Search Frictions in International Good Markets*

Clémence Lenoir† Julien Martin‡ Isabelle Mejean§

January 10, 2019

Abstract

This paper studies how search frictions in international good markets can distort competition between firms of heterogeneous productivity. We add bilateral search frictions between buyers and sellers in a Ricardian model of trade. Search frictions prevent buyers from identifying the most productive sellers which induces competitive distortions and benefits low-productivity firms at the expense of high-productivity ones. We use French firm-to-firm trade data and a GMM estimator to recover search frictions faced by French exporters at the product and destination level. They are found more severe in large and distant countries and for products that are more differentiated. In a counterfactual exercise, we show that reducing the level of search frictions leads to an improvement in the efficiency of the selection process because the least productive exporters are pushed out of the market while the export probability and the conditional value of exports increase at the top of the productivity distribution. As a consequence, the mean productivity of exporters increases significantly.

JEL Classification: F10, F11, F14, L15

Keywords: Firm-to-firm trade, Search frictions, Ricardian trade model, Structural estimation

---

*The paper has benefited from insightful comments by participants to the seminars at Warwick University, LSE, University of Cergy-Pontoise, Paris School of Economics, CREST, Yale-PhD seminar, GSIE, RIEF and ETSG. We are particularly indebted to Andrew Bernard, Tibor Besedes, Johannes Boehm, Edouard Chretien, Arnaud Costinot, Gregory Corcos, Jonathan Eaton, Etienne Lalé, Sam Kortum, Francis Kramarz, Thierry Mayer, Andreas Moxnes, Gianmarco Ottaviano, Sophie Osotimehin and Esteban Rossi-Hansberg who provided feedbacks at various stages of the project. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 714597).

†CREST-ENSAE, Email address: clemence.lenoir@ensae-paristech.fr.

‡UQAM, CREST, and CEPR, Email address: martin.julien@uqam.ca.

§CREST-Ecole Polytechnique and CEPR, Email address: isabelle.mejean@polytechnique.edu.

1 Introduction

Since the seminal paper by Melitz (2003), the trade literature has extensively discussed the self-selection of firms into exporting as an engine of welfare gains from trade.[1] While the Melitz model holds true on average, it is well-known that its performances are more mixed when it comes to explaining individual firms' export decisions (Eaton et al., 2011). This is especially true among the sub-sample of small and medium firms, as illustrated in Figure 1. Namely, the correlation between the export probability and the destination's market access, which is expected to be positive due to the self-selection of the most productive firms in difficult destinations, is as strong as .88 among the subsample of the 15% largest exporters (the red labels in Figure 1). However, the correlation is substantially lower, at .65, among relatively small exporters represented by the green labels.

This paper argues that adding search frictions into a Ricardian model of trade is an attractive way of accounting for such fuzzy export patterns at the bottom of the productivity distribution. Search frictions enable buyers to identify the most efficient suppliers which distorts the strength of competition in international markets in favor of low-productivity exporters. This can explain some of the randomness observed in the data as the strength of self-selection mechanisms is lowered in markets displaying high search frictions. In those markets, a relatively large share of low-productivity sellers enter which reduces the average productivity of exports.

Our model is a partial equilibrium version of Eaton and Kortum (2002) in which there is a large number of ex-ante homogeneous buyers of each variety located in each country. As in Eaton et al. (2018), each of these buyers meet with a random number of suppliers drawn from the overall distribution of potential producers of the variety. Conditional on their draw, buyers choose to interact with the lowest-cost supplier. Importantly, it is assumed that the random matching is systematically biased geographically due to heterogeneous bilateral search frictions. In practice, all sellers from, say, France face the same level of search frictions and thus have the same probability of meeting with a foreign buyer when exporting to a given destination. But the bilateral heterogeneity implies that French exporters face different frictions in different markets. Moreover, they compete within a market with firms originating

[1] See Melitz and Redding (2015) and the literature cited therein. A vast empirical literature has provided support to the theory. From this literature, it is well-known that exporters are on average more productive than non-exporters (Bernard and Jensen, 1995; Mayer and Ottaviano, 2008), and that the mean productivity of exporters is higher in more difficult destinations (Eaton et al., 2004). Episodes of trade liberalization have also been used to show how opening to international trade affects the within and between productivity of domestic firms (Fernandes, 2007; Topalova and Khandelwal, 2011; Pavcnik, 2002).

Figure 1: ***Export probability as a function of market potential, small and large exporters***

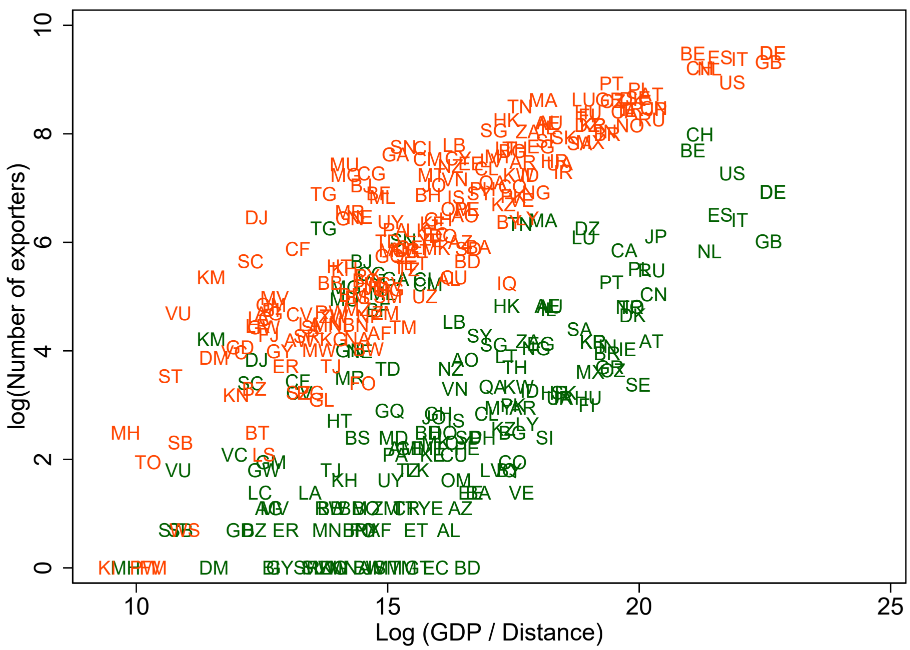

Notes: This figure correlates the stock of French exporters active in a given destination with the destination's market potential, for the top 15% and the bottom 15% of firms (red and green labels, respectively). The size of a firm is defined as the value of its worldwide exports. Source: French customs described in Bergounhon et al. (2018).

from other countries, which do not face the same level of frictions.

We first draw analytical predictions regarding how the magnitude of search frictions affects export patterns, at the product-level and for individual firms. At the product-level, bilateral exports are negatively correlated with the relative size of search frictions there. A larger share of a country's consumption is spent on goods produced in countries displaying low search frictions, everything else equal. From this point of view, search frictions are no different from other physical barriers to international trade studied in Eaton and Kortum (2002). At most, their introduction into the model can help explain why countries that are culturally closer trade more together.[2] More interesting are the model's predictions regarding individual

[2]It has long been recognized that physical barriers to international trade are not the only impediment to international trade. In the gravity literature, a common language or former colonial ties are well-known to contribute substantially to the model's explanatory power (Head and Mayer, 2014). More closely related to the interpretation we have of search frictions, Rauch (1999); Rauch and Trindade (2002) provide evidence that

firms' export patterns. Theoretically, the impact of frictions on export probabilities is indeed non-monotonic along the distribution of productivities. While more productive firms always suffer from more frictions, the impact is lower, or actually reversed, at the bottom of the distribution. The non-monotonicity is a direct consequence of the model's assumptions. In our framework, serving a given client abroad requires to i) meet with her and ii) be chosen as her partner, conditional on a meeting. While the meeting probability is constant across firms, the likelihood that a firm is chosen depends on its productivity, in relative terms with competing firms in the importer's random choiceset. The strength of competition is reduced in highly frictional destinations where each importer on average meets with a smaller number of potential partners. This tends to increase the chance that a low-productivity exporter ends up serving the firm.

This prediction is the key element of the model which we argue can explain fuzzy export patterns at the bottom of the distribution of exporters' size. By chance, even a poorly productive exporter can end up serving any foreign country, the probability that this happens being increased in more frictional destinations. The prediction is also what lets us estimate the frictions structurally, separately from other barriers to international trade. To this aim, we exploit firm-to-firm trade data covering the universe of French exporters and each of their individual client in the European Union.[3] Such data allow us to document a new dimension of heterogeneity among exporters, regarding the number of partners they serve in a given destination country. This number displays a significant degree of heterogeneity within a product and destination. On average, large firms tend to serve more buyers. Within a firm, the number of partners served in a destination is decreasing in distance from France but is systematically larger in destinations that have closer links to France through past bilateral migration flows.

Our estimation of search frictions exploits this heterogeneity. Namely, we use as empirical moment the dispersion across French exporters in terms of the *number* of partners they serve in a given destination. This heterogeneity, we argue, cannot be explained by physical trade barriers. In our model, the dispersion comes from search frictions affecting individual firms' export probabilities. More frictions reduce the dispersion across individual firms by dampening high-productivity firms' export premium. Since iceberg trade costs do not have such distorsive effect, exploiting this moment of the data is useful to recover search frictions

the stock of migrants from one origin in a country is significantly correlated with more bilateral trade. Their interpretation of this finding is that migrants help reduce information frictions characterizing international good markets.

[3]See Bernard et al. (2018b); Carballo et al. (2018) for examples of papers using similar data covering other exporting countries.

separately from other trade barriers.

Using this empirical moment and its theoretical counterpart, bilateral search frictions are estimated by the generalized method of moments for about 10,000 product and destination country pairs. In countries for which the product set is sufficiently comparable, the maximum degree of average frictions faced by French exporters is found in Greece and Finland while the less frictional country is Belgium. Search frictions are estimated to be stronger in differentiated product markets. Within a product, they are more pronounced in distant and more populated countries, while lower in countries where the population of French migrants is larger. Importantly, the estimated model is able to explain about one fifth of the heterogeneity observed in the data regarding the share of exporters serving a given number of importers in a destination. Given the simplicity of the model which relies on a single parameter to explain this heterogeneity, we find this measure of fit quite encouraging.

Once estimated the model can be used to run counterfactuals. Our main experiment consists in simulating the impact that a reduction in bilateral frictions with Greece to the level observed in Belgium would have on aggregate and firm-level export patterns. Results can be summarized as follows. First, a reduction in frictions with Greece, keeping everything else unchanged, explains a 7 percentage point increase in the market share of the median French product in Greece. This aggregate effect however hides a substantial impact on the allocation of resources across exporting firms. Namely, the export probability to Greece falls in the bottom 15% of the distribution, by 3.5 percentage points, on average, to reach less than 3.5%. At the top of the distribution instead, export probabilities increase from 65.3 to 84% among the top 15% productivity percentiles and from 70 to 92.5% among the top 5%. Within the sub-sample of exporters, a reduction in search frictions also reallocates market shares with the expected number of clients served by high-productivity firms increasing substantially. All in all, the mean productivity of exporters increases by 10 to 20% as a consequence of Greek importers being better able to identify the most productive French suppliers.

In comparison with other barriers to international trade, search frictions thus have important misallocative consequences. For this reason, reducing such frictions might be of especially strong policy relevance. It also comes with a cost for the least efficient firms that are likely to exit the market. Within the toolbox of export-promoting agencies, programs aimed at increasing the visibility of domestic sellers abroad can be an efficient tool for increasing export flows in a non-distortive way, especially if they target small but highly productive firms.[4]

[4]The French export promoting agency offers several programs, which are meant to help firms meet with foreign clients. The agency notably helps financing firms' participation to international trade fairs or organizing bilateral meetings with representatives of the sector in the destination country. See details on the agency's website, www.businessfrance.fr

Our paper is related to two strands of the literature. The first one concerns a number of recent contributions which have used similar firm-to-firm trade data to study the matching between exporters and importers in international markets (Bernard et al., 2018b; Carballo et al., 2018). The heterogeneity between exporters in terms of the number of buyers they serve is explained in models imposing an additional source of ex-ante heterogeneity, regarding the productivity or the preferences of importers in foreign markets. Our model instead displays ex-ante homogenous importers which ex-post heterogeneity is solely driven by the randomness in the matching process (Chaney, 2014; Eaton et al., 2018; McCallum and Krolikowski, 2018).[5] In term of modeling, our framework borrows from Eaton et al. (2018) by introducing random frictions in a Ricardian model of trade. Eaton et al. (2018) use this framework to study the interplay between trade and the labor share both at the firm and at the aggregate levels. Instead, we analyze the heterogeneous impact of frictions on high- vs low-productivity exporters. The role of search and information frictions in international markets is the topic of an older empirical and theoretical literature. Rauch (2001) thus explains the role of migrant networks in international markets by way of such frictions. More recently, Lendle et al. (2016), Bernard et al. (2018a), Akerman et al. (2018), and Steinwender (2018) provide evidence of such frictions being an important barrier to international trade, using various natural experiments of a decrease in information frictions, namely the launching of a telegraph line between London and New York in Steinwender (2018), the opening of the Japanese high-speed train (Shinkansen) in Japan in Bernard et al. (2018a), the adoption of broad band internet in Norwegian municipalities in Akerman et al. (2018), and the development of online markets in Lendle et al. (2016). This topic has also been studied theoretically in several recent contributions. In Allen (2014), information frictions hit the seller side of the economy; exporters ignore the potential price of their crops abroad, thus enter into a sequential search process. We instead introduce frictions on the demand side of the economy, with buyers having an imperfect knowledge of the supply curve. From this point of view, our model is closer to Dasgupta and Mondria (2018). Their model of inattentive importers assumes that buyers optimally choose how much to invest into information processing to discover potential suppliers. In comparison with theirs, our model is based on simpler assumptions since the search process is purely random. The tractability of this framework allows us to derive closed-form

[5] As will become clear from the presentation of the model, the fact that buyers are homogenous ex-ante implies that this side of the market is very stylized. In particular, the model will not reproduce a stylized fact which is extensively analyzed in Bernard et al. (2018b), namely that individual importers display a strong degree of heterogeneity in terms of the number of sellers they are connected to, ex-post. Because our purpose is to explain the fuzziness in *exporters*' participation to trade, we see this property of our model innocuous, although unrealistic.

solutions and estimate frictions structurally.[6] This tractability also allows us to emphasize the non-monotonic impact of search frictions at the individual level.[7]

The rest of the paper is organized as follows. In section 2, we present the data and stylized facts on firm-to-firm trade which we will later use to build and test the model. We most specifically focus on the number of buyers served by a given firm, and study how it varies across firms, products and destinations. Section 3 describes our theoretical model and derives analytical predictions regarding the expected number of clients that an exporter will serve in its typical destination. Section 4 explains how we estimate the magnitude of search frictions using a GMM approach. We also provide summary statistics on the estimated frictions and the model fit. Section 5 uses the estimated frictions in a counterfactual exercise to discuss how search frictions affect the allocation of resources across exporters. Finally, Section 6 concludes.

2 Data and stylized facts

2.1 Data

The empirical analysis is conducted using detailed export data covering the universe of French firms. The data are provided by the French Customs and are described in details in Bergounhon et al. (2018). The full data set covers all transactions that involve a French exporter and an importing firm located in the European Union, over 1995-2017. Our analysis focuses on data for 2007 but we checked that statistics are not sensitive to the choice of the reference year. Since the analysis is conducted at the product-level, we have no choice but to drop all transactions that are reported under the simplified declaration regime, for which the product category is not recorded. This concerns 10% of firms which overall exports in the European Union during the year is below 150,000 euros.[8]

[6]The cost of this tractability is an extreme degree of passivity of firms regarding frictions. In general, one would expect high-productivity firms to be willing to invest in advertising so that to increase their visibility in foreign markets (Arkolakis, 2010). Instead, our model assumes that large exporters have the same probability of meeting with a buyer as low-productivity ones. While this is not very realistic, our purpose is to explain what happens at the bottom of the distribution, whereby low-productivity firms display fuzzy export patterns. For these firms, the assumption that the meeting probability is taken as given seems more realistic. We discuss in Appendix A.3 the sensitivity of our results to this assumption.

[7]In our framework, the effect of frictions is ambiguous at the individual level but not at the aggregate level. See Petropoulou (2011) for a model where search frictions may have a non-monotonic impact on aggregate trade flows.

[8]One might be concerned that this selection bias our empirical analysis since the neglected small exporters are likely to display systematically different patterns of exports. While we cannot rule this out with certainty,

For each transaction, the data set records the identity of the exporting firm (its SIREN identifier), the identification number of the importer (an anonymized version of its VAT code), the date of the transaction (month and year), the product category (at the 8-digit level of the combined nomenclature) and the value of the shipment. It is also possible to link each exporter to its sector of activity using INSEE data. In the analysis, data are aggregated across transactions within a year, for each exporter-importer-hs6 product triplet. Such aggregation helps focus on the most important novelty in the data, which is the explicit identification of both sides of the markets; the exporter and its foreign partner. The product dimension will allow conditioning our results on the good being traded, as in the model. A "seller" will thus be an exporter of a specific product. This hypothesis comes down to redefining a French exporter as a single-product firm and neglecting any potential complementarity between products sold by the same firm.

