Search Frictions in International Goods Markets

Clémence Lenoir* Julien Martin† Isabelle Mejean‡

October 30, 2020

Abstract

This paper studies how frictions in the acquisition of new customers distort the allocation of resources across heterogeneous producers. We add bilateral search frictions in a Ricardian model of trade and use French firm-to-firm trade data to estimate search frictions faced by French exporters. Estimated frictions are more severe in large and distant countries and for differentiated products. A counterfactual reduction in the level of search frictions improves the efficiency of the selection process and increases the average productivity of exports, because the least productive exporters are pushed out of the market, whereas exports increase at the top of the productivity distribution.

JEL Classification: F10, F11, F14, L15

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

*CREST-ENSAE. Address: CREST 5, Av Henry Le Chatelier, 91120 Palaiseau, France. Email address: clemence.lenoir@ensae-paristech.fr.

†UQAM, CREST, and CEPR. Address: Université du Québec à Montréal, C.P. 8888, succursale A, Montréal, QC H3C 3P8, Canada. Email address: martin.julien@uqam.ca.

‡CREST-Ecole Polytechnique and CEPR. Address: CREST 5, Av Henry Le Chatelier, 91120 Palaiseau, France. Email address: isabelle.mejean@polytechnique.edu. Tel: +33170266824.

1 Introduction

Customer acquisition, which is central for firms' economic development, is subject to various forms of frictions, such as search frictions or asymmetric information.[1] Despite their prevalence in most product markets, the abundant literature on the sources of misallocation among heterogeneous producers has overlooked such frictions.[2] In this paper, we ask whether and how frictions in the acquisition of new customers distort the effectiveness of resource allocation across heterogeneous producers. We do so in the context of international goods markets, in which search frictions are likely pervasive, interact with other barriers to trade, and for which we have rich data to estimate search frictions.

We develop and estimate a model of firm-to-firm trade displaying frictions that affect the matching of firms and consumers in international markets. We discuss the consequences of these frictions for the efficiency of selection into exporting, and the size of the customer base, conditional on trade. Search frictions penalize disproportionally the most productive producers, and thus distort the allocation of resources across exporters within an origin country. We use this prediction of the model to develop a structural estimator of search frictions. Estimates recovered for a large cross-section of products and destination countries are used to quantify the impact of search frictions on the allocation of resources across heterogeneous producers.

The starting point is a Ricardian model of trade à la Eaton and Kortum (2002). Their model constitutes a useful benchmark to study the efficiency of selection into export activities, because it displays an extreme form of selection. The assumption of perfect substitutability between heterogeneous producers of the same good implies the ex-post distribution of active firms is degenerated with the most productive firms of each good from each country being the only ones eventually active ex post. Our model does

[1] Luttmer (2006) discusses the role of the customer base as a determinant of firms' size. Arkolakis (2010) focuses on the impact of penetration costs associated with acquiring additional consumers on trade patterns. The impact of frictions is studied in various recent contributions, in the trade and macroeconomic literatures. Perla (2019) shows how information frictions can impede customer acquisition and thus firms' growth. Gourio and Rudanko (2014) use statistics on the size of marketing expenditures at the firm level as evidence of search frictions in product markets and study their consequences for the dynamics of firms. In a business-to-business trade context, asymmetric information on market conditions (Allen, 2014) or the seller's reliability (Macchiavello and Morjaria, 2015) have been argued to affect firms' pricing and quantity decisions.

[2] See Hopenhayn (2014) for a review of the related literature. Among the distortions that are extensively discussed in the empirical and theoretical literature, one can cite regulations (Garicano et al., 2016), financial constraints (Midrigan and Xu, 2014), or - closer to our paper - information frictions (David et al., 2016).

not display such feature. Search frictions imply each consumer meets with a random number of potential suppliers from each origin country. As a consequence, consumers are not able to identify who is the most competitive producer for the good they are after and multiple heterogeneous producers can be active ex post, within the same origin country. Search frictions are thus a sufficient ingredient to generate inefficiencies in the selection of firms into exporting. The tractability of the model makes quantifying the magnitude of the inefficiency possible. We use these analytical predictions to estimate product-level search frictions using firm-to-firm trade data.

The matching process in our model displays random search. A discrete number of ex-ante homogeneous consumers in each market meet with a random number of heterogeneous producers of a perfectly substitutable good. Conditional on her random choice set, the consumer chooses the lowest-cost supplier. As in Eaton and Kortum (2002), iceberg trade costs interact with technological parameters in shaping Ricardian advantages. In the absence of search frictions, all consumers within a country would end up interacting with the firm that can offer the good at the lowest price thanks to a combination of high productivity and low trade costs. Because of search frictions, consumers end up interacting with different suppliers that serve the market at heterogeneous costs. The ex-post dispersion of suppliers' costs, and its implication for realized prices, is thus a sufficient statistics for the inefficiency of the matching process in this simple framework.[3]

In this framework, the likelihood that two firms meet is determined by the search friction parameter, which we assume is product- and country-pair specific. Search frictions can thus be considered as another source of bilateral trade costs, beyond iceberg trade costs.[4] In the model, both parameters have the same qualitative impact on bilateral trade flows at the product level. High search frictions and large iceberg costs

[3]A side benefit of the theoretical exercise is that it produces a Ricardian model of trade that has interesting predictions regarding export probabilities *at the firm level*. To recover these predictions, the key elements are a discrete number of producers of each good, as in Eaton et al. (2012), and the presence of search frictions. The randomness in the matching process implies even poorly productive firms have a strictly positive (but low) probability of exporting in any destination, despite the Ricardian forces.

[4]This assumption is consistent with evidence in the gravity literature, which uses dyadic proxies for information frictions and finds their impact on the geography of bilateral trade is significant. In this literature, a common language or former colonial ties are well-known to contribute substantially to the model's explanatory power (Head and Mayer, 2014). Rauch (1999) and Rauch and Trindade (2002) provide evidence that the stock of migrants from one origin in a country is significantly correlated with more bilateral trade. Their interpretation is that migrants help reduce information frictions characterizing international goods markets.

both reduce the market share of a given origin country in the destination. These frictions are not isomorphic though, because search frictions further distort the allocation of resources among exporters from a given origin in the destination country. The reason is that search frictions are unequally costly along the productivity distribution. High-productivity firms always suffer from high search frictions that prevent them from increasing the number of consumers in their portfolio, and thus their market share. The impact is lower, or actually reversed, at the bottom of the distribution. Frictions indeed reduce the number of competing firms in the importer's choice set, which increases the chance that a low-productivity exporter ends up serving the firm.[5] We exploit this prediction to identify search frictions in the data, conditional on other barriers to trade.[6] The structural estimator exploits the product-level dispersion in the number of buyers served by exporters from a given origin in a particular destination, conditional on the first moment of the distribution that depends on iceberg trade costs. In the 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. Because iceberg trade costs do not have such a distortive effect, exploiting this moment of the data is useful to recover search frictions separately from other trade barriers.

The empirical counterpart of this moment is computed using firm-to-firm trade data covering the universe of French exporters and each of their individual clients in the European Union, restricted to its 15 oldest members. Search frictions are estimated by a

[5]In the benchmark model, the log-supermodularity of export probabilities arises from the probability of being chosen conditional on a match varying along the distribution of productivities, whereas the matching probability is constant by assumption. Dickstein and Morales (2018), however, find large firms have more knowledge of foreign markets than small firms. Moreover, endogenous search effort would distort the distribution of matching probabilities towards high-productivity exporters. We study the robustness of our predictions to these possibilities. We show in the appendix that our model can handle matching probabilities that increase with the firm's productivity while the main result remains qualitatively unchanged. For the qualitative result to survive, the log-supermodularity needs to survive as well, which happens if the cross derivative of the matching probability with respect to search frictions is not too decreasing in productivity.

[6]The prediction can also be used to rationalize other patterns in trade data, namely, that "too many" firms are selling small amounts in a typical export destination and that firms do not enter markets according to an exact hierarchy. As discussed in Eaton et al. (2011), these patterns are not consistent with a Melitz (2003) model, which needs to be augmented with additional degrees of randomness to fit individual data. In our model, such patterns emerge endogenously. Due to search frictions, any firm always has a chance to enter a difficult market, even though, on average, it is more likely to serve a nearby than a remote country. But conditional on serving a difficult market, the likelihood that a low-productivity firm serves several consumers there is tiny, thus the prevalence of small export flows in the data. See also Eaton et al. (2018) for a more systematic discussion of export patterns under search frictions.

generalized method of moments for about 10,000 product and destination country pairs. The recovered distribution displays substantial dispersion, with most product markets featuring moderated search frictions, whereas a small number of product×destination pairs are found to be highly frictional. In the last quartile of the distribution, the probability of meeting zero consumers for a French firm willing to export is above 4% and can reach 50% in the last decile. The maximum degree of average frictions that French exporters face is found in Greece and Finland, and the less frictional country is Belgium. The country dimension, however, explains only 13% of the recovered variance. As expected, search frictions are estimated to be stronger in differentiated product markets. Within a product, they are more pronounced in distant and more populated countries, and lower in countries where the population of French migrants is larger. Importantly, the estimated model is able to fit the distribution of the number of buyers that exporters serve within a country and product. Frictions notably help explain the skewness of the distribution. They 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, this measure of fit is sizeable.

What are the distributional consequences of search frictions? Qualitatively, we can assess how distortive estimated search frictions are by studying their correlation with Ricardian comparative advantages. Intuitively, being unable to meet with foreign consumers is especially costly when exporters are in a strong competitive position in the market. To test this prediction of the model, we infer revealed comparative advantages at the product level using a model-consistent structural gravity equation applied to multilateral data. Revealed comparative advantages are then correlated with estimated search frictions. The empirical correlation between Ricardian comparative advantages and search frictions is positive; that is, search frictions are higher, on average, in sectors in which France has a comparative advantage. This correlation magnifies the distortive impact of search frictions.

To go beyond these qualitative results, we rely on counterfactuals. Our main experiment consists of 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 1.14 percentage point increase in the market share of French exporters in Greece, in the median product market. This relatively small aggregate effect 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 of productivities, from 7% to 5.7%, on average. At the top of the distribution, export probabilities instead increase from 61% to 82% among the top 15% productivity percentiles and from 66% to 90% among the top 5%. Within the sub-sample of exporters, a reduction in search frictions also reallocates market shares as the expected number of buyers served by high-productivity firms increases substantially. All in all, the mean productivity of exporters increases by 6% to 11% due to Greek importers being more likely to identify the most productive French suppliers. Such allocative gains are specific to search frictions. When we instead consider a reduction in iceberg trade costs, the results are very different. Namely, a reduction in iceberg costs that leads to the same change in product-level trade patterns as those generated by the drop in search frictions leads to a *decrease* in the average productivity of exporters. The reason is that a decrease in iceberg costs benefits disproportionately low-productivity firms.

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.[7]

Related literature. Our paper is related to different strands of the literature. The role of search and information frictions in international markets is the topic of an old empirical and theoretical literature. Rauch (2001) thus explains the role of migrant networks in international markets by way of such frictions. More recently, a series of papers 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.

[7]Business France, the French export-promoting agency, offers several programs that are meant to help firms meet with foreign clients. The agency notably helps financing firms' participation in international trade fairs or organizing bilateral meetings with representatives of the sector in the destination country. Using their data on the size of these programs, we show Business France indeed targets high-frictional markets, on average, as shown by the positive correlation between the intensity of their activities and our estimates of search frictions across destinations.

