product dimension. We include these firms as observations and all their shipments as sales.

Table 12 summarizes various dimensions of the data. Total exports in this sample aggregate to US$257 billion in 2005, capturing 78 percent of total French exports of finished goods to these destinations in the World Input-Output Database (WIOD), described below.[41]

For each destination n we calculate the number of French firms N_{nF} exporting there and the number of firms buying from them F_{nF}. Since we know the identity of individual buyers as well as sellers, we calculate the number French exporters with b buyers, $N_{nF}(b)$ and the number of importers with s French sellers $F_{nF}(s)$. We also construct the number of relationships R_{nF} between these French exporters and their buyers in destination n. We can also compute how much is sold to each buyer in destination n. Aggregating across buyers in n and exporters in France, we obtain the value of exports by French producers to n, X^P_{nF}, for $n \neq F$. Table 13 displays the basic statistics by destination. They form the observations for the regressions in Table 1.

To check consistency between the VAT data and the TEC data, we compare the two bilateral exporter series for France as the exporter. Figure 7 plots N_{nF} for 2012 from TEC against N_{nF} for 2005 from VAT.

Figure 7: Number of French exporters in 2005 and 2012

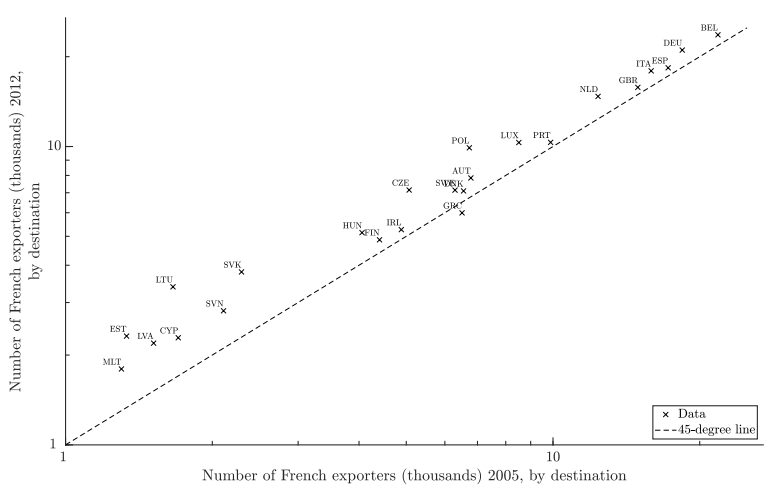

A.3 French Firm Labor-Share Data (DADS)

We obtain firm-level data on labor shares as follows. We consider only French manufacturing firms, merging administrative-origin tax data from firm-level balance-sheets from Fichier

[41] French Customs generates both the VAT data and the export data used in the WIOD. The VAT data follow a different protocol and have different coverage (excluding, for instance, exports to individuals or shipments in which either seller or buyer lacks a VAT identifier).

Table 13
2005 VAT Data, Final dataset

	R_{nF}	$\bar{x}_{nF}$ [1]	N_{nF}	$\bar{b}_{nF}$	F_{nF}	$\bar{s}_{nF}$	X^P_{nF} [2]	π_{nF}
Austria	23,668	131	6,784	3.5	12,267	1.9	3,103	0.019
Belgium	173,379	126	21,798	8.0	59,989	2.9	21,851	0.106
Cyprus	2,679	85	1,703	1.6	1,344	2.0	227	0.028
Czech Republic	9,919	245	5,069	2.0	4,796	2.1	2,430	0.021
Germany	196,830	207	18,415	10.7	103,120	1.9	40,809	0.028
Denmark	17,711	139	6,554	2.7	8,098	2.2	2,455	0.024
Spain	135,911	239	17,230	7.9	68,419	2.0	32,422	0.045
Estonia	1,872	96	1,334	1.4	973	1.9	180	0.020
Finland	10,527	145	4,407	2.4	4,813	2.2	1,531	0.014
United Kingdom	94,132	288	14,930	6.3	44,882	2.1	27,142	0.034
Greece	20,242	112	6,508	3.1	9,106	2.2	2,259	0.021
Hungary	7,161	295	4,057	1.8	3,520	2.0	2,112	0.027
Ireland	13,031	171	4,885	2.7	6,015	2.2	2,225	0.028
Italy	146,039	167	15,900	9.2	79,046	1.8	24,319	0.022
Lithuania	2,389	113	1,661	1.4	1,218	2.0	270	0.020
Luxembourg	23,314	50	8,509	2.7	6,603	3.5	1,168	0.079
Latvia	2,184	69	1,516	1.4	1,046	2.1	151	0.014
Malta	1,747	64	1,302	1.3	783	2.2	112	0.033
Netherlands	55,650	189	12,372	4.5	27,697	2.0	10,530	0.042
Poland	17,857	242	6,736	2.7	8,352	2.1	4,319	0.021
Portugal	38,837	107	9,888	3.9	17,176	2.3	4,150	0.035
Slovakia	3,468	210	2,296	1.5	1,677	2.1	727	0.018
Slovenia	3,243	458	2,109	1.5	1,690	1.9	1,487	0.061

Source: VAT Data Year: 2005
[1] $ thousands. [2] $ millions.

complet unifié de SUSE (Système unifié de statistiques d'entreprises) (FICUS)[42] with firm-level employment data from Declaration Annuelle des Données Sociales (DADS) for the year 2005.[43] The DADS data report the wage bill by the qualification level of workers. From these data we take the wage bills for unskilled and for skilled production workers. We divide each by total variable costs (total intermediate purchases, reported in FICUS, plus total payments to production labor), to make them shares in production costs. The results deliver the quantiles of the shares of each type of production labor, reported in Figures 5a and 5b, used to construct moments for estimation as described in Section 5.1.3.