Since we are interested in the extent of search frictions that an exporter faces in foreign markets, we restrict each exporter's product portfolio to products that represent at least 10% of export sales for at least one French seller in the firm's sector of activity.[9] This restriction substantially reduces the number of exporter $\times$ product pairs covered (by almost 50%) without having much of an impact on the aggregate value of exports (-8%), on the population of importers (-4%) and on the population of exporters (which is left unaffected).

Table A1 in appendix provides summary statistics on the number of sellers, buyers and products, by destination. In 2007, we have information on 44,255 French firms exporting to 572,536 individual importers located in the 26 countries of the European Union. Total exports by these firms amount to 216 billions euros. This represents 53% of France worldwide exports. Table A1 displays the number of individuals involved in each bilateral trade flow. Most of the

we believe that the bias should not be substantial based on evidence reported in Figure A.1. Namely, the distribution of sellers' degrees, which product-specific equivalent is used to compute the empirical moments in the estimation, is very similar in the whole sample and in the sample restricted to the 90% of exporters that declare a product category. While there are obviously more exporters with one buyer in the restricted sample, the difference is roughly proportional to the total number of such exporters in the whole dataset (bottom panel).

[9] The rationale for such restriction is that we see in the data firms selling many different products, some of which being relatively "close" to the firm's activity (say exports of wine in agricultural sectors) and others being hardly related to their main activity (e.g. export of glasses for wine producers). In this example, glasses are most probably side products which the firm sells to its customers while they buy some of its wine. While information frictions might be important to identify potential wine consumers, we shall not expect frictions in the market of glasses to affect the wine producer's ability to sell this product; such tied selling only depends on the firm's ability to meet with wine consumers. In practice, it is almost impossible to decide which products are tied and which are not. The statistical criterion that we use thus considers that a product which no firm in the sector sells in large enough quantities is probably tied and is thus removed from the sample.

time, the number of importers is larger than the number of exporters selling to this destination (Columns (1) and (2)). This suggests that the degree of exporters (number of importers they are connected to) is on average larger than the degree of importers (number of French exporters they interact with). This is even more true once we focus on product-specific trade flows as in Columns (4) and (5). Column (3) in Table A1 reports the number of exporter-importer pairs which are active in 2007 and Column (6) the number of exporter-importer-product triplets. These numbers are small in comparison with the number of *potential* relationships, equal to the number of active exporters times the number of importers. This suggests that the density of trade networks is low on average.

The firm-to-firm dataset is complemented with several product-level and aggregate variables used to run gravity regressions in Section 2.2. Distance data are taken from CEPII (Mayer and Zignago, 2011). We control for the market's overall demand using HS6-specific imports in the destination, less the demand for French goods. Multilateral import data are from ComTrade. Finally, information frictions are controlled for using the stock of migrants per origin and destination countries, taken from the UN database on Trends in International Migrant Stock. Following Rauch and Trindade (2002), the degree of information frictions between France and destination i is expected to be inversely related to the share of French citizens in the destination's population and the share of migrants from i in France.

2.2 Descriptive Statistics

As explained in Section 2.1, the most important novelty in our data is the identification of both sides of international trade flows, not only individual exporters but also their foreign clients in each destination. We now present stylized facts exploiting this dimension to characterize the nature of interactions between sellers and buyers engaged in international trade. The facts are later used to motivate the model's assumptions and back out a number of theoretical predictions.

Figure 2 shows the strong heterogeneity in the number of buyers per seller within a destination.[10] The left panel documents the share of sellers interacting with a given number of buyers while the right panel depicts their relative weight in overall exports. To illustrate the amount of heterogeneity across destination countries, Figure 2 displays the distribution obtained in the average European destination (circle points) as well as those computed for two specific destinations, which represent extreme cases around this average, namely Romania

[10] Remember that here and in the rest of the paper, a seller is identified by its siren number *and* the product category. The statistics underlying this graph is thus somewhat difference than in Bernard et al. (2018b), although the conclusion regarding the strong heterogeneity in exporters' number of clients holds in both cases.

Figure 2: ***Distribution of the number of buyers per seller, across exporters***

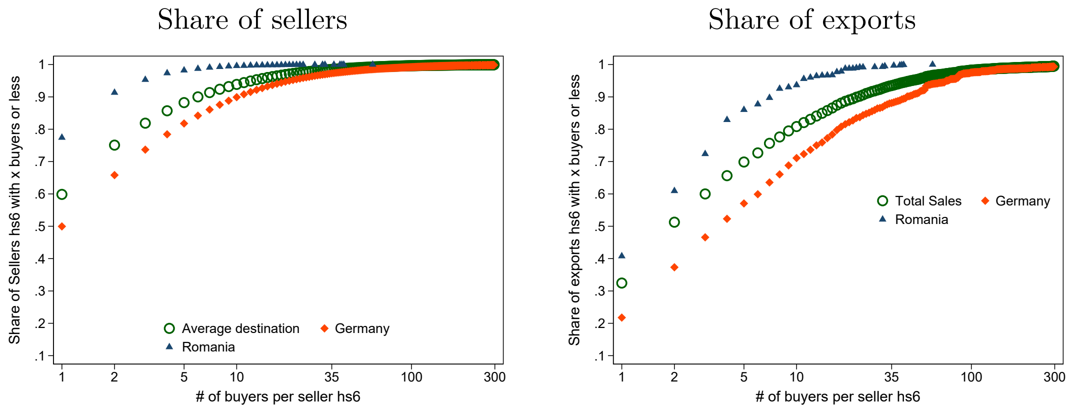

Notes: The figure displays the proportion of sellers (left panel) and the share of trade accounted for by sellers (right panel) that serve x buyers or less in a given destination, in 2007. A seller is defined as an exporter-HS6 product pair. The green circles correspond to the average across EU destinations. The blue triangles and red diamonds are respectively obtained from exports to Romania and Germany.

and Germany (triangle and diamond points, respectively).

In France's typical export market, 65% of sellers interact with a single buyer, and 90% with at most 5 buyers. At the other side of the spectrum, one percent of sellers interact with more than 100 buyers in the same destination. As shown by the right panel in Figure 2, sellers interacting with a single buyer in their typical destination account for about a third of French exports and are thus smaller than the average firm in the distribution. Still, 80% of trade is made up by sellers interacting with at most 10 buyers. From this, we conclude that French exports are dominated by sellers interacting with a small number of buyers.

Figure 2, circle points, hides a substantial amount of heterogeneity in the number of buyers per seller, across both sectors and destinations. The other two distributions depicted in Figure 2 illustrate the cross-country heterogeneity.[11] While the median degree of sellers is equal to just one buyer in all destination countries, the mean varies quite substantially, due to varying shares of sellers who manage to serve more clients. Such heterogeneity also exists across sectors, although perhaps less pronounced. A full variance decomposition however shows that more than 80% of the heterogeneity in the number of buyers served by a seller is within a sector and destination. This is the dimension of heterogeneity that the structural estimation uses to identify search frictions.

At this level, heterogeneity in terms of the number of buyers is significantly correlated with the seller's size, as measured by the worldwide value of the firm's exports. The conditional

[11]Table A2 in Appendix provides more systematic evidence based on the whole set of destination countries.

correlation coefficient is equal to .28 and size explains 37% of the within-variance. The positive correlation between a seller's size and the number of importers it is able to serve within a destination is consistent with evidence in Bernard et al. (2018b) and Carballo et al. (2018) based on similar data for other countries. In Bernard et al. (2018b), the heterogeneity in exporters' degrees is explained in a two-sided heterogeneity model in which importers of heterogeneous size can interact with several exporters. In our model instead, it is assumed that an importer is matched with a single seller, at a point in time. This is justified by another property of our data, which is that more than 89% of importers purchase a given product from a single French exporter. This explains that the mean degree of buyers which can be recovered from the comparison of columns (5) and (6) in Table A1 is very close to one in all destinations.[12]

Having documented new dimensions of heterogeneity in firm-to-firm trade data, we close this section with an empirical analysis using the gravity framework to show how the buyer margin affects the geography of French exports. Table 1 summarizes the results. The gravity equation is run at the product level (columns (1)-(4)) and within a firm (columns (5)-(7)). Bilateral trade is explained by distance to France, proxies for market size, namely the country's (product-specific) import demand and GDP per capita, and proxies for information frictions, the stock of natives from the destination country in France and the stock of French citizens in the destination.

Column (1) confirms the results found in the rest of the literature, namely that product-level bilateral trade is larger towards closer, bigger and wealthier destination markets. Moreover, it is positively correlated with the stock of migrants in France and the destination country which we interpret as information frictions having a negative impact on bilateral trade.[13] These results are also confirmed within a firm, in Column (5). As largely documented in the previous literature, see e.g. Bernard et al. (2007), gravity effects in international trade are attributable to the cross-country heterogeneity of bilateral trade flows at the intensive margin,

[12] While our model is consistent with this property of the data, it fails to take into account another property of the data, which Bernard et al. (2018b) analyze, namely that importers are heterogeneous in terms of the *number of products they import*, which also determines the number of exporters they are connected to. Because we work at the product-level, we implicitly assume that the same importer importing two products can be considered as two importers purchasing two different products. This assumption might be problematic if these buyers were able to enjoy economies of scale on search costs by purchasing the two products from the same exporter. This is not what happens in general, as shown by the very high correlation in the data between the number of sellers a buyer is connected to and its number of seller×product pairs.

[13] In comparison with a specification that does not control for information frictions, the impact of distance is reduced by about a third. This suggests that information frictions are correlated with distance from France in this sample.

Table 1: ***Product- and Firm-level gravity equations***

	Dependent Variable (all in log)						
	Product-level				Firm-level		
	Value of Exports	# Sellers	# Buyers per Seller	Mean export per Buyer-seller	Value of Exports	# Buyers	Exports per Buyer
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
log Distance	-1.079***	-0.448***	-0.303***	-0.327***	-0.390***	-0.269***	-0.121**
	(.076)	(.035)	(.025)	(.045)	(0.070)	(0.029)	(0.052)
log Import Demand	0.737***	0.226***	0.118***	0.393***	0.396***	0.162***	0.234***
	(.015)	(.007)	(.005)	(.009)	(.014)	(.007)	(.011)
log GDP per Capita	0.284***	0.157***	0.121***	0.007	0.112***	0.038***	0.074***
	(.041)	(.019)	(.013)	(.024)	(.024)	(.014)	(.018)
French Migrants	0.028***	0.044***	0.004	-0.020***	0.027***	0.019***	0.008**
	(.009)	(.005)	(.004)	(.005)	(.006)	(.003)	(.004)
Migrants in France	0.121***	0.046***	0.038***	0.038***	0.054***	0.027***	0.027***
	(.008)	(.003)	(.003)	(.004)	(.005)	(.002)	(.004)
Observations	60,720	60,720	60,720	60,720	577,237	577,237	577,237
R-squared	0.655	0.797	0.461	0.587	0.693	0.445	0.720
Fixed effects	Product	Product	Product	Product	Firm	Firm	Firm

Notes: Standard errors, clustered in the country dimension, in parentheses with ***, ** and * respectively denoting significance at the 1, 5 and 10% levels. "log Distance" is the log of the weighted distance between France and the destination. "log Import demand" is the log of the value of the destination's demand of imports for the hs6-product, less the demand addressed to France. "log GDP per capita" is the log-GDP per capita in the destination. "French migrants" is the number of French citizens in the destination country, per 1000 inhabitants. "Migrants in France" is the number of migrants from the destination in France, expressed as a stock per 1000 inhabitants in France. The dependent variable is either the log of product-level French exports in the destination (column (1)) or one of its components, namely the number of sellers involved in the trade flow (column (2)), the mean number of buyers they serve (column (3)) and the mean value of a seller-buyer transaction (column (4)). Column (5) uses as left-hand side variable the log of firm-level bilateral exports while columns (6) and (7) use one of its components, the number of buyers served (column (6)) or the value of exports per buyer (column (7)).

i.e. in terms of the mean shipment per firm, *and* at the extensive margin, in terms of the number of firms exporting. We confirm this result in Columns (2)-(4) and Columns (6)-(7), where bilateral trade flows are further decomposed into intensive and extensive components. Importantly, the buyer dimension of the data allows us to control for an additional source of extensive adjustments, namely the number of buyers in existing exporters' portfolio of clients (see also Bernard et al. (2018b) for a similar decomposition based on Norwegian data).[14] All margins of bilateral trade significantly contribute to the sensitivity of trade to gravity variables. In particular, the "buyer" extensive margin is responsible for 28% of the overall distance elasticity at the product-level, a number that jumps to 69% once gravity coefficients are identified within a firm.[15] Likewise, the buyer margin accounts for a substantial share of the overall impact of migrants. Our interpretation of this finding is that migrants help alleviate information frictions in international markets, which in turn facilitates the matching between exporters and buyers.

This analysis thus confirms previous results in the literature regarding the heterogeneity across exporting firms, in terms of the number of buyers they serve in a destination. This

[14]More specifically, the product-level decomposition used in Table 1, Columns (1)-(4), is based on the following decomposition:

$$\ln x_{pd} = \underbrace{\ln \#^S_{pd}}_{\#\ Sellers} + \underbrace{\ln \frac{1}{\#^S_{pd}} \sum_{s \in S_{pd}} \#^B_{spd}}_{\#\ Buyers\ per\ Seller} + \underbrace{\ln \frac{1}{\#^{SB}_{pd}} \sum_{s \in S_{pd}} \sum_{b \in B_{spd}} x_{sbpd}}_{Mean\ exports\ per\ Buyer-seller}$$

where x_{pd} denotes the value of French exports of product p in destination d which is the sum of all firm-to-firm transactions x_{sbpd}. S_{pd} is the set of the sellers serving this market and B_{spd} the set of the importers purchasing product p from seller s. $\#^S_{pd}$, $\#^B_{spd}$ and $\#^{SB}_{pd}$ respectively denote the number of sellers, the number of buyers seller s is connected to and the total number of active seller-buyer pairs in market pd.

Likewise, the decomposition of firm-level exports in Columns (5)-(7) of Table 1 is based on the following decomposition of trade into an extensive and an intensive terms:

$$\ln x_{spd} = \underbrace{\ln \#^B_{spd}}_{\#\ Buyers} + \underbrace{\ln \frac{1}{\#^B_{spd}} \sum_{b \in B_{spd}} x_{sbpd}}_{Mean\ exports\ per\ Buyer}$$

[15]Note that the contribution of the buyer margin is artificially low in the decomposition of product-level trade in Columns (1)-(4) because of the multicolinearity between the "seller" and 'buyer" extensive margins. If we instead work with this decomposition:

$$\ln x_{pd} = \ln \#^S_{pd} + \ln \#^B_{pd} + \ln \frac{\#^{SB}_{pd}}{\#^S_{pd} \times \#^B_{pd}} + \ln \frac{1}{\#^{SB}_{pd}} \sum_{s \in S_{pd}} \sum_{b \in B_{spd}} x_{sbpd}$$

which treats sellers and buyers symmetrically, the distance elasticity is found larger on the buyer than the seller margin (i.e. $\left|\frac{d \ln \#^B_{pd}}{d \ln Dist_d}\right| > \left|\frac{d \ln \#^S_{pd}}{d \ln Dist_d}\right|$).

number is systematically correlated with the size of the exporter. It also varies within a firm, across destinations, with on average less buyers served in distant destinations displaying more information frictions. In the next section, we build a model which is consistent with the main features of the data.

3 Model

The model is a partial equilibrium version of Eaton and Kortum (2002), extended to search frictions as in Eaton et al. (2018). The analysis is conducted at the level of a product, given factor prices and we do not aggregate across sectors. After having summarized the main assumptions, we derive a number of analytical predictions which are later used in the structural estimation.

3.1 Assumptions

The economy is composed of N countries indexed by $i = 1, ..., N$. In this economy, a single good is consumed and produced into perfectly substitutable varieties by a continuum of firms, some of which being inactive ex-post.

The supply side of the model is almost the same as in Eaton and Kortum (2002). Namely, there is a continuum of producers of the good in each country j, of measure $T_j z_{min}^{-\theta}$. Those firms produce with a constant-returns-to-scale technology using an input bundle which unit price c_j is taken as exogenous. The productivity of a firm s_j located in country j is independently drawn from a Pareto distribution of parameter θ and support $[z_{min}, +\infty[$. The measure of firms in country j that can produce with efficiency above z is thus:

$$\mu_j^Z(z) = T_j z^{-\theta}$$

In the rest of the analysis, firms will be designated by their productivity, with z_{s_j} being the realized productivity of firm s_j. The exporter-hs6 product pairs studied in Section 2 are the empirical counterpart of these firms. Heterogeneity across firms regarding the number of buyers they serve in a destination will later be explained by the underlying productivity heterogeneity.