(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). Several recent contributions have also studied this topic theoretically. Krolikowski and McCallum (2018) introduce random matching frictions in a Melitz (2003) type framework. In their model, matched producers trade with a single buyer within a country. We instead focus on the role search frictions in explaining heterogeneity in the number of buyers across sellers. Chaney (2014) develops a model of firms' network in international trade. In his model, ex-ante homogenous sellers face informational frictions and meet sellers through directed search and through their existing networks. In Allen (2014), information frictions also hit the seller side of the economy; exporters ignore the potential price of their crops abroad, and 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 buyers optimally choose how much to invest in information processing to discover potential suppliers. In comparison with theirs, our model is based on simpler assumptions about the search technology that is purely random in our case. Our model is instead richer on the modeling of the supply side as we allow for multiple heterogeneous producers in each origin country whereas they have a single firm per exporting country. The tractability of our framework allows us to derive closed-form solutions, estimate frictions structurally, and discuss the allocation of resources among heterogeneous sellers of the same origin country.[8,9]

We also contribute to a series of recent papers that 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; Eaton et al., 2018). The main stylized fact we document, exporters' heterogeneity in terms of the number of buyers they serve in a given destination, is robust across country datasets.[10] In Bernard et al.

[8]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 in order to increase their visibility in foreign markets (Arkolakis, 2010). Instead, our model assumes large exporters have the same probability of meeting with a buyer as low-productivity ones. We discuss in Appendix A.2 the sensitivity of our results to this assumption.

[9]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.

[10]Our analysis however displays a notable difference in comparison with the previous literature. Once we condition on a particular product being traded, we indeed show that more than 90% of importers in our data source a given product from a single French exporter. Instead, their overall degree, in our

(2018b); Carballo et al. (2018), the heterogeneity is studied in monopolistic competition models with two-sided heterogeneity. In Eaton et al. (2018) (EKKII hereafter) as in our paper, the matching of exporters and importers is governed by random search in a Ricardian framework. Whereas both papers use common tools, their scopes are quite different. Our main goal is to understand how frictions shape the distribution of product-level sales across French exporters. EKKII introduces frictions in a model of firm-to-firm trade to understand the link between aggregate trade and the labor share. These complementary objectives induce several differences in modeling, and in the empirical objects we examine. First, we extend the Ricardian analysis to the case of a discrete number of producers following the logic introduced in Eaton et al. (2012). Unlike EKKII, working with a discrete number of firms allows the model to have Eaton and Kortum (2002) as a limit case whenever there are no search frictions. Second, the relative simplicity of our model allows us to derive analytical predictions regarding the heterogeneity in export performances across firms within a product and destination market, whereas EKKII estimate frictions at the country level. We are thus able to describe cross-sectoral heterogeneity in search frictions. Third, firms in EKKII are both buyer and seller, and can choose to perform tasks in-house (with labor) or use inputs. This extra layer allows them to connect trade and labor macro-outcomes, which is not examined in our paper. We instead focus on misallocation across sellers at a more micro level (within product-level categories).

The introduction of a countable number of firms also relates our work to recent papers that examine trade patterns in models with a finite number of firms (Eaton et al., 2012; Gaubert and Itskhoki, 2018). Whereas in these papers, the coexistence of several firms in a given market is due to imperfect substituability of the varieties produced, we instead consider perfectly substitutable varieties that can co-exist in a market due to the combination of search frictions and the presence of multiple buyers. Contrary to us, Gaubert and Itskhoki (2018) solve their model in general equilibrium, by assuming that granular firms affect sectoral trade patterns but are atomistic with respect to aggregate prices. Likewise, it would be possible to insert our partial equilibrium structure in a GE framework using a continuum of sectors. As discussed at the end of section 3, such GE structure would not add much to the analysis, which aims to quantify the extent

data as in others, is often above one as importers tend to source several products from several French firms. Our analysis is agnostic about this particular dimension of heterogeneity, which we argue does not interact with search frictions as importers rarely import two products from the same exporter. Their decisions to source two different products are thus taken as independent.

to which product-level search frictions can distort the allocation of resources among exporters in that market.

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 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 that number 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 coefficients 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 detail in Bergounhon et al. (2018). The full dataset 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. Because the analysis is conducted at the product-level, we must drop all transactions that are reported under the simplified declaration regime, for which the product category is not recorded. This restriction concerns 10% of firms whose overall exports in the European Union during the year are below 150,000 euros.[11]

For each transaction, the dataset contains 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

[11] One might be concerned that this selection biases our empirical analysis, because the neglected small exporters are likely to display systematically different patterns of exports. Although we cannot rule out this possibility, we believe the bias should not be substantial, based on evidence reported in Figure A.1. Namely, the distribution of sellers' degrees, whose 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. Although the restricted sample obviously contains more exporters with one buyer, the difference is roughly proportional to the total number of such exporters in the whole dataset (bottom panel).

(at the 8-digit level of the combined nomenclature), and the value of the shipment. Linking each exporter to its sector of activity is also possible, using INSEE data. In the analysis, data are aggregated across transactions within a year, for each exporter-importer-hs6 product triplet. The product dimension allows 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. To be consistent with this assumption, we further 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, defined at the 4-digit level of the NAF-1993 nomenclature.[12] 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).[13]

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 billion euros. This number represents 53% of France worldwide exports. Table A1 displays the number of individuals involved in each bilateral trade flow. Most of the time, the number of importers is larger than the number of exporters selling to this destination (Columns (1) and (2)). Such asymmetry suggests 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). The discrepancy

[12] The rationale for such restriction is that we see in the data firms selling many different products, some of which are relatively "close" to the firm's activity (e.g., 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 probably side products that the firm sells to its customers while they buy some of its wine. Although information frictions might be important to identify potential wine consumers, we do not expect frictions in the glasses market 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, deciding which products are tied and which are not is almost impossible. The statistical criterion that we use thus considers that a product that no firm in the sector sells in large enough quantities is probably tied, and is thus removed from the sample.

[13] We intentionally kept in the sample all French firms, whether wholesalers or manufacturers. One may instead be tempted to drop wholesalers since their very presence in the data may indicate search frictions, which trade intermediaries help mitigate. There are three arguments that convinced us not to drop these firms from the sample. First, the precision of our estimator increases with the number of exporters. Second, we cannot identify intermediaries on the other side of the border. Dropping French wholesalers would have meant treating these two categories of intermediaries asymmetrically. Finally, data do not indicate that the number of buyers per seller is significantly different in the wholesale sector than in the rest of the economy, as illustrated in Figure A.2.

in average seller and buyer degrees is stronger 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 that are active in 2007, and column (6) the number of exporter-importer-product triplets. These numbers are an order of magnitude smaller than the number of *potential* relationships, equal to the number of active exporters times the number of importers. This finding suggests 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. 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 the CEPII-BACI database (Gaulier and Zignago, 2010). Finally, the stock of migrants per origin and destination countries, taken from the UN database on Trends in International Migrant Stock, is used as a proxy for information frictions. 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.

2.2 Descriptive Statistics

The most important novelty in firm-to-firm trade 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. These stylized facts are to a large extent consistent with facts uncovered from other data sources (Bernard et al., 2018b; Carballo et al., 2018), even though we systematically condition upon the particular product being traded while they do not.

Figure 1 shows the strong heterogeneity in the number of buyers per seller within a destination. The left panel documents the share of sellers interacting with a given number of buyers, and the right panel depicts their relative weight in overall exports. To illustrate the amount of heterogeneity across destination countries, Figure 1 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 and Germany (triangle and diamond points, respectively).

Figure 1: ***Distribution of the number of buyers per seller, across exporters***

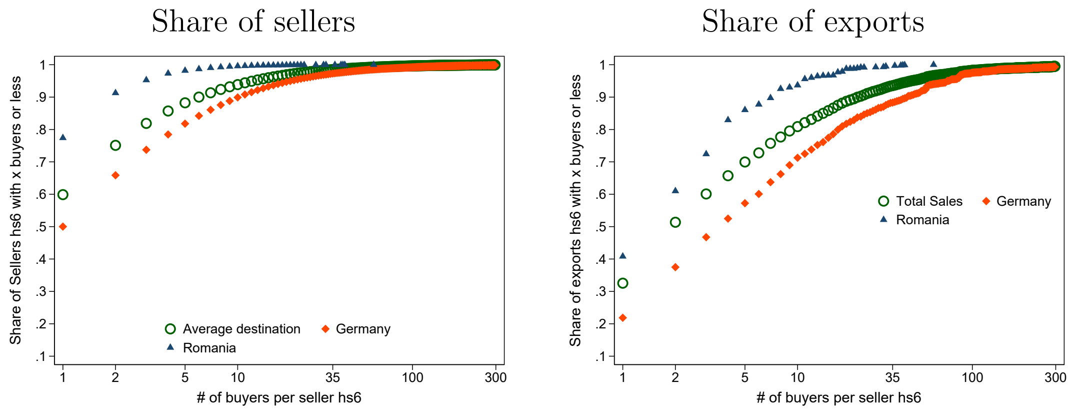

Notes: The figure displays the proportion of sellers (left panel) and the share of trade accounted for by sellers (right panel) that serve x buyers or fewer in a given destination, in 2007. A seller is defined as an exporter-HS6 product pair. The green circles correspond to the average across EU destinations.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

In France's typical export market, 65% of sellers interact with a single buyer, and 90% with at most five buyers. At the other side of the spectrum, 1% of sellers interact with more than 100 buyers in the same destination. As the right panel in Figure 1 shows, 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 of sellers interacting with at most 10 buyers. Based on such evidence, we conclude that French exports are dominated by sellers interacting with a small number of buyers.

Figure 1, 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 1 illustrate the cross-country heterogeneity.[14] Whereas the median degree of sellers is equal to just one buyer in all destination countries, the mean varies substantially, due to varying shares of sellers who serve more clients. Such heterogeneity also exists across sectors, although perhaps less pronounced. A full variance decomposition, however, shows more than 80% of the heterogeneity in the number of buyers served by a seller is within a product and destination. The structural estimation uses this dimension of heterogeneity 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.

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

The conditional correlation coefficient is equal to .28, and the 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, we assume an importer is matched with a single seller, at a point in time. This assumption 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. As a consequence, the mean degree of buyers that can be recovered from the comparison of columns (5) and (6) in Table A1 is very close to 1 in all destinations.[15]

We close this section with an empirical analysis using the gravity framework to show how the buyer margin correlates with 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 the stock of French citizens in the destination, relative to its population, as a proxy for information frictions. A country's import demand is defined as the total value of its worldwide imports, less the value of imports purchased from French firms (Source: CEPII-BACI).[16]

Column (1) confirms the results found in the rest of the literature, namely, that product-level bilateral trade is larger toward closer, bigger, and wealthier destination markets. Moreover, it is positively correlated with the stock of French migrants living in the destination country, which we interpret as information frictions having a negative

[15] Although 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 the same importer importing two products can be considered 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. In the data, we observe a very high correlation between the number of sellers a buyer is connected to and its number of seller$\times$product pairs, however. Such correlation rules out the importance of economies of scale for buyers' search costs.

[16] The stock of French migrants may be argued to be a poor proxy for information frictions since migrations and trade costs could themselves be correlated. The use of this variable is meant to provide indicative evidence of the role of information frictions but the rest of the analysis does not rely on this proxy.

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	-0.653***	-0.308***	-0.193***	-0.152***	-0.196***	-0.194***	-0.006
	(0.068)	(0.033)	(0.024)	(0.046)	(0.052)	(0.026)	(0.042)
log Import Demand	0.795***	0.238***	0.137***	0.419***	0.423***	0.171***	0.253***
	(0.015)	(0.008)	(0.006)	(0.010)	(0.013)	(0.007)	(0.011)
log GDP per Capita	-0.117***	-0.052**	0.018	-0.084***	-0.084***	-0.066***	-0.018
	(0.044)	(0.020)	(0.014)	(0.025)	(0.032)	(0.018)	(0.020)
log French Migrants	0.372***	0.206***	0.091***	0.075***	0.189***	0.103***	0.086***
	(0.020)	(0.008)	(0.007)	(0.012)	(0.015)	(0.008)	(0.011)
Observations	60,770	60,770	60,770	60,770	578,947	578,947	578,947
R-squared	0.648	0.788	0.441	0.588	0.691	0.439	0.720
Fixed effects	Product	Product	Product	Product	Firm	Firm	Firm

Notes: Standard errors, clustered in the country dimension,are 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 1,000 inhabitants. “Migrants in France” is the number of migrants from the destination in France, expressed as a stock per 1,000 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 variables the log of firm-level bilateral exports, whereas 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)).

impact on bilateral trade.[17] These results are also confirmed within a firm, in column (5). In columns (2)-(4) and columns (6)-(7), 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).[18] 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.[19] Likewise, the buyer margin accounts for a substantial share of the overall impact of migrants. Our interpretation of this finding is that migrants

[17]In comparison with a specification that does not control for information frictions, the impact of distance is reduced by about a third. This finding suggests information frictions are correlated with distance from France in this sample. Similar patterns arise when we instead use the number of migrants from the destination country living in France as a proxy for information frictions.

