

SEARCH FRICTIONS IN INTERNATIONAL GOODS MARKETS

Clémence Lenoir* Julien Martin[†] Isabelle Mejean[‡]

January 14, 2022

Abstract

This paper studies how frictions in the acquisition of new customers distort the allocation of activities 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 in foreign markets. Estimated coefficients display a strong degree of heterogeneity across countries and products, that correlates with various proxies for information frictions. 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.

[‡]Sciences Po and CEPR. Address: Sciences Po, Department of Economics, 28 rue des Saints Pères, 75006 Paris, France. Email address: isabelle.mejean@sciencespo.fr

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 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 sellers and buyers in international markets. We discuss the consequences of these frictions for the efficiency of selection into exporting, and the size of the firm's customer base, conditional on trade. By reducing the strength of competition, search frictions penalize the most productive producers, and thus distort the allocation of resources. As a consequence, the export premium of high-productivity firms is dampened in frictional product markets. We develop a structural estimator of search frictions that exploits firm-to-firm trade data. Estimates recovered for a large cross-section of products and destination countries are used to quantify the impact of search frictions on the selection of firms into export markets.

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: Ex-post, only the most-productive technology can eventually be exported. We introduce random search in the Ricardian framework. A discrete number of ex-ante homogeneous consumers in

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. [Drozd and Nosal \(2012\)](#) discuss how frictions affecting the building of market shares can explain the dynamics of international prices. [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.

²See [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](#)).

each market meet with a random number of heterogeneous producers of a perfectly substitutable good.³ Conditional on her random choice set, the consumer chooses the lowest-cost supplier. As in [Eaton and Kortum \(2002\)](#), iceberg trade costs and technological parameters shape Ricardian advantages. Search frictions however interact with these Ricardian comparative advantages to determine the geography of international trade: Conditional on comparative advantages, higher bilateral frictions dampen bilateral trade.⁴

Whereas search frictions and iceberg costs have the same qualitative impact on bilateral trade at product-level, search frictions further distort the allocation of resources among exporters of a given origin country. High search frictions reduce the average number of sellers met by any consumer thus dampening the strength of competition. Reduced competitive pressures benefit, in relative terms, low-productivity exporters: High-productivity firms always suffer from a high degree of bilateral frictions, but sufficiently low-productivity firms can instead display higher export propensities towards more frictional markets. This property is a distinctive feature of models displaying search frictions.⁵

We show how the model can identify search frictions structurally using firm-to-firm trade data. The structural estimator exploits the product-level dispersion in the customer base of exporters from a given origin in a particular destination. 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. Search frictions are estimated by a generalized method of moments for

³Unlike [Antràs et al. \(2017\)](#), our model abstracts from (ex-ante) buyers heterogeneity, and we assume there is no fixed cost of outsourcing. The selection of importers into different importing strategy is explained by the heterogeneity in the set of suppliers met by each buyer.

⁴This result 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. See, among others, [Head and Mayer \(2014\)](#), [Rauch \(1999\)](#) and [Rauch \(2001\)](#).

⁵We show in an appendix that this prediction survives an extension of the model in which high-productivity firms display relatively high matching probabilities provided the extension preserves the log-supermodularity of export probabilities with respect to search frictions and the firm's productivity. It is the case under realistic parametric assumptions.

about 12,000 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 \times 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. Search frictions are estimated to be stronger in more downstream and more input-specific product markets. Within a product, they are more pronounced in distant countries but lower in countries where the population of French migrants is larger, where citizens tend to speak a common language or where the degree of social connectedness is stronger. Importantly, the estimated model is able to fit the distribution of the number of consumers that exporters serve within a country and product, including the skewness of the distribution.

What are the distributional consequences of search frictions? We first show that estimated search frictions are on average larger in those product markets in which French firms have a Ricardian comparative advantage. This correlation magnifies the distortive impact of search frictions because it implies French firms have more difficulties matching with foreign buyers in those markets in which they would be in a strong competitive position with respect to foreign competitors, in the absence of search frictions. We then test and confirm the model's prediction that search frictions dampen the export premium of high-productivity exporters. Our results indicate that the export premium of firms in the top quartile of the distribution of sectoral domestic sales (resp. sectoral labor productivity) increases from 31.9 to 40.3% (resp. 23.8 to 30.3%) when moving from the ninth to the first decile of the distribution of estimated search frictions.

Counterfactual simulations show that a decrease in search frictions lead to an increase in average productivity through its non-monotonic impact along the distribution of firms' productivities, with low-productivity firms seeing their export probability reduced and high-productivity firms gaining in terms of both export probability and their expected number of buyers, conditional on export. Instead, a quantitatively comparable reduction in iceberg costs benefits disproportionately low-productivity firms and reduces average productivity.⁶

⁶This statement may seem to contradict [Melitz \(2003\)](#) result that a decrease in trade barriers improves the allocative efficiency by shifting resources from low to high-productivity firms. The discrepancy comes from the fact that [Melitz \(2003\)](#) focuses on multilateral trade liberalization, whereas our thought experiment considers a unilateral decrease in trade barriers. Reducing the cost of serving foreign markets without easing the entry of foreign exporters in the domestic economy increases the competitiveness of all domestic firms abroad, which benefit infra-marginal firms that can pass the

In comparison with other barriers to international trade, search frictions thus have important misallocative consequences. Reducing such frictions might thus 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 in Japan in [Bernard et al. \(2018a\)](#), the adoption of broad band internet in Norwegian municipalities in [Akerman et al. \(2018\)](#), and the development of online markets in [Lendle et al. \(2016\)](#). In a related work, [Chen and Wu \(2021\)](#) study the interplay between reputation and information frictions in the online trade of T-shirts.

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 of search frictions in explaining heterogeneity in firms' customer base. [Chaney \(2014\)](#) and [Allen \(2014\)](#), both develop models in which frictions hit the seller-side of the economy. We instead introduce frictions on the demand side, with consumers 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 threshold for profitable exports.

⁷Business 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.

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 test how the estimated frictions affect the selection of exporters into foreign markets.⁸

We also contribute to a series of recent papers that have used 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., 2020). 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.⁹ In Bernard et al. (2018b) and Carballo et al. (2018), the heterogeneity is studied in monopolistic competition models with two-sided heterogeneity. A distinctive feature of our model in comparison with theirs is the distortive effect of frictions that we confirm prevails in the data. The distortion also characterizes the model in Eaton et al. (2020), in which the matching of exporters and importers is also governed by random search in a firm-to-firm framework. Eaton et al. (2020) focus on how country-level frictions affect individual firms’ decisions to outsource some productive tasks, with an end-effect on the labor market. We instead focus on the impact of search frictions on the allocation of activities across exporters within narrowly defined foreign markets. Our empirical strategy allows us to remain flexible regarding the amount of heterogeneity of search frictions across products and destinations. To our knowledge, we are the first to provide systematic evidence of the export premium of high-productivity firms being dampened in frictional markets.¹⁰

The introduction of a countable number of firms also relates our work to recent

⁸In 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.

⁹Our 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 90% of importers in our data source a given product from a single French exporter. Instead, the overall number of French sellers they are connected to is often above one as importers tend to source several products from several French firms.

¹⁰Our paper also displays important differences in the modeling of the matching of sellers and buyers, in comparison with Eaton et al. (2020). To our knowledge, our paper is the first to extend the Ricardian analysis in Eaton and Kortum (2002) to a discrete setting. In doing so, we follow the logic introduced in Eaton et al. (2012) to study how a discrete number of sellers affect predictions of the Melitz (2003) model. Working with a discrete number of firms allows the model to have Eaton and Kortum (2002) as a limit case when search frictions become infinitely small. We discuss in Appendix A.2 how the comparison with Eaton and Kortum (2002) can help gather intuitions regarding the general equilibrium consequences of search frictions.

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 substitutability 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.

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, our proxy for a firm’s customer base, and study how that number varies across firms, products, and destinations. We notably study its correlation with proxies for search and information frictions. Section 3 describes our theoretical model and derives analytical predictions regarding the expected customer base that an exporter will serve in its typical destination, depending on its productivity and the level of search frictions. 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 data covering the universe of French exporting firms. The data are provided by the French Customs and are described in Bergounhon et al. (2018). The dataset covers each single transaction that involves a French exporter and an importing firm located in the European Union, in 2007. Using another reference year, such as 1997 or 2016, does not alter the results. Firms with annual export sales in Europe below 150,000 euros are allowed to fill a simplified form that does not contain information on the product category. The restriction concerns 31% of firms who cumulate less than 1% of aggregate exports.¹¹ For each transaction, the dataset contains the identity of the exporting firm (its SIREN identifier), an identifier for the importer (an anonymized version of its VAT code), the date of the

¹¹We show in Figure A.4 that the restriction is unlikely to affect our results. The distribution of sellers’ degrees, whose product-specific equivalent is used to compute the empirical moments in the estimation, is indeed almost identical in the whole sample and in the sample restricted to the 70% of exporters that declare a product category.

transaction (month and year), the product category (at the 8-digit level of the combined nomenclature), and the value of the shipment. From the firm identifier, we can recover the exporter’s sector of activity using data from the French statistical institute. In the rest of the analysis, data will be aggregated across transactions within a year, at the exporter-importer-hs6 product level. A unit of observation will thus be an exporter-importer-product triplet.

The analysis uses a sample restricted to a subset of each firm’s product portfolio that constitutes the core of the firm’s activity. Core products are defined as those that represent at least 10% of the firm’s export sales plus all products that constitute at least 10% of sales for at least one firm in the same 4-digit NAF sector. Information frictions are expected to be less of a problem for non-core products, that the firm sells occasionally. The restriction 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). The restriction thus helps focus the analysis on exporters that actually compete for serving foreign markets with their core products. We have also reproduced the estimation of search frictions on the full dataset with most results being qualitatively unchanged.

In 2007, we have information on 44,280 French firms exporting to 572,585 individual importers located in the 26 countries of the European Union. Total exports by these firms amount to 216 billion euros, which represents 53% of France worldwide exports. Table A2 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)), i.e. French exporters interact with more foreign partners than foreign buyers with French partners. The asymmetry is more pronounced once we focus on product-specific trade flows as in columns (4) and (5). Column (3) in Table A2 reports the number of active exporter-importer pairs, 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. 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 products. Multilateral import

data are from the CEPII-BACI database (Gaulier and Zignago, 2010). Finally, we use three alternative proxies for information frictions at country-level: the stock of French migrants in each destination, taken from the UN database on Trends in International Migrant Stock, the probability that individuals in France and the destination speak the same language (constructed by Melitz and Toubal (2014)) and a measure of social connectedness between France and the destinations, computed by Bailey et al. (2020) using anonymized Facebook data. The stock of French migrants is measured per thousand of inhabitants in the destination. Social connectedness is defined as the probability that two users in France and the destination country have a friendship link.

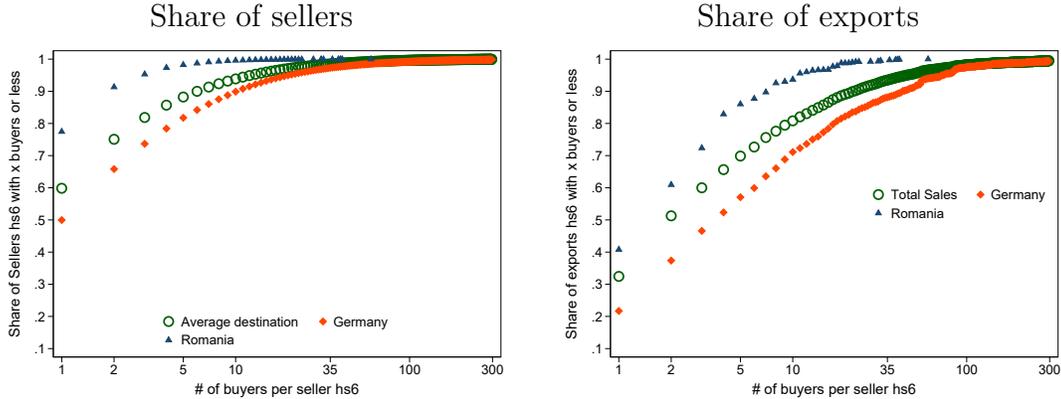
2.2 Descriptive Statistics

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. This information is of particularly high quality in the context of intra-EU trade as the corresponding data are collected for tax purposes and are thus exhaustive. 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 including customs data (Bernard et al., 2018b; Carballo et al., 2018, see, e.g.), and online trade data for a specific product (Chen and Wu, 2021). In comparison with these papers, we also show that the number of buyers in a firm’s portfolio is correlated with proxies for information frictions, conditional on other gravity variables.

Figure 1 shows the strong heterogeneity in the number of buyers per seller within a destination, our proxy for their customer base.¹² 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

¹²These facts are robust to alternatively defining the customer base as the number of buyers a seller interacts with over a three-year window, or the number of buyers interacting more than once with the seller over a 3-year window.

Figure 1: *Distribution of the number of buyers per seller*



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.

and diamond points, respectively).

In France’s typical export market, 60% of firms interact with a single buyer, and 88% with at most five buyers. At the other side of the spectrum, 1% of firms interact with more than 100 buyers in the same destination. As the right panel in Figure 1 shows, firms 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 firms interacting with at most 10 buyers. These numbers are not significantly different between wholesalers and other exporters (Figure A.5). Based on such evidence, we conclude that French exports are dominated by sellers interacting with a small number of buyers.