There are iceberg trade costs between countries. To serve market i with one unit of the good, firms from country j need to produce $d_{ij} > 1$ units. The cost of serving market i for a firm s_j is thus $\frac{c_j d_{ij}}{z_{s_j}}$. Given input prices and international trade costs, the measure of firms from j that can serve market i at a cost below p is $\mu_{ij}(p) = T_j \left(\frac{d_{ij} c_j}{p}\right)^{-\theta}$.

Summing over all producing countries gives the measure of firms which can serve country i at a cost below p, $\mu_i(p) = p^\theta \sum_{j=1}^{N} T_j(d_{ij}c_j)^{-\theta} = p^\theta \Upsilon_i$. As in Eaton and Kortum (2002), $\Upsilon_i = \sum_{j=1}^{N} T_j(d_{ij}c_j)^{-\theta}$ reflects "multilateral resistance" in country i and governs the country's price distribution: the higher Υ_i, the more competitors with low costs in this country.

In Eaton and Kortum (2002), the demand side of the model is summarized by the CES demand of a representative consumer in each country i. We depart from their framework and instead assume that each country is populated by a large number B_i of (ex-ante) homogeneous buyers, each one being willing to spend one unit of the numeraire into the good. Because of search frictions, each buyer b_i meets with a discrete number of suppliers, drawn into the distribution $\mu_i(p)$. Conditional on the subset of producers met, the buyer decides on which one to purchase from, by comparing the prices they offer. In the rest of the analysis, we will assume that producers price at their marginal cost, as in a perfect competition framework. As a consequence, buyer b_i chooses to purchase the good from the lowest-cost supplier who she met and pays the price:[16]

$$p_{b_i} = \underset{s_j \in \Omega_{b_i}}{\arg\min} \left\{ \frac{c_j d_{ij}}{z_{s_j}} \right\}$$

where Ω_{b_i} is the set of producers drawn by buyer b_i in the distribution $\mu_i(p)$.

The number of potential suppliers in the set Ω_{b_i} reflects the extent of search frictions in the economy. In a frictionless world, each buyer b_i would meet with all the firms in $\mu_i(p)$. Within a destination, all buyers would thus end up paying the same price for the homogenous good and the assumption of a representative consumer would be suitable. This is the assumption in Eaton and Kortum (2002), which generates an ex-post degenerated distribution of firms since only the lowest-cost suppliers are active in market i. We instead assume that the number of price quotes in Ω_{b_i} is a random variable.

Namely, the number of suppliers from j that buyer b_i meets follows a Poisson law of

[16] One might question the assumption of marginal cost pricing in a context of frictional good markets. We think of marginal cost pricing as the result of some "price-posting" process, a situation in which producers need to define their price ex-ante, before the matching process. Under such pricing rule, and because the extent of competition within the mass $\mu_i(p)$ is important, marginal cost pricing is an equilibrium outcome. Ex-post, the producer might however be willing to deviate from this pricing rule. An alternative would be to assume that firms drawn by a buyer b_i compete à la Bertrand. Under such assumption, buyer b_i would optimally match with the lowest cost supplier, as in the case of marginal cost pricing, but would be charged a price which would equal the marginal cost of the second lowest-cost supplier. Under the assumption of inelastic demands, this does not change the results. If demand was elastic, we would need to keep track of both the first and second lowest cost supplier while computing the ex-post value of trade. Since most of our results exploit predictions regarding who each buyer is matched with, while neglecting the price at which the good is purchased, we think of the problem as being irrelevant. For this reason, we will stick to the simpler assumption of marginal cost pricing.

parameter $\lambda_{ij} T_j z_{min}^{-\theta}$.[17] Likewise, the number of suppliers from j offering a price below p can be represented by a Poisson process of parameter $\lambda_{ij}\mu_{ij}(p)$. In the rest of the analysis, λ_{ij} is interpreted as an inverse measure of frictions, which we assume is specific to each country pair (and each product). A coefficient closer to one implies that buyers from i gather more information on potential suppliers in country j and are thus more likely to identify the most competitive ones.

Heterogeneity in the magnitude of search frictions across countries means that the subset of firms which a buyer meets is biased towards firms located in countries with which search frictions are lower, on average. Within an origin country, all producers however have the same probability of being drawn, no matter their productivity. This is the key assumption which will generate ex-post heterogeneity across buyers regarding the price they pay. Namely, lucky buyers will end up with a random choiceset Ω_{b_i} which contains low cost producers. As a consequence, they will pay the homogeneous good at a low price. At the other side of the border, even poorly productive sellers can end up serving a distant and frictional country, which happens if they are lucky enough to be drawn by an unlucky buyer which has no better choice than buying the good from this high cost producer.

As shown in Eaton et al. (2018), the assumption of Poisson draws into a Pareto distribution delivers a Weibull distribution for the minimum price at which a buyer b_i can purchase the good (see the proof in Appendix A.1):

$$G_i(p) = 1 - e^{-p^{\theta}\Upsilon_i \kappa_i \widetilde{\lambda}_i}$$

where $\widetilde{\lambda}_i = \dfrac{\sum\limits_{j=1}^{N} \lambda_{ij} T_j z_{min}^{-\theta}}{\sum\limits_{j=1}^{N} T_j z_{min}^{-\theta}}$ is the share of the overall mass of suppliers which buyers from i have access to, on average and $\kappa_i = \dfrac{\sum_{j=1}^{N} \frac{\lambda_{ij} T_j}{\sum_{j=1}^{N} \lambda_{ij} T_j} (d_{ij} c_j)^{-\theta}}{\sum_{j=1}^{N} \frac{T_j}{\sum_{j=1}^{N} T_j} (d_{ij} c_j)^{-\theta}}$ is a measure of how search frictions distort the distribution of the nationality of sellers that the buyer meets, in comparison with a frictionless world in which this dispersion only depends on the geography of costs. From this, it comes that, on average, buyers in country i pay a lower price if competition is high (larger $\kappa_i \Upsilon_i$) and search frictions are low (greater $\widetilde{\lambda}_i$). Finally, the probability of being matched with

[17] This modelling amounts to assuming that every producer located in j has a probability $\lambda_{ij} d\mu_{ij}(p)$ to be drawn by buyer b_i. Due to the assumption of independent draws, the number of suppliers from j that buyer b_i meets follows a binomial law of parameters $(T_j z_{min}^{-\theta}, \lambda_{ij})$, that can be approximated by a Poisson law of parameter $\lambda_{ij} T_j z_{min}^{-\theta}$. The approximation rests on the convergence in law of the binomial distribution towards the Poisson law, when the number of trials goes to infinity while the product $T_j z_{min}^{-\theta} \lambda_{ij}$ remains constant. We follow Eaton et al. (2018) by directly assuming that the number of suppliers met follows a Poisson process.

a low cost supplier is also increasing in θ, conditional on Υ_i: A less heterogeneous distribution of prices implies that even unlucky buyers will end up paying a price which is not very far from the lowest cost supplier they would have been able to reach in the absence of search frictions.

3.2 Analytical predictions

In this section, we first derive predictions regarding the magnitude of bilateral trade flows between any two countries. Such predictions help understand how search frictions modify the benchmark frictionless model in Eaton and Kortum (2002). We then derive predictions regarding export probabilities along the distribution of firms' productivities, which are later used to identify search frictions in the data, separately from other barriers to trade.

3.2.1 Aggregate trade

Under the assumption of inelastic demand, the share of country j's consumption which is imported from country i, denoted π_{ij}, is the sum of unitary demands aggregated across all buyers which interact with a seller from j divided by aggregate consumption:

$$\pi_{ij} = \frac{\sum_{b_i=1}^{B_i} \mathbb{1}\{s(b_i) = j\}}{B_i} = \mathbb{E}_{b_i}[\mathbb{1}\{s(b_i) = j\}]$$

where $\mathbb{1}\{s(b_i) = j\}$ is a dummy variable which is equal to one if buyer b_i ultimately chooses to purchase the good from a supplier from j and $\mathbb{E}_{b_i}[.]$ is the expectation operator, defined across buyers from i. Under the assumptions of the model, $\{\mathbb{1}\{s(b_i) = j\}\}_1^{B_i}$ are random variables which are independent and identically distributed. Using the law of large numbers, π_{ij} is thus equal to the expected value of $\mathbb{1}\{s(b_i) = j\}$, across buyers in i. It is the probability that the lowest cost supplier encountered by any buyer from i is located in country j.

In order to derive analytical predictions regarding π_{ij}, two assumptions are important. First, the assumption that firms' productivity is drawn into a Pareto distribution which shape parameter is homogeneous across source countries implies that, at any price p, the share of firms from j in the distribution $\mu_i(p)$ is constant. Second, the assumption that draws in this distribution follow a Poisson process means that this property subsists in the frictional world

(Eaton et al., 2012). Using these two assumptions, one can show that:[18]

$$\pi_{ij} \quad = \quad \frac{\lambda_{ij}\mu_{ij}(p)}{\sum_{j=1}^{N}\lambda_{ij}\mu_{ij}(p)} = \frac{T_j(d_{ij}c_j)^{-\theta}}{\Upsilon_i\kappa_i}\frac{\lambda_{ij}}{\widetilde{\lambda}_i} \tag{1}$$

The share of products from country j in destination i's final consumption depends on i) the relative competitiveness of its firms in comparison with the rest of the world, $\frac{T_j(d_{ij}c_j)^{-\theta}}{\Upsilon_i\kappa_i}$, and ii) the relative size of search frictions its firms encounter while serving market i, $\frac{\lambda_{ij}}{\widetilde{\lambda}_i}$.[19] The first determinant is almost identical to the formula derived in Eaton and Kortum (2002), though they derive it for the aggregate economy exploiting the law of large numbers across imperfectly substitutable varieties rather than across buyers within a product. It shows how the combined impact of technology and geography determines international trade flows in a Ricardian world. The key insight from our model is that search frictions can distort trade flows, in comparison with this benchmark. This is what the second term in equation (1) captures. Taking the derivative of equation (1) with respect to λ_{ij} yields Proposition 1 regarding the sensitivity of aggregate trade to search frictions:

Proposition 1. The market share of a country always increases following a reduction in bilateral frictions:

$$\frac{d\ln\pi_{ij}}{d\lambda_{ij}} = \frac{1-\pi_{ij}}{\lambda_{ij}} > 0, \ \forall \ \lambda_{ij} \ \in \ [0,1]$$

See the Proof in Appendix A.2.

To recover the intuition surrounding this result, first note that

$$\frac{d\ln\pi_{ij}}{d\lambda_{ij}} = \underbrace{\frac{d\ln\lambda_{ij}}{d\lambda_{ij}}}_{\text{Visibility channel}} - \underbrace{\left[\frac{d\ln\kappa_i}{d\lambda_{ij}} + \frac{d\ln\widetilde{\lambda}_i}{d\lambda_{ij}}\right]}_{\text{Competition channel}}$$

[18] We conjecture that having an iso-elastic demand would not affect this result. The reason is that, exactly as in Eaton and Kortum (2002), the ex-post price at which buyers purchase their good follows a distribution that does not depend on the origin of sellers offering these goods. For this reason, trade shares only depend on the probability that a supplier from a given origin supplies the good, exactly as in Eaton and Kortum (2002), whether demand is inelastic or iso-elastic.

[19] In this formula, κ_i can be interpreted as the distortion that frictions induce on the destination's multilateral resistance index. To see why, notice that:

$$\Upsilon_i\kappa_i = \frac{\sum_{j=1}^{N}T_j}{\sum_{j=1}^{N}\lambda_{ij}T_j}\sum_{j=1}^{N}\lambda_{ij}T_j\left(d_{ij}c_j\right)^{-\theta}$$

This term can be interpreted as an "ex-post" multilateral resistance index measuring how input costs, geographic barriers *and* frictions distort country i's effective state of technology towards technologies emanating from closer, cheaper and less frictional countries. This "ex-post" multilateral index is equal to the "ex-ante" multilateral index described in Eaton and Kortum (2002) when frictions disappear (i.e. $\Upsilon_i\kappa_i \rightarrow \Upsilon_i$ when $\lambda_{ij} \rightarrow 1, \ \forall i$). κ_i can be below or above 1 depending on how frictions correlate with countries' competitiveness.

The impact of a reduction in bilateral search frictions on aggregate trade flows can be decomposed into two channels. First a "Visibility" channel that captures the direct impact of search frictions on the likelihood that any exporter from j meets with a buyer from i. Second, a "Competition" channel which affects the likelihood that any seller is chosen for serving a buyer, conditional on meeting with her. The first effect is positive since lower search frictions increase the probability that any supplier from j will be drawn by any buyer from i. As shown in Appendix A.2, the "Competition" channel plays in the opposite direction since less frictions increase the strength of competition between French firms, conditional on being drawn, thus reducing the likelihood that any seller from j is chosen by a buyer from i. However, the "Visibility" effect always dominates the "Competition" channel at the product level. The reason is that the strength of the competition channel induced by a change in bilateral frictions exerts over the subsample of French firms competing for foreign buyers. Because of this, a reduction in search frictions unambiguously increases the exporting country's share in the destination's absorption. This is in line with the argument in Rauch (1999) that search frictions can contribute to reducing the magnitude of bilateral trade between more distant countries, if correlated with (physical and cultural) distance between countries.

Finally, note that the model is compatible with structural gravity. Namely, log-linearizing equation (1) implies:

$$\ln \pi_{ij} = FE_i + FE_j - \theta \ln d_{ij} + \ln \lambda_{ij} \tag{2}$$

where $FE_i \equiv \ln \Upsilon_i \kappa_i \widetilde{\lambda}_i$ and $FE_j \equiv \ln T_j (c_j)^{-\theta}$. The cross-sectional variation in bilateral trade flows can be explained by a full set of origin- and destination-country fixed effects and a number of bilateral variables correlated with the magnitude of trade frictions. In comparison with standard gravity-compatible models, the difference is that our model predicts physical trade barriers as well as information frictions to enter the gravity equation. A corollary is that predictions on product-level trade cannot be expected to help identify search frictions, separately from other barriers to trade, since both sources of frictions have the same qualitative impact on trade. We now explain why this is no longer true when studying the model's predictions regarding firm-level trade.

3.2.2 Firm-to-firm matching

Having derived predictions regarding the magnitude of aggregate trade flows, we now study the matching process between any two firms. Such predictions are new to our model and can be confronted to firm-to-firm trade data. Because we observe the universe of French exporters, and their clients abroad, we will take the point-of-view of individual sellers and derive predictions regarding the expected number of clients they can reach, in each destination.

Consider first the probability that a given supplier from j, France in our data, serves a buyer in i. In our framework, this probability decomposes into the probability that s_j meets with b_i times the probability that it is the lowest cost supplier, within b_i's random set:[20]

$$\begin{aligned}\rho_{s_j i} &= \mathbb{P}\left(s_j \in \Omega_{b_i}\right)\mathbb{P}\left(\min_{s'_k \in \Omega_{b_i}}\left\{\frac{c_k d_{ik}}{z_{s'_k}}\right\} = s_j\right) \\ &= \lambda_{ij} e^{-(c_j d_{ij})^{\theta} z_{s_j}^{-\theta}\Upsilon_i \kappa_i \widetilde{\lambda}_i} \end{aligned} \tag{3}$$

Because of the Poisson assumption, the probability of being drawn by a buyer is constant and only depends on the size of search frictions. More productive sellers however have a higher probability to end up serving any buyer from i because, conditional on being drawn, they have a higher chance to be the lowest cost supplier. And conditional on productivity, a seller has a higher chance to serve a buyer located in a market which can be served at a low cost (d_{ij} close to one), where competition is limited (Υ_i low) and which displays important frictions, on average ($\tilde{\lambda}_i$ small). These predictions are consistent with evidence presented in Section 2.2.

One can verify that, under some parameter restrictions, an increase in the meeting probability has an ambiguous impact on the probability of a seller to be chosen by a particular buyer. This leads us to Proposition 2:

Proposition 2. The impact of search frictions varies along the distribution of productivities, with high-productivity firms benefiting more, in terms of export performances, from a reduction in search frictions:

$$\frac{\partial \ln \rho_{s_j i}}{\partial \lambda_{ij}} = \frac{1}{\lambda_{ij}} - z_{s_j}^{-\theta} T_j \quad \text{and} \quad \frac{\partial^2 \ln \rho_{s_j i}}{\partial \lambda_{ij} \partial z_{s_j}} > 0. \tag{4}$$

For low-enough search frictions, an increase in λ_{ij} affects firms at the bottom of the distribution negatively, i.e.:

$$\frac{\partial \ln \rho_{s_{min} i}}{\partial \lambda_{ij}} < 0 \quad \text{if} \quad \lambda_{ij} > \frac{1}{z_{min}^{-\theta} T_j} \tag{5}$$

where $\rho_{s_{min} i}$ is the export probability in i of a firm with productivity z_{min}. .

See the Proof in Appendix A.3.