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

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

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} is the set of the importers purchasing product p from seller s. $\#^S_{pd}$, $\#^B_{spd}$, and $\#^{SB}_{pd}$ 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, respectively.

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:

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

[19]Note 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:

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

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

help alleviate information frictions in international markets, which in turn facilitates the matching between exporters and importers.

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

3 Model

This section presents a Ricardian model of firm-to firm trade with search frictions. 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 that we later use in the structural estimation.

3.1 Assumptions

The economy is composed of N countries indexed by $i = 1, ..., N$. The partial equilibrium analysis focuses on a single good produced into perfectly substitutable varieties.[20] As in Eaton et al. (2012), a discrete number of producers of the good are located in each country j. These firms produce with a constant-returns-to-scale technology using an input bundle whose 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 number of firms with productivity higher than z is the realization of a Poisson variable with parameter $T_j z^{-\theta}$. 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. We later use the underlying productivity heterogeneity to explain the dispersion across firms regarding the number of buyers they serve in a destination.

[20] It would be possible to plug the partial equilibrium model described here in a general equilibrium framework, by assuming that there is a continuum of such goods, that countries trade in equilibrium. Labor and goods market equilibria would then be used to solve for equilibrium factor prices consistent with balanced trade. Because the purpose of the model is to help identify search frictions in the data, and the data cover a single exporting country, this development is not necessary.

Iceberg trade costs exist between countries but no entry cost. 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 number of firms from j that can serve market i at a cost below p is a Poisson random variable of parameter $\mu_{ij}(p) = T_j \left(\frac{d_{ij}c_j}{p}\right)^{-\theta}$. Summing over all producing countries, the number of firms from any country in the world that can serve country i at a cost below p is distributed Poisson of parameter $\mu_i(p) = p^\theta \sum_{j=1}^N T_j(d_{ij}c_j)^{-\theta} = p^\theta \Upsilon_i$. 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 is, the more competitors with low costs are in this country.

We depart from the representative consumer's assumption used in most of the literature and instead assume each country is populated by a large number B_i of (ex-ante) homogeneous buyers, each one characterized by its own iso-elastic function. Because of search frictions, each buyer b_i meets with a random subset of the potential suppliers of the good, with each supplier from country j having a probability λ_{ij} of being drawn. Conditional on the subset of producers met, the buyer decides which one to purchase from, by comparing the prices they offer.

In the rest of the analysis, we assume 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:[21]

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

where Ω_{b_i} is the set of producers drawn by buyer b_i.

The number of potential suppliers in the set Ω_{b_i} reflects the extent of search frictions in the economy. In a frictionless world, for $\lambda_{ij} = 1 \ \forall \ (i, j)$, each buyer b_i would meet with

[21] One might question the assumption of marginal-cost pricing in a context of frictional goods 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 a pricing rule, and because the extent of competition across potential suppliers 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 firms drawn by a buyer b_i compete à la Bertrand. Under such an 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 that would equal the marginal cost of the second lowest-cost supplier. Because most of the results discussed here rely on the realization of the match rather than the value traded conditional on a match, most of our results would remain unchanged.

all suppliers. Within a destination, all buyers would thus end up paying the same price for the homogeneous good and the assumption of a representative consumer would be suitable. Eaton and Kortum (2002) make this assumption, which generates an ex-post degenerated distribution of firms because at most one, the lowest-cost supplier, is active in any market i. Our model does not display such degenerated ex-post distribution. Each buyer b_i meets with a random number of potential suppliers, drawn from a Poisson distribution of parameter $\sum_j \lambda_{ij} T_j z_{min}^{-\theta}$. Likewise, the number of suppliers from j (resp. from any country) offering a price below p can be represented by a Poisson process of parameter $\lambda_{ij}\mu_{ij}(p)$ (resp. $\sum_j \lambda_{ij}\mu_{ij}(p)$). Under this assumption, any supplier from j has a strictly positive probability of ending up serving market i. In the rest of the analysis, λ_{ij} is interpreted as an inverse measure of bilateral frictions. A coefficient closer to 1 implies buyers from i gather more information on potential suppliers in country j and are thus more likely to identify the most competitive one.

Given the property of the Poisson distribution, the minimum price at which a buyer b_i can purchase the good can be shown to follow a Weibull distribution:

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

where $\kappa_i \equiv \frac{\sum_j \lambda_{ij} T_j (c_j d_{ij})^{-\theta}}{\sum_j T_j (c_j d_{ij})^{-\theta}}$ measures the expected number of suppliers met, in relative terms with the number of suppliers met under no search frictions. κ_i can also be interpreted as a weighted average of bilateral search frictions, with the weights representative of the relative comparative advantage of the different origin countries in market i, that is, $\kappa_i = \sum_j w_{ij}\lambda_{ij}$ with $w_{ij} \equiv \frac{T_j (c_j d_{ij})^{-\theta}}{\sum_j T_j (c_j d_{ij})^{-\theta}}$.

As in Eaton and Kortum (2002), the ex-post distribution of prices in this economy depends on the strength of competition there, as measured by Υ_i, and the amount of heterogeneity in firms' prices, which is inversely proportional to θ. In comparison with standard Ricardian models, expected prices are, however, inflated by search frictions (because $\kappa_i < 1$). The presence of search frictions indeed implies buyers fail to identify the lowest-cost supplier in the whole distribution of potential producers. This lack of information is distortive, thus inflating the average price paid by consumers in country i. The size of this distortion is inversely related to κ_i. It is larger when search frictions λ_{ij} are negatively correlated with the country's comparative advantages as measured by $T_j(c_j d_{ij})^{-\theta}$. Intuitively, being unable to meet with all potential suppliers is all the more costly for consumers when search frictions increase the relative probability that they meet with poorly competitive firms. In the rest of the analysis, we thus refer to

κ_i as an inverse measure of the distortive impact of 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 us understand how search frictions modify the predictions of Ricardian models à la Eaton and Kortum (2002). We then derive predictions regarding export probabilities along the distribution of firms' productivities, which we later use to identify search frictions in the data, separately from other barriers to trade.

3.2.1 Aggregate trade

In this model, the share of country j's (product-level) consumption that is imported from country i, denoted π_{ij}, is the expected value of goods purchased by buyers that end up interacting with a supplier from j, normalized by aggregate consumption:

$$\pi_{ij} = \mathbb{E}\left[\sum_{b_i=1}^{B_i} I^{(1)}_{b_ij} \frac{X_{b_i}}{X_i}\right],$$

where $I^{(1)}_{b_ij}$ is a dummy variable equal to 1 if the lowest-cost supplier met by b_i originates from country j, and X_{b_i} and X_i, respectively, denote the demand expressed by buyer b_i and market i. Properties of the Poisson distribution imply the probability of the lowest-cost supplier being located in j is constant and independent of b_i. Trade shares simplify into

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

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, $\dfrac{T_j(d_{ij}c_j)^{-\theta}}{\Upsilon_i}$, and (ii) the relative size of search frictions its firms encounter while serving market i, $\dfrac{\lambda_{ij}}{\kappa_i}$. The first ratio is 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. The impact of search frictions is captured by the second term in equation (1). 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.1.

The intuition for this result is straightforward. As search frictions decrease, the likelihood that an exporter from j meets with a buyer from i increases. If parameters governing the rest of the world are left unchanged, the market share of country j in destination i increases. The elasticity is below 1, however, because the improved visibility of exporters is somewhat compensated by an increase in competitive pressures attributable to buyers from i meeting a larger number of exporters from j, on average. A reduction in search frictions unambiguously increases the exporting country's share in the destination's absorption, which is in line with the argument in Rauch (1999) that search frictions can contribute to reducing the magnitude of bilateral trade.

Finally, note 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 $FE_i \equiv \ln\Upsilon_i\kappa_i$ 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 d_{ij} as well as information frictions λ_{ij} to enter the gravity equation.[22] A corollary is that predictions on product-level trade cannot be expected

[22]The estimation of the gravity equation presented in equation (2) can notably be used to estimate the elasticity of trade to iceberg costs (θ) provided the identification restriction is not violated by some correlation between search frictions and iceberg costs. For instance, the identification strategy in Caliendo and Parro (2015) that uses tariffs to instrument iceberg costs and exploits a transformation of equation (2) involving two-way trade within country triplets continues to deliver unbiased trade

to help identify search frictions, separately from other barriers to trade, because both sources of frictions have the same qualitative impact on 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 novel to our model and can be used together with firm-to-firm trade data to estimate search frictions. Because we observe the universe of French exporters, and their clients abroad, we 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:[23]

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

By assumption, the probability of being drawn by a buyer is constant and only depends on the size of bilateral search frictions. More productive sellers, however, have a higher probability of ending up serving any buyer from i because, conditional on being drawn, they have a higher chance of being the lowest-cost supplier. And conditional on productivity, a seller has a higher chance of serving a buyer located in a market that can be served at a low cost (d_{ij} close to one), where competition is limited (Υ_i low), and that displays highly distortive average search frictions (κ_i small). These predictions are consistent with evidence presented in section 2.2.

elasticities provided the identification restriction in their equation (23) holds true for the residual components of *both* iceberg costs and search frictions. More precisely, the estimation of θ is unbiased if the asymmetric bilateral components of d_{ij} and λ_{ij} are uncorrelated with tariffs. On the other hand, the identification strategy proposed by Eaton and Kortum (2002), which instruments iceberg costs by observed price differentials recovered from the UN International Comparison Program, is trickier to interpret in the context of frictional good markets. The reason is that information frictions are a barrier to arbitrage, so that relative prices may no longer be bounded up by trade costs (Allen, 2014).

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

The probability in equation (3) is log-supermodular in bilateral search frictions and firms' productivity. Search frictions do not equally affect firms at different points of the productivity distribution. Under some parameter restrictions, one can further show that reducing search frictions improves export prospects for high-productivity firms while reducing low-productive firms' export probability. These results are summarized in 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:

$$\frac{\partial \ln \rho_{ij}(z)}{\partial \lambda_{ij}} = \underbrace{\frac{\partial \ln \lambda_{ij}}{\partial \lambda_{ij}}}_{\text{Visibility channel}} - \underbrace{\frac{\partial \left(c_j d_{ij}\right)^{\theta} z^{-\theta} \kappa_i \Upsilon_i}{\partial \lambda_{ij}}}_{\text{Competition channel}} = \frac{1}{\lambda_{ij}} - z^{-\theta} T_j \quad \text{and} \quad \frac{\partial^2 \ln \rho_{ij}(z)}{\partial \lambda_{ij} \partial z} > 0 \tag{4}$$

High-productivity firms always benefit from a reduction in frictions:

$$\lim_{z \to +\infty} \frac{\partial \ln \rho_{ij}(z)}{\partial \lambda_{ij}} = \frac{1}{\lambda_{ij}} > 0.$$

For low-enough search frictions, an increase in λ_{ij} instead has a negative impact on firms at the bottom of the distribution; that is,

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

where $\rho_{ij}(z_{min})$ is the export probability in i of a firm from j with productivity z_{min}. .

See the Proof in Appendix A.2.