Our structural estimation of search frictions exploits the heterogeneity across sellers, within a product and destination. At this level, heterogeneity in terms of the number of buyers is significantly correlated with the seller’s size (Bernard et al., 2018b; Carballo et al., 2018). In our data, the coefficient of correlation between the log of the firm’s worldwide exports and the log of the number of partners served in a particular destination is equal to .28. Alone, the firm’s size explains 37% of the within-variance. Whereas there is strong heterogeneity in the number of partners served by French exporters, about 90% of foreign buyers purchase a given product from a single French exporter (Figure A.6). As a consequence, the mean degree of buyers that can be recovered from the comparison of columns (5) and (6) in Table A2 is very close to 1 in all

destinations.¹³

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)). Our dataset covers French exports to 26 EU countries. As a consequence, standard gravity variables such as distance, the border effect and the common language dummy end up strongly correlated. This explains that our specification has less control variables than a typical gravity regression based on multilateral data. Namely, bilateral trade is explained by distance to France, two proxies for market size, the country’s (product-specific) import demand and GDP per capita, and a proxy for search and information frictions, namely the likelihood that citizens in France and the destination speak the same language taken from Melitz and Toubal (2014). Table A1 reproduces the same regressions using two alternative proxies for information frictions.

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. Trade is also positively correlated with the language proximity of the origin and destination countries, our baseline proxy for frictions. The effect is robust to the choice of a proxy for information frictions, as shown in Table A1. These results are also confirmed within a firm, in column (5). Information frictions have a significant impact on exports in our sample: A one standard deviation higher probability of speaking the same language increases trade flows by 6%.

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 examine the buyer extensive margin, as measured by the number of buyers in each exporter’s portfolio of clients (see also Bernard et al. (2018b) for a similar decomposition based on Norwegian data). All margins of bilateral trade significantly contribute to the sensitivity of trade to gravity variables. In particular,

¹³Although 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. This heterogeneity is illustrated in Figure A.6 that compares the number of sellers an importer is connected to, within and across products. Our focus is on the matching of sellers and buyers within a product, and we thus use the fact the vast majority of buyers interact with a single seller at the product level to justify the model’s many-to-one matching structure. As discussed in Fontaine et al. (2021), it is straightforward to account for multi-product importers using the same setting by assuming buyers combine individual products within their production function.

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	-0.920*** (0.068)	-0.449*** (0.031)	-0.237*** (0.023)	-0.234*** (0.048)	-0.306*** (0.055)	-0.196*** (0.028)	-0.110** (0.045)
log Import Demand	0.847*** (0.014)	0.257*** (0.006)	0.152*** (0.005)	0.438*** (0.010)	0.444*** (0.010)	0.193*** (0.007)	0.252*** (0.009)
log GDP per Capita	0.173*** (0.036)	0.141*** (0.018)	0.090*** (0.011)	-0.058** (0.023)	0.023 (0.028)	-0.004 (0.017)	0.027 (0.019)
Proba Common Language	2.540*** (0.397)	2.205*** (0.186)	0.836*** (0.141)	-0.501** (0.195)	1.358*** (0.206)	0.999*** (0.110)	0.359** (0.144)
Observations	66,335	66,335	66,335	66,335	633,136	633,136	633,136
R-squared	0.630	0.767	0.414	0.578	0.684	0.425	0.715
Fixed effects	Product	Product	Product	Product	Firm	Firm	Firm

Notes: Standard errors, clustered in the country×HS2 chapter 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. ”Proba Common Language” is the probability that a French citizen and a citizen from the destination country speak the same language, as computed by Melitz and Toubal (2014). 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)):

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

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 . $\#_{pd}^S$, $\#_{spd}^B$, and $\#_{pd}^{SB}$ 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.

Column (5) uses the log of firm-level bilateral exports as left-hand-side variable, 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)):Likewise, the decomposition of firm-level exports in columns (5)-(7) of Table 1 is based on the following decomposition of trade into an extensive and an intensive terms:

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

the “buyer” extensive margin is responsible for 26% of the overall distance elasticity at the product level, a number that jumps to 64% once gravity coefficients are identified within a firm.¹⁴ Likewise, the buyer margin accounts for a substantial share of the overall impact of information frictions: Around one third when the gravity equation is estimated at the product-level and between 54 and 73% at firm-level, depending on the chosen proxy for information frictions, i.e. the language proximity in Table 1, the stock of French migrants in the destination or the degree of social connectedness in Table A1. Our interpretation of this finding is that a high probability of speaking the same language, a large stock of French migrants in the destination or a high degree of social connectedness 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. To alleviate notations, we neglect the product dimension until necessary, although all parameters can be understood as being potentially product-specific. After having summarized the main assumptions, we derive a number of analytical predictions that we later use in the structural estimation. We then discuss possible extensions of the model and alternative theoretical frameworks.

¹⁴Note the contribution of the buyer margin is artificially low in the decomposition of product-level trade in columns (1)-(4) because of the multicollinearity between the “seller” and “buyer” extensive margins. If we instead work with this decomposition:

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

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

3.1 Assumptions

The economy is composed of N countries indexed by $i = 1, \dots, N$. The partial equilibrium analysis focuses on a single good produced into perfectly substitutable varieties. 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 w_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 $[z, +\infty[$. 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.

The model has bilateral iceberg trade costs but no fixed 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 equal to $\frac{w_j d_{ij}}{z_{s_j}}$. 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 $\mu_{ij}(p) = T_j \left(\frac{d_{ij} w_j}{p} \right)^{-\theta}$. 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 $\mu_i(p) = p^\theta \sum_{j=1}^N T_j (d_{ij} w_j)^{-\theta} = p^\theta \Upsilon_i$. As in [Eaton and Kortum \(2002\)](#), $\Upsilon_i = \sum_{j=1}^N T_j (d_{ij} w_j)^{-\theta}$ reflects “multilateral resistance” in country i : the higher Υ_i is, the more competitors with low costs can serve the country.

We depart from the representative consumer’s assumption used in most of the literature and instead assume each country is populated by a finite number B_i of (ex-ante) homogeneous buyers, each one characterized by its own iso-elastic demand function:

$$c_{b_i} = p_{b_i}^{-\sigma} \bar{X}_i, \quad \sigma > 1$$

where c_{b_i} is the quantity of the good purchased by buyer b_i given the price p_{b_i} she is offered and a demand shifter \bar{X}_i that we assume is shared across all buyers within a market. 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.

To simplify the analysis, we will assume that producers price at their marginal cost. As a consequence, buyer b_i chooses to purchase the good from the lowest-cost supplier who she met and pays the price:

$$p_{b_i} = \arg \min \left\{ \frac{w_j d_{ij}}{z_{s_j}}; s_j \in \Omega_{b_i}; \forall j = 1, \dots, N \right\},$$

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 all suppliers. Within a destination, all buyers would thus end up paying the same price for the homogeneous good and the model would collapse into a representative buyer's Ricardian setup à la [Eaton and Kortum \(2002\)](#). The frictionless equilibrium has a distribution of firms that is degenerate ex-post with only the most efficient technology being eventually active. Our model does not display such a degenerate ex-post distribution. Each buyer b_i meets with a random number of potential suppliers, drawn from a Poisson distribution of parameter $\sum_j \lambda_{ij} T_j z^{-\theta}$. 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:¹⁵

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

where $\kappa_i \equiv \frac{\sum_j \lambda_{ij} T_j (w_j d_{ij})^{-\theta}}{\sum_j T_j (w_j d_{ij})^{-\theta}}$ measures the expected number of suppliers met, in relative terms with respect to the maximum number of suppliers that would compete 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 : $\kappa_i = \sum_j \omega_{ij} \lambda_{ij}$ with $\omega_{ij} \equiv \frac{T_j (w_j d_{ij})^{-\theta}}{\sum_j T_j (w_j d_{ij})^{-\theta}}$.

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,

¹⁵All analytical details are postponed to Appendix [A.1](#).

which is inversely proportional to θ . In comparison with the frictionless equilibrium, expected prices are systematically inflated by search frictions (since $\kappa_i < 1$). 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 the expected number of suppliers met, κ_i . It is larger when λ_{ij} and $T_j(w_j d_{ij})^{-\theta}$ are negatively correlated i.e. when search frictions are high in markets where the country has a Ricardian comparative advantage. 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 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.

3.2.1 Product-level trade

As we demonstrate in [Appendix A.1](#), our model inherits the property of the [Eaton and Kortum \(2002\)](#) model that explains the geography of trade by the probability for a given unit of consumption to be sourced in a particular origin country. Namely, the ex-post share of country j 's (product-level) consumption that is imported from country i , denoted π_{ij} , can be shown to be exactly equal to the ex-ante expected probability that any buyer b_i ends up interacting with a supplier from j :

$$\pi_{ij} = \mathbb{E} \left[\mathbb{1}_{b_i j}^{(1)} \right],$$

where $\mathbb{1}_{b_i j}^{(1)}$ is a dummy variable equal to one if the lowest-cost supplier met by b_i originates from country j and $\mathbb{E}(\cdot)$ is the expectation operator. Properties of the Poisson distribution imply the probability of the lowest-cost supplier being located in j is con-

stant and independent of b_i . Trade shares thus simplify into

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

The share of country i 's absorption of the product that is sourced from country j thus depends on (i) the relative competitiveness of its firms in comparison with the rest of the world, $\frac{T_j(d_{ij}w_j)^{-\theta}}{\Upsilon_i}$, and (ii) the relative size of search frictions its firms encounter while serving market i , $\frac{\lambda_{ij}}{\kappa_i}$. The first ratio is the formula derived in [Eaton and Kortum \(2002\)](#). 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} implies:

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

i.e. the market share of a country always increases following a reduction in bilateral frictions.

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

$$\ln \pi_{ij} = FE_i + FE_j - \theta \ln d_{ij} + \ln \lambda_{ij}, \quad (2)$$

where $FE_i \equiv \ln \Upsilon_i \kappa_i$ and $FE_j \equiv \ln T_j (w_j)^{-\theta}$. 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.¹⁶ A corollary is that predictions on product-level trade cannot be expected 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

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 customers abroad, we take the point of view of individual sellers and derive predictions regarding the expected number of customers they can reach, in each destination.

Consider the probability that a given supplier from j , France in our data, serves a buyer in i . Given buyers are ex-ante homogeneous, multiplying the probability by B_i directly delivers the expected number of buyers served by an exporter, which we measure in the data. In our framework, this probability is the product of 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:

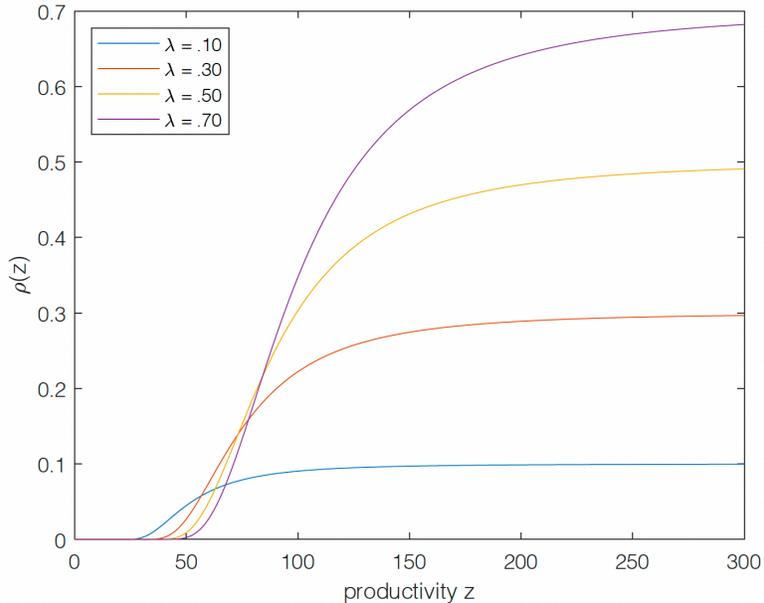
$$\begin{aligned} \rho_{ij}(z_{s_j}) &= \mathbb{P}(s_j \in \Omega_{b_i}) \mathbb{P}\left(s_j : \min\left\{\frac{w_k d_{ik}}{z_{s'_k}}; s'_k \in \Omega_{b_i}\right\} = \frac{w_j d_{ij}}{z_{s_j}}\right) \\ &= \lambda_{ij} e^{-(w_j d_{ij})^\theta z_{s_j}^{-\theta} \Upsilon_i \kappa_i} \end{aligned} \quad (3)$$

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 average cost $w_j d_{ij}$, where competition is limited (Υ_i low), and that displays highly distortive average search frictions (κ_i small). These predictions are

¹⁶In the trade literature, equation (2) is used to estimate the elasticity of trade to iceberg costs (θ). Equation (2) shows unbiased estimates of trade elasticities requires instruments for iceberg costs that are uncorrelated with search frictions.

consistent with evidence presented in section 2.2.

Figure 2: *Probability of serving a buyer as a function of the seller's productivity*



Notes: This figure illustrates how the probability of serving a buyer varies with the seller's productivity, for four different values of bilateral frictions.

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. This property of the model is illustrated in Figure 2 which shows the probability of a match $\rho_{ij}(z_{s_j})$ as a function of the firm's productivity, for four alternative values of the bilateral meeting probability. Consistent with empirical evidence, the probability of serving a foreign buyer is increasing in the firm's productivity. However, the slope of the relationship is reduced when search frictions increase (for a lower value of λ). 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 1:

Proposition 1. 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:

$$\frac{\partial \ln \rho_{ij}(z)}{\partial \lambda_{ij}} = \underbrace{\frac{\partial \ln \lambda_{ij}}{\partial \lambda_{ij}}}_{\text{Visibility channel}} - \underbrace{\frac{\partial (w_j d_{ij})^\theta z^{-\theta} \kappa_i \Upsilon_i}{\partial \lambda_{ij}}}_{\text{Competition channel}} = \frac{1}{\lambda_{ij}} - z^{-\theta} T_j \quad (4)$$

and

$$\frac{\partial^2 \ln \rho_{ij}(z)}{\partial \lambda_{ij} \partial z} > 0$$

High-productivity firms always benefit from a reduction in frictions (an increase in the meeting probability λ_{ij}):

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

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,

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

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

See the Proof in Appendix A.3.