The ambiguous impact of more bilateral search frictions (a lower meeting probability λ_{ij}) on the probability to serve a particular buyer conditional on the level of productivity again explains by the opposite impact of the visibility and competition channels:

[20]Since buyers are ex-ante homogeneous, the probability is the same for all buyers b_i located in country i.

$$\frac{\partial \ln \rho_{s_j i}}{\partial \lambda_{ij}} = \underbrace{\frac{\partial \ln \lambda_{ij}}{\partial \lambda_{ij}}}_{\text{Visibility channel}} - \underbrace{\frac{\partial \left(c_j d_{ij}\right)^{\theta} z_{s_j}^{-\theta} \kappa_i \Upsilon_i \widetilde{\lambda}_i}{\partial \lambda_{ij}}}_{\text{Competition channel}}$$

On the one hand, a decrease in search frictions through the "visibility" channel increases the likelihood that seller s_j will serve any buyer in country i as it enhances its probability to meet with the buyer. On the other hand, conditional on being drawn, less bilateral search frictions means than s_j faces fiercer competition from other domestic suppliers. This reduces the probability that it is the lowest-cost supplier met by any particular buyer, especially if the seller's productivity is low. For high-productivity sellers, the visibility channel dominates and they always benefit from a reduction in search frictions. For these firms, the main impediment to their export development is a lack of visibility in foreign markets. For low-productivity sellers instead, the competition channel is stronger which explains that their privately optimal value of the meeting probability, defined as the level of λ_{ij} which maximizes their export probability, is low. If frictions are not too strong, the competition channel dominates the visibility channel at the bottom of the productivity distribution and sufficiently low-productivity sellers benefit from more frictions.[21]

A corollary of Proposition 2 is that the export premium of high-productivity firms is affected by the level of frictions:

$$\ln \frac{\rho_{\bar{z}_j i}}{\rho_{\underline{z}_j i}} = \frac{\lambda_{ij}}{\pi_{ij}} T_j \left(\underline{z}_j^{-\theta} - \bar{z}_j^{-\theta} \right) \tag{6}$$

where $\rho_{\bar{z}_j i}$ and $\rho_{\underline{z}_j i}$ denote export probabilities in country i of a firm from j with a high-productivity $\bar{z}_j$ and a low productivity $\underline{z}_j$, respectively. Equation (6) is positive which reflects the fact that, everything else being equal, high productivity firms are more likely to serve any buyer in country i. However, it is decreasing in $\frac{\pi_{ij}}{\lambda_{ij}} = \frac{\pi_{ij}^{\lambda_{ij}=1}}{\kappa_i \widetilde{\lambda}_i}$, a measure of the distorsive impact of frictions in market i (where $\pi_{ij}^{\lambda_{ij}=1}$ denotes the market share of country j in market i, in the absence of frictions, i.e. when $\lambda_{ij} \to 1, \ \forall i$). In high frictional markets, buyers meet with a small number of sellers, on average. This reduces the strength of competition, thus the competitive advantage of high-productivity exporters. Because of this, the export premium of high-productivity firms is smaller. This feature of the model can help rationalize

[21] While the analytical results crucially rely on the size of the visibility channel being independent of firms' productivity, we argue in Appendix A.3 that the result is more general than this. In particular, we discuss the case in which the probability of a meeting is increasing in firms' productivity, as high-productivity firms are arguably less likely to suffer from a lack of visibility abroad. We argue that this should not overcome our result as long as the cross derivative of the meeting probability with respect to λ_{ij} and z_{s_j} is not too negative.

the randomness in small firms' export patterns. It is also what lets us identify search frictions in the data, separately from other barriers to international trade. Indeed, while the export premium of high-productivity firms is reduced in highly frictional countries, it is exhacerbated in countries featuring high iceberg trade costs, i.e. $\frac{d \ln \frac{\rho_{\bar{z}_j i}}{\rho_{\underline{z}_j i}}}{d \ln d_{ij}} > 0$. Using the heterogeneity in export performances across firms allows identifying search frictions, separately from iceberg trade costs.

Since all buyers play independently from each other, equation (3) immediately delivers an analytical expression for the expected number of buyers served in country i, conditional on the location and productivity of the seller. Namely, the expected number of clients in i of a seller s_j, which is also the expected value of exports, is:

$$\mathbb{E}[B_{s_j i}|z_{s_j} > z_{min}] = \lambda_{ij} e^{-(c_j d_{ij})^{\theta} z_{s_j}^{-\theta} \Upsilon_i \kappa_i \widetilde{\lambda}_i} B_i$$

where $B_{s_j i}$ denotes the number of buyers from i in s_j's portfolio of clients. Again, more productive sellers are expected to serve more buyers in each destination, a prediction which is consistent with evidence in Figure 2. In our framework, this relationship comes from more productive sellers being more likely to be chosen by any buyer. This differentiates us from Carballo et al. (2018) and Bernard et al. (2018b) who also rationalize the relationship between a firm's productivity and the number of buyers it serves in a destination, though with quite different arguments.[22]

4 Estimation

In this section, we first justify the moments used to estimate search frictions, independently from other barriers to international trade. We then describe the GMM estimator and its implementation, with details postponed to Appendix B. Finally, we discuss the results.

4.1 Moment choice

Results in Section 3.2.2 provide insights on the *expected* number of buyers in each destination. The randomness of the matching process however generates dispersion around this mean. To

[22] In Carballo et al. (2018), more productive exporters serve more consumers in each destination because they can produce and sell products further away from their "core segment", thus reaching a wider set of heterogeneous buyers. In Bernard et al. (2018b), the heterogeneity comes from more productive exporters being able to serve a larger range of less productive buyers in presence of match-specific fixed costs. Both papers need to introduce another source of ex-ante heterogeneity, between buyers. We instead assume buyers to be ex-ante homogeneous and attribute all the ex-post heterogeneity to the random meeting process.

confront the model with the data, we thus derive the probability that seller s_j has *exactly* M buyers in country i, conditional on its productivity. Given the independence of draws, one can show that it follows a binomial law of parameters B_i and $\rho_{s_j i}$:

$$\mathbb{P}(B_{s_j i} = M | z_{s_j} > z_{min}) = C_{B_i}^M \rho_{s_j i}^M (1 - \rho_{s_j i})^{B_i - M}$$

Integrating over the distribution of productivities gives the expected measure of firms from j with exactly $M > 0$ buyers in i (See details in Appendix A.4)

$$h_{ij}(M) = \frac{\pi_{ij}}{\lambda_{ij}} \frac{1}{M} I_{\lambda_{ij}}(M, B_i - M + 1) \tag{7}$$

where $I_a(b, c) = \frac{B(a;b,c)}{B(b,c)}$ denotes the regularized incomplete beta function.

Equation (7) shows that the mass of firms serving a given number of clients is decreasing in M, which is consistent with evidence in Section 2.2. In our model, this comes from the independence of matches: The probability that a given seller is drawn by a large number of buyers shrinks rapidly when the number of buyers increases. The shape of $h_{ij}(M)$ is also a function of λ_{ij}. Conditional on π_{ij} and B_i, one can use the predicted value for $h_{ij}(M)$ and its counterpart in the data to recover a structural estimate for λ_{ij}, for each product and destination.[23]

Once normalized by the total measure of firms in the market ($T_j z_{min}^{-\theta}$) to recover a convergent moment, equation (7) can be used to estimate search frictions. We decided not to use this exact moment, though, because of its empirical sensitivity to distance, that potentially reflects the impact of other physical trade barriers on the firm-level stock of partners within a destination. This sensitivity is illustrated in Table 2 which shows the correlation between various transformations of the empirical moment, and distance from France, used as a proxy for iceberg trade costs.[24] The correlation between the number of firms with exactly M buyers in a destination and the distance to this destination is negative and strongly significant. This is consistent with evidence in Section 2.2 that French sellers tend to serve less partners, if any at all, in more distant countries. This result should be expected from the model, since the π_{ij} component in equation (7) is negatively correlated with iceberg trade costs d_{ij} which can reasonably be assumed to be increasing in distance. In principle, this can be controlled for using readily available data for those trade shares.

[23] Since our dataset only covers exporters located in France, the j country will always be France and we will use the heterogeneity across destinations and sectors to recover a distribution of estimated parameters.

[24] For practical reasons detailed below, we restrict our attention to four values for $h_{ij}(M)$, corresponding to the bottom of the distribution of sellers' degrees.

Table 2: ***Correlation between various empirical moments and distance from France***

	log Distance	Std Dev.	Adjusted R-squared
Dependent Variable			
# sellers with:			
1 buyer	-15.92***	(1.48)	.698
2 buyers	-5.82***	(.556)	.536
3 buyers	-3.21***	(.362)	.417
4 buyers	-1.98***	(.252)	.335
# sellers (in relative terms with respect to the sellers with 1 buyer) with:			
2 buyers	.020***	(.009)	.339
3-4 buyers	-.026***	(.008)	.373
5+ buyers	-.123***	(.021)	.412
Variance of the relative shares of sellers:			
across M	.002	(.010)	.211
across M, controlling for migrants	-.016	(.014)	.211
	coef. on migrants: -.008** (.003)		

Notes: Robust standard errors, clustered at the country level, in parentheses with ***, ** and * respectively denoting significance at the 1, 5 and 10% levels. The last regression uses as right-hand side variables the (log of) distance from France and the stock of migrants.

What is not explained by the model is that the correlation with distance survives when these shares are further normalized by the destination-specific proportion of sellers with one buyer, i.e. when the empirical counterpart of $\frac{h_{ij}(M)}{h_{ij}(1)}$ is correlated with distance, as in the second panel of Table 2. In principle, the normalization shall neutralize the impact of trade shares, thus of iceberg trade costs. While the correlation may be explained by search frictions being correlated with distance, there might also be other channels through which iceberg trade costs affect the ratios, which the model does not encompass but the data reveal. To avoid that this pollutes our estimates of search frictions, we decided to use an alternative moment which is not affected by distance to France and is thus more likely to help us extract from the data information on pure search frictions.

The moment chosen exploits information on the *dispersion* in the number of buyers served by sellers serving the same destination with the same product. Namely, the theoretical moment is defined as the variance in the $\frac{h_{ij}(M)}{h_{ij}(1)}$ ratios:

$$Var_{ij}\left(\lambda_{ij}\right) = \frac{1}{B_i - 1}\sum_{M=2}^{B_i}\left(\frac{h_{ij}(M)}{h_{ij}(1)} - \frac{1}{B_i - 1}\sum_{M=2}^{B_i}\frac{h_{ij}(M)}{h_{ij}(1)}\right)^2 \tag{8}$$

This moment is related to the curvature of the distribution of sellers' number of partners represented in Figure 2 (left panel) and is positively correlated with λ_{ij}. Intuitively, less frictions reduce the mass of exporters serving a small number of buyers while increasing the density at high values of M. This tends to increase the variance in equation (8). We use this property to identify search frictions.

In theory, the dispersion can be calculated across $B_i - 1$ ratios. However, these ratios do not convey a lot of relevant information since they are almost all equal to zero above a certain level of M.[25] For this reason, we decided to restrict our attention to the variance computed over three empirically relevant $\frac{h_{ij}(M)}{h_{ij}(1)}$ ratios, namely $M = \{2, [3, 4], [5, B_i]\}$, $M = \{2, 3, [4, B_i]\}$ or $M = \{[2, 3], [4, 5], [6, B_i]\}$ depending on the product and destination. As documented in the last panel of Table 2, the empirical counterpart of the moment in equation (8) is not correlated with distance from France. However, it is negatively correlated with the stock of French migrants in the destination, our proxy for information frictions. This moment is thus a good candidate for estimating search frictions independently from other physical trade barriers. The intuition behind this empirical finding can directly be drawn from the model. Physical trade barriers affect French exporters in a homogeneous way along the whole distribution. They

[25] As shown in Figure 2 (left panel), most of the variance in the number of buyers served by French exporters is indeed found at values for $B_{s_j i}$ below 10 and thus using all the individual moments regarding the number of firms with $B_{s_j i} > 10$ clients would be inefficient. Moreover, this would artificially reduce the dispersion in the data, in a way that is not independent from B_i.

reduce their relative competitiveness with respect to foreign competitors. However, search frictions do not impact small and large exporters in the same way, as discussed in Proposition 2. Small exporters "benefit" from high search frictions because they reduce competition exerted by other French firms on these low-competitive sellers. Instead, large exporters suffer from the lack of visibility induced by large enough frictions. This non-monotonic relationship is the key reason why the dispersion in sellers' degrees is informative on the size of search frictions in our framework.

4.2 Estimation strategy

We estimate search frictions with a Generalized Method of Moments. As just explained, we focus on the theoretical moment defined in equation (8) which conditional on B_i solely depends on λ_{ij}. The empirical counterpart of this theoretical moment is observed in our data:

$$\widehat{Var_{ij}} = Var\left(\frac{\sum_{s_j=1}^{S_j} \mathbb{1}\{B_{s_j i} = m_1\}}{\sum_{s_j=1}^{S_j} \mathbb{1}\{B_{s_j i} = 1\}}, \frac{\sum_{s_j=1}^{S_j} \mathbb{1}\{B_{s_j i} = m_2\}}{\sum_{s_j=1}^{S_j} \mathbb{1}\{B_{s_j i} = 1\}}, \frac{\sum_{s_j=1}^{S_j} \mathbb{1}\{B_{s_j i} = m_3\}}{\sum_{s_j=1}^{S_j} \mathbb{1}\{B_{s_j i} = 1\}}\right) \tag{9}$$

where $\mathbb{1}\{B_{s_j i} = M\}$ is an observed dummy equal to one when firm s_j has exactly M buyers in destination i and m_1, m_2 and m_3 respectively denote the first, second and third element of $M = \{2, [3, 4], [5, B_i]\}$, $M = \{2, 3, [4, B_i]\}$ or $M = \{[2, 3], [4, 5], [6, B_i]\}$.

As explained in Appendix B.1, the following convergence result applies:

$$\sqrt{S_j}\left(\widehat{Var_{ij}} - Var_{ij}(\lambda_{ij})\right) \xrightarrow[S_j \to +\infty]{\mathcal{D}} \mathcal{N}(0, \Omega_{ij}(\lambda_{ij})) \tag{10}$$

where $\Omega_{ij}(\lambda_{ij})$ is the variance of $\widehat{Var_{ij}}$.[26]

Using the convergence result, it is possible to identify λ_{ij} uniquely. Indeed, $Var_{ij}(\lambda_{ij}) - \widehat{Var_{ij}} = 0$ has a unique solution on $[0, 1]$. The minimization program writes as follows:

$$\min_{\lambda_{ij}} [Var_{ij}(\lambda_{ij}) - \widehat{Var_{ij}}]' \Omega_{ij}^{-1}(\lambda_{ij}) [Var_{ij}(\lambda_{ij}) - \widehat{Var_{ij}}] \tag{11}$$

Note that $\Omega_{ij}(\lambda_{ij})$ is the optimal matrix of weights as defined in Appendix B.1. Moreover, with an Asymptotic Least Squares estimation methodology, the estimated variance of estimated frictions writes:

$$\widehat{\Sigma}_{\lambda_{ij}} = \left[\frac{\partial Var_{ij}(\widehat{\lambda_{ij}})}{\partial \lambda_{ij}}' \Omega_{ij}^{-1}(\hat{\lambda_{ij}}) \frac{\partial Var_{ij}(\widehat{\lambda_{ij}})}{\partial \lambda_{ij}}\right]^{-1}$$

[26] $\Omega_{ij} = \nabla g(\lambda_{ij}) \Sigma_{ij} \nabla' g(\lambda_{ij})$ where g is the variance function and Σ_{ij} is the variance-covariance matrix of the random variables $\mathbb{1}\{B_{s_j i} = M\}$ for $M = m_1, m_2, m_3$.

In the rest of the analysis, we focus on sellers from one single country, $j = France$ and buyers from each European country. Search frictions are estimated independently for each product and destination. With a targeted moment which has an analytical formula, the implementation is straightforward. The only practical difficulty concerns the measurement of S_j and B_i in the data. Indeed, the theoretical moment in (8) is a function of λ_{ij} and B_i such that we need to measure the population of buyers in each destination country and sector. Moreover, the total number S_j of potential suppliers of a $hs6$ reference is needed to compute both the optimal weights entering the objective function and the asymptotic variance of the estimator (see details in Appendix B.1).

We recover measures of the population of buyers in each destination country and sector using predictions of the model regarding trade shares. Under the assumptions of the model, π_{ij} is both the share of goods from j in country i's total consumption and the ratio of the number of buyers from i buying their consumption from a seller in j divided by the total number of buyers in i ($\pi_{ij} = B_{ij}/B_i$). π_{ij} can easily be recovered from sectoral bilateral trade and absorption data.[27] B_{ij} is directly observed into our data. Based on this, one can recover a value of B_i for each destination and sector.[28]

Information on the number of *potential* suppliers by $hs6$ product is not available in any administrative dataset. To proxy S_j for each product, we exploit information on the universe of French firms recovered from the INSEE-Ficus database and the sector of activity they belong to. All firms belonging to a sector in which at least one firm makes 10% of its exports in a product are considered potential suppliers of the product. Atalay et al. (2014) use a comparable strategy to proxy the number of firms susceptible of purchasing a firm's output.