The ambiguous impact of more bilateral search frictions (a lower meeting probability λ_{ij}) on the probability of serving a particular buyer conditional on the level of productivity can be explained by the opposite impact of the visibility and competition channels. On the one hand, a decrease in search frictions increases the likelihood that seller s_j will serve any buyer in country i as it enhances its probability of meeting with the buyer ("visibility" channel). On the other hand, conditional on being drawn, less bilateral search friction means s_j faces fiercer competition from other domestic suppliers. As a consequence, the probability that it is the lowest-cost supplier met by any particular buyer is reduced, 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 so that the expected number of sellers from j that buyers from i meet is above 1 ($\lambda_{ij} z_{min}^{-\theta} T_j > 1$), the competition channel dominates the visibility channel at the bottom of the productivity distribution, and sufficiently low-productivity sellers benefit from more frictions.[24]

Proposition 2 shares some similarity with results in Dasgupta and Mondria (2018) who also establish a correlation between the distribution of trade and search frictions. The objects of interest in the two papers are however quite different. We study the distribution of bilateral trade along the distribution of heterogeneous producers, which we map to the firm-to-firm trade data. Dasgupta and Mondria (2018) instead study how imports are allocated among various source countries when importers face information frictions, which they map to bilateral trade flows. The nature of frictions also differs between their paper and ours. In Dasgupta and Mondria (2018), uncertainty is about the distribution of prices across origin countries and frictions benefit low expected costs countries because importers put more weight on priors when information is more expensive. In our model, buyers have full information about seller prices, conditional on meeting them. High search costs benefit high-cost sellers because they insulate them from competition.

A consequence of the non-monotonic impact of frictions along the productivity distribution is that the export premium of high-productivity firms is affected by the level of frictions:

$$\ln \frac{\rho_{ij}(\bar{z})}{\rho_{ij}(\underline{z})} = (c_j d_{ij})^{\theta} \Upsilon_i \kappa_i \left(\underline{z}^{-\theta} - \bar{z}^{-\theta} \right), \tag{6}$$

where $\rho_{ij}(\bar{z})$ and $\rho_{ij}(\underline{z})$ denote export probabilities in country i of a firm from j with a high-productivity $\bar{z}$ and a low-productivity $\underline{z}$, 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 increasing in κ_i, which is

[24] Although the analytical results crucially rely on the size of the visibility channel being independent of firms' productivity, we argue in Appendix A.2 that the result is more general. 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. Such systematic correlation between the firm's productivity and its chance to meet with a foreign buyer does 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.

consistent with the idea that more distortive search frictions reduce the competitive advantage of high-productivity firms. In markets displaying high and distortive search frictions, buyers meet with a small number of relatively low competitive firms, on average. As a consequence, the strength of competition is reduced and the export premium of high-productivity firms is smaller.

Whereas the export premium of high-productivity firms is reduced in highly frictional countries, it is exacerbated in countries featuring high iceberg trade costs; that is,

$$\frac{d \ln \frac{\rho_{ij}(\bar{z})}{\rho_{ij}(\underline{z})}}{d \ln d_{ij}} > 0.$$

The reason is that a decrease in iceberg costs improves the relative competitiveness of all French firms - but the gain in competitiveness is stronger for low-productivity ones. Low-productivity firms are thus more likely to serve the buyer, conditional on a match. The heterogeneity in export performances across firms is thus an interesting moment to exploit for the identification of search frictions.

Because 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:

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

where $B_{ij}(z_{s_j})$ 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 that is consistent with evidence in Figure 1. In our framework, this relationship comes from more productive sellers being more likely to be chosen by any buyer.

3.2.3 Discussion

Extending our model to general equilibrium should be possible by assuming that there is a continuum of individual products in the economy, so that the randomness that we emphasize at the product-level does not have aggregate consequences (see Gaubert and Itskhoki, 2018). There are two reasons why one may want to account for general equilibrium effects (GE hereafter). First, GE effects would allow us to account for the potential feedback effect of a drop in trade frictions on relative prices. In the context of our model, such aggregate price adjustments in turn affect the selection of firms into exporting, through domestic firms' relative competitiveness abroad. Importantly, such adjustment would not affect the strength of competition among domestic firms, which we argue is a key component explaining the non-monotonic impact of frictions on domestic firms. Moreover, these price effects are likely to be small quantitatively in our counterfactual experiment that focuses on change in bilateral frictions in a (small) destination. Second, GE effects could help us discuss the welfare impact of trade frictions on foreign consumers. However, to make a quantitative evaluation of such effects, one would need information on bilateral frictions faced by non-French exporters to this destination, which cannot be computed from our dataset. For these reasons we focus on the reallocation of exports across French firms and abstract from GE forces.

Another natural question is the extent to which the theoretical predictions later used to identify search frictions could be rationalized differently under alternative assumptions. We discuss this point into more details in appendix A.3, in the context of a Melitz-type model. We sketch a partial equilibrium version of a monopolistic competition model that introduces market penetration costs *à la* Arkolakis (2010) in the discrete version of the Melitz model proposed by Eaton et al. (2012). We then study how variations in both variable and fixed trade costs affect the number of buyers served by each exporter in such framework. We show that low-productivity firms benefit more from a decrease in variable and fixed costs than high-productivity ones. This is in contrast with what happens when search frictions decrease in our model. Interestingly, we can also prove that the theoretical moment used in the estimation is independent from all trade costs in this context.

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 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 $\rho_{ij}(z_{s_j})$:

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

Integrating over the expected distribution of productivities gives the expected number of firms from j with exactly $M > 0$ buyers in i (see details in Appendix A.4):[25]

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

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

Equation (7) shows the expected number of firms serving a given number of clients is decreasing in M, which is consistent with evidence in section 2.2. This property 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. Because our dataset only covers exporters located in France, the j country will always be France, and we use the heterogeneity across destinations and sectors to recover a distribution of estimated parameters.

Once normalized by the expected number 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, which 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.[26] The correlation be-

[25]Integrating over the expected distribution of productivities amounts to neglecting additional distortions induced by the assumption of a discrete number of French suppliers. With a discrete and finite number of French suppliers, the ex-ante Pareto distribution of productivities does not exactly coincide with the ex-post distribution of productivities. We neglect this discrepancy and derive a distribution of the number of buyers per firm, whose shape solely depends on search frictions. This assumption is innocuous as long as the number of potential suppliers of the product is large enough, which is the case in practice in the data.

[26]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.

tween the number of firms with exactly M buyers in a destination and distance to the destination is negative and strongly significant. This finding is consistent with evidence in section 2.2 that French sellers tend to serve fewer partners, if any at all, in more distant countries. This result should be expected from the model, as the π_{ij} component in equation (7) is negatively correlated with iceberg trade costs d_{ij}, which are likely to be increasing in distance. In principle, the correlation can be controlled for using readily available data for those trade shares.

Another option is to normalize the expected number of firms with M buyers with the destination-specific proportion of sellers with one buyer, i.e. compute the theoretical moment $\frac{h_{ij}(M)}{h_{ij}(1)}$ and compare it with its empirical counterpart. In theory, this convergent moment is useful to identify search frictions as it varies monotonically with λ_{ij} (see Figure A.3 in Appendix). Moreover, several ratios can be combined to identify precisely search frictions along a wide range of possible values. Unfortunately, the corresponding empirical moments are still correlated with distance, which the model does not explain (see the second panel of Table 2). In principle, the normalization should neutralize the impact of trade shares, and thus of iceberg trade costs. A correlation between search frictions and distance may explain this result. However, iceberg trade costs may also affect the ratios through other channels, which the model does not encompass but the data reveal. To prevent such correlation from polluting our estimates of search frictions, we use an alternative moment that 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 1 (left panel). As illustrated in the simulations reported in Figure 2, this moment is also correlated positively with λ_{ij} and is thus useful for identification. Intuitively, fewer frictions reduce the expected number of exporters serving a small number of buyers while increasing the density at high values of M, thus increasing the variance in equation (8). As shown in the third panel of Table 2, the empirical

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

Dependent Variable	log Distance	Std Dev.	Adjusted R-squared
# sellers with:			
1 buyer	-16.04***	(1.48)	.697
2 buyers	-5.87***	(.559)	.534
3 buyers	-3.23***	(.363)	.416
4 buyers	-2.00***	(.253)	.333
# sellers (in relative terms with respect to the sellers with 1 buyer) with:			
2 buyers	.020***	(.008)	.343
3-4 buyers	-.027***	(.008)	.374
5+ buyers	-.123***	(.021)	.410
Variance of the relative shares of sellers:			
across M	.001	(.010)	.211
across M, controlling for migrants	-.016	(.014)	.212
	coef. on migrants: -.008** (.003)		

Notes: Robust standard errors, clustered at the country level, are 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.

counterpart of this moment is not correlated with distance. On the other hand, it is negatively correlated with the stock of French migrants in the destination, our proxy for information frictions. Finally, we discuss in appendix A.3 whether this moment could capture other forms of trade frictions in the context of a Melitz-type model. In this model, the dispersion across sellers in the number of buyers served reflects the shape of the underlying productivity distribution. The ex-post correlation of our estimates with external proxies for search frictions corroborates our interpretation of the data.

In theory, the dispersion can be calculated across $B_i - 1$ ratios. However, these ratios do not convey a lot of relevant information, because they are almost all equal to 0 in the data, above a certain level of M.[27] 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. For consistency, the moments in Figure 2 use the same

[27]As shown in Figure 1 (left panel), most of the variance in the number of buyers served by French exporters is indeed found at values for $B_{ij}(z_{s_j})$ below 10. Using all the individual moments regarding the number of firms with $B_{ij}(z_{s_j}) > 10$ clients would thus be inefficient and would artificially reduce the dispersion in the data, in a way that is not independent from B_i.

Figure 2: ***Correlation between the variance of the*** $h(M)/h(1)$ ***ratios and the value of search frictions***

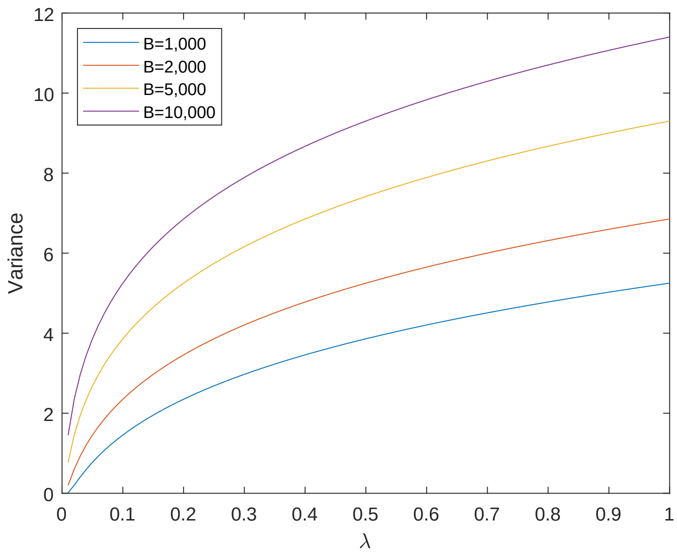

Notes: This figure shows the theoretical relationship between the underlying value of search frictions (λ, x-axis) and the variance of the $h(M)/h(1)$ ratios, i.e. the theoretical moment used to identify search frictions. The relationship is derived conditional on the underlying number of buyers (B) and using three ratios, namely $\frac{h(2)}{h(1)}$, $\frac{h(3)+h(4)}{h(1)}$ and $\frac{\sum_{m=5}^{B} h(m)}{h(1)}$.

convention, although we have checked that the variance recovered from all possible ratios is also increasing in λ_{ij}.

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:

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

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

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 $\widehat{Var_{ij}}$.[28] Using the convergence result, identifying λ_{ij} uniquely is possible. With an asymptotic least squares estimation strategy, the estimated variance of estimated frictions writes

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

with $\Omega_{ij}(\lambda_{ij})$ the optimal matrix of weights defined in Appendix B.1.