[42]See https://www.insee.fr/fr/information/2407173

[43]We exclude wholesalers since accounting conventions for them aren't compatible with those for manufacturing firms.

A.4 Sectoral Data (WIOD)

We use these data both to construct absorption and trade share measures for 2012 and 2005 used in Sections 2 and 5, and to obtain service and good sector preference and production share parameters used in Section 5.1.4.

A.4.1 Trade Shares

Our trade measures ignore the world outside the EU. From rows of the WIOD we calculate the flow X_{ni}^P of goods from the manufacturing and wholesale sectors in source i to all buyers in destination n.[44] We calculate trade shares as:

$$\pi_{ni} = \frac{X_{ni}^P}{X_n^P},$$

where X_n^P is total absorption of the output of producers in EU destinations n:

$$X_n^P = \sum_i X_{ni}^P, \quad n, i \in EU.$$

To maintain consistency with the French firm-to-firm data, for 2005 we replace X_{nF}^P $(n \neq F)$ from WIOD with the value obtained by summing the VAT data across French exporters. In both the WIOD and VAT data, our measures of absorption exclude retail markups.

A.4.2 Preference and Production Shares

We now turn to the preference and production shares for the goods and services sectors used in Section 5.1.1 We define the goods sector (superscript G) as the sum of manufacturing, wholesale, and retail. The services sector (superscript S) is the sum of the remaining sectors.

Preference Parameters We measure total final consumption X_n^F in WIOD, separating goods $X_n^{F,G}$ and services $X_n^{F,S}$. We include all purchases from domestic or foreign sources (including from outside the EU) by households and the government (combining it with households). We treat investment as intermediate purchases. We compute household spending shares as:

$$\alpha_n^G = \frac{X_n^{F,G}}{X_n^F}; \quad \alpha_n^S = \frac{X_n^{F,S}}{X_n^F}.$$

[44] EU destinations represented 58 percent of French exports of these goods worldwide in 2005.

Production Parameters We need to adjust the WIOD data to reflect our treatments of investment and of trade outside the EU.

a. adjustments WIOD provides measures of each sector's output, $X_n^{I,G}$ and $X_n^{I,S}$, used for investment, but does not indicate the sector spending on investment. To fold investment spending into intermediates, we apportion each sector's contribution to investment according to its contribution to intermediates. Our measures of sectoral value added are consequently:

$$\begin{aligned}
\tilde{W}_n^G &= Y_n^G - X_n^{GG} - X_n^{SG} - \frac{X_n^{GG}}{X_n^{GG} + X_n^{SG}} X_n^{I,G} - \frac{X_n^{GS}}{X_n^{GS} + X_n^{SS}} X_n^{I,S}; \\
\tilde{W}_n^S &= Y_n^S - X_n^{SG} - X_n^{SS} - \frac{X_n^{SG}}{X_n^{GG} + X_n^{SG}} X_n^{I,G} - \frac{X_n^{SS}}{X_n^{GS} + X_n^{SS}} X_n^{I,S}.
\end{aligned}$$

We scale our measures of sectoral value added so that total value added in the EU equals total final spending by the EU, since we ignore the world outside the EU. Our scaling factor is:

$$\varsigma = \frac{\sum_{n' \in EU} X_{n'}^F}{\sum_{n' \in EU} \left(\tilde{W}_{n'}^G + \tilde{W}_{n'}^S \right)},$$

which turns out to be 0.955. We set $W_n^G = \varsigma \tilde{W}_n^G$, $W_n^S = \varsigma \tilde{W}_n^S$, and $W_n = W_n^G + W_n^S$. With our scaling, the trade deficit of any country n within the EU is $D_n = X_n^F - W_n$ and the EU as a whole has balanced trade with the rest of the world.

We need to decompose country-level deficits into trade imbalances in goods and trade imbalances in services. Since we don't model trade in services, we treat service deficits as transfers. We take trade deficits in goods from our matrix of bilateral trade flows:

$$D_n^G = \sum_{i \neq n} X_{ni}^P - \sum_{k \neq n} X_{kn}^P. \tag{42}$$

We calculate trade imbalances in the service sector as $D_n^S = D_n - D_n^G$.

b. services In the model, services are produced with intermediate goods and labor, thus netting out intermediate services. In the data, we remove services intermediates from WIOD's gross services output $\tilde{Y}^S$ to obtain our model-consistent measure of services output:

$$Y_n^S = \tilde{Y}_n^S - \tilde{X}_n^{SS}.$$

Since labor is the only source of value added, the input shares for the services sector are:

$$\beta_n^{S,L} = \frac{W_n^S}{Y_n^S}$$

and

$$\beta_n^{SG} = 1 - \beta_n^{S,L},$$

our Cobb-Douglas share parameters.

c. goods Our goods sector sums manufacturers, wholesalers, and retailers. In our model manufacturers and wholesalers (collectively producers) sell goods to each other, to retailers, and abroad. Retailers sell final goods to consumers and intermediates to the service sector. Contrary to WIOD's convention, we treat the cost of goods sold by wholesalers and retailers as intermediate expenditures and hence part of their gross production.