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 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 ($\lambda_{ij} z^{-\theta} T_j > 1$), the competition channel dominates the visibility channel at the bottom of the productivity distribution, and sufficiently low-productivity

sellers benefit from more frictions.¹⁷

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

$$\begin{aligned} \ln \frac{\rho_{ij}(z^H)}{\rho_{ij}(z^L)} &= (w_j d_{ij})^\theta \Upsilon_i \kappa_i \left((z^L)^{-\theta} - (z^H)^{-\theta} \right) \\ &= \frac{\lambda_{ij}}{\pi_{ij}} T_j \underline{z}^{-\theta} \left[\left(\frac{z^L}{\underline{z}} \right)^{-\theta} - \left(\frac{z^H}{\underline{z}} \right)^{-\theta} \right], \end{aligned} \quad (6)$$

where $\rho_{ij}(z^H)$ and $\rho_{ij}(z^L)$ denote export probabilities in country i of a firm from j with a high-productivity z^H and a low-productivity z^L , 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 and λ_{ij} , which is 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. We provide evidence of this distortive impact of frictions in Section 5.

Before concluding, it is important to note that the export premium of high-productivity firms is exacerbated in countries featuring high iceberg trade costs; that is,

$$\frac{d \ln \frac{\rho_{ij}(z^H)}{\rho_{ij}(z^L)}}{d \ln d_{ij}} > 0.$$

The reason is that an increase in iceberg costs deteriorates the relative competitiveness of all French firms - but the competitiveness loss is stronger for low-productivity ones. Low-productivity firms thus become less likely to serve the buyer, conditional on a match. Whereas search frictions and iceberg costs have the same qualitative impact on product-level trade, their impact on individual firms' export probabilities is instead different. This discrepancy explains that the heterogeneity in export performances

¹⁷Proposition 1 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 and nature of trade frictions are however quite different in ours and their models. The non-monotonicity discussed in Proposition 1 also has a similar flavor than lemma 1 in [Ornelas et al. \(2021\)](#), where contract incompleteness affects disproportionately the most productive suppliers.

across firms is useful to identify search frictions separately from iceberg costs.

3.3 Discussion

In this sub-section, we discuss the robustness of our results to various assumptions. We also compare our model with alternative frameworks used in the literature to describe the matching of sellers and buyers in international markets.

Possible extensions. In section 3.1, we have assumed marginal cost pricing within each buyer’s random choicest. In section A.4, we show that the model’s main properties are robust to assuming that firms Bertrand compete. The reason is that the distribution of markups recovered under Bertrand competition is invariant to the nationality of either the supplier or the buyer, exactly as in the Bernard et al. (2003) extension of the Eaton and Kortum (2002) model. As a consequence, the geography of trade discussed in Section 3.2.1 is the same as in the baseline model and so is the expression for export probabilities derived in Section 3.2.2.

A more crucial assumption is that the probability of meeting with a seller is invariant along the productivity distribution. Arguably, high-productivity firms may benefit from more visibility, and we shall expect the lack of visibility to matter mostly for relatively small firms.¹⁸ In Section A.5, we show that the distortive effect of search frictions does not necessarily disappear if we relax the assumption that meeting probabilities are constant along the productivity distribution and instead allow high-productivity firms to benefit from more visibility. As in the baseline case, an upward shift in the distribution of meeting probabilities increases the likelihood that any firm meets with a buyer while in the meantime strengthening competition, conditional on a match. In comparison with the baseline case, the competition effect should in general be stronger because the additional sellers met in the less frictional equilibrium are more productive, on average. However, the visibility channel now varies along the distribution of productivities. Whereas the outcome of the model obviously depends on the exact functional form linking meeting probabilities and sellers’ productivity, we show that under a reasonable parametric assumption, the model still displays the log-supermodularity of $\rho_{ij}(z_{s_j})$ in λ_{ij} and z_{s_j} . This property is summarized by the comparison of Figures 2 and A.1 that show how export probabilities vary along the productivity distribution in both

¹⁸This intuition is confirmed by estimates in Eaton et al. (2020) who use a more flexible matching technology and estimate low-cost firms to be more visible in international markets, on average.

models. Whereas increasing meeting probabilities rise the relative probability of high-productivity firms to serve any particular foreign buyer, a reduction in the mean level of search frictions still improves the relative export performances of high-productivity firms, a consequence of the distortive impact of frictions.

General equilibrium. Appendix A.2 discusses how the model could be incorporated into a general equilibrium framework. Such extension can be achieved by assuming that there is a continuum of products as in [Atkeson and Burstein \(2008\)](#) and [Gaubert and Itskhoki \(2018\)](#). Equilibrium factor prices are then pinned down by good market equilibrium conditions. For similar reasons as in [Eaton et al. \(2020\)](#), such extension can bring the model closer to the class of models described in [Arkolakis et al. \(2012\)](#), with the additional complication implied by the rich product heterogeneity assumed in our model.¹⁹ Whereas numerically solving this model is beyond the scope of this paper, the general equilibrium effect of a change in frictions is unlikely to affect the predictions described earlier, at least qualitatively. Unlike search frictions, wage adjustments indeed affect all firms symmetrically.

Comparison with alternative models. Beside the model’s robustness to alternative assumptions, another natural question is the extent to which the theoretical predictions later used to identify search frictions could be rationalized in a completely different model of the matching between sellers and buyers. In Appendices A.6, A.7 and A.8, we discuss three alternative frameworks. In Appendix A.6, we first discuss the properties of a discrete choice model in which buyers display heterogeneous preferences for horizontally differentiated varieties offered by sellers originating from various countries. As notably discussed in [Head and Mayer \(2014\)](#), such model can deliver a gravity structure, exactly as a Ricardian framework does. In Appendix A.6, we set up such model and show that the product-level prediction for the geography of trade is similar in this and our baseline model, if one interprets λ_{ij} as a dyadic preference parameter shaping the (Pareto) distribution of consumers’ valuation for varieties produced in a particular country. Such model has nothing to say about the heterogeneity across sellers in their ability to serve a particular market, however. As this heterogeneity is at the core of our identification strategy, we can rule out that consumers’ heterogeneous

¹⁹Quantifying the general equilibrium impact of search frictions would notably require to estimate product-level bilateral search frictions faced by producers from the rest of the world, which is not possible based on our empirical strategy.

preferences, alone, shape the empirical moment used to identify search frictions.²⁰

In appendix A.7, we then describe a partial equilibrium version of a 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. In Appendix A.8, we develop a model of two-sided heterogeneity *à la* Bernard et al. (2018b). The supply side of the model is again taken from Eaton et al. (2012) but buyers are assumed heterogeneous in terms of the level of their demand. Sellers choose the number and identity of buyers they want to serve, which generates negative assortative matching between buyers and sellers. Both models can rationalize the stylized facts in Section 2.2, most notably the increasing relationship between a firm’s number of buyers and its productivity. None of these models however displays the distortive trade friction that is central in our analysis. In these models, reducing trade frictions, whether the fixed penetration cost in the model of appendix A.7 or the fixed cost of matching with one more buyer in appendix A.8, increases the probability of export everywhere along the productivity distribution (see Figures A.2 and A.3). This is in contrast with what we have in our model, in which reducing search frictions has a non-monotonic impact on low- and high-productive firms (Figure 2). We provide evidence of such non-monotonicity in Section 5, that we argue is difficult to explain in alternative contexts.

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. To simplify notation and since the empirical analysis solely uses data on French exporters, the index for the country of origin (j in Section 3) is now neglected and in-

²⁰The same is true of a slightly modified version of the frictionless Eaton and Kortum (2002) model in which productivity would be drawn from a Fréchet distribution displaying a bilateral scale parameter as in: $F_{ij}(z) = Pr[Z_{ij} \leq z] = e^{-T_{ij}z^{-\theta}}$ where Z_{ij} denotes the random productivity drawn by sellers in j to serve market i . In such model, the distribution of prices in country i becomes $G_i(p) = 1 - e^{-p^\theta \sum_j T_{ij}(w_j d_{ij})^{-\theta}}$ which is the same expression as in our baseline model for $T_{ij} = \lambda_{ij}T_j$. In this variant of the Eaton and Kortum (2002) model, the ex-post distribution of active exporters within each country is still degenerated though, whereas our model exploits predictions for export probabilities along the distribution of firms’ productivity.

dividual firms are just identified by their productivity z . On the other hand, we will introduce the product dimension k that was neglected until now. λ_i^k will thus denote the level of frictions faced by French producers of product k , in market i . In the data, destination countries are all located in the European Union.

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 a given French exporter 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^k and $\rho_i^k(z)$:

$$\mathbb{P}(B_i^k(z) = M | z > \underline{z}) = C_{B_i^k}^M \rho_i^k(z)^M (1 - \rho_i^k(z))^{B_i^k - M}.$$

Integrating over the expected distribution of productivities gives the expected number of French exporters with exactly $M > 0$ buyers in i (see details in Appendix A.9):²¹

$$h_i^k(M) = \frac{\pi_i^k}{\lambda_i^k} \frac{1}{M} I_{\lambda_i^k}(M, B_i^k - M + 1), \quad (7)$$

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 $h_i^k(M)$ is also a function of λ_i^k . Conditional on π_i^k and B_i^k , one can use the predicted value for $h_i^k(M)$ and its counterpart in the data to recover a structural estimate for λ_i^k , for each product and destination.

Once normalized by the expected number of firms in the market ($T^k \underline{z}^{-\theta}$) to recover a

²¹Integrating 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.

convergent moment, equation (7) can be used to estimate search frictions. Empirically, this moment however varies with distance, which potentially reflects the impact of other physical trade barriers on a firm’s customer base in 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.²² The correlation between 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_i^k component in equation (7) is negatively correlated with iceberg trade costs d_i^k , 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 $\frac{h_i^k(M)}{h_i^k(1)}$ and compare it with its empirical counterpart. In theory, this convergent moment is useful to identify search frictions as it varies monotonically with λ_i^k (see Figure A.7 in Appendix). Moreover, several ratios can be combined to identify precisely search frictions along a wide range of possible values. In the data, 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 correlated with 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 $\frac{h_{ij}^k(M)}{h_{ij}^k(1)}$ ratios:

$$Var_i^k(\lambda_i^k) = \frac{1}{B_i^k - 1} \sum_{M=2}^{B_i^k} \left(\frac{h_i^k(M)}{h_i^k(1)} - \frac{1}{B_i^k - 1} \sum_{M=2}^{B_i^k} \frac{h_i^k(M)}{h_i^k(1)} \right)^2. \quad (8)$$

²²For practical reasons detailed below, we restrict our attention to four values for $h_i^k(M)$, corresponding to the bottom of the distribution of sellers’ degrees.

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

Dependent Variable	log Distance	Std Dev.	Adjusted R-squared
# sellers with:			
1 buyer	-15.92***	((1.51)	.698
2 buyers	-5.89***	(.575)	.535
3 buyers	-3.24***	(.374)	.417
4 buyers	-2.00***	(.261)	.334
# sellers (in relative terms with respect to the sellers with 1 buyer) with:			
2 buyers	.021**	(.009)	.339
3-4 buyers	-.027***	(.008)	.372
5+ buyers	-.121***	(.021)	.408
Variance of the relative shares of sellers:			
across M	.002	(.011)	.210
across M , controlling for search frictions	-.010	(.014)	.210
	coef. on search frictions: -0.086** (.041)		

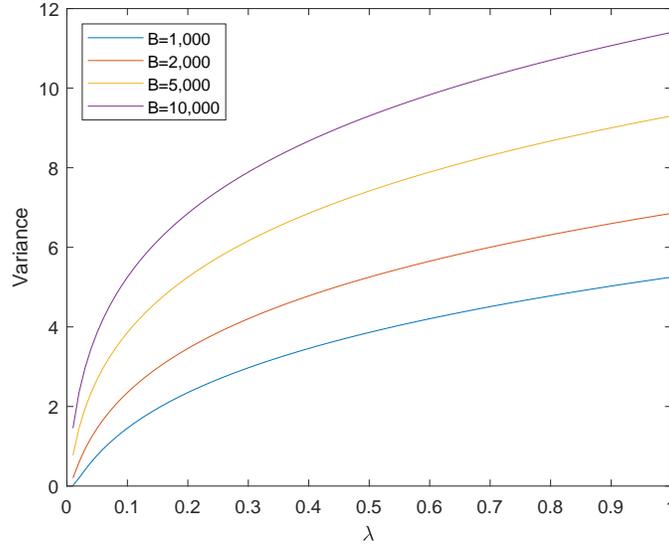
Notes: Robust standard errors, clustered at the country level, are in parentheses, with ***, **, and * respectively denoting significance at the 1%, 5% and 10% levels. The last regression uses as right-hand-side variables the (log of) distance from France and the probability of citizens in France and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

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 3, this moment is also correlated positively with λ_i^k and is thus useful for identification. Intuitively, when frictions get very high, the probability of a seller reaching more than one buyer approaches zero, resulting in a variance that approaches zero. As shown in the third panel of Table 2, the empirical counterpart of this moment is not correlated with distance. On the other hand, it is negatively correlated with our proxy for search frictions recovered from the probability of citizens speaking the same language. From that point of view, this moment is a good candidate to identify search frictions in the data.²³

In theory, the dispersion can be calculated across $B_i^k - 1$ ratios. However, these ratios do not convey a lot of relevant information, because they are almost all equal

²³These results are the only ones that are not robust to using the full data sample instead of restricting each exporter's portfolio to its core products. Because the invariance of the empirical moment to distance is important to insure search frictions are identified separately from other barriers to trade, we have decided to use the restricted sample in the analysis.