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 that 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 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.[29] B_{ij} is directly observed in our data. Based on this, one can recover a value of B_i for each destination and sector.[30]

[28] $\Omega_{ij} = \nabla g(\lambda_{ij})\Sigma_{ij}\nabla' g(\lambda_{ij})$, 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 $M = m_1, m_2, m_3$.

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

[30] 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 issue, we winsorized the number of potential buyers at 20,000, i.e., $B_i = min\left\{20,000; \frac{B_{iF}}{\pi_{iF}}\right\}$. This constraint is binding for 13% of product×country pairs.

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 for the number of firms susceptible to 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 recover estimated values for the meeting probabilities. Because the minimization program 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.

4.3 Results

Summary statistics. Search frictions are estimated at the (product×country) level for a total of 10,427 λ_{ij} parameters, among which 10,402 are statistically significant. To get meaningful comparisons, we restrict our analysis to countries where we have at least 200 estimated parameters. With this restriction, we keep 9,855 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 sellers from country j that a given buyer in country i would meet, on average. We see an important level of dispersion in these probabilities. Indeed, 10% of product-country pairs have a meeting probability below .01%, wherease 10% have a meeting probability above 1.7%. A basic variance decomposition exercise shows 13% of the dispersion in our friction parameters is driven by the destination-country dimension, 43% is product-specific, and the remaining 45% is within a product×country.

Determinants of search frictions. 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 product fixed effects. In column (3), we remove the product fixed effects to include a measure of product differentiation. In column (4), we focus on the role of product characteristics and thus control for country fixed effects. The results show market size (measured by population) and physical distance

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

	Meeting Probability λ_{ij} (en %)	Probability of Meeting 0 Buyer $(1-\lambda_{ij})^{B_i}$ (en %)
Mean	0.76	12.0
Percentile 10	0.01	0.00
Percentile 25	0.06	0.00
Percentile 50	0.21	0.02
Percentile 75	0.69	4.37
Percentile 90	1.79	56.74
# Observations	9,855	9,855

Notes: The first column in this table presents summary statistics on the λ_{ij} coefficients, estimated by country $\times$ 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.

are positively correlated with frictions. The positive correlation between frictions and market size suggests the search process is easier when economic activity is spatially concentrated. Spillovers in the search process might explain this correlation.[31] Whereas, search frictions are higher in large markets, the probability of meeting a buyer increases with market size because large markets are populated with more buyers. Note 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 with transportation costs. Our findings show distance further affects trade flows by impeding the search process between buyers and foreign sellers. As expected, search frictions are found to be lower in countries where French migrants are more numerous (though the effect is not always significant at conventional levels). This finding is consistent with the view that migrants convey information on their origin country, thus reducing information frictions. Finally, the results show search frictions are higher for more differentiated products (according to the Rauch classification). This finding is consistent with the view that the search process

[31]The smaller probability of meeting in large market is also consistent with the existence of congestion frictions (see, eg. Eaton et al., 2018).

is easier for products traded in organized markets.

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

	(1)	(2)	(3)	(4)
		Dep. Var: $\ln(\lambda_{ij})$		
ln distance		-0.527*	-0.809***	
		(0.248)	(0.190)	
ln population	-0.478***	-0.508***	-0.208***	
	(0.071)	(0.056)	(0.046)	
ln French migrants	0.179*	0.072	0.218***	
	(0.088)	(0.076)	(0.057)	
Rauch dif.			-0.382***	-0.384***
			(0.080)	(0.081)
Fixed Effects	Product	Product	No	Country
Observations	9,855	9,855	9,855	9,855
R-squared	0.580	0.585	0.103	0.115

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

Search frictions across countries. Next, we investigate what these estimated matching probabilities imply in terms of the probability that a given French exporter meets with zero buyers in each destination, which is positively linked to the extent of frictions. Because the meeting process is a binomial, this probability is equal to $(1-\lambda_{ij})^{B_i}$, with B_i being 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 buyers in a destination is 12%. This number, however, hides a lot of heterogeneity. In more than 50% of country and sector pairs, the probability is below 1%. At the other side of the distribution, 10% of country$\times$sector pairs display high frictions, with French exporters having more than a 56% chance of meeting with no buyer there. Figure 3 compares these probabilities, on average across destinations.[32] 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, Greece,

[32]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 we thus chose to present this measure in relative terms with respect to Germany.

Figure 3: ***Comparison of frictions that French exporters face across countries***

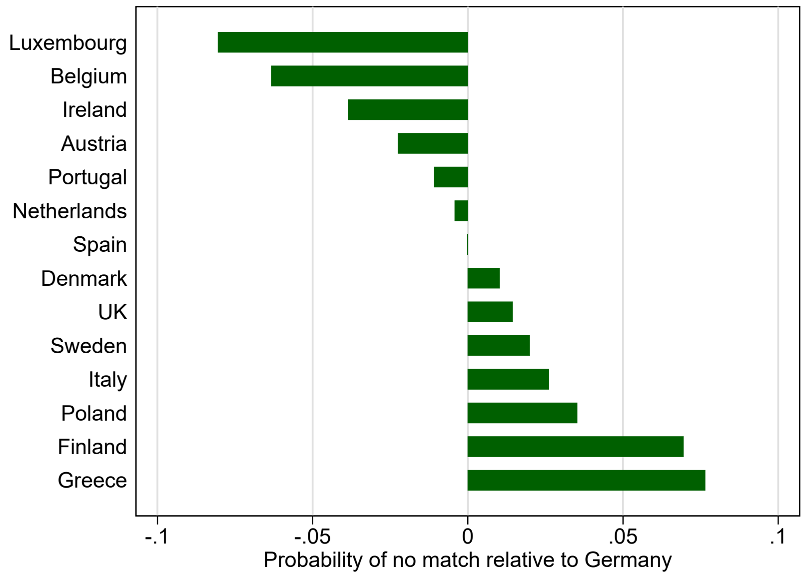

Notes: Mean probability of meeting with zero buyers across countries, in relative terms with respect to Germany.

Finland, and Poland, three countries which are relatively distant from France along several metrics, display the highest no match probabilities, on average.

Search frictions and export-promoting policies. We now discuss another piece of external validity that relies on data provided to us by Business France, the French export-promoting agency. Among the policies that France implements to support export activities, Business France has a program that is particularly well suited to address information frictions. Namely, they finance for some selected firms their participation in foreign trade fairs. The announced objective of such a program is to help them "gain visibility and meet buyers at an event abroad." Business France provided to us a list of firms and the location of the trade fair financed over 2005-2007. We use these data to compute a measure of the intensity of export- promoting activities by country. The measure is computed as the ratio of the share of firms supported in a destination (relative to the share of firms supported in EU15 destinations) relative to the share of that destination in French total exports to the EU15. The higher the ratio, the higher the intensity of the effort exerted by Business France in a market. To the extent that Business France has some privileged knowledge on the extent of information frictions in

Figure 4: ***Correlation of search frictions with the export-promoting policies***

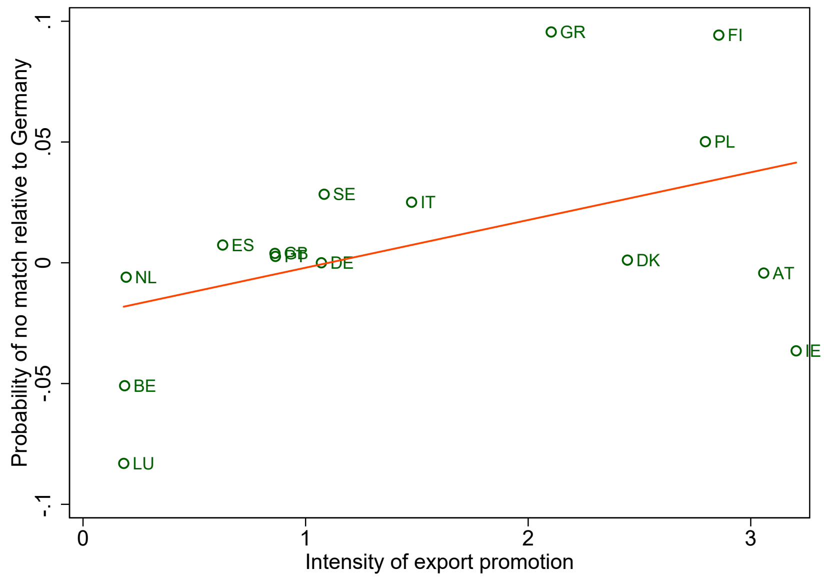

Notes: The graph is a scatter plot of a measure of the activities of the French export-promoting agency against the log of estimated search frictions (averaged across products). The measure of export-promoting activities is computed as the share of French firms supported through a fair trade in a given destination relative to the share of that destination in French exports. The higher the ratio, the higher the intensity of bilateral export-promoting activities.

various destinations, this intensity should be positively correlated with the magnitude of estimated frictions.[33] Figure 4 plots this measure against the country-level estimated probability of no match presented in Figure 3. Efforts of the export-promoting agency are higher in markets that display higher search frictions. This finding suggests that, on average, the agency focuses on the "right" markets by this metrics. The fit is not perfect though, and more effort could seemingly be directed to Greece or Finland, and less toward Ireland or Austria.

Test of empirical predictions. Another way to assess the validity of our estimates is to confront the model's predictions to the data. Proposition 1 unambiguously shows an increase in bilateral search frictions within a product category between France and

[33] One may worry that the correlation could be reversed if the policy was sufficiently effective so that targeted markets ended up less frictional. We argue that this is unlikely to be the case because the program is relatively young and modest in size. Over 2005-2007, 2,302 firms benefited from such help to meet with potential partners in the EU15.

a trade partner should lead to a reduction in French market shares. We thus regress the logarithm of French market shares (computed by destination-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 and the share of French migrants in the destination. As expected, bilateral distance is an impediment to French exports, whereas 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 estimated search frictions. The results in columns (2) to (4) show French market shares are lower for product-destination pairs that exhibit a higher level of search frictions. This finding 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 percentage is sizeable.

Model Fit. Having shown our estimates of search frictions correlate with observables in a theory-consistent way, we now evaluate the model's ability to reproduce key features of the data. We use our parameter estimates to simulate the expected number of sellers interacting with 0 to 10 buyers within a destination market. We then simulate the cumulated distribution of sellers' number of buyers in a destination, and compare it with the data.[34] Figure 5 reports the observed and predicted CDFs for the 15 countries in our sample. A visual inspection shows 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 expected number of sellers serving one client in our set of moments. Interestingly, our simple model captures quite well the share of sellers serving a single buyer within a destination, that is, the fit is good regarding the

[34]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.

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

	(1)	(2)	(3)	(4)
	Dep. Var: ln French Market Share			
ln distance	-0.749***	-0.711***		-0.736***
	(0.221)	(0.177)		(0.221)
ln French migrants	0.365***	0.262***		0.367***
	(0.063)	(0.049)		(0.064)
ln Meeting proba		0.255***	0.352***	
		(0.017)	(0.056)	
Proba no meeting				-0.008***
				(0.002)
Fixed Effects	Product	Product	Product	Product
Observations	9,855	9,855	9,855	9,808
R-squared	0.659	0.745	0.610	0.662

Notes: The dependent variable is the log of France's 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, and 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. It 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, are in parentheses with ***, **, and * respectively denoting significance at the 1%, 5% and 10% levels.

curvature of CDFs *and* their intercept.[35] Although 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. 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, whereas column (3) has country and product fixed effects. The R^2 of the first regression is .19, suggesting 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 the correlation is valid within countries across products as well as across products within countries.[36]

[35]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.

[36]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.