To adjust WIOD's measure of goods sector output to include the cost of goods sold by wholesalers and retailers, we measure goods sector output by adding up all the inputs used by manufacturers, wholesalers, and retailers:

$$Y_n^G = W_n^G + X_n^{GG} + X_n^{GS}.$$

To calculate X_n^{GS} we decompose service output as:

$$\alpha_n^S X_n^F + X_n^{GS} - D_n^S = Y_n^S.$$

Substituting back into the expression for goods sector output, we get:

$$Y_n^G = W_n^G + X_n^{GG} + (Y_n^S - \alpha_n^S X_n^F + D_n^S).$$

Input shares in the goods sector are thus:

$$\beta_n^{G,L} = \frac{W_n^G}{Y_n^G},$$

$$\beta_n^{GG} = \frac{X_n^{GG}}{Y_n^G},$$

and:

$$\beta_n^{GS} = 1 - \beta_n^{G,L} - \beta_n^{GG}.$$

We now turn to how the retail sector fits into our accounting framework. Since goods intermediates are all supplied by producers, producer output in source country i is the sum of sales of intermediate goods in all destinations n:

$$Y_i^P = \sum_n \pi_{ni} X_n^P,$$

where X_n^P, absorption of the output of producers, is the same, in our accounting framework, as X_n^{GG}, the goods sector's use of goods intermediates. Retail output, which includes its cost of goods sold, is:

$$\begin{aligned} Y_n^R &= Y_n^G - Y_n^P \\ &= X_n^{F,G} + X_n^{SG} + X_n^{GG} - D_n^G - Y_n^P \\ &= \alpha_n^G X_n^F + \beta_n^{SG} Y_n^S + \beta_n^{GG} Y_n^G - D_n^G - Y_n^P \\ &= \alpha_n^G X_n^F + \beta_n^{SG} Y_n^S + X_n^P - D_n^G - Y_n^P \\ &= \alpha_n^G X_n^F + \beta_n^{SG} Y_n^S. \end{aligned}$$

d. type-specific labor shares We calibrate type-specific labor shares by sector, $\beta_n^{G,l}$ and $\beta_n^{S,l}$, for each type of labor $l \in \Omega^L = \{p, s, t\}$ from WIOD's Social and Economic Accounts (SEA) for 2005 and 2009.[45] The WIOD SEA classifies workers into six skill categories based on levels of educational attainment. We group their lowest two categories into primary (p), their middle two into secondary (s), and their top two into tertiary (t). For the goods sector we pool data from the manufacturing, wholesale, and retail sectors with the remaining sectors comprising services.

We measure labor compensation in each sector, S and G, going to skill-type $l \in \{p, s, t\}$, which we denote by $W_n^{S,l}$ and $W_n^{G,l}$. We construct type-specific labor shares, using the overall labor shares calibrated above:

$$\beta_n^{S,l} = \left(\frac{W_n^{S,l}}{\sum_{l'} W_n^{S,l'}} \right) \beta_n^{S,L}.$$

[45] We use data for 2009 instead of 2012, as it's the most recent year with labor compensation disaggregated by educational attainment.

and

$$\beta_n^{G,l} = \left(\frac{W_n^{G,l}}{\sum_{l'} W_n^{G,l'}} \right) \beta_n^{G,L}$$

For S and for G, except for France in 2005, we use WIOD SEA's data on labor compensation by level of education to construct these measures. To be consistent with data on the distribution of labor shares, for the French goods sector in 2005, we use data from DADS rather than from WIOD SEA to create $W_F^{G,l}$, treating unskilled production workers (from DADS) as primary workers and skilled production workers (from DADS) as secondary workers.

B Model Derivations

Here we describe the solution of (17) to obtain the parameters of the cost distribution, conditional on wages, and the equilibrium conditions (44) used to obtain wages.

B.1 Conditions for a Finite Constant

The constant term that cancels out of the cost distribution is:

$$\begin{aligned} g_i(m) &= \prod_k \left(\int_0^\infty e^{-x_1} ... \int_0^\infty e^{-x_{m-1}} \int_0^\infty e^{-x_m} \left(\sum_{\omega=1}^m x_\omega^{-\frac{\sigma-1}{\theta(1-\gamma)}} \right)^{\frac{\theta\beta_{k,i}}{\sigma-1}} dx_m dx_{m-1} ... dx_1 \right) \\ &= \prod_k \left(\int_0^\infty e^{-x_1} ... \int_0^\infty e^{-x_{m-1}} \int_0^\infty e^{-x_m} \left[\left(\sum_{\omega=1}^m x_\omega^{-\frac{\sigma-1}{\theta(1-\gamma)}} \right)^{-\frac{\theta(1-\gamma)}{\sigma-1}} \right]^{-\frac{\beta_{k,i}}{1-\gamma}} dx_m dx_{m-1} ... dx_1 \right) \\ &\leq \prod_k \left(\int_0^\infty e^{-x_1} ... \int_0^\infty e^{-x_{m-1}} \int_0^\infty e^{-x_m} \underline{x}^{-\beta_{k,i}/(1-\gamma)} dx_m dx_{m-1} ... dx_1 \right), \end{aligned} \tag{43}$$

where $\underline{x} = \min\{x_1, x_2, ..., x_m\}$. This minimum has distribution:

$$\Pr[X \leq x] = 1 - \prod_\omega \Pr[x_\omega > x] = 1 - e^{-mx},$$

allowing us to write the last term in (43) as:

$$\begin{aligned}\prod_k \left(\int_0^\infty m e^{-m\underline{x}} \underline{x}^{-\beta_{k,i}/(1-\gamma)} d\underline{x}\right) &= \prod_k \left(m^{\beta_{k,i}/(1-\gamma)} \int_0^\infty e^{-y} y^{-\beta_{k,i}/(1-\gamma)} dy\right) \\ &= \prod_k \left(m^{\beta_{k,i}/(1-\gamma)} \Gamma\left(1 - \frac{\beta_{k,i}}{1-\gamma}\right)\right).\end{aligned}$$

An implication is that $g_i(m)$ is finite if $\beta_{k,i} < 1 - \gamma$ for all k.