Figure 3: *Correlation between the variance of the $h(M)/h(1)$ ratios and the value of the meeting probability*



Notes: This figure shows the theoretical relationship between the underlying meeting probability (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 $\frac{\sum_{m=5}^B h(m)}{h(1)}$.

to 0 in the data, above a certain level of M .²⁴ For this reason, we decided to restrict our attention to the variance computed over three empirically relevant $\frac{h_i^k(M)}{h_i^k(1)}$ ratios, namely, $M = \{2, [3, 4], [5, B_i^k]\}$, $M = \{2, 3, [4, B_i^k]\}$ or $M = \{[2, 3], [4, 5], [6, B_i^k]\}$ depending on the product and destination. For consistency, the moments in Figure 3 use the same convention, although we have checked that the variance recovered from all possible ratios is also increasing in λ_i^k . Finally, in unreported results, we have estimated search frictions using the variance over four rather than three ratios, namely $M = \{2, [3, 4], [5, 6, 7], [8, B_i^k]\}$. The correlation between the baseline estimates and estimates recovered from the alternative definition is high, at .75, thus suggesting our strategy is not too sensitive to the precise moments aggregated within the variance operator.

²⁴As 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 .

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^k , solely depends on λ_i^k . The empirical counterpart of this theoretical moment is observed in our data:

$$\widehat{Var}_i^k = Var \left(\frac{\sum_{z=1}^{S^k} \mathbb{1}\{B_i^k(z) = m_1\}}{\sum_{z=1}^{S^k} \mathbb{1}\{B_i^k(z) = 1\}}, \frac{\sum_{z=1}^{S^k} \mathbb{1}\{B_i^k(z) = m_2\}}{\sum_{z=1}^{S^k} \mathbb{1}\{B_i^k(z) = 1\}}, \frac{\sum_{z=1}^{S^k} \mathbb{1}\{B_i^k(z) = m_3\}}{\sum_{z=1}^{S^k} \mathbb{1}\{B_i^k(z) = 1\}} \right), \quad (9)$$

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

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

$$\sqrt{S^k} \left(\widehat{Var}_i^k - Var_i^k(\lambda_i^k) \right) \xrightarrow[S^k \rightarrow +\infty]{\mathcal{D}} \mathcal{N}(0, \Omega_i^k(\lambda_i^k)) \quad (10)$$

where $\Omega_i^k(\lambda_i^k)$ is the variance of \widehat{Var}_i^k .²⁵ Using the convergence result, identifying λ_i^k uniquely is possible. With an asymptotic least squares estimation strategy, the estimated variance of estimated frictions writes

$$\widehat{\Sigma}_{\lambda_i^k} = \left[\frac{\partial Var_i^k(\widehat{\lambda}_i^k)}{\partial \lambda_i^k} (\Omega_i^k)^{-1} (\widehat{\lambda}_i^k) \frac{\partial Var_i^k(\widehat{\lambda}_i^k)}{\partial \lambda_i^k} \right]^{-1},$$

with $\Omega_i^k(\lambda_i^k)$ the optimal matrix of weights defined in Appendix B.1.

With a targeted moment that has an analytical formula, the implementation is straightforward. The only practical difficulty concerns the measurement of S^k and B_i^k in the data. Indeed, the theoretical moment in (8) identifies λ_i^k conditional on B_i^k . Moreover, the total number S^k 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

²⁵ $\Omega_i^k = \nabla g(\lambda_i^k) \Sigma_i^k \nabla' g(\lambda_{ij})$, where g is the variance function and Σ_i^k is the variance-covariance matrix of the random variables $\mathbb{1}\{B_i^k(z) = M\}$ for $M = m_1, m_2, m_3$.

sector using predictions of the model regarding trade shares. Under the assumptions of the model, π_i^k is both the share of French products in country i 's absorption of product k and the ratio of the number of buyers from i buying their consumption from a French producer divided by the total number of buyers in i ($\pi_i^k = B_{iF}^k/B_i^k$). π_i^k can easily be recovered from sectoral bilateral trade and absorption data.²⁶ B_{iF}^k is observed in our data. Based on this, one can recover a value of B_i^k for each destination and sector.²⁷

Information on the number of *potential* suppliers by *hs6* product is not available in any administrative dataset. We measure S^k based on information on the number of firms in each sector available from the INSEE-Repertoire Siren database. 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_i^k to select the algorithm's starting point for each country and product.

4.3 Results

Summary statistics. Search frictions are estimated at the (product \times country) level for a total of 12,631 λ_i^k parameters, among which 12,599 are statistically significant at the 5% level. Table 3, first column, provides summary statistics on the estimated parameters. Remember that in the model, the λ_i^k coefficient is defined as the share of sellers from France 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 .00%, whereas 10% have a meeting probability

²⁶We use bilateral trade flows from the CEPII-BACI database ([Gaulier and Zignago, 2010](#)) and production data from the World Input-Output Database. π_i^k is defined as the ratio of bilateral trade from France to country i over absorption in country i .

²⁷In sectors and countries in which the market share of French firms is very low, our empirical strategy implies very high values for B_i^k , above a million firms. Such high values might artificially bias our estimation of λ_i^k down. To avoid this issue, we winsorized the number of potential buyers at the 95th percentile of each country-specific distribution. The winsorizing is relatively homogenous across sectors, except for three exceptions. The share of winsorized product \times country pairs is slightly larger in chapter 49 (Printed books), 69 (Ceramic products) and 71 (Precious and semi-precious stones), for which we measure a relatively low share of foreign products in absorption despite the number of French exporters being significant.

above 2.51%. The second column in Table 3 provides statistics regarding the probability that a French firm willing to export meets with zero buyer in the destination, which is equal to $(1 - \lambda_i^k)^{B_i^k}$ in the context of our model. The median probability of meeting zero buyer is .03% but increases to 4% and 55% for the top 25% and top 10% of product×country pairs, respectively. These results are found robust over time, using 1997 and 2016 data as references.

Table 3: *Summary statistics on estimated coefficients*

	Meeting Probability	Probability of Meeting 0 Buyer	Number of Buyers
	λ_i^k (en %)	$(1 - \lambda_i^k)^{B_i^k}$ (en %)	B_i^k
Mean	0.36	11.8	5,997
Percentile 10	0.00	0.00	293
Percentile 25	0.09	0.00	705
Percentile 50	0.35	0.03	1,920
Percentile 75	1.08	4.38	5,514
Percentile 90	2.51	55.30	15,068
# Observations	12,631	12,631	12,631

Notes: The first column in this table presents summary statistics on the λ_i^k coefficients, estimated by country × hs6 product. The second column summarizes the subsequent probabilities that a French exporter meets with no buyer in the destination computed as $(1 - \lambda_i^k)^{B_i^k}$ for each country and product. Statistics on the number of potential buyers are reported in the third column.

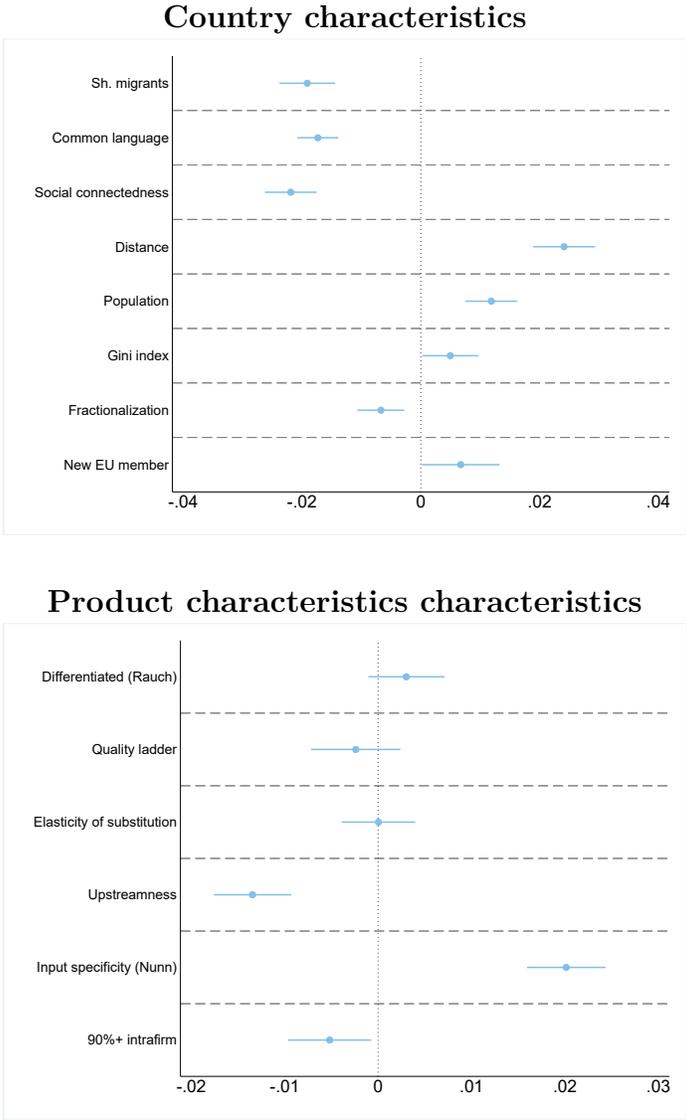
Correlates of search frictions. In Figure 4, we examine how the estimates correlate with different country and product attributes. The top panel focuses on country-specific attributes. Each point in the figure shows the result of a univariate regression of the zero-match probability on the corresponding (standardized) country-level variable, conditional on product fixed effects. As expected, our estimates correlate negatively with the three proxies for information frictions, namely the stock of French migrants in the destination, the index of social connectedness and the share of citizens from France and the destination that speak a common language. Because we cannot rule out the influence of other forces affecting the empirical moment used in the structural estimation, we also correlate our estimates with additional variables. First, our estimates are

positively correlated with distance, which means that French exporters have more difficulties meeting foreign buyers located in distant countries. Zero-match probabilities are also positively correlated with the size of the population in the destination country. This correlation is consistent with the presence of congestion effects in large markets (Eaton et al., 2020). It may also reveal the influence of consumer heterogeneity, a determinant of trade that our model neglects but has been discussed extensively in the trade literature.²⁸ As discussed in Appendix A.8, consumers' heterogeneity can indeed affect the matching of sellers and buyers in international markets, and thus the moment used for identification. To investigate this possibility further, we correlate our estimated zero-match probabilities with two proxies for the extent of consumers' heterogeneity, namely the Historical Index of Ethnic Fractionalization (Drazanova, 2020) and a Gini index of income inequalities. Zero-match probabilities are reduced in countries that are more heterogeneous in terms of cultural preferences, but they are higher in more unequal countries. Whether frictions increase with consumer heterogeneity or whether our estimates of frictions capture income heterogeneity is hard to assess at this stage. In any case, the correlation between our estimates and bilateral variables such as distance or common proxies for search frictions suggest that estimated frictions capture more than consumer heterogeneity. Last, frictions tend to be stronger in new EU member states which display more recent trade ties with the rest of the European Union.

We then examine the correlation between estimated frictions and product-level characteristics. The bottom panel of Figure 4 reports the results of separate regressions of zero-match probabilities on each of these characteristics and country fixed effects. We first examine whether frictions are correlated with horizontal product differentiation as measured by Rauch (2001) or by estimates of the elasticity of substitution (Imbs and Mejean, 2015), and with vertical differentiation as measured by Khandelwal (2010)'s quality ladder. These measures of differentiation are not systematically tied to our measures of search frictions. However, we find a strong and positive correlation between input specificity (Nunn, 2007) and search frictions, which suggests that frictions are higher among product categories featuring stronger investments in specific inputs. Last, search frictions are found lower among more upstream product categories and for products that are mostly traded intra-firm.

²⁸The correlation with country size may also emerge if firms are constrained in their capacity to serve a large number of buyers.

Figure 4: *Correlates of the no-match probabilities with product and country attributes*



Notes: This figure shows the estimated coefficient recovered from a regression that has the probability of zero match on the left-hand side of the estimated equation and the variable described on the left of the graph as right-hand side variable. All explanatory variables are standardized. Results of the regressions summarized in the top panel control for product fixed effects whereas country fixed effects are used as control in the bottom panel. The spikes correspond to 95% confidence interval around the estimated coefficients.

Test of empirical predictions. To assess the validity of our estimates, we confront the model’s predictions to the data. Section 3.2.1 unambiguously shows an increase in bilateral search frictions within a product category between France and a trade partner should lead to a reduction in French exports. We test this prediction of the model in Table 4. We first regress the logarithm of French exports (computed by destination-product pair) on standard gravity variables, namely distance, market size, and income per capita. All specifications include product fixed effects to capture differences in French comparative advantages across product categories. In column (2), we add our estimates of search frictions. Consistent with the model’s prediction, estimated search frictions are negatively correlated with French exports. This negative correlation is robust to the inclusion of other proxies for information frictions in Columns (3)-(5). Controlling for search frictions does not significantly influence the coefficient on distance (col. (1)-(2)).