Figure 5: ***Model fit: Distribution of sellers' degrees***

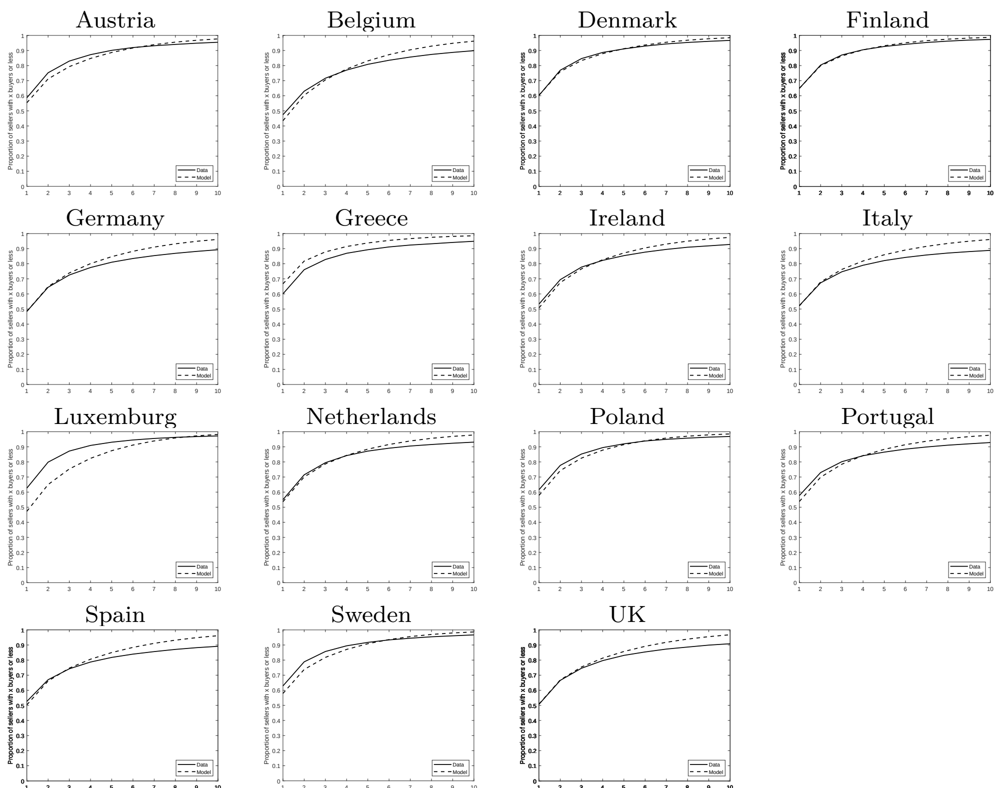

Notes: Observed and predicted CDF of sellers' numbers of buyers, by country. Predicted CDF are obtained using the model's definition of $h_{ij}(M)$, at the country and product level, before aggregating across products using information on the relative number of producers of each good in France.

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

	Dep.Var.: Empirical share of one buyer		
	(1)	(2)	(3)
Predicted share	.295***	.276***	.177***
	(.006)	(.005)	(.005)
Constant	.391***		
	(.003)		
# obs	10,427	10,427	10,059
Fixed Effects	No	Country	Country Product
R-squared	.194	.276	.579

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 are in parentheses, with *** denoting significance at the 1% level.

5 Implications for the allocative efficiency

Having shown our methodology delivers convincing estimates of bilateral search frictions French exporters face, we now ask what such frictions imply for the efficiency of selection into exporting. The analysis proceeds in two steps. We first study the correlation of frictions with Ricardian comparative advantages, which the theoretical analysis has shown matters for the level of distortion induced by frictions. In a second step, we quantify the efficiency loss using a counterfactual analysis.

5.1 Search frictions and Ricardian comparative advantage

As explained in section 3, the strength of search-induced distortions depends on how they correlate with comparative advantages. Intuitively, search frictions are all the more distortive if they hit firms that would, on average, display strong comparative advantages in the frictionless economy. We now investigate whether this is the case in the data, using cross-sectoral measures of revealed comparative advantages and the dispersion in estimated frictions, across products.

Revealed comparative advantages are measured using a strategy inspired from Costinot et al. (2012). Exploiting the gravity structure of the model, equation (2) can be used to recover a statistical decomposition of bilateral exports into its different variance components:

$$\ln \pi_{ijk} = FE_{ik} + FE_{jk} + FE_{ij} + \varepsilon_{ijk}, \tag{11}$$

where we now explicitly introduce the product dimension k. π_{ijk} thus measures the share of producers from country j in country i's consumption of product k.

The exporter-product fixed effect FE_{jk} in equation (11) absorbs the impact of Ricardian technological advantages that affect a country's sales in all export destinations (i.e., $\ln T_{jk} c_{jk}^{-\theta_k}$ using the model's notations). A negative correlation between revealed comparative advantages and estimated search frictions would thus indicate search frictions are distortive, in the sense that they are high in those product markets in which French exporters have a Ricardian comparative advantage.[37] We test this theory in the

[37] In such statistical decomposition, the fixed effect also captures all components of trade frictions that are common across destinations (i.e. $\ln \frac{1}{N} \sum_i \lambda_{ijk} d_{ijk}^{-\theta_k}$). Such components can mechanically create a positive correlation with estimated search frictions, as our strategy allows to recover the absolute level of λ_{ijk}, that can eventually have a common component across destination markets. A positive correlation between the estimated exporter-product fixed effect and estimated search frictions would thus be difficult to interpret. Our results suggest the correlation is negative, which in the model needs to come from a negative covariance of $T_{jk} c_{jk}^{-\theta_k}$ and λ_{ijk}.

Figure 6: ***Correlation of search frictions with comparative advantages***

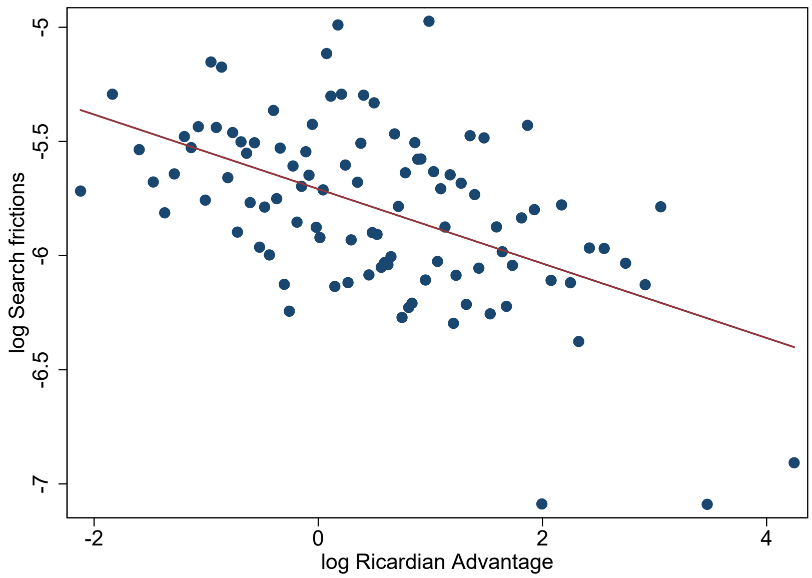

Notes: The graph is a binned scatter plot of the log of revealed comparative advantages measured for each hs6 product, using equation (11) against the log of estimated search frictions (averaged across destinations).

data by first estimating revealed comparative advantages using equation (11) and the CEPII-BACI multilateral trade database available at the product level for 2007.

Results suggest a strong, negative correlation between revealed comparative advantages and search-friction parameters λ_{ijk}. The unconditional correlation (in logs) is equal to -.10, and the correlation across products within a destination is equal to -.15 and is significant at the 1% level. The correlation is illustrated in Figure 6, which shows how average frictions per product correlate with France's comparative advantages recovered from equation (11). In the figure, the high level of heterogeneity is smoothed by grouping the two measures into bins. The negative correlation is consistent with search frictions faced by French firms in Europe being distortive, because they penalize more those sectors in which French firms have a comparative advantage.

5.2 Quantifying the efficiency loss induced by search frictions

Results illustrated in Figure 6 indicate search frictions are distortive, on average. Although the result is qualitatively interesting, it says little about the quantitative impact of these distortions. To recover such numbers, we now turn to a counterfactual analysis. Namely, we simulate a drop in frictions that French exporters face, keeping parameters governing the rest of the world competitiveness unchanged. We then compute the impact of the reduction in bilateral frictions on the mean productivity of exports.

5.2.1 Methodology

Throughout the exercise, we focus on the Greek market, identified as the most frictional country in our data, on average. Using this benchmark, we simulate how French exporters' behaviors would adjust if the level of bilateral frictions decreased in this destination, to the mean observed in the less frictional country in our sample, which is Belgium.[38] In practice, this means 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 $\lambda^c_{ij} = \hat{\lambda}_{ij} \times 4.5$, where 4.5 is the mean ratio of search frictions estimated for Belgium and Greece, conditional on product characteristics).

Reducing search frictions without moving any other parameter of the model has a positive impact on French firms' expected market share in Greece that can be computed using the result in proposition 1. Under our counterfactual calibration, the median (resp. maximum) gain in market shares is relatively small, +1.14 (resp. +14) percentage points. Gains in market shares are not the most important criteria to evaluate the benefit of reducing search frictions, though. On top of increasing the overall competitiveness of French exporters in the destination, such a policy has additional consequences on the allocation of resources across exporting firms.

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 a Pareto distribution of productivity, $\left(\frac{z}{z_{min}}\right)^{-\theta}$ is the expected share of firms with productivity above z. The estimated value of λ_{ij} is taken as a benchmark, and

[38]Luxembourg is actually found to be slightly less frictional than Belgium, on average, in Figure 3, but the number of products underlying this average is lower than for Belgium. Moreover, Figure 5 shows the model fit for Luxembourg is not as good.

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 expected number of potential suppliers in country j (France in our experiment). We calibrate this object to fit the data regarding large firms' export premium in any given product market:[39]

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

Given observed λ_{ij} and π_{ij}, we can calibrate $T_j z_{min}^{-\theta}$ 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 to the former is our measure of the product- and destination-specific export premium.[40] 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 to be invariant to the counterfactual shift in search frictions.

Armed with the calibrated expected number 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 $(B_i \rho_{ij}(z))$, for each percentile.

[39] 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.

[40] 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. Because the model is not consistent with a negative export premium, we have no choice but to discard the corresponding products from the counterfactual analysis.

5.2.2 Results

Figure 7, 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 up along 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 68% among the 1% most 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; that is, some low-productivity firms are evicted from the Greek market while higher-productivity firms enter. Less than 40% of firms benefit from the reduction in frictions in this experiment. The winners are firms at the top of the productivity distribution whereas export probabilities decrease for firms in the first 6 deciles of the productivity distribution. Because of this asymmetry, the mean productivity of exporters improves significantly, by about 10%.[41]

The right-hand side of Figure 7 illustrates how the reduction in frictions further affects the allocation of resources, by reallocating market shares across exporting firms, at the intensive (buyer) margin. In the data as in the counterfactual, the expected number of partners for firms at the bottom of the distribution is equal to one, because exporting at that level of productivity implies being lucky enough to meet with a buyer that has not met any more competitive firm, which is unlikely to happen. In this graph, the difference between the data and the counterfactual mostly appears at the

[41]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 11.42%, a value that is reduced to 6.63% for $\theta = 5$.