B.2 Computing the Cost Distribution

Given $w_{k,i}$ and B_i we can compute the Υ_n's by iterating on (17), repeated here for convenience:

$$\Upsilon_n = \sum_i \lambda_{ni} d_{ni}^{-\theta} T_i \prod_k \left(\frac{\lambda_k}{1-\gamma} B_i^{-\varphi} \Upsilon_i^{1-\gamma} + w_{k,i}^{-\theta(1-\gamma)}\right)^{\beta_{k,i}/(1-\gamma)},$$

for $n = 1, ..., \mathcal{N}$.

Define $\boldsymbol{\upsilon} = [\ln \Upsilon_1, \ln \Upsilon_2, ..., \ln \Upsilon_{\mathcal{N}}]$ which is the fixed point of

$$\boldsymbol{\upsilon} = F(\boldsymbol{\upsilon}).$$

The mapping F has n'th element:

$$F_n(y) = \ln\left[\sum_i \exp\left(A_{ni} + \sum_{k=1}^K \frac{\beta_{k,i}}{1-\gamma} \ln\left(u_{k,i} e^{(1-\gamma)y_i} + w_{k,i}^{-\theta(1-\gamma)}\right)\right)\right],$$

where

$$A_{ni} = \ln \lambda_{ni} d_{ni}^{-\theta} T_i - \theta \beta_{0,i} \ln w_{0,i}$$

and

$$u_{k,i} = \frac{\lambda_k}{1-\gamma} B_i^{-\varphi}.$$

We verify that F satisfies Blackwell's conditions for a contraction. For monotonicity, it's apparent that if $x \leq y$ then $F_n(x) \leq F_n(y)$ for each $n = 1, ..., \mathcal{N}$. For discounting, consider

$a > 0$ so that, for each $n = 1, ..., \mathcal{N}$:

$$
\begin{aligned}
F_n(y+a) &= \ln \sum_i \exp\left(A_{ni} + \sum_{k=1}^{K} \frac{\beta_{k,i}}{1-\gamma} \ln\left(u_{k,i} e^{(1-\gamma)(y_i+a)} + w_{k,i}^{-\theta(1-\gamma)}\right)\right) \\
&= \ln \sum_i \exp\left(A_{ni} + \sum_{k=1}^{K} \frac{\beta_{k,i}}{1-\gamma} \left[(1-\gamma)a + \ln\left(u_{k,i} e^{(1-\gamma)y_i} + e^{-(1-\gamma)a} w_{k,i}^{-\theta(1-\gamma)}\right)\right]\right) \\
&= \ln \sum_i \exp\left(A_{ni} + (1-\beta_{0,i})a + \sum_{k=1}^{K} \frac{\beta_{k,i}}{1-\gamma} \ln\left(u_{k,i} e^{(1-\gamma)y_i} + e^{-(1-\gamma)a} w_{k,i}^{-\theta(1-\gamma)}\right)\right) \\
&\leq \ln \sum_i \exp\left(A_{ni} + (1-\underline{\beta}_0)a + \sum_{k=1}^{K} \frac{\beta_{k,i}}{1-\gamma} \ln\left(u_{k,i} e^{(1-\gamma)y_i} + w_{k,i}^{-\theta(1-\gamma)}\right)\right) \\
&= \ln \sum_i \exp\left(A_{ni} + \sum_{k=1}^{K} \frac{\beta_{k,i}}{1-\gamma} \ln\left(u_{k,i} e^{(1-\gamma)y_i} + w_{k,i}^{-\theta(1-\gamma)}\right)\right) + \ln e^{(1-\underline{\beta}_0)a} \\
&= F_n(y) + (1-\underline{\beta}_0)a,
\end{aligned}
$$

where $\underline{\beta}_0 = \min_i \{\beta_{0,i}\}$. Boundedness follows from positing a common upper bound $\overline{y} < \infty$ and lower bound $\underline{y} > -\infty$ on each dimension of y and showing that evaluating the right-hand-side at these bounds yields a left-hand-side outcome within them.

B.3 Market Clearing

To solve for wages w_i^l and final spending X_i^F, begin with the uses of goods and services output in country i in final demand and as inputs into the goods and services sector:

$$
\begin{aligned}
Y_i^G &= \alpha_i^G X_i^F + \beta_i^{GG} Y_i^G + \beta_i^{SG} Y_i^S - D_i^G \\
Y_i^S &= \alpha_i^S X_i^F + \beta_i^{GS} Y_i^G - D_i^S
\end{aligned}
$$

to get:

$$
\begin{aligned}
Y_i^G &= \frac{\left[(1-\alpha_i^S) + \beta_i^{SG}\alpha_i^S\right] X_i^F - D_i^G}{1 - \beta_i^{GG} - \beta_i^{SG}\beta_i^{GS}} \\
Y_i^S &= \alpha_i^S X_i^F + \beta_i^{GS} \frac{\left[(1-\alpha_i^S) + \beta_i^{SG}\alpha_i^S\right] X_i^F - D_i^G}{1 - \beta_i^{GG} - \beta_i^{SG}\beta_i^{GS}}.
\end{aligned}
$$