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. Based on this, we can compute the cumulated distribution of sellers’ number of buyers in a market, and compare it with the data.²⁹ Figure 5 summarizes the results obtained when pooling all product×country pairs. Country-specific figures are reproduced in the appendix, Figure A.8. A visual inspection shows the model performs relatively well in the left-hand side of the distribution, but the right-tail is fatter in the data than in the model. The λ_i^k 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 reproduces almost perfectly the share of sellers serving a single buyer within a destination, that is, the fit is good regarding the curvature of CDF *and* its intercept. 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 5. Instead of aggregating across products within countries, we predict the share of sellers serving one buyer for each product-country pair. Table 5 reports

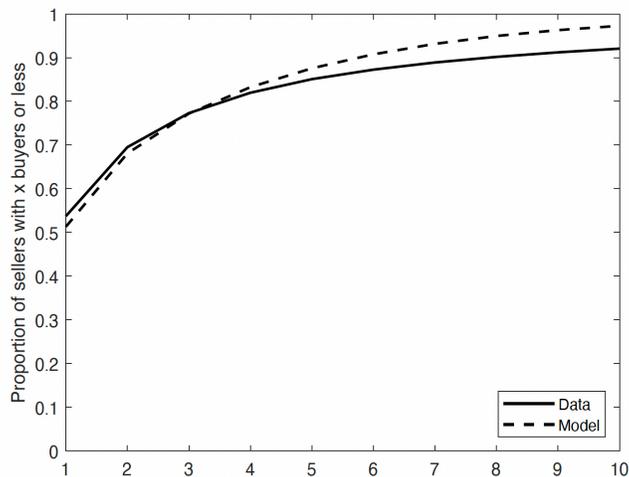
²⁹More precisely, we use the estimated λ_i^k coefficients to predict the share of exporters serving a given number of buyers, in each destination and product. These shares are then aggregated across products and countries using information on the relative number of suppliers of each product in France.

Table 4: *Search frictions and French market shares*

	(1)	(2)	(3)	(4)	(5)
Dep. Variable: log of product-level exports					
log Distance	-0.932*** (0.0852)	-0.920*** (0.0837)	-0.691*** (0.0978)	-0.466*** (0.0823)	-0.283*** (0.0847)
log Import demand	0.736*** (0.0214)	0.737*** (0.0214)	0.763*** (0.0223)	0.837*** (0.0214)	0.841*** (0.0209)
log GDP per capita	-0.437*** (0.0640)	-0.435*** (0.0637)	-0.378*** (0.0616)	-0.388*** (0.0531)	-0.465*** (0.0585)
Proba no match		-0.183*** (0.0518)	-0.174*** (0.0511)	-0.147*** (0.0489)	-0.165*** (0.0478)
Common language			1.229*** (0.352)		
Social connectedness				0.270*** (0.0274)	
Share migrants					0.273*** (0.0220)
Observations	12,247	12,247	12,247	12,247	12,247
R-squared	0.794	0.795	0.796	0.802	0.808

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. “Proba no match” is the estimated probability of the seller meeting with zero buyer in the destination. “Common language” is the probability that citizens from France and the destination speak a common language. “Social connectedness” is the social connectedness between France and destinations as measured by [Bailey et al. \(2020\)](#) using anonymized data from Facebook. “Share migrants” is the number of French citizens in the destination country, per 1,000 inhabitants. The dependent variable is the log of product-level bilateral exports.

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



Notes: The figure compares the CDF of the distribution of exporters' degree, in the data and the estimated model. The CDF is first simulated at the country \times product level based on estimated parameters, and then aggregated across country \times product pairs.

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 implies that our simple model accounts for 16% of the dispersion in the share of sellers serving a single buyer. The correlation with the predicted shares is highly significant in the three specifications, which confirms the correlation is valid within countries across products as well as across products within countries. Quantitatively similar results are obtained when we investigate the model's ability to explain the share of sellers with two or three buyers.

5 The distortive impact of search frictions

Having explained how firm-to-firm trade data can be used to recover estimates for search frictions, we now turn to the paper's main question, namely how such frictions distort the selection of firms in international markets.

Table 5: *Model fit: Share of one-buyer sellers*

	Dep.Var.: Empirical share of one buyer		
	(1)	(2)	(3)
Predicted share	0.284***	0.263***	0.171***
	(.006)	(.006)	(.005)
Constant	.392***		
	(.003)		
# obs	12,631	12,631	12,247
Fixed Effects	No	Country	Country Product
R-squared	.164	.245	.558

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.

5.1 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 it is the case in the data, using cross-sectoral measures of revealed comparative advantages and the dispersion in estimated frictions, across products.

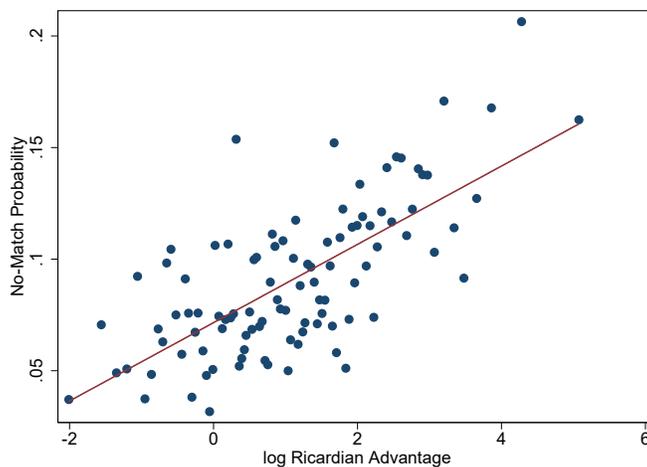
Revealed comparative advantages are measured using a strategy inspired from [Costinot et al. \(2012\)](#). Exploiting the gravity structure of the model, equation (2) can be used to recover a statistical decomposition of bilateral exports into its different variance components:

$$\ln \pi_{ij}^k = FE_i^k + FE_j^k + FE_{ij} + \varepsilon_{ij}^k, \quad (11)$$

where we now explicitly introduce the product dimension k . π_{ij}^k thus measures the share of producers from country j in country i 's consumption of product k . In this equation, the exporter-product fixed effect FE_j^k absorbs the impact of Ricardian technological advantages that affect a country's sales in all export destinations. A positive correlation between this term and estimated search frictions is thus indicative of magnified distortions, i.e. search frictions that are particularly high in those product markets in which French exporters have a Ricardian comparative advantage.

To test whether it is the case in the data, we first estimate equation (11) using the CEPII-BACI multilateral trade database available at the product level for 2007. Estimated revealed comparative advantages for France are then correlated with the

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 mean value of the no-match probability (averaged across destinations within a product).

product-specific average of estimated search frictions. Results shown in Figure 6 show a strong positive correlation between France’s comparative advantages and the average probability of the seller meeting with zero buyer in the destination, our proxy for the magnitude of search frictions. Note that the correlation is strongly positive, at 21%, despite the fact estimated revealed comparative advantages also absorb the mean level of dyadic trade frictions in such statistical decomposition (the mean of $\ln \lambda_{ij}^k (d_{ij}^k)^{-\theta^k}$ across destinations within a product and origin country). For this reason, there is a mechanical negative correlation between the estimated fixed effect and the zero-match probability that plays against the correlation recovered in Figure 6.³⁰ The positive 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.

³⁰The negative correlation would be reinforced if meeting probabilities were endogenous. If firms could invest in decreasing frictions, they should invest in markets in which they have a comparative advantage. The positive correlation recovered from the data suggests that this force, if it exists, is not strong enough to counteract the unconditional relationship between search frictions and comparative advantages.

5.2 Search frictions, productivity, and export performances

A consequence of the distortive impact of frictions, discussed in Section 3.2.2, is that search frictions also affect selection mechanisms *across* sellers *within* an origin country. In the frictionless benchmark, only the most efficient firms can expect serving foreign markets. In the frictional model instead, low-productivity firms have a strictly positive probability of serving at least one buyer in each destination and the probability is all the higher since frictions are severe. The mapping between the productivity and the export performances of French firms in export markets is suggestive of the size of competitive distortions induced by search frictions.

Table 6: *Impact of search frictions on the relative export performances of heterogeneously productive firms*

	Dep.Var.: ln firm-level bilateral exports					
	(1)	(2)	(3)	(4)	(5)	(6)
ln Domestic Sales	0.216*** (.011)	0.221*** (.011)				
- × Proba No match		-0.031*** (.012)				
1 Top Quartile Sectoral Sales			0.403*** (.033)			
- × Proba No match			-0.132** (.057)			
ln L Productivity				0.284*** (.022)	0.288*** (.023)	
- × Proba No match					-0.030 (.026)	
1 Top Quartile Sectoral L Prod.						0.303*** (.038)
- × Proba No match						-0.102** (.043)
Observations	451,625	451,625	451,625	433,535	433,535	433,535
R-squared	0.235	0.235	0.217	0.220	0.220	0.217
Fixed effects	Product	Product	Product	Product	Product	Product
	-Country	-Country	-Country	-Country	-Country	-Country

Notes: Standard errors, clustered in the firm dimension, are in parentheses, with ***, ** and * respectively denoting significance at the 1%, 5% and 10% levels. The dependent variable is the log of the firm’s exports to the destination country. The right-hand side variables include a proxy for the firm’s productivity, and its interaction with the level of frictions (“Proba No Match”). We use four alternative proxies for the firm’s productivity: the value of its domestic sales (“ln Domestic Sales”), a dummy variable equal to one if the firm’s domestic sales fall in the top quartile of the firm’s sector-specific distribution (“1 Top Quartile Sectoral Sales”), the firm’s apparent labor productivity computed as the ratio of value added over the number of employees (“ln L productivity”) and a dummy equal to one if the firm falls into the top quartile of the sector-specific distribution of labor productivities (“1 Top Quartile Sectoral L Prod.”).

To test this prediction of the model, we leverage upon external balance-sheet data on French firms to measure the relative productivity of exporters. We use two measures of productivity, the size of domestic sales and the firm’s apparent labor productivity. We then estimate heterogeneity in export performances across firms within a product and destination and how it varies depending on the strength of estimated search frictions. According to our model, we shall see the relative export performance of high-productivity firms being dampened in more frictional markets.

Results are summarized in Table 6. Columns (1)-(3) use domestic sales to characterize firms’ heterogeneity while columns (4)-(6) use the firm’s apparent labor productivity. Our measure of a firm’s export performance is based on the value of annual exports but we have reproduced the same analysis using the number of buyers instead and results are qualitatively unchanged. All regressions control for destination \times product fixed effects so that identification is across firms within the same market. We first confirm that, on average in our data, larger and more productive firms display better performances than smaller and less productive firms engaged in the same market (columns (1) and (4)). The remaining columns show that the export premium of large firms is dampened in markets that we estimate are more frictional. Indeed, the coefficient on the interaction between the firm’s productivity and the estimated no-match probability is always negative and significant, except when productivity is measured by the log of the firm’s labor productivity without controlling for the heterogeneity across sectors. In quantitative terms, the export premium of firms in the top quartile of the distribution of sectoral domestic sales (resp. sectoral labor productivity) increases from 31.9 to 40.3% (resp. 23.8 to 30.3%) when moving from the first to the ninth decile of the distribution of estimated no-match probabilities.³¹

5.3 Quantifying the efficiency loss induced by search frictions

To conclude the analysis, we run a counterfactual exercise that helps digging into the mechanism of the distortion just discussed. Namely, we simulate a drop in frictions that French exporters face, keeping parameters governing competition from the rest of the

³¹This empirical exercise also provides additional external validity to our interpretation of estimated parameters in terms of search frictions. As discussed in Section 3.3, one may be tempted to interpret the empirical moment used to identify search frictions in the context of alternative models of the matching of sellers and buyers in international markets. The role of buyers’ heterogeneity, which the model neglects, is an open question since a model of two-sided heterogeneity can a priori replicate any matching pattern in the data. However, the results in Table 6 are difficult to rationalize without relying on the distortion mechanism that is specific to our model.

world unchanged. The heterogeneous impact of the counterfactual reduction in search frictions induces composition effects that affect the mean productivity of exports and the export premium of high-productivity firms. In practice, our simulations exploit the prediction of the model summarized in equation (6). Given estimated and counterfactual frictions λ_i^k , observed and counterfactual market shares π_i^k and a calibrated mass of potential exporters $T_j^k z^{-\theta}$, the equation can be used to compute the consequences of reducing search frictions at different percentiles of the (Pareto) distribution.