Using information on the number of potential sellers and buyers in each country and destination plus the information on the number of buyers in each seller's portfolio, one can directly solve the program in (11) and recover estimated values for the meeting probabilities. Since the minimization program in (11) is somewhat sensitive to the initial value, we use a grid search algorithm over 200 values of λ_{ij} to select a country and product-specific starting point.

[27] We use bilateral trade flows from the CEPII-BACI database (Gaulier and Zignago, 2010) and production data from Prodcom. π_{ij} is defined as the ratio of trade from j to i over absorption in country i.

[28] In sectors and countries in which the market share of French firms is very low, our empirical strategy implies very high values for B_i, above a million firms. Such high values might artificially bias our estimation of λ_{ij} down. To avoid this, we winsorized the number of potential buyers at 20,000, i.e. $B_i = min\left\{20,000; \frac{B_{iF}}{\pi_{iF}}\right\}$.

4.3 Results

Summary statistics. Search frictions are estimated at the (sector×country) level for a total of 9,860 λ_{ij} parameters, among which 9,838 are significant at the 5% level. To get meaningful comparisons, we restrict our analysis to countries where we have at least 200 estimated parameters. With this restriction, we keep 9,304 λ_{ij} parameters covering 15 countries.

Table 3, first column, provides summary statistics on the estimated parameters. Remember that, in the model, the λ_{ij} coefficient is defined as the share of the (continuum of) sellers from country i a given buyer in country j would meet, on average. We see an important level of dispersion in these probabilities. Indeed, ten percent of product-country pairs have a meeting probability below .01%; while 10 percent of the pairs have a meeting probability above 2.3%. A basic variance decomposition exercise shows that 13% of the dispersion in our friction parameters is driven by the destination country dimension, 43% is sector-specific, and the remaining 55% is within a sector×country.

Table 3: ***Summary statistics on estimated coefficients***

	Meeting Probability λ_{ij} (en %)	Probability of Meeting 0 Buyer $(1-\lambda_{ij})^{B_i}$ (en %)
Mean	1.04	12
Percentile 10	.01	.00
Percentile 25	.07	.00
Percentile 50	.28	.02
Percentile 75	.91	4.21
Percentile 90	2.33	5.58
# Observations	9,304	9,304

Notes: The first column in this table presents summary statistics on the λ_{ij} coefficients, estimated by country × hs6 product. The second column summarizes the subsequent probabilities that a French exporter meets with no buyer in the destination computed as $(1-\lambda_{ij})^{B_i}$ for each country and product.

In Table 4, we examine how the estimates relate to different country and product characteristics. Columns (1) and (2) focus on country characteristics, controlling for HS6-product fixed effects. In column (3), we remove the product fixed effects to include a measure of prod-

Table 4: ***Correlates of bilateral search frictions***

	(1)	(2)	(3)	(4)
		Dep. Var: $\ln(\lambda_{ij})$		
ln distance		-0.525*	-0.711***	
		(.255)	(.229)	
ln population	-0.494***	-0.523***	-0.258***	
	(.068)	(.057)	(.053)	
French migrants	0.220**	0.112	0.263***	
	(.090)	(.074)	(.061)	
Rauch dif.			-0.214***	-0.211***
			(.055)	(.053)
Fixed Effects	Product	Product	No	Country
Observations	9,304	9,304	8,705	8,705
R-squared	0.557	0.561	0.118	0.136

Notes: Robust standard errors, clustered at the country level, in parentheses with ***, ** and * respectively denoting significance at the 1, 5 and 10% levels.

uct differentiation. In column (4), we focus on the role of product characteristics and thus control for country fixed effects. The results show that market size (measured by population) and physical distance are positively correlated with frictions. The positive correlation between frictions and market size suggests that the search process is easier when economic activity is spatially concentrated. This might be explained by spillovers in the search process. While, search frictions are higher in large markets, the probability to meet a buyer increases with market size as there are more buyers in large markets. It is worth noting that distance has a negative impact on frictions, even though the moments used to estimate frictions are not correlated with distance (see the last panel in Table 2). The impact of distance on trade flows is often associated to transportation costs, our findings show that distance further impacts trade flows by impeding the search process between buyers and foreign sellers. As expected, search frictions are found lower in countries where French migrants are more numerous (though the effect is not always significant at conventional levels). This is consistent with the view that migrants convey information on their origin country, thus reducing information frictions. Finally, the results show that search frictions are higher for more differentiated products (according to the Rauch classification). This is consistent with the view that the search process is easier for products traded in organized markets.

While the λ_{ij} parameters are interesting to investigate, they are not easy to interpret.

More easy to interpret is the implication of these estimates in terms of the probability that a given French exporter meets with zero buyer in each destination, which is positively linked with the extent of frictions. Since the meeting process is a binomial, this probability is equal to $(1-\lambda_{ij})^{B_i}$, with B_i the number of consumers in country i. The distribution of probabilities over all country and hs6 product pairs is summarized in the second column of Table 3. On average, the probability of meeting with zero buyer in a destination is 12%. This number however hides a lot of heterogeneity. In 25% of country and sector pairs, the probability is below one percent. At the other side of the distribution, 10% of country×sector pairs display high frictions, with French exporters having almost 6% of chances of meeting with no buyer there. Figure 3 compares these probabilities, on average across destinations.[29] Belgium and Luxembourg, two countries contiguous to France with a high share of French speakers, are found to display low levels of search frictions for French sellers, on average. At the other side of the distribution, no match probabilities are found the largest, on average, in Greece, Finland, and Poland, three countries which are relatively distant from France along several metrics.

Another way of assessing the validity of our estimates is to confront the model's predictions to the data. Proposition 1 unambiguously shows that an increase in bilateral search frictions within a product category between France and a trade partner should lead to a reduction in French market shares. We thus regress the logarithm of French market shares (computed by destination-HS6 product pair) on our estimates of search frictions. We further control for other trade barriers, namely the share of French migrants and bilateral distance between France and the destination country. We also include product-fixed effects in all specifications to capture differences in French comparative advantages across product categories.

The results are presented in Table 5. Because we focus here on the subsample of products and destinations for which frictions are estimated, Column (1) first shows how market shares in this sample correlate with distance, the share of French migrants in the destination, and the share of migrants from the destination country in France. As expected, bilateral distance is an impediment to French exports while migrant networks foster bilateral trade. In column (2), we include our estimates of bilateral search frictions. In column (3), we include only the bilateral search frictions. Finally, column (4) controls for the probability of no meeting instead of the raw measure of search frictions. The results in columns (2) to (4) show that

[29] As the probability of no match has a product dimension, we measure the country-specific probability of no match by regressing this probability on product and country fixed effects. The product fixed-effects control for sectoral composition effects. The country-fixed effects allow us to compare the probability of no match across countries. One cannot estimate all the fixed effects and thus choose to present this measure in relative terms with respect to Germany.

Table 5: ***Search frictions and French market shares***

	(1)	(2)	(3)	(4)
	Dep. Var: ln French Market Share			
ln distance	-0.194*	-0.233**		-0.180*
	(.104)	(.116)		(.104)
French migrants	0.070***	0.052***		0.070***
	(.009)	(.010)		(.009)
Migrants in France	0.064***	0.075***		0.065***
	(.007)	(.008)		(.007)
ln Meeting proba		0.115***	0.130***	
		(.013)	(.016)	
Proba no meeting				-0.211***
				(.064)
Fixed Effects	Product	Product	Product	Product
Observations	8,778	8,778	8,778	8,778
R-squared	0.624	0.638	0.557	0.625

Notes: The dependent variable is the log of France market share in the destination, by product (π_{ij} using the model's notations). French migrants is the share of French migrants in the destination while migrants in France is the share of migrants from country i in France. Meeting proba is the estimated coefficient λ_{ij}. Proba no meeting is the probability that a French exporter does not meet any buyer in the destination country. Is is computed as $(1-\lambda_{ij})^{B_i}$ where B_i is the number of buyers in country i. Robust standard errors, clustered at the country level, in parentheses with ***, ** and * respectively denoting significance at the 1, 5 and 10% levels.

Figure 3: ***Comparison of frictions faced by French exporters across countries***

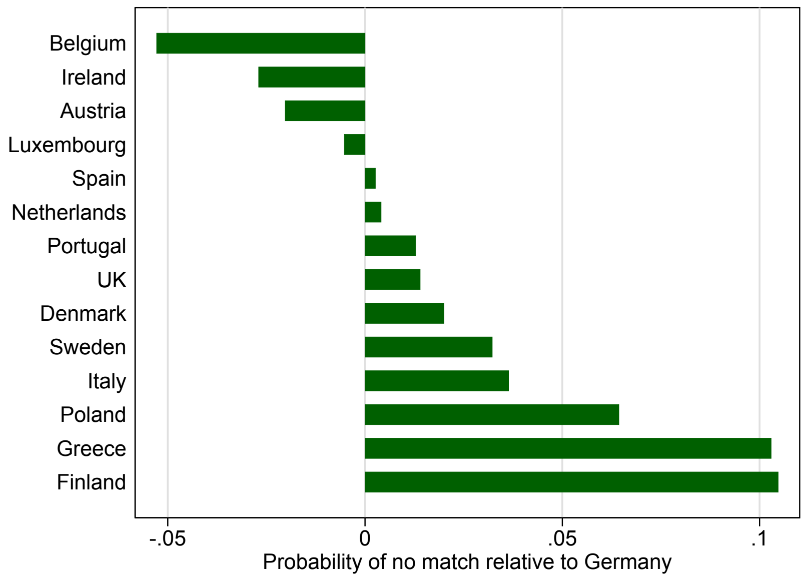

Notes: Mean probability of meeting with zero buyer across countries, in relative terms with respect to Germany.

French market shares are lower for product-destination pairs that exhibit a higher level of search frictions. This is consistent with Proposition 1. Alone, search frictions can explain as much as 55.7% of the variance in market shares across destinations within a product. This is sizable.

Model Fit. Having shown that our estimates of search frictions correlate with observables in a theory-consistent way, we now evaluate the ability of the model to reproduce key features of the data. We use our parameter estimates to simulate the mass of sellers interacting with zero to ten buyers within a destination market. Based on this, it is possible to predict the cumulated distribution of sellers' number of buyers in a destination, and compare it with the data.[30] Figure 4 reports the observed and predicted CDFs for the 15 countries in our sample. A visual inspection shows that the model nearly matches the distribution in most destinations. The parameters are estimated from the dispersion in the stock of buyers across French sellers serving the same destination. For reasons detailed in section 4.1, we do not consider the mass of sellers serving one client in our set of moments. Interestingly, our simple

[30]More precisely, we use the estimated λ_{ij} coefficients to predict the share of exporters serving a given number of buyers, in each destination and product. These shares are then aggregated across products using information on the relative number of suppliers of each product in France.

Figure 4: ***Model fit: Distribution of sellers' degrees***

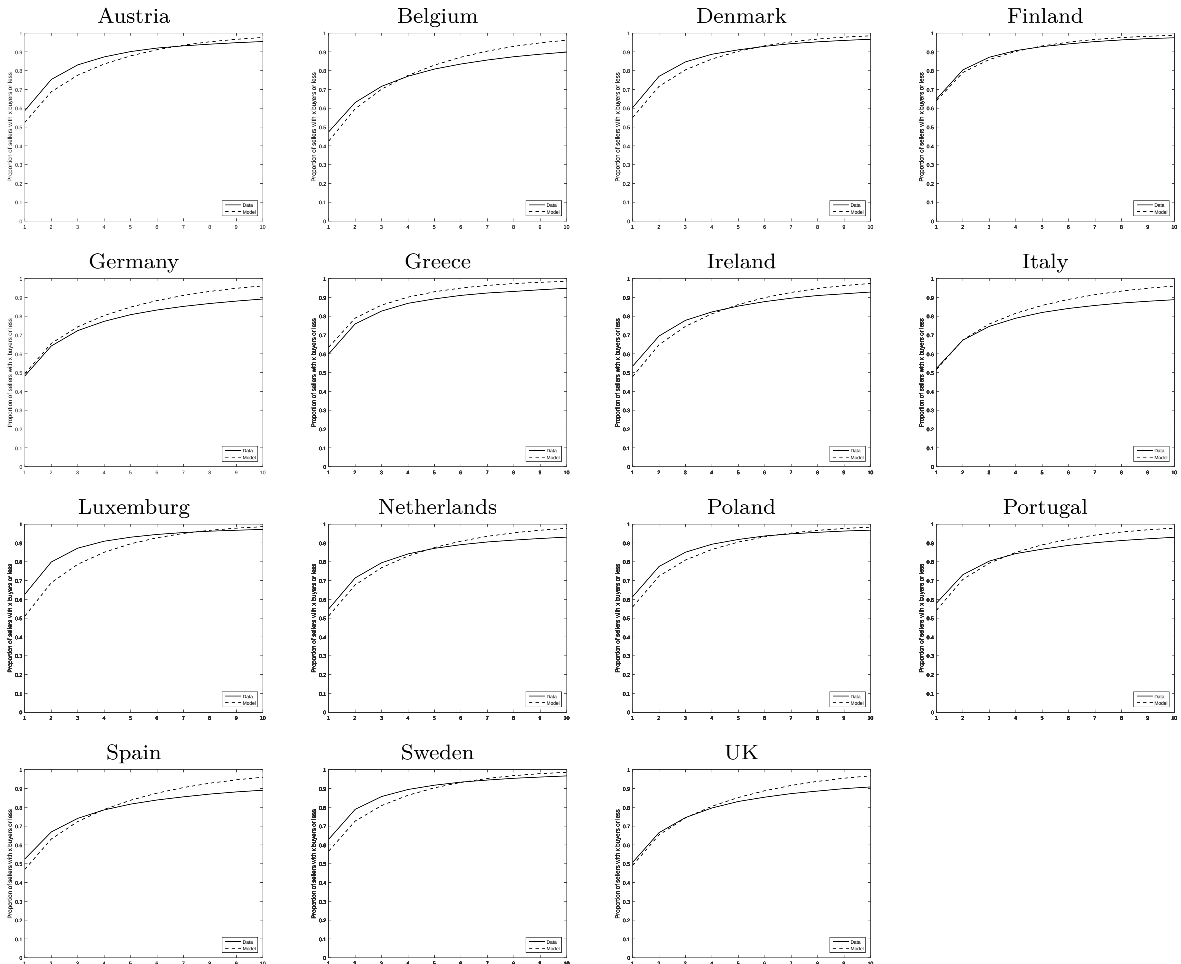

Notes: Observed and predicted CDF of sellers' numbers of buyers, by country. Predicted CDF are obtained using the model's definition of $h_{ij}(M)$, at the country and product level, before aggregating across products using information on the relative number of producers of each good in France.

Table 6: ***Model fit: Share of one-buyer sellers***

	Dep.Var.: Empirical share of one buyer		
	(1)	(2)	(3)
Predicted share	.286***	.265***	.163***
	(.006)	(.006)	(.006)
Constant	.391***		
	(.003)		
# obs	9,304	9,304	8,960
Fixed Effects	No	Country	Country Product
R-squared	.186	.230	.478

Notes: The predicted share of sellers with one buyer is calculated as $h_{ij}(1)/\sum_{M=1}^{B_i} h_{ij}(M)$. Robust standard errors in parentheses with *** denoting significance at the 1% level.

model captures quite well the share of sellers serving a single buyer within a destination, i.e. the fit is good regarding the curvature of CDFs *and* their intercept.[31] While the first moment is targeted in our estimation, the second is not.

The ability of the model to match the share of sellers serving a single buyer is further evaluated in Table 6. Instead of aggregating across products within countries, we predict the share of sellers serving one buyer for each product-country pair where we have estimates of frictions. Table 6 reports the correlation between the observed and predicted shares. In the first column, we report the unconditional correlation. In column (2), country fixed effects are introduced while column (3) has country and product fixed effects. The R^2 of the first regression is .19, suggesting that our simple model accounts for one-fifth of the dispersion in the share of sellers serving a single buyer. The correlation is highly significant in the three specifications which shows that the correlation is valid within countries across products as well as across products within countries.[32]

[31]One country for which we underestimate the share of sellers having a single buyer is Luxembourg. A possible reason for this poor performance is that the market share of French firms in Luxembourg is somewhat mismeasured due to bilateral trade data in BACI recording exports towards Belgium and Luxembourg together.

[32]We have run similar regressions considering the share of sellers with two buyers and with three buyers. The fit between the predicted and observed shares is very comparable.

5 Counterfactual Results

Having shown that our methodology delivers convincing estimates of bilateral search frictions faced by French exporters, we now use them to run a counterfactual analysis. The exercise is meant to quantify the extent to which search frictions contribute to explaining the randomness in export behaviors discussed in the introduction. We also provide estimates of the extent of the distortion induced by this particular form of barriers to trade.