Also using $X_i^P = \beta_i^{GG} Y_i^G$ and $Y_i^P = X_i^P - D_i^G$, equations (22) and (23) become:

$$\frac{\beta_i^{GG}\left[\left(1-\beta_i^{S,L}\alpha_i^S\right)X_i^F - \tilde{D}_i\right]}{\beta_i^{G,L}+\beta_i^{S,L}\beta_i^{GS}} - D_i^G = \sum_n \pi_{ni}\frac{\beta_n^{GG}\left[\left(1-\beta_n^{S,L}\alpha_n^S\right)X_n^F - \tilde{D}_n\right]}{\beta_n^{G,L}+\beta_n^{S,L}\beta_n^{GS}}$$

$$Y_i^l = \frac{\left[\alpha_i^S\left(\beta_i^{S,l}\beta_i^{G,L}-\beta_i^{S,L}\beta_i^{G,l}\right)+\beta_i^{G,l}+\beta_i^{S,l}\beta_i^{GS}\right]X_i^F - (\beta_i^{G,l}+\beta_i^{S,l}\beta_i^{GS})\tilde{D}_i}{\beta_i^{G,L}+\beta_i^{S,L}\beta_i^{GS}} - \beta_i^{S,l}D_i^S, \quad (44)$$

where $Y_i^l = w_i^l L_i^l$ and $\tilde{D}_i = D_i^G + (1-\beta_i^{S,L})D_i^S$. Note, of course, that labor incomes Y_i^l, trade shares π_{ni} (through Ξ_i), and labor shares in the goods sector $\beta_i^{G,l}$ (through $\varpi_{k,i}$) embody the solution for the w_i^l.

B.4 Deriving the Price Index

Like producers, retailers perform tasks of each type k with each task having cost $c_{\omega,k}$ drawn independently from the distribution $G_{k,i}(c)$ given by (14). Consider first retailers with m tasks of each type k. We can derive their price index from the expression:

$$\begin{aligned}
\left[P_i^R(m)\right]^{1-\sigma'} &= p(m)F_i^R E\left[\left(\prod_k\left(\sum_{\omega=1}^m c_{\omega,k}^{-(\sigma-1)}\right)^{-\beta_{k,i}/(\sigma-1)}\right)^{1-\sigma'}\right] \\
&= p(m)F_i^R\prod_k\int_0^\infty\ldots\int_0^\infty\left(\sum_{\omega=1}^m c_\omega^{-(\sigma-1)}\right)^{-\beta_{k,i}(1-\sigma')/(\sigma-1)} dG_{k,i}(c_1)\ldots dG_{k,i}(c_m) \\
&= p(m)F_i^R g_i^R(m)\Xi_i^{(\sigma'-1)/\theta},
\end{aligned}$$

where

$$g_i^R(m) = \prod_k\left(\int_0^\infty e^{-x_1}\ldots\int_0^\infty e^{-x_{m-1}}\int_0^\infty e^{-x_m}\left(\sum_{\omega=1}^m x_\omega^{-(\sigma-1)/[\theta(1-\gamma)]}\right)^{\beta_{k,i}/(\sigma-1)} dx_m dx_{m-1}\ldots dx_1\right).$$

The overall retail price index is:

$$P_i^R = \left[\sum_m\left(P_i^R(m)\right)^{1-\sigma'}\right]^{1/(1-\sigma')} = \left[\sum_m p(m)F_i^R g_i^R(m)\Xi_i^{(\sigma'-1)/\theta}\right]^{1/(1-\sigma')} = g_i^R\Xi_i^{-1/\theta},$$

where:

$$g_i^R = \left(\sum_m p(m) F_i^R g_i^R(m) \right)^{1/(1-\sigma')}.$$

Note that σ' enters the price index only through the constant g_i^R so drops out of price changes.

C Estimation

We divide our estimation procedures into those based on the French data and those based on the EU firm-to-destination data.

C.1 Estimation using French Data

We use these data to estimate the vector of parameters:

$$\Theta = \left\{ \gamma, \left\{ p(m) | m \in \Omega^M \right\}, \beta_s, \sigma \right\}$$

using French firm-to-firm and labor share data for 2005, as described in Sections 5.1.2 and 5.1.3. Our estimation procedure is to find parameter values to match moments from the model (either computed directly or simulated) with moments from the data.

C.1.1 Constructing Data Moments

The moments taken from the data are:

1. **Distribution of buyers per French seller.** We form bins $\tilde{b}$ for the number of buyers per French seller:

$$\tilde{b} \in \Omega^B = \{1, 2, [3,4], [5,8], [9,16], [17,32], [33,\infty)\}.$$

For each European destination $n \neq F$ we calculate $N_{nF}(\tilde{b})/N_{nF}$, the number of French sellers in bin $\tilde{b}$ as fraction of total French sellers to that destination, creating 24 sets of 7 moments, each denoted $m_n(1)$.

2. **Distribution of French sellers per buyer.** We form bins $\tilde{s}$ for the number of French sellers per buyer:

$$\tilde{s} \in \Omega^S = \{1, 2, [3,4], [5,8], [9,16], [17,32], [33,64], [65,128], [129,\infty)\}$$

For each European destination $n \neq F$ we calculate $F_{nF}(\tilde{s})/F_{nF}$, the number of buyers in bin $\tilde{s}$ as fraction of total buyers from French sellers in that destination, creating 24 sets of 9 moments, each denoted $m_n(2)$.

3. **Distribution of production labor shares.** For unskilled u and skilled s production workers we form bins according to the percentiles of the share of labor of that type in French manufacturing firms. We assign a French firm with labor shares β^u and β^s to its appropriate bins for each type of labor. Our moments are the fraction of firms in each bin. We divide firms into bins based on the qth percentiles of labor share where q takes the values:

$$\begin{aligned} q_u &\in \{0.50, 0.55, ..., 0.95, 0.99, 1\} \\ q_s &\in \{0.25, 0.30, ..., 0.95, 0.99, 1\}, \end{aligned}$$

where, since many firms have zero labor shares, we've combined the lower percentiles. We denote the 12 moments for β_u by $m_u(3)$ and the 17 moments for β_s by $m_s(3)$.