We simulate a homogenous upward shift in the value of meeting probabilities, that is scaled by the observed heterogeneity in estimated parameters. Namely, we multiply all parameters by 4.2, which is the ratio of estimated coefficients at the first and second quartiles of the distribution. Results are summarized in Figure 7 which shows the probability of export (left panel) and the average number of buyers conditional on exporting (right panel), at various percentiles of the productivity distribution, in the data and the counterfactual world. Increasing meeting probabilities has a non-monotonic impact along the distribution of firms' productivities, with low-productivity firms seeing their export probability reduced and high-productivity firms gaining in terms of both export probability and their expected number of buyers, conditional on export. On average across product markets, the probability of exports is reduced for firms below the 68th productivity percentile whereas the increase in the expected number of buyers is above one after the 59th percentile, and above 5 in the top 15% of the distribution. Together, these results imply that the mean productivity of exporting firms increases following the drop in search frictions, by 5 to 10%.³²

³²By definition, the mean productivity of exporters writes:

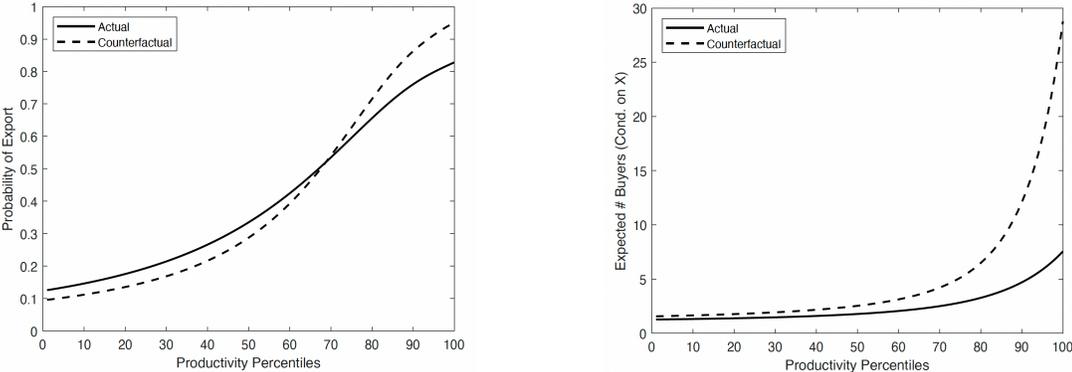
$$\mathbb{E}(Z|Export) = \frac{\int_z^{+\infty} z f(z) \mathbb{P}(Export|z) dz}{\int_z^{+\infty} f(z) \mathbb{P}(Export|z) dz}$$

where $f(z) = \frac{\theta z^{\theta-1}}{z^{\theta+1}}$ is the density of z and $\mathbb{P}(Export|z) = 1 - (1 - \rho_i^k(z))^{B_i^k}$ is the probability of exporting conditional 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:

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

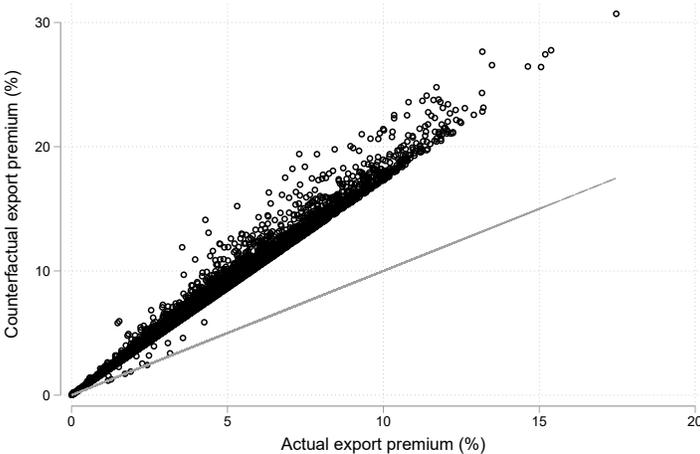
where the ^c superscript refers to the counterfactual state. After discretizing the productivity space in percentiles, this formula can be used, together with a calibrated value for θ , to recover the change in the mean productivity of exporters. For $\theta = 3$, the overall productivity improvement is found to be 9.34%, a value that is reduced to 5.43% for $\theta = 5$.

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



Notes: The figure shows the probability of serving at least one buyer (left panel) and the expected number of buyers conditional on exporting (right panel), in the average market of French firms, along the (product-specific) productivity distribution. The solid lines correspond to the actual equilibrium. The dashed lines are the counterfactual with reduced search frictions. The counterfactual is computed assuming that all meeting probabilities are shifted up, by a factor that corresponds to the ratio of estimated frictions at the median and first quartile of the distribution of estimated parameters. Results are then aggregated across products and countries using information on the relative number of firms in each product market.

Figure 8: *Impact of the drop in search frictions on product-level export premia*



Notes: The figure shows the actual and counterfactual export premium of firms at the 90th percentile relative to firms at the 25th percentile of the product-specific distribution of productivities, by product and destination. In the counterfactual exercise, the value of the meeting probability is multiplied by the ratio of estimated probabilities at the second and the first quartiles of the distribution. The solid line corresponds to the 45-degree line.

These numbers however hide a strong degree of heterogeneity across products and destinations, that we summarize in Figure 8. For each product \times country pair, we compute the export premium of firms at the 90th percentile of the productivity distribution, in comparison with firms at the 25th percentile. The export premium is computed in terms of the expected number of buyers in a particular destination. We then compare the actual value of the export premium with its counterfactual value in a world with lower search frictions. As expected, the export premium always increases in this simulation. The magnitude of the reallocation of export market shares from low to high productivity firms however varies across markets, from close to 0 to more than 300%. In our setting, a homogenous increase in meeting probabilities has more of an impact in product markets in which the initial market share of French firms is the largest such as the markets for live fish, fertilisers or iron and steel, most notably in Belgium and Luxembourg or in Estonia.

For comparison purposes, we ran another counterfactual exercise in which iceberg costs, instead of search frictions, are reduced product-by-product, 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 counterfactual experiment just described. Moving from the actual to this counterfactual equilibrium induces a substantial increase in export probabilities for French firms, from 40 to 70% on average. However, this increase in export probabilities does not induce an efficiency gain as in the case of a reduction in search frictions. The drop in iceberg costs actually benefits low-productivity firms, in relative terms. 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 more firms suffering from a lack of competitiveness. As a result, the average productivity of exporters decreases in this counterfactual experiment, with a drop of between 9 and 14%.

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 all exporters in foreign markets. 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. This heterogeneous impact of frictions along the productivity distribution is a distinctive feature of our model. In highly frictional markets, the export premium of high-productivity firms is lowered and the export probability of small and medium firms increased.

We provide direct evidence of such distortion in our data. Bilateral search frictions are first 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 size of its customer base 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 the destination. Intuitively, more frictional markets induce more distortions, 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 correlate with country and product attributes in a theoretically consistent way.

The estimated frictions are especially large in product markets where French firms have a comparative advantage, on average. Moreover, we provide evidence that the export premium of high-productivity firms is systematically reduced in markets that we estimate are more frictional. Using a counterfactual experiment, we show that reducing search frictions can generate sizeable gains in terms of the efficiency of selection mechanisms into export. Shifting meeting probabilities up by an amount that corresponds to the ratio of these probabilities at the second and first quartile of the distribution

reduces export probabilities up to the 68th percentile of the distribution of productivities while increasing export probabilities and the customer base of high-productivity exporters. The average productivity of exporters rises as a result, with an increase ranging from 5% to 10%. A comparable drop in iceberg costs would instead reduce the mean productivity of exporters, by about 10%.

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.
- Anderson, S. P., de Palma, A., and Thisse, J.-F. (1992). *Discrete Choice Theory of Product Differentiation*. MIT Press.
- Antràs, P., Fort, T. C., and Tintelnot, F. (2017). The Margins of Global Sourcing:

- Theory and Evidence from US Firms. *American Economic Review*, 107(9):2514–2564.
- Arkolakis, C. (2010). Market Penetration Costs and the New Consumers Margin in International Trade. *Journal of Political Economy*, 118(6):1151 – 1199.
- Arkolakis, C., Costinot, A., and Rodríguez-Clare, A. (2012). New trade models, same old gains? *American Economic Review*, 102(1):94–130.
- Atalay, E., Hortacsu, A., and Syverson, C. (2014). Vertical Integration and Input Flows. *American Economic Review*, 104(4):1120–1148.
- Atkeson, A. and Burstein, A. (2008). Pricing-to-Market, Trade Costs, and International Relative Prices. *American Economic Review*, 98(5):1998–2031.
- Bailey, M., Gupta, A., Hillenbrand, S., Kuchler, T., Richmond, R. J., and Stroebel, J. (2020). International Trade and Social Connectedness. NBER Working Papers 26960, National Bureau of Economic Research, Inc.
- Bergounhon, F., Lenoir, C., and Mejean, I. (2018). A guideline to French firm-level trade data.
- Bernard, A. B., Eaton, J., Jensen, J. B., and Kortum, S. (2003). Plants and productivity in international trade. *American Economic Review*, 93(4):1268–1290.
- 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.

- Chen, M. X. and Wu, M. (2021). The Value of Reputation in Trade: Evidence from Alibaba. *The Review of Economics and Statistics*, 103(5):857–873.
- 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.
- Costinot, A. and Rodriguez-Clare, A. (2014). Trade Theory with Numbers: Quantifying the Consequences of Globalization. In Gopinath, G., Helpman, ., and Rogoff, K., editors, *Handbook of International Economics*, volume 4 of *Handbook of International Economics*, chapter 0, pages 197–261. Elsevier.
- 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.
- Drazanova, L. (2020). Introducing the Historical Index of Ethnic Fractionalization (HIEF) Dataset: Accounting for Longitudinal Changes in Ethnic Diversity. *Journal of Open Humanities Data*, 6(1):6.
- Drozd, L. A. and Nosal, J. B. (2012). Understanding International Prices: Customers as Capital. *American Economic Review*, 102(1):364–395.
- Eaton, J. and Kortum, S. (2002). Technology, Geography, and Trade. *Econometrica*, 70(5):1741–1779.
- Eaton, J. and Kortum, S. (2010). *Technology in the Global Economy: A Framework for Quantitative Analysis*.
- Eaton, J., Kortum, S., and Kramarz, F. (2020). 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.

- Fontaine, F., Martin, J., and Mejean, I. (2021). Frictions and adjustments in firm-to-firm trade.
- 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 24807, National Bureau of Economic Research, Inc.
- 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.
- Imbs, J. and Mejean, I. (2015). Elasticity Optimism. *American Economic Journal: Macroeconomics*, 7(3):43–83.
- Khandelwal, A. (2010). The Long and Short (of) Quality Ladders. *Review of Economic Studies*, 77(4):1450–1476.
- 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, J. and Toubal, F. (2014). Native language, spoken language, translation and trade. *Journal of International Economics*, 93(2):351–363.
- 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.
- Nunn, N. (2007). Relationship-Specificity, Incomplete Contracts, and the Pattern of Trade. *The Quarterly Journal of Economics*, 122(2):569–600.
- Ornelas, E., Turner, J. L., and Bickwit, G. (2021). Preferential trade agreements and global sourcing. *Journal of International Economics*, 128(C).
- 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.
- 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 Theoretical appendix

A.1 Analytical details

Under the model's assumptions, the number of suppliers from j offering a price below p is drawn from a Poisson distribution of parameter $\lambda_{ij}\mu_{ij}(p) = \lambda_{ij}T_j(d_{ij}w_j)^{-\theta}p^\theta = \lambda_{ij}v_{ij}p^\theta$ where we define $v_{ij} \equiv T_j(d_{ij}w_j)^{-\theta}$ to alleviate notations. Likewise, the number of suppliers from any country offering a price below p is drawn from a Poisson distribution of parameter $\sum_j \lambda_{ij}\mu_{ij}(p) = \sum_j \lambda_{ij}T_j(d_{ij}w_j)^{-\theta}p^\theta = \kappa_i \Upsilon_i p^\theta$.

Preliminary results: The assumption of Poisson draws makes it possible to derive a number of useful properties for the distribution of prices offered to buyers in country i . The following theorem characterizes the joint distributions of prices offered to a particular buyer, when we note $P_i^{(n)}$ the n 'th lowest price offer received by a buyer in i and $P_{ij}^{(n)}$ the n 'th lowest price offer received by a buyer in i from a seller in j .³³

Theorem 1. The joint density of $P_i^{(n)}$ and $P_i^{(n+1)}$ is:

$$g_{i,n,n+1}(p_n, p_{n+1}) = \frac{\theta^2}{(n-1)!} (\kappa_i \Upsilon_i)^{n+1} p_n^{\theta n-1} p_{n+1}^{\theta-1} \exp[-\kappa_i \Upsilon_i p_{n+1}^\theta]$$

for $0 < p_n \leq p_{n+1} < \infty$. The marginal density of $P_i^{(n)}$ is:

$$g_{i,n}(p) = \frac{\theta}{(n-1)!} (\kappa_i \Upsilon_i)^n p^{\theta n-1} \exp[-\kappa_i \Upsilon_i p^\theta]$$

for $0 < p < \infty$.

Likewise, the joint density of $P_{ij}^{(n)}$ and $P_{ij}^{(n+1)}$ is:

$$g_{ij,n,n+1}(p_n, p_{n+1}) = \frac{\theta^2}{(n-1)!} (\lambda_{ij}v_{ij})^{n+1} p_n^{\theta n-1} p_{n+1}^{\theta-1} \exp[-\lambda_{ij}v_{ij}p_{n+1}^\theta]$$

for $0 < p_n \leq p_{n+1} < \infty$ while the marginal density of $P_{ij}^{(n)}$ is:

$$g_{ij,n}(p) = \frac{\theta}{(n-1)!} (\lambda_{ij}v_{ij})^n p^{\theta n-1} \exp[-\lambda_{ij}v_{ij}p^\theta]$$

for $0 < p < \infty$.

³³Here, we closely follow the steps in [Eaton and Kortum \(2010\)](#), chapter 4.

Proof. Under the model's assumptions, the distribution of P_i given $P_i \leq \bar{p}$ is:

$$F(p|\bar{p}) = \begin{cases} \left(\frac{p}{\bar{p}}\right)^\theta & \text{if } p \leq \bar{p} \\ 1 & \text{if } p > \bar{p} \end{cases}$$

The probability that a price is less than p_n is $F(p_n|\bar{p})$ and the probability that a price is more than p_{n+1} is $(1 - F(p_{n+1}|\bar{p}))$. Hence, if the buyer has met with m sellers with price below \bar{p} , the probability that n are lower than p_n and $m - n$ are greater than p_{n+1} is:

$$Pr \left[P_i^{(n)} \leq p_n, P_i^{(n+1)} \geq p_{n+1} | m \right] = \binom{n}{m} F(p_n|\bar{p})^n (1 - F(p_{n+1}|\bar{p}))^{m-n}$$

Taking the negative of the cross-derivative of this expression with respect to p_n and p_{n+1} gives the joint density of $P_i^{(n)}$ and $P_i^{(n+1)}$, conditional on m :

$$g_{i,n,n+1}(p_n, p_{n+1} | \bar{p}, m) = \frac{m! F(p_n|\bar{p})^{n-1} (1 - F(p_{n+1}|\bar{p}))^{m-n-1} F'(p_n|\bar{p}) F'(p_{n+1}|\bar{p})}{(n-1)!(m-n-1)!}$$

for $p_{n+1} \leq p_n$ and $m \leq n+1$. For $m < n+1$, $g_{i,n,n+1}(p_n, p_{n+1} | \bar{p}, m) = 0$.