Figure 7: ***Probability of exporting to Greece and expected number of buyers conditional on export, along the productivity distribution: Actual versus counterfactual***

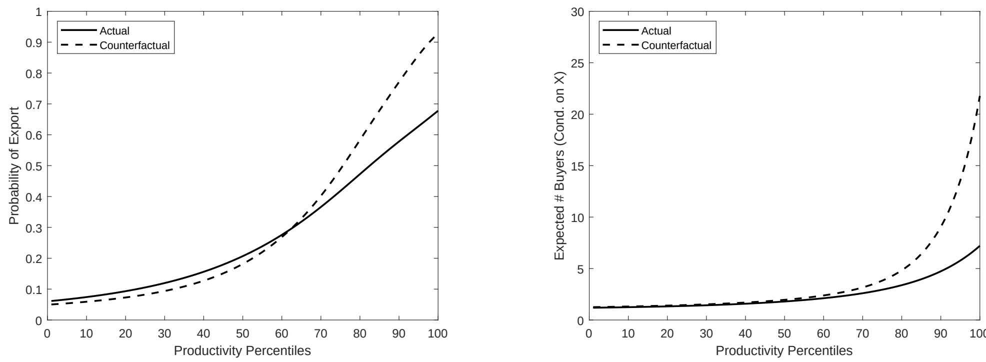

Notes: The graphs plot the probability of export to Greece (left panel) and the expected number of partners, conditional on exporting (right panel), conditional on the firm's position in the productivity distribution. The solid lines correspond to the actual equilibrium, and the dotted lines are the counterfactual. Export probabilities and the expected number of exporters are both calculated at the product level following the strategy described in section 5.2.1, before being aggregated across products using information on the relative number of firms in each product market.

right tail of the distribution as the expected number of clients of high productivity firms is substantially higher in the less frictional world. For the mean exporter at the 75^{th} percentile of its sector's productivity distribution, the expected number of partners increases from 2.9 to 3.8. At the 90^{th} percentile, the effect is more pronounced, with the expected number of clients shifting from 4.7 to 9.0. Finally, in the last percentile, the impact is substantial, with the expected number of clients increasing from 6.7 to 19.7.

Taken together, these two effects lead to an increase in high-productivity firms' export premium in the counterfactual equilibrium. In the data, firms at the 90^{th} percentile export 43 times more than firms at the 25^{th} percentile, in expectation in the median product market (the inter-quartile range being [15 144]). In the counterfactual equilibrium, export premia are an order of magnitude larger, reaching 901 for the median product and an inter-quartile range of [144 9,155]. The reason the effect is massive is that many firms in the 25^{th} percentile no longer exports in the counterfactual, as can be seen in Figure 7, 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, raising for high-productivity firms, in comparison with less productive exporters (Figure 7, right panel).

For comparison purposes, we ran another counterfactual exercise in which iceberg costs, instead of search frictions, were reduced product-by-product in the Greek market, keeping everything else unchanged. Because both parameters are not directly comparable, the counterfactual is calibrated such that the change in product-level trade shares is the same as in the main counterfactual experiment just described. Moving from the actual to this counterfactual equilibrium induces a substantial increase in export probabilities for French firms. In the aggregate, the export probability increases from 27.35% to 47.39%. However, this increase in export probabilities does not induce an efficiency gain as in the case of a reduction in search frictions. Actually, export probabilities increase the most at the bottom of the distribution of productivities, meaning the drop in iceberg costs benefits, in relative terms, low-productivity firms. The reason is that decreased iceberg costs push down the relative price offered by French relative to other countries' firms, thus increasing the likelihood of being the lowest-cost supplier conditional on a match. This competitiveness gain over non-French exporters benefits low-productive firms more, that is, those suffering the most from a lack of competitiveness. As a consequence, the mean productivity of exporters decreases in this counterfactual experiment, by 8% to 13%.

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. Such frictions thus benefit low-productivity firms, whereas they reduce the export probability and expected exports at the top of the distribution.

6 Conclusion

This paper shows how search frictions in international goods 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 number of potential suppliers of a perfectly substitutable good. The model combines two barriers to international trade. Physical (iceberg) trade costs reduce the competitiveness of exporters in foreign markets, in a way that is homogeneous 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. Although 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. We argue this heterogeneity is the key reason 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, 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 to be more severe in large and distant countries and for products that are more differentiated.

The estimated frictions are distortive. They are especially large in product markets where French firms have a comparative advantage, on average. A counterfactual analysis allows quantification of the size of the distorsion. 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. Therefore, 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.

Acknowledgments. The paper has benefited from insightful comments by participants at the seminars at the Graduate Institute in Geneva, Bocconi University, 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, Adrien Bilal, Johannes Boehm, Thomas Chaney, Edouard Chretien, Arnaud Costinot, Gregory Corcos, Jonathan Eaton, François Fontaine, Etienne Lalé, Sam Kortum, Francis Kramarz, Kalina Manova, Thierry Mayer, Andreas Moxnes, Mathieu Parenti, Gianmarco Ottaviano, Sophie Osotimehin, Frédéric Robert-Nicoud, 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).

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., 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.

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

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.

Costinot, A., Donaldson, D., and Komunjer, I. (2012). What Goods Do Countries Trade? A Quantitative Exploration of Ricardo's Ideas. *Review of Economic Studies*, 79(2):581–608.

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

David, J. M., Hopenhayn, H. A., and Venkateswaran, V. (2016). Information, Misallocation, and Aggregate Productivity. *The Quarterly Journal of Economics*, 131(2):943–1005.

Dickstein, M. J. and Morales, E. (2018). What do Exporters Know? *The Quarterly Journal of Economics*, 133(4):1753–1801.

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

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.

Garicano, L., Lelarge, C., and Reenen, J. V. (2016). Firm Size Distortions and the Productivity Distribution: Evidence from France. *American Economic Review*, 106(11):3439–3479.

Gaubert, C. and Itskhoki, O. (2018). Granular comparative advantage. NBER Working Papers No. 24807.

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.

Gourio, F. and Rudanko, L. (2014). Customer Capital. *Review of Economic Studies*, 81(3):1102–1136.

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.

Hopenhayn, H. A. (2014). Firms, Misallocation, and Aggregate Productivity: A Review. *Annual Review of Economics*, 6(1):735–770.

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

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.

Luttmer, E. G. J. (2006). Consumer search and firm growth. Working Papers 645, Federal Reserve Bank of Minneapolis.

Macchiavello, R. and Morjaria, A. (2015). The value of relationships: Evidence from a supply shock to kenyan rose exports. *American Economic Review*, 105(9):2911–45.

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

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

Midrigan, V. and Xu, D. Y. (2014). Finance and misallocation: Evidence from plant-level data. *The American Economic Review*, 104(2):422–458.

Perla, J. (2019). A model of product awareness and industry life cycles. Working paper.

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.

A Appendix: Proof of analytical results

A.1 Proof of proposition 1

Start with the definition of trade shares:

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

implying

$$\frac{d \ln \pi_{ij}}{d\lambda_{ij}} = \frac{1}{\lambda_{ij}} - \frac{1}{\kappa_i} \frac{d\kappa_i}{d\lambda_{ij}}.$$

Using

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

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

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

A.2 Proof of proposition 2

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

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

Depending on the current level of frictions (λ_{ij}), the expected number of firms in country j ($T_j z_{min}^{-\theta}$) and the position of the firm in the productivity distribution ($\left(\frac{z}{z_{min}}\right)^{-\theta}$), the derivative can be positive or negative. It is more positive for high values of z. At the limit, $\lim_{z \to +\infty} \frac{\partial \ln \rho_{ij}(z)}{\partial \lambda_{ij}} = \frac{1}{\lambda_{ij}}$. 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 $\lambda_{ij} > \frac{1}{z_{min}^{-\theta} T_j}$, that is, if frictions are not too strong so that buyers in expectation meet with at least one seller from j, a strictly positive mass of firms exists whose export probability decreases when search frictions are reduced: $\dfrac{\partial \ln \rho_{ij}(z_{min})}{\partial \lambda_{ij}} < 0$, where $\rho_{ij}(z_{min})$ 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:

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

These contrasted results are the key reason 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. By 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. Although functional forms obviously matter to obtain the analytical predictions, we argue this result applies more generally whenever

$$
\frac{d^2 \rho_{ij}(z)}{d\lambda_{ij} dz} > 0 \quad \text{and} \quad \frac{d^2 \rho_{ij}(z)}{d d_{ij} dz} > 0.
$$

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 $\lambda_{ij}(z) = f(\lambda_{ij}, z)$ with $\frac{df(\lambda_{ij},z)}{d\lambda_{ij}} > 0$ and $\frac{df(\lambda_{ij},z)}{dz} > 0$; that is, high-productivity firms meet with more buyers. In such a model

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

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 because a reduction in frictions implies 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; that is, $\frac{d^2 f(\lambda_{ij}, z_{s_j})}{d\lambda_{ij} dz_{s_j}}$ 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.3 A model of buyer acquisition under monopolistic competition

A natural question is the extent to which the moment exploited in our estimation could capture the impact of other determinants of trade, in an alternative model. Whereas our model is Ricardian in nature, an alternative interpretation of the buyer margin can be done in the context of an imperfect competition model *à la* Melitz (2003), as notably done by Bernard et al. (2018b); Carballo et al. (2018). In this section, we develop such a model using a structure and notations comparable to those used in our model to ease the comparison. The model introduces market penetration costs *à la* Arkolakis (2010) in the discrete version of the Melitz model proposed by Eaton et al. (2012). As in the paper's model, we abstract from any general equilibrium effects.

We start with the supply side structure used in our model, that features a discrete and random number of producers that are heterogeneous in their productivity. Remember that under our assumptions, borrowed from Eaton et al. (2012), the number of sellers from j that display a productivity above s is the realization of a Poisson variable with parameter $T_j z^{-\theta}$. Given exogenous input costs c_j and iceberg costs d_{ij} the number of firms serving market i at a cost below c is itself a Poisson variable of parameter $\mu_{ij}(c) = T_j \left(\frac{d_{ij} c_j}{c}\right)^{-\theta}$.

In the Ricardian framework, worldwide firms compete to serve market i with the same perfectly substitutable variety, which triggers prices towards marginal costs.[42] In the monopolistic competition variant, we instead follow Eaton et al. (2012), and assume

[42] Without search frictions, the outcome of such competition are prices exactly equal to the marginal cost of production. In a frictional equilibrium, this may not be the case since sellers and buyers could negotiate ex post. If prices are not pre-set, the outcome should be a price that is above the marginal cost but below the monopolistic competition price. We expect it to depend on the seller's bargaining power, a function of how low its marginal cost is, in comparison with other sellers that the buyer has met. Since our model's predictions do not rely on variables that are affected by prices, we do not derive equilibrium prices in the main text.

that each seller offers a differentiated variety and faces a demand which is isoelastic. Equilibrium prices are then a constant mark-up over marginal costs:

$$p_{ij}(z_{s_j}) = \frac{\sigma}{\sigma - 1}\frac{d_{ij}c_j}{z_{s_j}}$$

$p_{ij}(z_{s_j})$ is the price set by s_j in country i, which is uniform across buyers within a destination if the residual demand elasticity is itself homogeneous. We assume this is the case and denote $\sigma > 1$ this elasticity.

In Eaton et al. (2012), sellers face a representative consumer in each market i and decide whether to serve the market or not, depending on the size of some fixed export cost F_{ij}. To introduce the buyer margin, we instead assume that i) sellers can serve a discrete number B_i of homogeneous buyers in the destination and ii) the fixed cost of exporting is increasing in the number of buyers served. Namely, the residual real demand expressed by buyer b_i is assumed to be:

$$q_{b_i}(p) = p^{-\sigma}\tilde{q}_i.$$

It is thus homogeneous, decreasing in prices and shifted up and down by some real demand shifter $\tilde{q}_i$.[43] On top of iceberg costs, firms are also assumed to pay a fixed cost for exporting, which is increasing in the share of the market served:

$$F_{ij}\left(B_{ij}(z_{s_j})\right) = F_{ij} \times \frac{1 - \left(1 - \frac{B_{ij}(z_{s_j})}{B_i}\right)^{1-1/\lambda}}{1 - 1/\lambda}$$

where F_{ij} is a positive parameter and $\lambda > 0$ measures the increasing cost of reaching a larger fraction of potential buyers.

Solving for the seller's optimal number of buyers served implies:

$$\frac{B_{ij}(z_{s_j})}{B_i} = Max\left\{0; 1 - \left(\frac{p_{ij}(z_{s_j})^{1-\sigma}}{\sigma}\frac{B_i\tilde{q}_i}{F_{ij}}\right)^{-\lambda}\right\}$$

[43] In general equilibrium, $\tilde{q}_i$ would be a function of the CES ideal price index faced by the buyer, which could potentially be heterogeneous across buyers due to differences in the variety of goods they have access to. In Arkolakis (2010) and the literature that followed, this possibility is ruled out using a law of large numbers argument. We implicitly rely on the same argument to simplify the analysis.