C.1.2 Computing Model Moments

In parallel to our data moments, given a parameter vector Θ, we compute the following sets of moments from our model:

1. **Distribution of buyers per French seller.** We use equation (36), summing over $b \in \tilde{b}$, to form $N_{nF}(\tilde{b})/N_{nF}$. We invert (28) for $i = F$, given data on $\overline{b}_{nF}$, to recover $\tilde{\lambda}_{nF}$. We compute integrals numerically. This step delivers the $\hat{m}_n(1, \Theta)$, the model analogue of $m_n(1)$.

2. **Distribution of French sellers per buyer.** We use equation (37), summing over $s \in \tilde{s}$, to form $F_{nF}(\tilde{s})/F_{nF}$. Since (37) builds on (34) we need to compute the binomial probability $\Pr[s_{k,nF} = s_k|m]$. For large m we use the Poisson approximation to the binomial (with the same mean $m\pi_{nF}\varpi_{k,n}$). This step delivers the $\hat{m}_n(2, \Theta)$, the model analogue of $m_n(2)$.

3. **Simulating labor-share distributions.** To compute the distributions of the shares of unskilled u and skilled s production workers of French firms, we simulate $10,000$ firms. For each firm we proceed in four steps:

(a) We draw the number of tasks of each type k from the distribution $p(m)$, $m \in \Omega^M$.

(b) For each type of task, the fraction of spending devoted to task ω is:

$$\pi_{k,\omega} = \beta_k \frac{c_{k,\omega}^{-(\sigma-1)}}{\sum_{\omega'=1}^{m} c_{k,\omega'}^{-(\sigma-1)}},$$

where the task-specific costs have distribution (14). Since $\Phi_{k,n} c_{k,\omega}^{\theta(1-\gamma)}$ has a unit exponential distribution, we draw $x_{k,\omega}$ from that parameter-free distribution to compute:

$$\frac{x_{k,\omega}^{-(\sigma-1)/(\theta(1-\gamma))}}{\sum_{\omega'=1}^{m} x_{k,\omega'}^{-(\sigma-1)/(\theta(1-\gamma))}} = \frac{c_{k,\omega}^{-(\sigma-1)}}{\sum_{\omega'=1}^{m} c_{k,\omega'}^{-(\sigma-1)}}.$$

(c) Each task is carried out by the firm's own workers according to the outcome of a Bernoulli trial with probability of success $1-\varpi_{1,F}$ for unskilled tasks and $1-\varpi_{2,F}$ for skilled tasks.

(d) Combining steps b and c we aggregate across the tasks of each type to obtain the firm-level share of costs for unskilled β^u and skilled β^s production workers.

(e) We assign the firm to its appropriate percentile bin from step 3 of Section C.1.1. The fraction of firms in each bin delivers $\hat{m}_u(3,\Theta)$ and $\hat{m}_s(3,\Theta)$, the model analogues of $m_u(3)$ and $m_s(3)$.

C.1.3 Method-of-Moments Estimation

Stacking our 413 moments we form the vector of residuals between data and model:

$$y(\Theta) = \begin{bmatrix} y_1(1;\Theta) \\ \vdots \\ y_{24}(1;\Theta) \\ y_1(2;\Theta) \\ \vdots \\ y_{24}(2;\Theta) \\ y_u(3;\Theta) \\ y_s(3;\Theta). \end{bmatrix} = \begin{bmatrix} m_1(1) - \hat{m}_1(1;\Theta) \\ \vdots \\ m_{24}(1) - \hat{m}_{24}(1;\Theta) \\ m_1(2) - \hat{m}_1(2;\Theta) \\ \vdots \\ m_{24}(2) - \hat{m}_{24}(2;\Theta) \\ m_u(3) - \hat{m}_u(3;\Theta) \\ m_s(3) - \hat{m}_s(3;\Theta). \end{bmatrix}$$

We treat $y_n(1;\Theta)$ and $y_n(2;\Theta)$ as generated by sampling error in the data moment and $y_u(3;\Theta)$ and $y_s(3;\Theta)$ as generated by simulation error in the model moment. Our moment condition

is:

$$E[y(\Theta_0)] = 0$$

where Θ_0 is the true value of Θ. We thus seek a Θ that achieves:

$$\hat{\Theta} = \arg\min_{\Theta} : \{y(\Theta)'\mathbf{W}y(\Theta)\}$$

where $\mathbf{W}$ is a weighting matrix.

The weighting matrix $\mathbf{W}$ is block diagonal with each of the 50 sets of moments constituting a block. Each block is the Moore-Penrose inverse of the variance-covariance matrix of the corresponding set of moments.