The number m of price quotes is drawn from a Poisson distribution with parameter $\kappa_i \Upsilon_i \bar{p}^\theta$. The expectation of the joint distribution unconditional on m is thus:

$$\begin{aligned} g_{i,n,n+1}(p_n, p_{n+1} | \bar{p}) &= \sum_{u=0}^{\infty} \frac{\exp[-\kappa_i \Upsilon_i \bar{p}^\theta] (\kappa_i \Upsilon_i \bar{p}^\theta)^u}{u!} g_{i,n,n+1}(p_n, p_{n+1} | \bar{p}, m) \\ &= \frac{F(p_n|\bar{p})^{n-1} (\kappa_i \Upsilon_i \bar{p}^\theta)^{n+1} \exp[-\kappa_i \Upsilon_i \bar{p}^\theta F(p_{n+1}|\bar{p})] F'(p_n|\bar{p}) F'(p_{n+1}|\bar{p})}{(n-1)!} \\ &= \frac{\theta^2}{(n-1)!} (\kappa_i \Upsilon_i)^{n+1} p_n^{\theta n-1} p_{n+1}^{\theta-1} \exp[-\kappa_i \Upsilon_i p_{n+1}^\theta] \end{aligned}$$

which is the expression in Theorem 1 for $\bar{p} \rightarrow \infty$. The marginal density comes immediately from:

$$g_{i,n}(p) = \int_p^\infty g_{i,n,n+1}(p, p_{n+1}) dp_{n+1}$$

□

Theorem 1 thus characterizes the joint distribution of each pair of adjacent order statistics, in the overall subset of offers received by a buyer in i and in the subset of

offers originating from j . These distributions solely depend on θ , $\kappa_i \Upsilon_i$ and $\lambda_{ij} \nu_{ij}$.

Another useful property of random variables described by the marginal distribution in Theorem 1 is summarized in Theorem 2:

Theorem 2. For each order n , the b 'th moment ($b > -\theta n$) is:

$$E \left[\left(P_i^{(n)} \right)^b \right] = (\kappa_i \Upsilon_i)^{-b/\theta} \frac{\Gamma[(\theta n + b)/\theta]}{(n-1)!}$$

where $\Gamma(\alpha) = \int_0^\infty y^{\alpha-1} e^{-y} dy$ is the gamma function.

Likewise,

$$E \left[\left(P_{ij}^{(n)} \right)^b \right] = (\lambda_{ij} \nu_{ij})^{-b/\theta} \frac{\Gamma[(\theta n + b)/\theta]}{(n-1)!}$$

Proof. First consider $k = 1$:

$$\begin{aligned} E \left[\left(P_i^{(1)} \right)^b \right] &= \int_0^\infty p^b g_{i,1}(p) dp \\ &= \int_0^\infty (\kappa_i \Upsilon_i) \theta p^{\theta+b-1} \exp[-\kappa_i \Upsilon_i p^\theta] dp \end{aligned}$$

We now use a change of variable: $v = \kappa_i \Upsilon_i p^\theta$:

$$\begin{aligned} E \left[\left(P_i^{(1)} \right)^b \right] &= \int_0^\infty \left(\frac{v}{\kappa_i \Upsilon_i} \right)^{b/\theta} \exp[-v] dv \\ &= (\kappa_i \Upsilon_i)^{-b/\theta} \Gamma \left(\frac{\theta + b}{\theta} \right) \end{aligned}$$

which is defined for $\theta + b > 0$.

More generally:

$$E \left[\left(P_i^{(n)} \right)^b \right] = \int_0^\infty p^b g_{i,n}(p) dp$$

□

Weibull distribution of prices: Based on Theorem 1, one can recover the distribution of the n 'th lowest price $P_i^{(n)}$:

$$F_{i,n}(p) \equiv Pr[P_i^{(n)} \leq p] = 1 - \sum_{u=0}^{n-1} \frac{(\kappa_i \Upsilon_i p^\theta)^u}{u!} \exp[-\kappa_i \Upsilon_i p^\theta]$$

As is necessary for a cumulative distribution, $F_{i,n}(p)$ approaches 1 when p tends to infinity and $F'_{i,n}(p) = g_{i,n}(p)$.

In particular, the lowest price is distributed:

$$F_{i,1}(p) \equiv Pr[P_i^{(1)} \leq p] = 1 - \exp[-\kappa_i \Upsilon_i p^\theta]$$

which is the Weibull distribution used in the main text.

Likewise, the distribution of the lowest price received from j is:

$$F_{ij,1}(p) \equiv Pr[P_{ij}^{(1)} \leq p] = 1 - \exp[-\lambda_{ij} \nu_{ij} p^\theta]$$

Bilateral trade probabilities: Section 3.2.1 analyzes the share of buyers from i purchasing the product from country j . When the number of buyers is large enough, the share is also the expected value of $\mathbb{1}_{b_{ij}}^{(1)}$, a dummy variable that is equal to 1 if the lowest price received by buyer b_i originates from country j .

To derive this expected value, we first derive the probability that the lowest-cost seller from j that buyer b_i meets is the lowest cost supplier overall, conditional on a price p . By definition, this probability is equal to the probability that all lowest-cost sellers from a different country offer a price above p :

$$\prod_{j' \neq j} [1 - F_{ij',1}(p)] = \exp \left[-p^\theta \sum_{j' \neq j} \lambda_{ij'} \nu_{ij'} \right]$$

Integrating over p gives the probability that the lowest-cost supplier met from j is the lowest-cost supplier met:

$$\begin{aligned} E \left[\mathbb{1}_{b_{ij}}^{(1)} \right] &= \int_0^\infty \exp \left[-p^\theta \sum_{j' \neq j} \lambda_{ij'} \nu_{ij'} \right] dF_{ij,1}(p) \\ &= \frac{\lambda_{ij} \nu_{ij}}{\kappa_i \Upsilon_i} [1 - F_{i,1}(p)]_0^\infty \\ &= \frac{\lambda_{ij} \nu_{ij}}{\kappa_i \Upsilon_i} \end{aligned}$$

Bilateral trade shares: Under iso-elastic preferences, the nominal demand expressed by a buyer b_i is a function of the lowest price received, at the power $1 - \sigma$:

$$p_{b_i} c_{b_i} = \left(P_{b_i}^{(1)} \right)^{1-\sigma} \bar{X}_i$$

The expected value of bilateral imports from j to i can thus be written as the

expected value of individual purchases, across buyers in i that end up purchasing the good from j , i.e. conditional on the lowest-cost supplier from j that the buyer has met offering a price below the lowest-cost supplier of any other country:

$$\begin{aligned}\mathbb{E} [p_{b_i c_{b_i}} | \mathbb{1}_{b_{i,j}}^{(1)} = 1] &= \bar{X}_i \int_0^\infty p^{1-\sigma} \exp \left[-p^\theta \sum_{j' \neq j} \lambda_{ij'} v_{ij'} \right] dF_{ij,1}(p) \\ &= \bar{X}_i \theta \lambda_{ij} v_{ij} \int_0^\infty p^{\theta-\sigma} \exp \left[-p^\theta \kappa_i \Upsilon_i \right] dp\end{aligned}$$

To derive bilateral trade shares, this expression must be compared with the expected value of individual purchases, irrespective of the source country:

$$\begin{aligned}\mathbb{E} [p_{b_i c_{b_i}}] &= \bar{X}_i \int_0^\infty p^{1-\sigma} dF_{i,1}(p) \\ &= \bar{X}_i \theta \kappa_i \Upsilon_i \int_0^\infty p^{\theta-\sigma} \exp \left[-p^\theta \kappa_i \Upsilon_i \right] dp \\ &= \bar{X}_i (\kappa_i \Upsilon_i)^{\frac{\sigma-1}{\theta}} \Gamma \left(\frac{\theta + \sigma - 1}{\theta} \right)\end{aligned}$$

Taking the ratio of the two terms and simplifying gives:

$$\begin{aligned}\pi_{ij} &= \frac{\mathbb{E} [p_{b_i c_{b_i}} | \mathbb{1}_{b_{i,j}}^{(1)} = 1]}{\mathbb{E} [p_{b_i c_{b_i}}]} \\ &= \frac{\lambda_{ij} v_{ij}}{\kappa_i \Upsilon_i}\end{aligned}$$

As in [Eaton and Kortum \(2002\)](#), trade shares are fully summarized by the probability that any supplier from j ends up serving market i . The reason is that, conditional on the identity of the seller, the distribution of prices offered to buyers in i is the same whatever the origin of the seller. In this context, trade shares only depend on the likelihood that a seller from j is the lowest-cost supplier met by a buyer from i . In our model, the probability depends on country j 's comparative advantage in market i (v_{ij}/Υ_i) and the relative size of frictions (λ_{ij}/κ_i).

As discussed in the text, the semi-elasticity of this trade share with respect to the bilateral search parameter is unambiguously positive:

$$\begin{aligned}\frac{d \ln \pi_{ij}}{d \lambda_{ij}} &= \frac{1}{\lambda_{ij}} - \frac{1}{\kappa_i} \frac{d \kappa_i}{d \lambda_{ij}} \\ &= \frac{1 - \pi_{ij}}{\lambda_{ij}} > 0\end{aligned}$$

A.2 Plugging the model into a general equilibrium structure

The analysis developed in section 3 describes the matching equilibrium in a sector whose production costs are taken as exogenous as in partial equilibrium. We now show how this structure can be plugged into a general equilibrium framework. In this general equilibrium structure, the assumptions in the main text describe the matching of consumers and producers within a particular sector k . Following [Atkeson and Burstein \(2008\)](#) and [Gaubert and Itskhoki \(2018\)](#), the aggregate impact of the discreteness at product level is neglected by assuming the economy displays a continuum of products. A consumer in country i consumes a CES bundle of products:

$$c_{b_i} = \left[\int_0^1 (c_{b_i}^k)^{\frac{\sigma-1}{\sigma}} dk \right]^{\frac{\sigma}{\sigma-1}}$$

with σ the elasticity of substitution between product-level consumptions. Consumption at product level follows the assumptions in section 3. Each consumer b_i meets with a random number of potential suppliers for variety k , chooses the lowest-cost supplier met and pays the price:

$$p_{b_i}^k = \arg \min \left\{ \frac{w_j d_{ij}^k}{z_{s_j}}; s_j \in \Omega_{b_i}^k \right\},$$

where $\Omega_{b_i}^k$ now denotes the set of producers of variety k met by buyer b_i and d_{ij}^k is the product-specific iceberg cost. The price of the input bundle is instead assumed homogenous across products and we will now interpret it as the wage rate: Firms produce out of labor with a constant returns to scale technology and labor is perfectly mobile across sectors.

To solve the model, it is assumed that individual consumers maximize aggregate consumption based on the expected price index, i.e. they neglect the aggregate impact of the randomness in the matching process:

$$\begin{cases} \max_{\{c_{b_i}^k\}_{k \in [0,1]}} & \left[\int_0^1 (c_{b_i}^k)^{\frac{\sigma-1}{\sigma}} dk \right]^{\frac{\sigma}{\sigma-1}} \\ \text{s.t.} & \mathbb{E} \left[\int_0^1 p_{b_i}^k c_{b_i}^k dk \right] \leq \frac{R_i}{B_i} \end{cases}$$

R_i is the country's aggregate income that we assume is shared equally across buyers. In equilibrium $R_i = w_i B_i + \Pi_i$ where Π_i denotes aggregate profits, assumed to be distributed lump-sum to all consumers.

Under these assumptions, the demand addressed to the lowest cost supplier met by consumer b_i writes:

$$p_{b_i}^k c_{b_i}^k = \left(\frac{p_{b_i}^k}{P_i} \right)^{1-\sigma} \frac{R_i}{B_i}$$

where $P_i = \mathbb{E} \left[\int_0^1 (p_{b_i}^k)^{1-\sigma} dk \right]^{\frac{1}{1-\sigma}}$ is the expected price index in country i .

This demand function is consistent with the assumption in the main text under the notation $\bar{X}_i \equiv \frac{R_i}{B_i} P_i^{\sigma-1}$. As shown in Appendix A.1, aggregating these demand functions across buyers within a product implies:

$$\begin{aligned} X_i^k &= B_i \mathbb{E} [p_{b_i}^k c_{b_i}^k] = B_i \bar{X}_i (\kappa_i^k \Upsilon_i^k)^{\frac{\sigma-1}{\theta}} \Gamma \left(\frac{\theta + \sigma - 1}{\theta} \right) \\ \pi_{ij}^k &= \frac{X_{ij}^k}{X_i^k} = \frac{T_j^k (d_{ij}^k w_j)^{-\theta k}}{\Upsilon_i^k} \frac{\lambda_{ij}^k}{\kappa_i^k} \end{aligned}$$

where the notations are the same as in the main text except that we explicitly introduce the product dimension that was neglected to alleviate notations.

The model is closed using the trade balance condition. For each pair of countries i and j , we have:

$$\int_0^1 \pi_{ij}^k dk = \int_0^1 \pi_{ji}^k dk$$

These conditions for all pairs of countries define a system of equations that can be used to solve for equilibrium factor prices $\{w_j\}_{j=1\dots N}$. Solving the model numerically requires estimates for all parameters, most notably the whole vector of search frictions $\lambda_{ij}^k, \forall k \in [0, 1], i = 1, \dots, N$ and $j = 1 \dots N$. In the main text, we propose a strategy to estimate λ_{ij}^k from firm-to-firm trade data. Because we have access to data for French exporters only, we are unable to recover the whole set of parameters necessary to solve the model in general equilibrium.