From this, it comes:

$$\begin{aligned}
\frac{\partial \ln B_{ij}(z_{s_j})}{\partial z_{s_j}} &= \lambda \left[1 - \frac{B_{ij}(z_{s_j})}{B_i}\right] \frac{\sigma - 1}{z_{s_j}} > 0 \\
\frac{\partial \ln B_{ij}(z_{s_j})}{\partial d_{ij}} &= \lambda \left[1 - \frac{B_{ij}(z_{s_j})}{B_i}\right] \frac{1 - \sigma}{d_{ij}} < 0 \\
\frac{\partial \ln B_{ij}(z_{s_j})}{\partial F_{ij}} &= \lambda \left[1 - \frac{B_{ij}(z_{s_j})}{B_i}\right] \frac{-1}{F_{ij}} < 0
\end{aligned}$$

It is thus easily verified that the sensitivity to both iceberg trade costs and fixed costs is higher for low-productivity firms. This is in line with the prediction of our model.

One may wonder what parameter of this model could be estimated from the moment used in our estimation. The moments are the share of sellers with M buyers relative to the share of sellers with a single buyer. As there is a direct mapping between productivity and the number of buyers served, these ratios are mostly driven by the shape of the Pareto distribution and can hardly be used to back out (fixed and variable) trade frictions.

A.4 Expected number of firms serving M buyers

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

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

Using the following change of variable,

$$\rho_{ij}(z) = \lambda_{ij} e^{-\frac{\lambda_{ij}}{\pi_{ij}} T_j z^{-\theta}},$$

one can show that

$$h_{ij}(M) = \frac{\pi_{ij}}{\lambda_{ij}} C_{B_i}^{M} \int_{\rho_{ij}(z_{min})}^{\lambda_{ij}} \rho_{ij}(z)^{M-1} (1 - \rho_{ij}(z))^{B_i - M} d\rho_{ij}(z),$$

where $\rho_{ij}(z_{min})$ is the probability of the least productive firm in j serving a buyer in i.

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

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

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

Using properties of the Beta function, notice that

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

Then, the regularized incomplete beta function is

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

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:

$$h_{ij}(M) = \frac{\pi_{ij}}{\lambda_{ij}} \frac{1}{M} \Big(I_{\lambda_{ij}}(M, B_i - M + 1) - I_{\rho_{ij}(z_{min})}(M, B_i - M + 1) \Big).$$

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

$$\lim_{\rho_{ij}(z_{min}) \to 0} I_{\rho_{ij}(z_{min})}(M, B_i - M + 1) = \lim_{\rho_{ij}(z_{min}) \to 0} \int_0^{\rho_{ij}(z_{min})} \rho_{ij}(z)^{M-1} (1 - \rho_{ij}(z))^{B_i - M} d\rho_{ij}(z) = 0.$$

Using this property, 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 work with the following convergent moments as auxiliary parameters:

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

that is, the proportion of firms from j having exactly M buyers in destination i.[44] We first show the empirical counterparts of these auxiliary parameters are normally distributed. Then, 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 $\left[\mathbb{1}\{B_{ij}(z_{s_j}) = M\}\right]_{s_j \in S_j}$, the vector of dummy variables that equal 1 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 $M \in [0, B_i]$.[45] The central limit theorem implies

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

[44] 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.

[45] Independence comes from the fact that sellers are independent from each other. Note this assumption could be relaxed because 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 their productivity in the same distribution and face the same degree of search frictions.

where

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

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

We then consider the function

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

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 an estimate of λ_{ij} based on $g(.)$ is asymptotically normal:

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

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

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

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

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

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

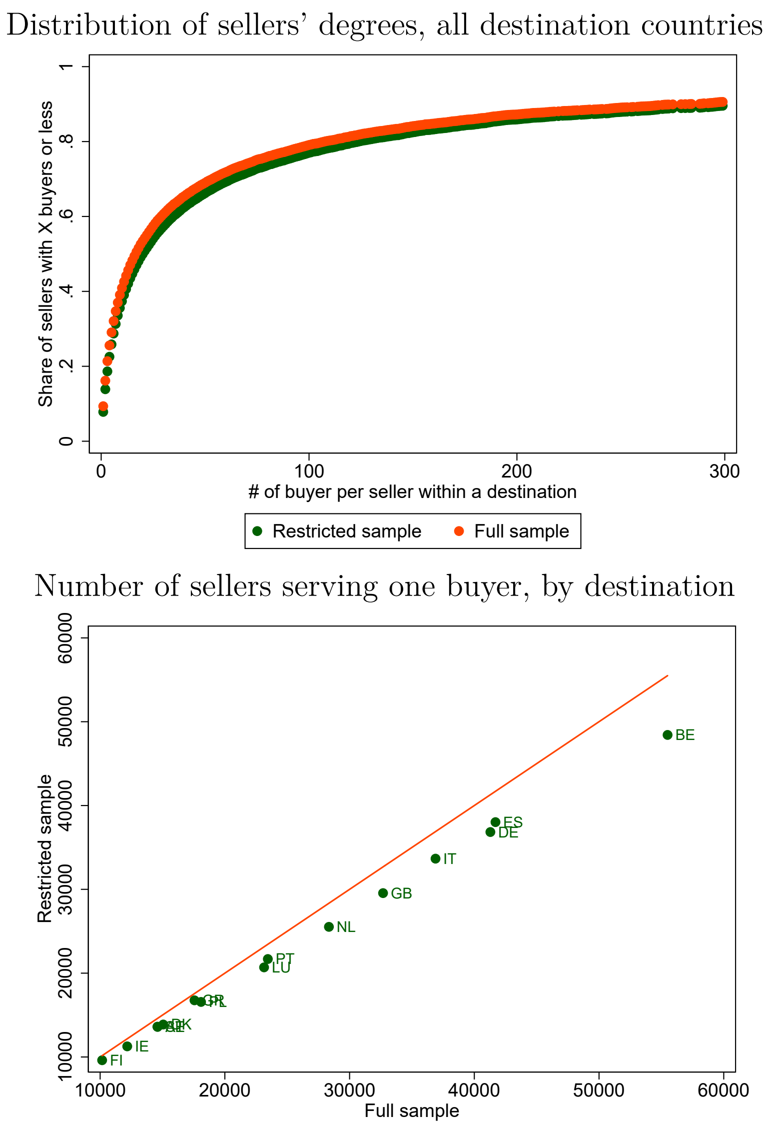

Notes: This figure compares the number of buyers per seller, in the whole dataset and in the estimation dataset, restricted to the 90% of exporters that declare the product category of their exports ("Restricted sample"). The top panel compares the distributions of sellers' degrees, where a firm's degree 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 (1)	Median (2)	p75 (3)	Sh. with 1 buyer (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.

Figure A.2: ***Number of buyers per seller, Wholesalers versus the rest of the economy***

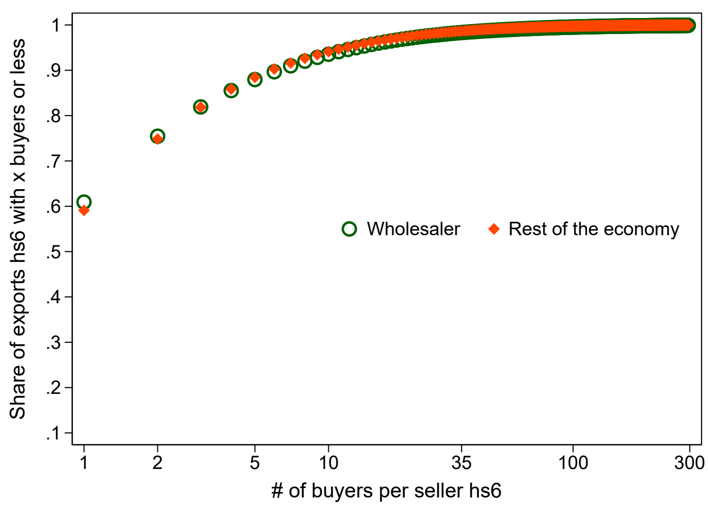

Notes: This figure compares the number of buyers per seller, in the wholesaler sector and in the rest of the economy.

Figure A.3: ***Identification power of the theoretical moments***

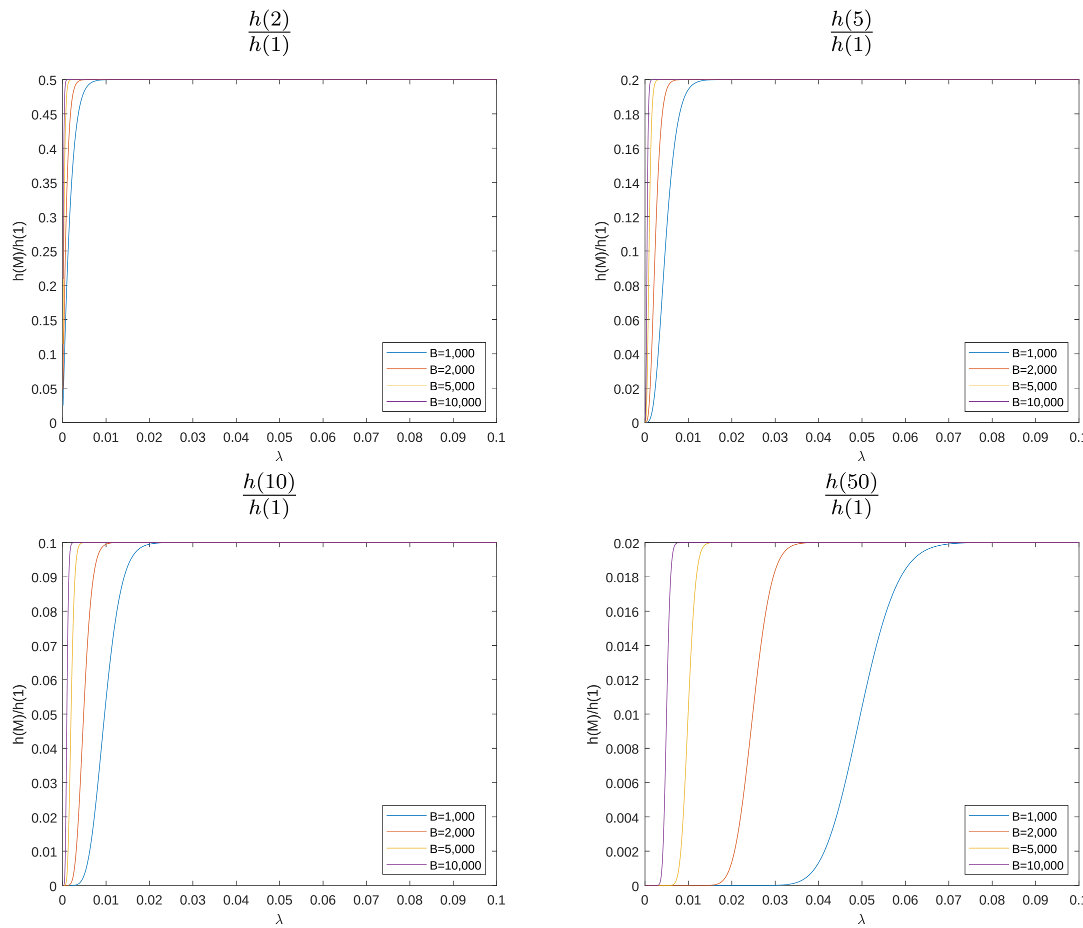

Notes: This figure shows the theoretical relationship between the underlying value of search frictions (λ, x-axis) and the share of firms with M buyers in the destination, in relative terms with respect to the expected number of firms with one buyer ($h(M)/h(1)$, y-axis). The relationship is derived conditional on the underlying number of buyers (B) and for various values of M, using the formula in equation (7).