5.1 Methodology

Throughout the exercise, we focus on the Greek market, identified as the second most frictional country in our data, on average.[33] Using this benchmark, we simulate how French exporters' behaviors would adjust would the level of bilateral frictions decrease in this destination, to the mean observed in the less frictional country in our sample, which is Belgium. In practice, this means that we compute expected export behaviors, in each product and in the aggregate, under the actual (estimated) search parameter ($\hat{\lambda}_{ij}$) and in a counterfactual in which the product-specific parameter is shifted up by the average difference in estimated frictions between Belgium and Greece (i.e for $\lambda^c_{ij} = \hat{\lambda}_{ij} \times 5.85$ where 5.85 is the mean difference in search frictions estimated for Belgium and Greece, conditional on product characteristics).

The distorsive impact of frictions is emphasized by comparing the impact of the counterfactual at various points of the productivity distribution. Using equations (1) and (3), the probability of serving a buyer in country i, conditional on a level of productivity z, writes:

$$\rho_{ij}(z) = \lambda_{ij} e^{-\frac{\lambda_{ij}}{\pi_{ij}} T_j z_{min}^{-\theta} \left(\frac{z}{z_{min}}\right)^{-\theta}} \tag{12}$$

Under Pareto, $\left(\frac{z}{z_{min}}\right)^{-\theta}$ is the share of firms with productivity above z, and is straightforward to calculate if z is interpreted as a particular percentile of the productivity distribution. The estimated value of λ_{ij} is taken as benchmark, and shifted up in the counterfactual state of the economy. Likewise, the trade share π_{ij} is observed in the benchmark and can be recovered in the counterfactual equilibrium using the formula in Proposition 1. The only unobservable component in this expression is thus $T_j z_{min}^{-\theta}$ which stands for the overall mass of potential suppliers in country j (France in our experiment). We decided to calibrate this object so as

[33] As shown in Figure 3, the mean probability of meeting with zero buyer is slightly larger for Finland than for Greece. We however decided to take the Greek market as reference because the cross-section of product-specific estimates is larger.

to fit the data regarding large firms' export premium in any given product market:[34]

$$\ln \frac{\rho_{\bar{z}_j i}}{\rho_{\underline{z}_j i}} = \frac{\lambda_{ij}}{\pi_{ij}} T_j z_{min}^{-\theta} \left[\left(\frac{\underline{z}_j}{z_{min}} \right)^{-\theta} - \left(\frac{\bar{z}_j}{z_{min}} \right)^{-\theta} \right]$$

Given observed λ_{ij} and π_{ij}, it is possible to calibrate $T_j z_{min}^{-\theta}$ so as to fit observed export premiums at different points of the productivity distribution. In practice, we use data on the apparent labor productivity of French firms, by sector, to assign each exporter to a productivity percentile. For each product and destination, we then compute the ratio of mean exports among firms below the 50th percentile in their sector and among firms above the 80th. The ratio of the later over the former is our measure of the product- and destination-specific export premium.[35] It is used to recover a calibrated value of $T_j z_{min}^{-\theta}$, for each product and destination. Consistent with the model, this object is assumed invariant to the counterfactual shift in search frictions.

Armed with the calibrated mass of firms in each sector and destination, the observed trade shares and the estimated search frictions, one can recover an estimate of $\rho_{ij}(z)$ for each percentile of the (Pareto) productivity distribution, and from this estimate the probability of exporting $(1-(1-\rho_{ij}(z))^{B_i})$ and the mean value of exports $(B_j \rho_{ij}(z))$, for each percentile. We now present results, focusing on two counterfactual exercises that help quantify the distortive impact of search frictions.

5.2 Results

Figure 5, left panel, shows how the probability of a firm exporting to Greece evolves along the productivity distribution, in the data (solid line) and in the counterfactual (dotted line). As expected, exporting to Greece is increasingly likely when moving right to the productivity distribution. In the equilibrium calibrated to actual data, less than 7% of firms in the first percentile serve at least one Greek client against more than 70% among the 1% most

[34] In the context of our model the export premium of large firms is the same whether expressed in terms of their relative probability of serving a given buyer, in terms of their expected number of buyers or in terms of the expected value of their exports. In the data, we use export premia recovered from average exports at different points of the productivity distribution. Results are qualitatively the same if we use instead information on firms' number of partners.

[35] The export premium is undefined in about 15% of product×destination pairs, either because we do not observe any firm in one of the two quantiles of the distribution used as reference or, in rare instances, because the recovered export premium is negative, i.e. low-productivity firms are found to export more on average than high-productivity firms. For Greece, negative export premia are found in 12 hs6 products out of 404. When the export premium is computed based on the export probability (instead of the mean value of exports), the number of negative premia falls to 2 out of 404 products. Since the model is not consistent with a negative export premium, we have no choice but to discard the corresponding products from the counterfactual analysis.

productive firms. More interesting is the model's prediction regarding the impact of shifting search frictions down, to the average level observed in Belgium. In this counterfactual less frictional Greek market, the export probability decreases at the bottom of the distribution while increasing at the top, i.e. some low-productivity firms are evicted from the Greek market while higher-productivity firms enter. The selection of firms is actually strong since only the last quartile of the distribution benefits from the reduction in frictions in this experiment. As a consequence, the overall export probability falls from 29.9% to 28.8% but the mean productivity of exporters improves, by 10 to 20%.[36]

To further document the impact of less frictions on the allocation of resources, Figure 6 shows the export premium that firms at the 90th percentile enjoy in comparison with competitors at the 25th percentile, in the benchmark and the counterfactual equilibrium. In the benchmark equilibrium, the simulated export premium matches exactly what is observed in the data for each product. It is equal to 40 for the median product, i.e. firms at the 90th percentile in this product market export 40 times more than firms at the 25th percentile, in expectations. In the counterfactual equilibrium (y-axis), export premia are an order of magnitude larger, to reach 3,800 for the median product. The reason why the effect is massive is obviously because many firms in the 25th percentile no longer exports in the counterfactual, as can be seen in Figure 5, left panel. The impact of low-productivity firms being evicted from the Greek market is further amplified by the value of exports, conditional on exporting, which also raises for high-productivity firms, in comparison with less productive exporters (Figure 5, right panel). Interestingly, this reinforcing force is entirely driven by large firms increasing their market share, in the counterfactual equilibrium in comparison with the benchmark. At the bottom of the distribution, the expected number of buyers conditional on exporting does

[36] By definition, the mean productivity of exporters writes:

$$\mathbb{E}(Z|Export) = \frac{\int_{z_{min}}^{+\infty} z f(z)\mathbb{P}(Export|z)dz}{\int_{z_{min}}^{+\infty} f(z)\mathbb{P}(Export|z)dz}$$

where $f(z) = \frac{\theta z_{min}^{\theta}}{z^{\theta+1}}$ is the density of z and $\mathbb{P}(Export|z)$ is the probability of exporting conditionally on z. After some simplifications, the change in the productivity of exporters in the counterfactual state of the economy, in relative terms with the benchmark, becomes:

$$\frac{\mathbb{E}^c(Z|Export)}{\mathbb{E}(Z|Export)} = \left[\int_{z_{min}}^{+\infty} \frac{\left(\frac{z}{z_{min}}\right)^{-\theta}\mathbb{P}(Export|z)}{\int_{z_{min}}^{+\infty}\left(\frac{z}{z_{min}}\right)^{-\theta}\mathbb{P}(Export|z)dz}\frac{\mathbb{P}^c(Export|z)}{\mathbb{P}(Export|z)}dz\right]\frac{\int_{z_{min}}^{+\infty}\left(\frac{z}{z_{min}}\right)^{-\theta-1}\mathbb{P}^c(Export|z)dz}{\int_{z_{min}}^{+\infty}\left(\frac{z}{z_{min}}\right)^{-\theta-1}\mathbb{P}(Export|z)dz}$$

where the ^{c} superscript refers to the counterfactual state. After discretizing the productivity space in percentiles, this formula can be used, together with a calibrated value for θ, to recover the change in the mean productivity of exporters. For $\theta = 3$, the overall productivity improvement is found to be 17.25%, a value which is reduced to 9.9% for $\theta = 5$.

Figure 5: ***Probability of exporting to Greece and expected number of buyers conditional on export, along the productivity distribution: Actual versus counterfactual***

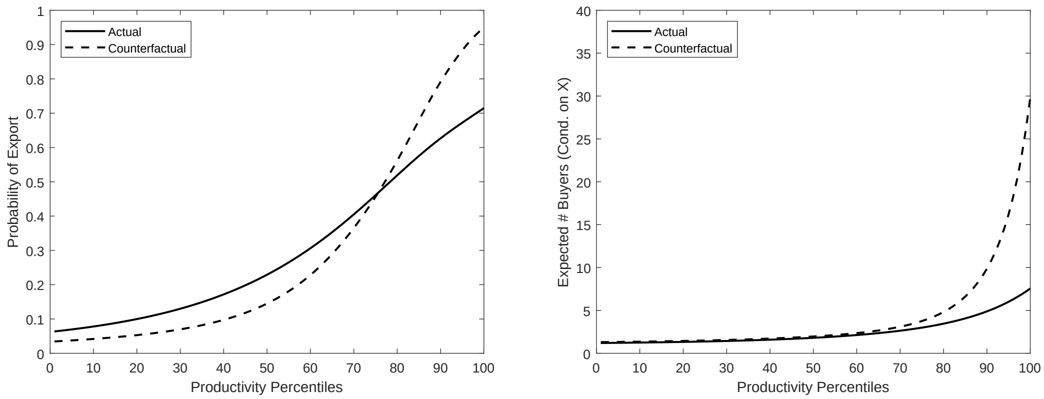

Notes: The graphs plot the probability of export to Greece (left panel) and the expected number of partners, conditional on exporting (right panel), conditional on the firm's position in the productivity distribution. The solid lines correspond to the actual equilibrium while the dotted lines are the counterfactual. Export probabilities and the expected number of exporters are both calculated at the product level following the strategy described in Section 5.1, before being aggregated across products using information on the relative number of firms in each product market.

not change much, and stays very close to one in many product markets and for firms in the first half of the distribution, roughly speaking. But the reduction in frictions has a strong positive effect on the expected number of clients for exporters at the top of the distribution. For the mean exporter at the 75th percentile of its sector's productivity distribution, the expected number of partners increases from 3.0 to 3.8. At the 90th percentile, the effect is more pronounced, with the expected number of clients shifting from 4.9 to 9.9. Finally, in the last percentile, the impact is substantial with the expected number of clients increasing from 7.6 to 29.8.

All in all, these results confirm the quantitatively important role of frictions. In comparison with standard barriers to international trade, they distort competition among potential exporters. This benefits, in relative terms, to low-productivity firms while reducing the export probability and expected exports at the top of the distribution.

6 Conclusion

This paper shows how search frictions in international good markets can distort competition between firms of heterogeneous productivity. We develop a Ricardian model of trade in which buyers in each market meet with a random subset of potential suppliers of a perfectly substitutable good. The model combines two barriers to international trade. Physical (iceberg)

Figure 6: ***Export premium (90th/25th percentiles) by product: Actual versus counterfactual***

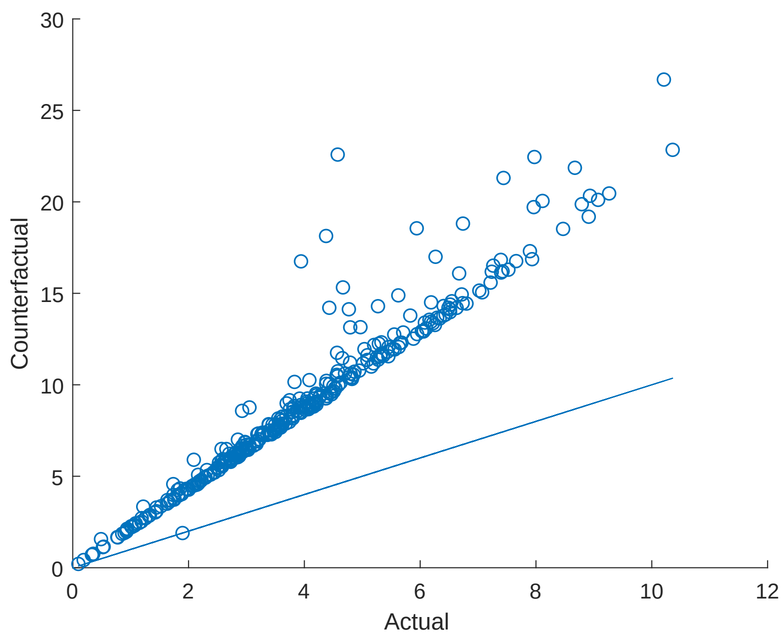

Notes: The graph plots the (log of the) export premium of firms at the 90th percentile in comparison with firms at the 25th percentile, by product, in the data (x-axis) and in the counterfactual (y-axis). The straight line is the 45-degree line.

trade costs reduce the competitiveness of exporters in foreign markets, in a way which is homogenous across firms. Instead, bilateral search frictions reduce the likelihood that any exporter will meet with a foreign consumer but also decrease competitive pressures, conditional on having met with a potential buyer. The relative strength of these two forces varies along the distribution of firms' productivity. While high-productivity firms always suffer from a lack of visibility in foreign markets, low-productivity firms can sometimes benefit from high search frictions because, conditional on having met with a buyer, these frictions reduce the strength of competition thus increasing the chances that the firm will be chosen to serve the buyer. This heterogeneity, we argue, is the key reason why search frictions can help explain the randomness in small and medium firms' export patterns that we observe in firm-level data. In highly frictional markets, the export premium of high-productivity firms is lowered and the export probability of small and medium firms increased.

Bilateral search frictions are estimated structurally using firm-to-firm trade data at the product and destination level. For each French firm and each product it sells, we can measure the number of clients it serves in a particular destination. In the model and in the data, heterogeneity across firms in this number is explained by firms' heterogeneous productivity and the magnitude of search frictions in this particular destination. Intuitively, more frictional markets induce more distorsions, which reduces the export premium of high-productivity firms. We use this property of the model to structurally recover a measure of search frictions, for each product and destination. Estimated frictions are found more severe in large and distant countries and for products that are more differentiated.

A counterfactual analysis allows quantifying the size of the distorsion induced by search frictions. When we simulate the impact of reducing the level of search frictions, in the most frictional country to the mean level observed in the least frictional one, we estimate substantial selection effects. Increasing the meeting probability between French sellers and Greek buyers on average pushes the least productive exporters out of the market while substantially increasing the export probability and the conditional value of exports for firms in the last quartile of the productivity distribution. Because of this, the mean productivity of exporters increases, by 10 to 20%, and their export premium is substantially increased.

The distortive impact of search frictions can rationalize a number of active policies used by export-promoting agencies. In a frictional world, any policy instrument that can help high-productivity firms that suffer from a lack of visibility abroad meet with foreign buyers induces aggregate productivity gains. Such policies may however hurt low-productivity exporters.

References

Akerman, A., Leuven, E., and Mogstad, M. (2018). Information Frictions, Internet and the Relationship between Distance and Trade. Memorandum 1/2018, Oslo University, Department of Economics.

Allen, T. (2014). Information Frictions in Trade. *Econometrica*, 82:2041–2083.

Arkolakis, C. (2010). Market Penetration Costs and the New Consumers Margin in International Trade. *Journal of Political Economy*, 118(6):1151 – 1199.

Atalay, E., Hortacsu, A., and Syverson, C. (2014). Vertical Integration and Input Flows. *American Economic Review*, 104(4):1120–1148.

Bergounhon, F., Lenoir, C., and Mejean, I. (2018). A guideline to French firm-level trade data.

Bernard, A. B. and Jensen, J. B. (1995). Exporters, Jobs, and Wages in U.S. Manufacturing: 1976-1987. *Brookings Papers on Economic Activity*, 26(1995 Micr):67–119.

Bernard, A. B., Jensen, J. B., Redding, S. J., and Schott, P. K. (2007). Firms in International Trade. *Journal of Economic Perspectives*, 21(3):105–130.

Bernard, A. B., Moxnes, A., and Saito, Y. U. (2018a). Production Networks, Geography and Firm Performance. *Journal of Political Economy*.

Bernard, A. B., Moxnes, A., and Ulltveit-Moe, K. H. (2018b). Two-sided heterogeneity and trade. *Review of Economics and Statistics*, 100(3):424–439.

Carballo, J., Ottaviano, G., and Volpe Martincus, C. (2018). The Buyer Margins of Firms' Exports. *Journal of International Economics*, 112(1):33–49.

Chaney, T. (2014). The Network Structure of International Trade. *American Economic Review*, 104(11):3600–34.

Dasgupta, K. and Mondria, J. (2018). Inattentive Importers. *Journal of International Economics*, 112(1):150–165.

Eaton, J. and Kortum, S. (2002). Technology, Geography, and Trade. *Econometrica*, 70(5):1741–1779.

Eaton, J., Kortum, S., and Kramarz, F. (2004). Dissecting Trade: Firms, Industries, and Export Destinations. *American Economic Review*, 94(2):150–154.