Sampling error implies the variance-covariance matrices:

$$\begin{aligned} V(m_n(1)) &= \frac{1}{N_{nF}}\left[\text{diag}(\hat{m}_n(1:\Theta)) - \hat{m}_n(1:\Theta)\hat{m}_n(1:\Theta)'\right] \\ V(m_n(2)) &= \frac{1}{F_{nF}}\left[\text{diag}(\hat{m}_n(2:\Theta)) - \hat{m}_n(2:\Theta)\hat{m}_n(2:\Theta)'\right]. \end{aligned}$$

Simulation error implies the variance covariance matrices:

$$\begin{aligned} V(\hat{m}_u(3,\Theta)) &= \frac{1}{2000}\left[\text{diag}(m_u(3)) - m_u(3)m_u(3)'\right] \\ V(\hat{m}_s(3,\Theta)) &= \frac{1}{2000}\left[\text{diag}(m_s(3)) - m_s(3)m_s(3)'\right]. \end{aligned}$$

C.1.4 Standard Errors

Table 14 displays the parameter estimates and their standard errors, with and without biased bootstrap correction. The mean across bootstraps is greater than the point estimates for φ, σ and some $p(m)$ for large m, indicating that the bootstrap procedure induces a positive bias for these parameters. An explanation is that, as we randomly sample French exporters from the VAT data with replacement in a setting with many-to-many matching, buyers with many French sellers are overrepresented in the bootstrapped samples. This decreases F_{nF}, the total number of European importers from France, and shifts the empirical distribution of French sellers per buyer upwards. The estimated probability distribution of tasks shifts towards larger values of m and the estimated buyer congestion parameter is also higher.

Table 14
Estimates and Bootstrap Results

	Estimates[1]	Mean across bootstraps	Standard error (without correction)	Standard error (with correction)	Error reduction through bias correction
γ	0.34	0.38	0.038	0.026	-29.8%
φ	0.34	0.33	0.020	0.019	-7.88%
σ	2.64	2.72	0.150	0.105	-29.8%
β_2	0.38	0.38	0.032	0.029	-10.2%
λ_1	1.65	1.66	0.062	0.062	-0.79%
λ_2	0.35	0.34	0.062	0.062	-0.79%
$p(m)$					
1	0.049	0.048	0.008	0.008	-1.50%
4	0.53	0.503	0.049	0.040	-17%
16	0.37	0.364	0.050	0.050	0.006%
64	0.048	0.075	0.037	0.027	-29.2%
256	0.005	0.007	0.003	0.003	-11.4%
1024	0.002	0.003	0.001	0.0008	-22.6%
4096	0.0006	0.001	0.0004	0.0003	-26.1%

Notes: Standard errors are across 25 estimates of the model each using moments computed from a separate bootstrap sample. Each bootstrap sample draws exporters with replacement from the original data, keeping the number of exporters the same. We present standard errors without and with bias correction. The last column shows the impact of bias correction.

[1] As reported in Tables 2, 3, and 4

C.2 Estimation using EU Firm-to-Destination Data

As described in Section 5.1.4, we use the EU firm-to-destination data to estimate the bilateral matching intensities λ_{ni} and the buyer congestion parameter φ.

In step 1 we regress our initial estimate of the total number of firms (producers plus retailers), $\ln \hat{F}_n$, on market size, $\ln X_n$. Figure 8 shows the relationship and the regression slope. The fitted values, multiplied by $2 \times \overline{m}$, become our estimates of the number of buyers B_n in market n.

In step 3 we perform a PPML regression of our initial estimate of buyers per seller $\hat{\bar{b}}_{ni}$ on source i and destination n effects and the log of the distance between i and n. Table 15 reports the results. The fitted values become our estimates of $\bar{b}_{ni}$.

Figure 8: Estimate of total number of firms and market size

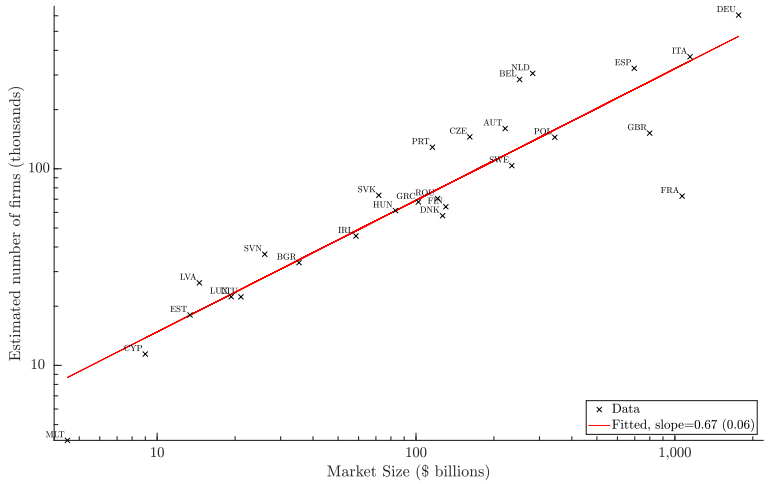

Table 15
Buyers per Seller Projection

Dependent variable		$\hat{\bar{b}}_{ni} - 1$
Log(Distance)		-0.64
		(0.07)
Years since EU entry[1]		
	6 years	-0.61
		(0.20)
	9 years	-0.45
		(0.16)
	18 years	-0.17
		(0.16)
	27 years	-0.17
		(0.17)
	32 years	0.00
		(excluded category)
	40 years	-0.10
		(0.27)
	55 years	-0.33
		(0.20)
Constant		6.58
		(0.53)
N		701
R^2		0.67

Notes: Estimated using Poisson Pseudo-Maximum Likelihood. We include destination and source fixed effects. We have 701 observations rather than 702 (27*26) as Cyprus has no exporter to Luxembourg in 2012.
[1]For each country pair, we calculate the number of years elapsed between the more recent EU entry date of the two and our reference year of 2012.

D Counterfactuals

We use total GDP across countries as our numéraire:

$$Y = \sum_{i=1}^{\mathcal{N}} Y_i = 1.$$

Hence, where x' denotes the counterfactual value of any magnitude x, the counterfactual values of GDP satisfy:

$$\sum_{i=1}^{\mathcal{N}} Y_i' = 1.$$

In addition to exogenous parameters (which include deficits), we also condition on initial values for trade shares, π_{ni}, outsourcing shares, $\varpi_{k,i}$, measures of buyers, B_i, and each type of labor's contribution to GDP, $Y_i^l = w_i^l L_i^l$.