Eaton, J., Kortum, S., and Kramarz, F. (2011). An Anatomy of International Trade: Evidence From French Firms. *Econometrica*, 79(5):1453–1498.

Eaton, J., Kortum, S., and Kramarz, F. (2018). Firm-to-Firm Trade: Imports, Exports, and the Labor Market.

Eaton, J., Kortum, S. S., and Sotelo, S. (2012). *International trade: Linking micro and macro*, volume II of *Advances in Economics and Econometrics: Tenth World Congress*. Applied economics edition.

Fernandes, A. M. (2007). Trade policy, trade volumes and plant-level productivity in Colombian manufacturing industries. *Journal of International Economics*, 71(1):52–71.

Gaulier, G. and Zignago, S. (2010). BACI: International Trade Database at the Product-Level. The 1994-2007 Version. CEPII Working Papers 2010-23, CEPII.

Gouriéroux, C., Monfort, A., and Trognon, A. (1985). Moindres carrés asymptotiques. *Annales de l'INSEE*, 58:91–122.

Head, K. and Mayer, T. (2014). *Gravity Equations: Workhorse,Toolkit, and Cookbook*, volume 4 of *Handbook of International Economics*, chapter 0, pages 131–195. Elsevier.

Lendle, A., Olarreaga, M., Schropp, S., and Vézina, P.-L. (2016). There goes gravity: how eBay reduces trade costs. *The Economic Journal*, 126:406–441.

Mayer, T. and Ottaviano, G. (2008). The Happy Few: The Internationalisation of European Firms. *Intereconomics: Review of European Economic Policy*, 43(3):135–148.

Mayer, T. and Zignago, S. (2011). Notes on CEPII's distances measures : the GeoDist Database. CEPII Working Papers 2011-25, CEPII Working Paper.

McCallum, A. and Krolikowski, P. (2018). Goods-market Frictions and International Trade.

Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6):1695–1725.

Melitz, M. J. and Redding, S. J. (2015). New Trade Models, New Welfare Implications. *American Economic Review*, 105(3):1105–1146.

Pavcnik, N. (2002). Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants. *Review of Economic Studies*, 69(1):245–276.

Petropoulou, D. (2011). Information costs, networks and intermediation in international trade. Globalization Institute Working Papers 76, Federal Reserve Bank of Dallas.

Rauch, J. E. (1999). Networks versus markets in international trade. *Journal of International Economics*, 48(1):7–35.

Rauch, J. E. (2001). Business and Social Networks in International Trade. *Journal of Economic Literature*, 39(4):1177–1203.

Rauch, J. E. and Trindade, V. (2002). Ethnic chinese networks in international trade. *Review of Economics and Statistics*, 84(1):116–130.

Steinwender, C. (2018). Information frictions and the law of one price:" When the States and the Kingdom became United". *American Economic Review*, 108:657–696.

Topalova, P. and Khandelwal, A. (2011). Trade Liberalization and Firm Productivity: The Case of India. *The Review of Economics and Statistics*, 93(3):995–1009.

A Appendix: Proof of analytical results

A.1 Weibull distribution of minimum price

To derive the distribution of the minimum price drawn by a buyer in country i, start with the probability of paying a price above p, conditional on the number of price quotes in the buyer's random choiceset:

$$\begin{aligned}
\mathbb{P}\left[\underset{s_j\in\Omega_{b_i}}{\text{Min}}\left(\frac{c_j d_{ij}}{z_{s_j}}\right) > p \mid D_{b_i} = d\right] &= \prod_{s_j\in\Omega_{b_i}} \mathbb{P}\left[\left(\frac{c_j d_{ij}}{z_{s_j}}\right) > p \mid s_j \in \Omega_{b_i}\right] \\
&= \left[1 - \mathbb{P}\left(\left(\frac{c_j d_{ij}}{z_{s_j}}\right) < p\right)\right]^d
\end{aligned}$$

where D_{b_i} is the number of prices in buyer b_i's random choiceset Ω_{b_i}. Here, we use the fact that all price quotes are drawn independently by buyer b_i and are independent from each other under the assumption of marginal cost pricing. $\mathbb{P}\left(\left(\frac{c_j d_{ij}}{z_{s_j}}\right) < p\right)$ represents the probability that a randomly drawn price is cheaper than price p in country i. This implies:

$$\mathbb{P}\left[\underset{s_j\in\Omega_{b_i}}{\text{Min}}\left(\frac{c_j d_{ij}}{z_{s_j}}\right) > p \mid D_{b_i} = d\right] = \left[1 - \frac{p^\theta \sum_{j=1}^N \lambda_{ij} T_j (d_{ij} c_j)^{-\theta}}{\sum\limits_{j=1}^N \lambda_{ij} T_j z_{min}^{-\theta}}\right]^d$$

Integrating over all possible random numbers of price quotes gives the unconditional probability of paying a price above p:

$$\begin{aligned}
\mathbb{P}\left[\underset{s_j\in\Omega_{b_i}}{\text{Min}}\left(\frac{c_j d_{ij}}{z_{s_j}}\right) > p\right] &= \sum_{d=0}^{+\infty} \mathbb{P}\left(\underset{s_j\in\Omega_{b_i}}{\text{Min}}\left(\frac{c_j d_{ij}}{z_{s_j}}\right) > p | D_{b_i} = d\right) \mathbb{P}(D_{b_i} = d) \\
&= \sum_{d=0}^{+\infty} \left[1 - \frac{p^\theta \sum_{j=1}^N \lambda_{ij} T_j (d_{ij} c_j)^{-\theta}}{\sum\limits_{j=1}^N \lambda_{ij} T_j z_{min}^{-\theta}}\right]^d \left[\frac{\left(\sum\limits_{j=1}^N \lambda_{ij} T_j z_{min}^{-\theta}\right)^d e^{-\sum\limits_{j=1}^N \lambda_{ij} T_j z_{min}^{-\theta}}}{d!}\right] \\
&= e^{-p^\theta \Upsilon_i \kappa_i \widetilde{\lambda}_i}
\end{aligned}$$

where $\Upsilon_i = \sum_{j=1}^N T_j (d_{ij} c_j)^{-\theta}$ is the multilateral resistance index, $\widetilde{\lambda}_i = \frac{\sum\limits_{j=1}^N \lambda_{ij} T_j z_{min}^{-\theta}}{\sum\limits_{j=1}^N T_j z_{min}^{-\theta}}$ is the "mean" level of frictions in country i and $\kappa_i = \frac{\sum_{j=1}^N \frac{\lambda_{ij} T_j}{\sum_{j=1}^N \lambda_{ij} T_j} (d_{ij} c_j)^{-\theta}}{\sum_{j=1}^N \frac{T_j}{\sum_{j=1}^N T_j} (d_{ij} c_j)^{-\theta}}$ is a measure of how matching frictions distort the distribution of the nationality of sellers that the buyer meets, in comparison with a frictionless world in which this dispersion only depends on the geography of costs. The probability for the minimum price encountered to be below p is thus the exponential of the total measure of firms whose price is below p in country i times the proportion of those which will be encountered on average. Using $\widetilde{\lambda}_i \kappa_i = \frac{\sum_{j=1}^{N} \lambda_{ij} T_j (d_{ij} c_j)^{-\theta}}{\sum_{j=1}^{N} T_j (d_{ij} c_j)^{-\theta}}$, one can also interpret the probability in relative terms with the frictionless EK case in which all buyers would meet with the whole distribution of potential sellers. In comparison with EK, search frictions have a distorsive impact on prices which is all the stronger since they increase the relative probability for a buyer to meet with a high cost firm (i.e. if $Corr(\lambda_{ij}, T_j (d_{ij} c_j)^{-\theta}) < 0$).

Based on this, the distribution of the lower price encountered by a particular buyer b_i in country i has the following Weibull cumulated distribution function:

$$G_i(p) = 1 - e^{-p^{\theta} \Upsilon_i \kappa_i \widetilde{\lambda}_i}$$

A.2 Proof of proposition 1

Start with the definition of trade shares:

$$\pi_{ij} = \frac{\lambda_{ij}}{\kappa_i \widetilde{\lambda}_i} \frac{T_j (d_{ij} c_j)^{-\theta}}{\Upsilon_i}$$

implying:

$$\frac{d \ln \pi_{ij}}{d \lambda_{ij}} = \frac{1}{\lambda_{ij}} - \frac{1}{\kappa_i \widetilde{\lambda}_i} \frac{d \kappa_i \widetilde{\lambda}_i}{d \lambda_{ij}}$$

Using

$$\kappa_i \widetilde{\lambda}_i = \frac{\sum_{j=1}^{N} \frac{\lambda_{ij} T_j}{\sum_{j=1}^{N} \lambda_{ij} T_j} (d_{ij} c_j)^{-\theta}}{\sum_{j=1}^{N} \frac{T_j}{\sum_{j=1}^{N} T_j} (d_{ij} c_j)^{-\theta}} \frac{\sum\limits_{j=1}^{N} \lambda_{ij} T_j z_{min}^{-\theta}}{\sum\limits_{j=1}^{N} T_j z_{min}^{-\theta}} = \frac{\sum_{j=1}^{N} \lambda_{ij} T_j (d_{ij} c_j)^{-\theta}}{\sum_{j=1}^{N} T_j (d_{ij} c_j)^{-\theta}}$$

the derivative of $\kappa_i \widetilde{\lambda}_i$ with respect to λ_{ij} is just $\frac{T_j (d_{ij} c_j)^{-\theta}}{\sum_{j=1}^{N} T_j (d_{ij} c_j)^{-\theta}} = \frac{\pi_{ij} \kappa_i \widetilde{\lambda}_i}{\lambda_{ij}}$. This finally implies:

$$\frac{d \ln \pi_{ij}}{d \lambda_{ij}} = \frac{1 - \pi_{ij}}{\lambda_{ij}} > 0$$

A.3 Proof of proposition 2

The sensitivity of export probabilities to search frictions can be assessed through the following derivative:

$$
\begin{aligned}
\frac{\partial \ln \rho_{s_j i}}{\partial \lambda_{ij}} &= \underbrace{\frac{\partial \ln \lambda_{ij}}{\partial \lambda_{ij}}}_{\text{Visibility channel}} + \underbrace{\frac{\partial \ln e^{-(c_j d_{ij})^\theta z_{s_j}^{-\theta} \kappa_i \Upsilon_i \widetilde{\lambda}_i}}{\partial \lambda_{ij}}}_{\text{Competition channel}} \\
&= \frac{1}{\lambda_{ij}} - (d_{ij} c_j)^\theta z_{s_j}^{-\theta} \Upsilon_i \frac{d\kappa_i \widetilde{\lambda}_i}{d\lambda_{ij}} \\
&= \frac{1}{\lambda_{ij}} - z_{s_j}^{-\theta} T_j
\end{aligned}
$$

Depending on the current level of frictions (λ_{ij}), the overall mass of firms in country j ($T_j z_{min}^{-\theta}$) and the position of s_j in the productivity distribution ($\left(\frac{z_{s_j}}{z_{min}}\right)^{-\theta}$), the derivative can be positive or negative. It is all the more positive since z_{s_j} is high with, at the limit, $\lim_{z_{s_j} \to +\infty} \frac{\partial \ln \rho_{s_j i}}{\partial \lambda_{ij}} = \frac{1}{\lambda_{ij}}$. Instead, low-productivity sellers' export probability is less sensitive to frictions, and can even be negatively affected by a decrease in frictions. Namely, if the level of frictions is such that $\lambda_{ij} > \frac{1}{z_{min}^{-\theta} T_j}$, i.e. if frictions are not too strong, there is a strictly positive mass of firms which export probability decreases when search frictions are reduced: $\frac{\partial \ln \rho_{s_{min} i}}{\partial \lambda_{ij}} < 0$ where $\rho_{s_{min} i}$ denotes the export probability of the least productive firm.

The sensitivity of export probabilities to iceberg trade costs is instead unambiguously negative, less so for more productive sellers:

$$
\begin{aligned}
\frac{\partial \ln \rho_{s_j i}}{\partial d_{ij}} &= -(c_j d_{ij})^\theta z_{s_j}^{-\theta} \Upsilon_i \kappa_i \widetilde{\lambda}_i \left[\frac{\theta}{d_{ij}} + \frac{\partial \ln \Upsilon_i}{\partial d_{ij}} + \frac{\partial \ln(\kappa_i \widetilde{\lambda}_i)}{\partial d_{ij}} \right] \\
&= -\frac{\theta}{d_{ij}} (c_j d_{ij})^\theta z_{s_j}^{-\theta} \Upsilon_i \kappa_i \widetilde{\lambda}_i (1 - \pi_{ij}) < 0
\end{aligned}
$$

These contrasted results are the key reason why search frictions and iceberg costs can be identified separately in firm-level export patterns in this model. Larger iceberg trade costs decrease the probability of serving any buyer in the destination, less so for more productive sellers. In contrast, more search frictions are more costly for high-productivity firms, in relative terms. This distorsive effect of search frictions is a direct consequence of the competition channel. While functional forms obviously matter to obtain the analytical predictions, we argue that this result applies more generally whenever:

$$
\frac{d^2 \rho_{s_j i}}{d\lambda_{ij} dz_{s_j}} > 0 \quad \text{and} \quad \frac{d^2 \rho_{s_j i}}{dd_{ij} dz_{s_j}} > 0
$$

In particular, one may wonder whether imposing the same meeting probability to all firms, whatever their productivity, is a key driver of the result. An alternative would be a model in which the meeting probability takes the form: $\lambda_{is_j} = f(\lambda_{ij}, z_{s_j})$ with $\frac{df(\lambda_{ij}, z_{s_j})}{d\lambda_{ij}} > 0$ and $\frac{df(\lambda_{ij}, z_{s_j})}{dz_{s_j}} > 0$, i.e. high-productivity firms meet with more buyers. In such model:

$$\frac{d^2 \rho_{s_j i}}{d\lambda_{ij} dz_{s_j}} = \left[\frac{\rho_{s_j i}}{\lambda_{ij}} \frac{d^2 f(\lambda_{ij}, z_{s_j})}{d\lambda_{ij} dz_{s_j}} + \frac{\rho_{s_j i}}{\mathbb{P}()} \frac{d^2 \mathbb{P}\left(\min_{s'_k \in \Omega_{b_i}} \left\{ \frac{c_k d_{ik}}{z_{s'_k}} \right\} = s_j \right)}{d\lambda_{ij} dz_{s_j}} \right]$$

As in the benchmark case, the second term is likely to be negative and increasing in z_{s_j}. The second derivative should be larger than in the benchmark since a reduction in frictions implies that the typical buyer in i meets with more sellers and the additional sellers met are more productive, on average. From this point of view, the competitive channel is even more distorsive in this case. However, a reduction in frictions also affects the relative meeting probabilities at different points of the distribution, i.e. $\frac{d^2 f(\lambda_{ij}, z_{s_j})}{d\lambda_{ij} dz_{s_j}}$ might no longer be zero. From this, it comes that the distortive impact of frictions is likely to show up in this model as well, whenever the cross derivative of the meeting probability with respect to λ_{ij} and z_{s_j} is not too negative.

A.4 Expected mass of firms serving M buyers

Integrating the probability of having exactly M buyers along the distribution of productivities gives the expected measure of firms from j with exactly M buyers in i:

$$h_{ij}(M) = -\int_{z_{min}}^{+\infty} C_{B_i}^M \rho_{s_j i}^M (1 - \rho_{s_j i})^{B_i - M} d\mu_j^Z(z)$$

Using the following change of variable:

$$\rho_{s_j i} = \lambda_{ij} e^{-\frac{\lambda_{ij}}{\pi_{ij}} T_j z_{s_j}^{-\theta}}$$

one can show that:

$$h_{ij}(M) \quad = \quad \frac{\pi_{ij}}{\lambda_{ij}} C_{B_i}^M \int_{\rho_{s_{min} i}}^{\lambda_{ij}} \rho_{s_j i}^{M-1} (1 - \rho_{s_j i})^{B_i - M} d\rho_{s_j i}$$

where $\rho_{s_{min} i}$ is the probability of the least productive firm in j to serve a buyer in i.

If we assume that $M > 0$ we can recognize a function of the family of the Beta function:

$$h_{ij}(M) = \frac{\pi_{ij}}{\lambda_{ij}} C_{B_i}^M \left(B(\lambda_{ij}, M, B_i - M + 1) - B(\rho_{s_{min} i}, M, B_i - M + 1) \right)$$

with $B(\lambda_{ij}, M, B_i - M + 1) = \int_0^{\lambda_{ij}} \rho_{s_j,i}^{M-1}(1-\rho_{s_j,i})^{B_i-M} d\rho_{s_j,i}$ being the incomplete beta function.

Using properties of the Beta function, notice that :

$$\begin{aligned} B(M, B_i - M + 1) &= \frac{\Gamma(M)\Gamma(B_i - M + 1)}{\Gamma(M + B_i - M + 1)} = \frac{\Gamma(M)\Gamma(B_i - M + 1)}{\Gamma(B_i + 1)} \\ &= \frac{(M-1)!(B_i - M)!}{B_i!} = \frac{1}{M}\frac{(M)!(B_i - M)!}{B_i!} \\ &= \frac{1}{M}\frac{1}{C_{B_i}^M} \end{aligned}$$

Then, the regularized incomplete beta function is :

$$I_{\lambda_{ij}}(M, B_i - M + 1) = \frac{B(\lambda_{ij}, M, B_i - M + 1)}{B(M, B_i - M + 1)} = B(\lambda_{ij}, M, B_i - M + 1) C_{B_i}^M M$$

Now, we can rewrite the expression for the mass of suppliers from j with M buyers in i with the help of the regularized incomplete beta function:

$$h_{ij}(M) = \frac{\pi_{ij}}{\lambda_{ij}}\frac{1}{M}\left(I_{\lambda_{ij}}(M, B_i - M + 1) - I_{\rho_{s_{min}i}}(M, B_i - M + 1)\right)$$

Finally, note that if $\rho_{s_{min}i}$ goes to 0, $I_{\rho_{s_{min}i}}(M, B_i - M + 1)$ goes to 0 as well:

$$\lim_{\rho_{s_{min}i}\to 0} I_{\rho_{s_{min}i}}(M, B_i - M + 1) = \lim_{\rho_{s_{min}i}\to 0} \int_0^{\rho_{s_{min}i}} \rho_{s_j,i}^{M-1}(1 - \rho_{s_j,i})^{B_i - M} d\rho_{s_j,i} = 0$$

Using this, one recovers equation (7) in the text:

$$h_{ij}(M) = \frac{\pi_{ij}}{\lambda_{ij}}\frac{1}{M} I_{\lambda_{ij}}(M, B_i - M + 1)$$

B Details on the empirical strategy

B.1 Distribution of the Auxiliary Parameter

We will work with the following convergent moments as auxiliary parameters:

$$\theta_{ij}(\lambda_{ij}, M) = \frac{h_{ij}(M)}{\sum\limits_{M=0}^{B_i} h_{ij}(M)} = \frac{1}{M}\frac{I_{\lambda_{ij}}(M, B_i - M + 1)}{\int_0^{\lambda_{ij}} \frac{(1-\rho_{s_j i})^{B_i}}{\rho_{s_j i}} d\rho_{s_{ji}} + \sum\limits_{M=1}^{B_i} \frac{1}{M} I_{\lambda_{ij}}(M, B_i - M + 1)} \tag{13}$$

i.e. the proportion of firms from j having exactly M buyers in destination i.[37] We first show that the empirical counterparts of these auxiliary parameters are normally distributed. Then

[37] Here and in the rest of the section, the number B_i of buyers in country i is treated as known. Section 4.2 explains how we measure it in the data.

we apply the delta-method to work with the moment we chose to identify λ_{ij}. Finally, we discuss the asymptotic distribution of our estimator of λ_{ij}.

In line with our theoretical framework we note $\left[\mathbb{1}\{B_{s_j i} = M\}\right]_{s_j \in S_j}$ the vector of dummy variables which equal one whenever a firm in the sample has exactly M buyers in country i. The vector is of size S_j, the number of observations in the sample under consideration. The dummies are independent and identically distributed random variables of mean $\theta_{ij}(\lambda_{ij}, M)$ and of variance $\sigma^2_{ij}(M)$. This is true for all $M \in [0, B_i]$.[38] The Central Limit Theorem implies:

$$\sqrt{S_j}\left(\hat{\theta}_{ij} - \theta_{ij}(\lambda_{ij})\right) \underset{S_j \to +\infty}{\xrightarrow{\mathcal{D}}} \mathcal{N}_B(0, \Sigma_{ij}) \tag{14}$$

where

$$\hat{\theta}_{ij} = \begin{pmatrix} \frac{\sum_{s_j=1}^{S_j} \mathbb{1}\{B_{s_j i}=1\}}{S_j} \\ \frac{\sum_{s_j=1}^{S_j} \mathbb{1}\{B_{s_j i}=2\}}{S_j} \\ \ldots \\ \frac{\sum_{s_j=1}^{S_j} \mathbb{1}\{B_{s_j i}=B_i\}}{S_j} \end{pmatrix} \quad \text{and} \quad \theta_{ij}(\lambda_{ij}) = \begin{pmatrix} \frac{h_{ij}(1)}{\sum_{M=0}^{B_j} h_{ij}(M)} \\ \frac{h_{ij}(2)}{\sum_{M=0}^{B_i} h_{ij}(M)} \\ \ldots \\ \frac{h_{ij}(B_i)}{\sum_{M=0}^{B_i} h_{ij}(M)} \end{pmatrix}$$

respectively denote the vector of empirical and auxiliary parameters and Σ_{ij} is the variance-covariance matrix of the B_i random variables $\mathbb{1}\{B_{s_j i} = M\}$, for $M \in \{1..., B_i\}$.

We then consider the function

$$\begin{array}{cccc} g : & \mathbb{R}^{B_i} & \mapsto & \mathbb{R} \\ & \begin{pmatrix} \theta_{ij}(\lambda_{ij}, 1) \\ \theta_{ij}(\lambda_{ij}, 2) \\ \ldots \\ \theta_{ij}(\lambda_{ij}, B_i) \end{pmatrix} & \to & Var\left(m_1 = \frac{\theta_{ij}(\lambda_{ij}, 2)}{\theta_{ij}(\lambda_{ij}, 1)}, m_2 = \frac{\sum_{M=3}^{6} \theta_{ij}(\lambda_{ij}, M)}{\theta_{ij}(\lambda_{ij}, 1)}, m_3 = \frac{\sum_{M=7}^{B_i} \theta_{ij}(\lambda_{ij}, M)}{\theta_{ij}(\lambda_{ij}, 1)}\right) \end{array}$$

where $Var(.)$ is the variance operator. g is derivable and verifies the property $\nabla g(\theta_{ij}(\lambda_{ij})) \neq 0$. Using the Delta-Method, one can show that an estimate of λ_{ij} based on $g(.)$ is asymptotically

[38] Independence comes from the fact that sellers are independent from each other. Note that this assumption could be relaxed since we could eventually use a version of the central limit theoreim based on weak dependence conditions. They are identically distributed ex-ante as sellers draw there productivity in the same distribution and face the same degree of search frictions.

normal:

$$\sqrt{S_j}[g(\hat{\theta}_{ij}) - g(\theta_{ij}(\lambda_{ij}))] \underset{S_j \to +\infty}{\xrightarrow{\mathcal{D}}} \mathcal{N}\left(\begin{pmatrix} 0 \end{pmatrix}, \Omega(\theta_{ij}(\lambda_{ij})) = \nabla' g(\theta_{ij}(\lambda_{ij}))\Sigma_{ij}\nabla g(\theta_{ij}(\lambda_{ij}))\right) \quad (15)$$

where $\nabla g(\theta_{ij}(\lambda_{ij}))$ is of dimension $[B_i, 1]$ and is defined as

$$\begin{pmatrix} \dfrac{\partial g}{\partial \theta_{ij}(\lambda_{ij},1)} = -\frac{2}{3}\sum_{p=1}^{3} \frac{(m_p - \bar{m})m_p}{\theta_{ij}(\lambda_{ij},1)} \\ \dfrac{\partial g}{\partial \theta_{ij}(\lambda_{ij},2)} = \dfrac{2}{3}\frac{m_1 - \bar{m}}{\theta_{ij}(\lambda_{ij},1)} \\ \dfrac{\partial g}{\partial \theta_{ij}(\lambda_{ij},3)} = \dfrac{2}{3}\frac{m_2 - \bar{m}}{\theta_{ij}(\lambda_{ij},1)} \\ \ldots \\ \dfrac{\partial g}{\partial \theta_{ij}(\lambda_{ij},6)} = \dfrac{2}{3}\frac{m_2 - \bar{m}}{\theta_{ij}(\lambda_{ij},1)} \\ \dfrac{\partial g}{\partial \theta_{ij}(\lambda_{ij},7)} = \dfrac{2}{3}\frac{m_3 - \bar{m}}{\theta_{ij}(\lambda_{ij},1)} \\ \ldots \\ \dfrac{\partial g}{\partial \theta_{ij}(\lambda_{ij},B_i)} = \dfrac{2}{3}\frac{m_3 - \bar{m}}{\theta_{ij}(\lambda_{ij},1)} \end{pmatrix}$$

with $\bar{m} = \frac{1}{3}\sum_{p=1}^{3} m_p$.

In practice, our estimation is implemented in two steps. First we use an estimation of the $\Omega(\widehat{\theta}_{ij})$ weight matrix using our observations $\nabla g(\hat{\theta}_{ij})$ and $\widehat{\Sigma_{ij}}$. Second, with the $\hat{\lambda}_{ij}$ estimated in the first step we re-run our estimation with $\Omega(\theta(\hat{\lambda}_{ij}))$.

As proved in Gouriéroux et al. (1985), the variance of the GMM estimator of λ_{ij} is:

$$\Sigma_{\lambda_{ij}} = \left[\frac{\partial g(\theta_{ij}(\lambda_{ij}))}{\partial \lambda_{ij}} \Omega(\theta_{ij}(\lambda_{ij}))^{-1} \frac{\partial g(\theta_{ij}(\lambda_{ij}))}{\partial \lambda_{ij}}\right]^{-1}$$

with

$$\begin{aligned} \frac{\partial g(\theta_{ij}(\lambda_{ij}))}{\partial \lambda_{ij}} &= \tfrac{2}{3}(m_1 - \bar{m})\frac{\partial \theta_{ij}(\lambda_{ij},2)/\theta_{ij}(\lambda_{ij},1)}{\partial \lambda_{ij}} \\ &\quad + \tfrac{2}{3}(m_2 - \bar{m})\sum_{M=3}^{6} \frac{\partial \theta_{ij}(\lambda_{ij},M)/\theta_{ij}(\lambda_{ij},1)}{\partial \lambda_{ij}} \\ &\quad + \tfrac{2}{3}(m_3 - \bar{m})\sum_{M=7}^{B_i} \frac{\partial \theta_{ij}(\lambda_{ij},M)/\theta_{ij}(\lambda_{ij},1)}{\partial \lambda_{ij}} \end{aligned}$$

Figure A.1: ***Number of buyers per seller, full and restricted sample***

Distribution of sellers' degrees, all destination countries

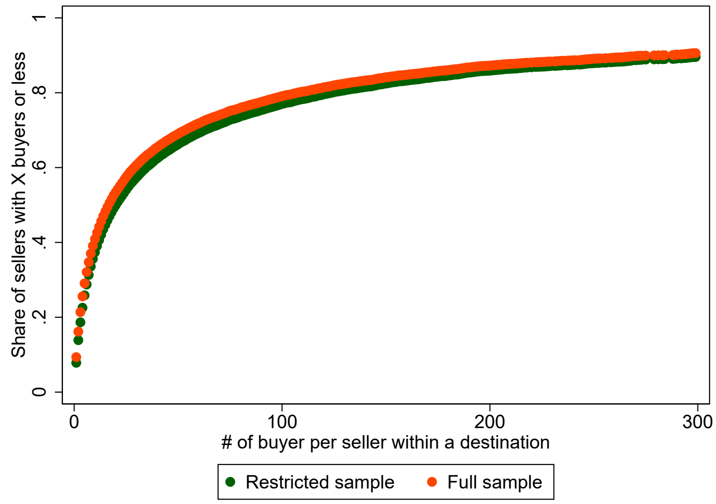

Number of sellers serving one buyer, by destination

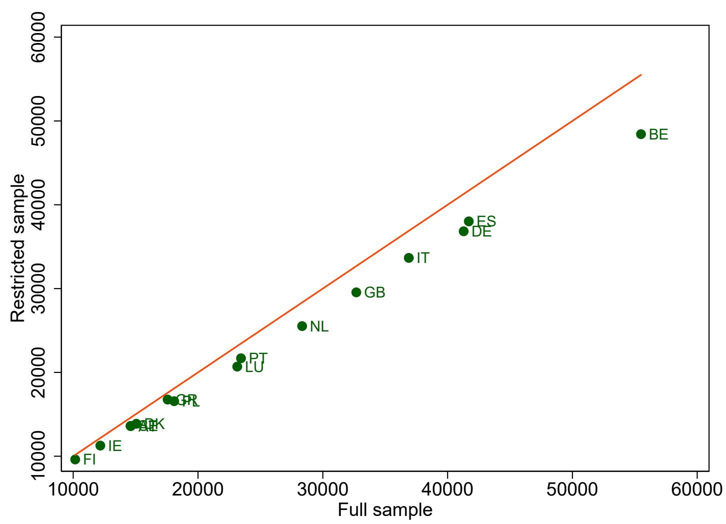

Notes: This figure compares the number of buyers per seller, in the whole dataset and in the estimation dataset, restricted to the 90% of exporters that declare the product category of their exports ("Restricted sample"). The top panel compares the distributions of sellers' degrees, where a firm's degree is computed as the total number of buyers it serves in a given destination. The bottom panel compares the number of exporters declaring to serve one buyer in a given destination, in the full sample (x-axis) and the restricted sample (y-axis). The red line is the 45-degree line.

Table A1: ***French sellers and EU buyers, 2007***

	Number of			*Number of*		
	Exporters	Importers	Pairs	Exporter-HS6	Importer-HS6	Triplets
	(1)	(2)	(3)	(4)	(5)	(6)
Overall	44,255	572,536	1,260,001	184,435	2,390,249	2,879,448
Austria	8,205	14,035	28,128	21,393	52,916	61,478
Belgium	29,468	71,271	214,070	97,415	379,490	482,960
Bulgaria	2,294	2,287	3,657	5,747	6,886	7,630
Cyprus	2,362	1,627	3,735	7,252	8,342	10,041
Czech Republic	6,846	6,117	13,196	16,544	21,491	25,192
Denmark	8,356	8,832	20,846	21,105	37,411	46,574
Estonia	1,802	1,235	2,494	5,230	5,477	6,358
Finland	5,257	5,167	11,592	13,704	21,924	26,046
Germany	24,641	117,935	236,536	73,735	391,424	462,759
Greece	7,792	11,261	25,412	26,054	55,601	68,533
Hungary	5,375	4,437	9,554	12,912	16,309	18,670
Ireland	6,351	6,670	16,265	17,938	38,169	49,297
Italy	20,123	95,864	183,238	63,494	375,681	438,393
Latvia	2,063	1,355	2,948	5,895	6,060	7,430
Lithuania	2,913	1,853	4,698	7,235	7,306	9,891
Luxembourg	10,734	7,652	28,566	31,379	54,959	70,251
Malta	1,781	930	2,552	4,709	4,715	5,781
Netherlands	16,442	33,637	69,833	43,548	131,420	157,913
Poland	9,733	12,857	30,230	24,687	43,482	52,631
Portugal	11,648	19,676	42,925	35,073	95,385	113,477
Romania	5,036	4,855	9,502	12,499	16,446	18,416
Slovakia	3,272	2,306	5,003	7,345	8,078	9,400
Slovenia	2,842	2,227	4,389	7,516	8,634	9,751
Spain	21,633	77,592	159,636	70,410	359,825	419,895
Sweden	7,682	10,198	20,391	20,212	39,315	45,462
UK	18,892	50,660	110,605	55,276	203,503	255,219

Notes: This table gives the number of exporters, importers, exporter-importer pairs, exporter-HS6 product pairs, importer-HS6 product pairs, and importer-exporter-HS6 products triplets involved in a given bilateral trade flow. The data are for 2007 and are restricted to transactions with recorded CN8-products.

Table A2: ***Number of buyers per seller across destination countries***

	Mean	Median	p75	Sh. with 1 buyer
	(1)	(2)	(3)	(4)
Austria	2.3	1	2	67%
Belgium	4.3	1	3	54%
Bulgaria	1.2	1	1	87%
Cyprus	1.3	1	1	82%
Czech Republic	1.4	1	1	79%
Denmark	2.2	1	2	68%
Estonia	1.2	1	1	87%
Finland	1.7	1	2	74%
Germany	5.0	1	3	55%
Greece	2.2	1	2	68%
Hungary	1.3	1	1	82%
Ireland	2.6	1	2	67%
Italy	5.0	1	3	59%
Latvia	1.2	1	1	87%
Lithuania	1.3	1	1	83%
Luxembourg	1.8	1	2	70%
Malta	1.2	1	1	87%
Netherlands	3.3	1	2	61%
Poland	1.7	1	2	74%
Portugal	2.8	1	2	67%
Romania	1.3	1	1	81%
Slovenia	1.3	1	1	82%
Slovakia	1.3	1	1	85%
Spain	4.2	1	3	59%
Sweden	2.0	1	2	67%
United Kingdom	3.9	1	3	59%
Across countries	12.6	2	8	39%

Notes: Columns (1)-(3) respectively report the mean, median, and third quartile number of buyers per seller in each destination. Column (4) gives the share of sellers having a unique buyer. A seller is defined as an exporter-HS6 product pair. The data are for 2007 and are restricted to transactions with recorded CN8-products.