Our applications concern equilibrium responses to exogenous changes in bilateral matching frictions (from $\{\lambda_{ni}\}$ to $\{\lambda_{ni}'\}$) and in iceberg trade costs (from $\{d_{ni}\}$ to $\{d_{ni}'\}$).[46] Solving for the new equilibrium requires solving four systems of equations simultaneously to obtain $\{\hat{Y}_i^l, \hat{w}_i^l, \hat{B}_i, \hat{\Upsilon}_n\}$.

The first systems two involve writing (44) in terms of changes as:

$$\frac{\beta_i^{GG}\left[\left(1-\beta_i^{S,L}\alpha_i^S\right)(Y_i'+D_i)-\tilde{D}_i\right]}{\beta_i^{G,L\,\prime}+\beta_i^{S,L}\beta_i^{GS}} - D_i^G = \sum_n \pi_{ni}' \frac{\beta_n^{GG\,\prime}\left[\left(1-\beta_n^{S,L}\alpha_n^S\right)(Y_n'+D_n)-\tilde{D}_n\right]}{\beta_n^{G,L\,\prime}+\beta_n^{S,L}\beta_n^{GS}}$$

$$Y_i^{l\,\prime} = \frac{\left[\alpha_i^S\left(\beta_i^{S,l}\beta_i^{G,L\,\prime}-\beta_i^{S,L}\beta_i^{G,l\,\prime}\right)+\beta_i^{G,l\,\prime}+\beta_i^{S,l}\beta_i^{GS}\right](Y_i'+D_i)-(\beta_i^{G,l\,\prime}+\beta_i^{S,l}\beta_i^{GS})\tilde{D}_i}{\beta_i^{G,L\,\prime}+\beta_i^{S,L}\beta_i^{GS}} - \beta_i^{S,l}D_i^S,$$

where $Y_i^{l\,\prime} = Y_i^l \hat{w}_i^l$, $\pi_{ni}' = \pi_{ni}\hat{\pi}_{ni}$, $\beta_i^{G,L\,\prime} = \sum_l \beta_i^{G,l\,\prime}$, and $Y_i' = \sum_l Y_i^{l\,\prime}$.

The third system writes (38) in terms of changes as:

$$B_i' = K\bar{m}\int_0^\infty \left(1 - e^{-\sum_n \lambda_{ni}'(B_n')^{1-\varphi/(1-\gamma)}\tilde{\eta}_n'\left(\lambda_{ni}'(B_n')^{-\varphi/(1-\gamma)}y/\pi_{ni}'\right)}\right) dy + K\bar{m}F_i^R,$$

where $\tilde{\eta}_n'(x)$ is given in (30) with $\varpi_{k,n}'$ in place of $\varpi_{k,n}$.

[46] Our methodology would allow us to consider changes in other exogenous parameters, such as technology (from $\{T_i\}$ to $\{T_i'\}$).

The fourth system writes (17) in terms of changes as:

$$\hat{\Upsilon}_n = \sum_i \pi_{ni} \hat{\lambda}_{ni} \hat{d}_{ni}^{-\theta} \prod_{k=0}^{K} \left(\varpi_{k,i} \hat{B}_i^{-\varphi} \hat{\Upsilon}_i^{1-\gamma} + (1 - \varpi_{k,i}) \, \hat{w}_{k,i}^{-\theta(1-\gamma)} \right)^{\beta_k/(1-\gamma)}, \tag{45}$$

The solution requires calculating:

$$\begin{aligned}
\hat{\pi}_{ni} &= \frac{\hat{\lambda}_{ni} \hat{d}_{ni}^{-\theta} \left(\prod_{k=0}^{K} \hat{\Phi}_{k,i}^{\beta_{k,i}/(1-\gamma)} \right)}{\hat{\Upsilon}_n} \\
\hat{\Phi}_{k,i} &= \varpi_{k,i} \hat{B}_i^{-\varphi} \hat{\Upsilon}_i^{1-\gamma} + (1 - \varpi_{k,i}) \, \hat{w}_{k,i}^{-\theta(1-\gamma)} \quad k = 1, 2 \\
\varpi'_{k,i} &= \varpi_{k,i} \frac{\hat{B}_i^{-\varphi} \hat{\Upsilon}_i^{1-\gamma}}{\hat{\Phi}_{k,i}} \quad k = 1, 2 \\
\hat{\Phi}_{0,i} &= \hat{w}_{0,i}^{-\theta(1-\gamma)} \\
\hat{w}_{0,i} &= (\hat{w}_i^t)^{\beta^{G,t}/(\beta_i^{G,t} + \beta_i^{GS})} (\hat{P}_i^S)^{\beta_i^{GS}/(\beta_i^{G,t} + \beta_i^{GS})} \\
\hat{P}_i^S &= \left(\prod_{k=0}^{K} \hat{\Phi}_{k,i}^{\beta_{k,i}/(1-\gamma)} \right)^{-\beta_i^{SG}/\theta} \prod_l (\hat{w}_i^l)^{\beta_i^{S,l}}
\end{aligned}$$

where $\beta_i^{G,p\,\prime} = \beta_{1,i} \left(1 - \varpi'_{1,i}\right)$, $\beta_i^{G,s\,\prime} = \beta_{2,i} \left(1 - \varpi'_{2,i}\right)$, $\hat{w}_{1,i} = \hat{w}_i^p$, and $\hat{w}_{2,i} = \hat{w}_i^s$.