The general equilibrium extension discussed in this section however gives intuition for how search frictions affect welfare in general equilibrium. The gravity structure at product-level makes it possible to compare our theoretical framework with other trade models displaying structural gravity, most notably the multi-sector extension of Eaton and Kortum (2002) in Caliendo and Parro (2015). As discussed in Costinot and Rodriguez-Clare (2014), a gravity structure, together with other common micro and macro restrictions, defines a general class of trade models which welfare predictions can

be summarized using a simple formula that solely involves trade shares and measures of trade elasticities. As long as the search frictions we introduced in the model are technological constraints that the planner faces, so that the planner's solution coincides with the decentralized equilibrium, the results in [Costinot and Rodriguez-Clare \(2014\)](#) are likely to apply in our model as in the class of multi-sector models they discuss. By distorting the geography of trade, search frictions will thus affect the welfare benefits from trade.

A.3 Proof of proposition 2

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

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

Depending on the current level of frictions (λ_{ij}), the expected number of firms in country j ($T_j \underline{z}^{-\theta}$) and the position of the firm in the productivity distribution ($\left(\frac{z}{\underline{z}}\right)^{-\theta}$), the derivative can be positive or negative. It is more positive for high values of z . At the limit, $\lim_{z \rightarrow +\infty} \frac{\partial \ln \rho_{ij}(z)}{\partial \lambda_{ij}} = \frac{1}{\lambda_{ij}}$. 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 $\lambda_{ij} > \frac{1}{T_j \underline{z}^{-\theta}}$, 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: $\frac{\partial \ln \rho_{ij}(\underline{z})}{\partial \lambda_{ij}} < 0$, where $\rho_{ij}(\underline{z})$ denotes the export probability of the least productive firm.

This non-monotonicity is to be compared with the sensitivity of export probabilities to iceberg trade costs, which 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.

A.4 Alternative market structure assumption

Results in the main text rely on the assumption that firms price at their marginal cost. The randomness induced by matching frictions however increases the market power of the lowest-cost supplier met by a particular buyer. Even if the firm has announced a price at its marginal cost, it has an incentive to deviate ex-post and exploit its market power. As we now show, the model is flexible enough to handle more realistic price strategies.

Suppose that firms in each buyer's random choicest compete à la Bertrand. Under Bertrand competition, the lowest-cost producer ends up setting the price which is just sufficient to beat the second lowest-cost supplier, unless this price is above the monopoly price.³⁴ As in the baseline discussed in the main text, each buyer ends up purchasing the product from the lowest-cost supplier she has met. The price that she pays is however equal to the marginal cost of the second lowest-cost supplier.

In this setting, bilateral trade probabilities are the same as in the baseline case. The value of trade, conditional on a match, is however a function of the distribution of the second lowest-cost supplier, through the lowest-cost supplier's pricing function. As demonstrated in [Bernard et al. \(2003\)](#), endogenous mark-ups do not distort the geography of trade, and thus trade shares are still equal to bilateral trade probabilities. The reason is that the distribution of markups is the same whatever the origin of the firm setting this markup: within a destination, no source sells at systematically higher markups. This result is summarized in Theorem 3.

³⁴In the set-up under study, the monopoly markup is $\frac{\sigma}{\sigma-1}$ for $\sigma > 1$.

Theorem 3. Under Bertrand competition, the distribution of markups set to buyers in country i writes:

$$H_i(m) = P[M_i \leq m] = 1 - m^{-\theta}$$

with M_i the ratio of the price set by the lowest cost supplier over its marginal cost, $P_i^{(2)}/P_i^{(1)}$ if we keep the notations used in Appendix A.1.

Proof. The distribution of the second to the first lowest cost $P_i^{(2)}/P_i^{(1)}$, conditional on $P_i^{(2)} = p_2$ is:

$$\begin{aligned} \mathbb{P}\left[\frac{P_i^{(2)}}{P_i^{(1)}} \leq m \mid P_i^{(2)} = p_2\right] &= \mathbb{P}\left[P_i^{(1)} \geq \frac{p_2}{m} \mid P_i^{(2)} = p_2\right] \\ &= 1 - \mathbb{P}\left[P_i^{(1)} \leq \frac{p_2}{m} \mid P_i^{(2)} = p_2\right] \\ &= 1 - \int_0^{\frac{p_2}{m}} \frac{g_{i,1,2}(p, p_2)}{g_{i,2}(p_2)} dp \\ &= 1 - \left(\frac{p_2}{m}\right)^\theta \\ &= 1 - (m)^{-\theta} \end{aligned}$$

The unconditional distribution $H_i(m)$ comes immediately. \square

Whereas the baseline model sticks to the assumption of marginal cost pricing, the result in Theorem 3 shows that the main conclusions would be left unchanged if we instead assumed Bertrand competition among the random choicest of firms met by a particular buyer. The pricing assumption instead has consequences for the dynamics of trade adjustment to shocks, that is studied into more details in [Fontaine et al. \(2021\)](#).

A.5 Increasing meeting probabilities

One may wonder whether imposing the same meeting probability to all firms, whatever their productivity, is a key driver of the result. An alternative approach would assume the meeting probability to be increasing in the firm's productivity. Such increasing relationship would for instance emerge under endogenous search effort. In a reduced-form set-up, this assumption would imply that the meeting probability for a firm of productivity z in country j that seeks to serve market i can be summarized by

$$\lambda_{ij}(z) = f(\lambda_{ij}, z)$$

with $\frac{df(\lambda_{ij}, z)}{d\lambda_{ij}} > 0$ and $\frac{df(\lambda_{ij}, z)}{dz} > 0$, i.e. high-productivity firms meet more buyers on average but more structural frictions reduce meeting probabilities at each point of the productivity distribution.

Under such assumption, the probability for a firm with productivity z_{s_j} to serve a buyer in i is still equal to the probability of a match times the probability of being the lowest cost supplier, conditional on this match. However, the cross-derivative of $\rho_{ij}(z_{s_j})$ with respect to the (exogenous) search parameter and the firm's productivity now takes a more complicated form:

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

As in the benchmark case, the second term is likely to be negative and increasing in z_{s_j} . The second derivative of the probability of serving the buyer conditional on a match with respect to λ_{ij} and z_{s_j} is expected larger than in the baseline, however. The reason is that 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 distortive 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.

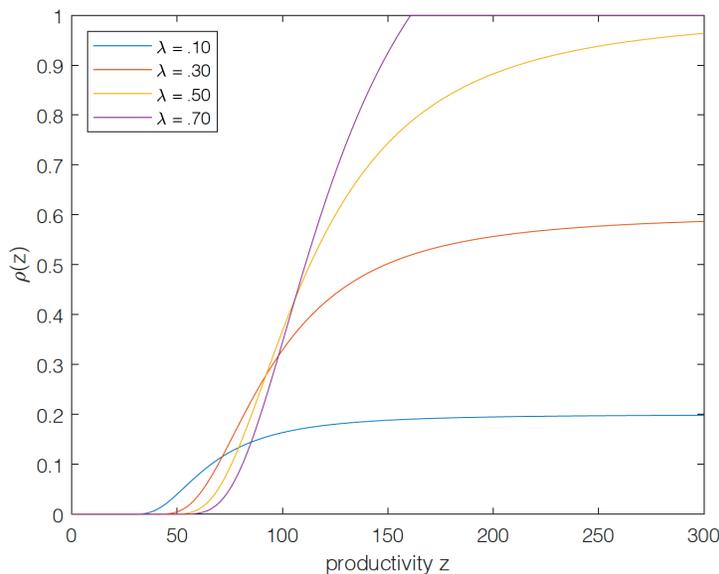
Figure A.1 illustrates the impact of varying the meeting probability under a specific parametric assumption, which can be compared with Figure 2 in the benchmark case. Namely, we simulate the model assuming:

$$\lambda_{ij}(z_{s_j}) = 2\lambda_{ij} \left[1 - \left(\frac{z_{s_j}}{\underline{z}} \right)^{-\theta} \right] \quad (12)$$

Under this parametric assumption, the expected meeting probability is equal to λ_{ij} , as in the baseline, but now varies between 0 and $2\lambda_{ij}$ along the productivity distribution. Assuming the meeting probability to be increasing in the seller's productivity mechanically increases the likelihood that a high-productivity seller will end up serving a foreign buyer. As a consequence, the probability of a match is larger at the top of

the distribution in the extended model than in the benchmark case. Whereas the level probabilities are different in the baseline and extended models, the extended model still displays the log-supermodularity in λ_{ij} and z_{sj} , which we exploit in the empirical section. For this reason, we believe that our empirical strategy is not threatened by the strong parametric assumption imposed to recover analytical results.

Figure A.1: *Probability of serving a buyer as a function of the seller's productivity under increasing meeting probabilities*



Notes: This figure illustrates how the probability of serving a buyer varies with the seller's productivity, for four different values of bilateral frictions.

A.6 A model of heterogeneous consumers

In this section, we discuss the predictions of another class of models which the trade literature often compares with Ricardian models à la [Eaton and Kortum \(2002\)](#) or the monopolistic competition structure in [Melitz \(2003\)](#), namely models displaying horizontal differentiation and heterogeneous consumers. As discussed in [Anderson et al. \(1992\)](#) and [Head and Mayer \(2014\)](#), discrete choice models of demand for horizontally differentiated varieties can deliver a gravity structure under the adequate parametric assumptions. In this section, we show how it is possible to interpret our predictions

in the context of this class of models. We also show that the moment used to identify search frictions in Section 4 is likely to be orthogonal to the parameters of such a model.

Consider a model in which buyers display heterogeneous preferences with regard to the varieties produced by the (discrete) set of worldwide producers. More specifically, let us denote $\psi_{b_i s_j}$ a random variable characterizing the preference of buyer b_i with respect to variety s_j . Following the literature, we will assume that the preference parameters are independently drawn into a Pareto distribution of shape θ :

$$\mathbb{P}(\psi_{b_i s_j} \geq \underline{\psi}) = \begin{cases} \left(\frac{\psi}{\underline{\psi}}\right)^{-\theta} & \text{if } \psi \geq \underline{\psi} \\ 1 & \text{if } \psi < \underline{\psi} \end{cases}$$

Let us further assume that consumers' preferences are systematically biased towards varieties produced in some countries. Namely, the number of varieties from j that deliver a preference parameter above ψ in country i is assumed to follow a Poisson distribution of parameter $T_j \lambda_{ij} \psi^{-\theta}$. T_j can be interpreted as the mean quality of varieties produced in j whereas λ_{ij} measures a dyadic preference bias.

Finally, suppose as in the paper's baseline model that the cost of serving market i from j is equal to $w_j d_{ij}$ and that the revenue of a consumer in country i is equal to $x_i \equiv \frac{R_i}{B_i}$. Under these assumptions, the indirect utility recovered from the consumption of variety s_j writes:³⁵

$$V_{b_i s_j} = \frac{x_i}{w_j d_{ij}} \psi_{b_i s_j}$$

Using the properties of the Poisson distribution, it is straightforward to show that the number of varieties from j (resp. from any country) delivering an indirect utility above v is distributed Poisson of parameter $v^{-\theta} c_i^\theta T_j \lambda_{ij} (w_j d_{ij})^{-\theta}$ (resp. $v^{-\theta} c_i^\theta \sum_j T_j \lambda_{ij} (w_j d_{ij})^{-\theta}$).

In this model, the probability that a buyer b_i chooses a variety produced in country j , conditional on the variety delivering indirect utility above v , writes:

$$\mathbb{P}\left(V_{b_i}^{(1)} \text{ originates from } j \mid V_{b_i}^{(1)} \geq v\right) = \frac{v^{-\theta} c_i^\theta (w_j d_{ij})^{-\theta} T_j \lambda_{ij}}{v^{-\theta} c_i^\theta \sum_j (w_j d_{ij})^{-\theta} T_j \lambda_{ij}}$$

As the probability is independent from v and homogenous across buyers, it is also equal to the probability that any buyer from i purchases the variety from j , which is also the

³⁵As in the paper's model, it is implicitly assumed that the seller s_j serves buyer b_i at marginal cost. Assuming instead that the seller exploits her competitive advantage to price at the second lowest preference-adjusted marginal cost would complicate the analysis although results regarding the selection of firms into exporting would be left unaffected.

trade share in this model in which the value of trade is homogenous across buyers in expectation:

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

where $\Upsilon_i \equiv \sum_j (w_j d_{ij})^{-\theta} T_j$ and $\kappa_i = \frac{\sum_j (w_j d_{ij})^{-\theta} T_j \lambda_{ij}}{\sum_j (w_j d_{ij})^{-\theta} T_j}$.

Based on bilateral trade shares, it is thus not possible to discriminate our model, that combines Ricardian comparative advantages with search frictions, and a model that displays (biased) heterogeneous preferences. As discussed in the main text though, this prediction of the model is not sufficient to identify search frictions anyway. Indeed, the geography of trade involves two dyadic components, d_{ij} and λ_{ij} , which cannot be separated in the data based on predicted and observed trade shares. What the discussion in this section adds is that, even if we were able to control for iceberg costs, using this prediction of the model would not be desirable because the same structural equation arises from a completely different model in which λ_{ij} interprets as a preference parameter instead of a search friction.

Our identification strategy does not exclusively relies on the gravity structure of the model though. Instead, we exploit the model's prediction regarding *individual* trade patterns. Namely, the moment used for identification is based on the expected number of firms from a given country that serve exactly M buyers in i :

$$h_{ij}(M) = \int C_{B_i}^M \rho_{ij}(s_j)^M (1 - \rho_{ij}(s_j))^{B_i - M} f(s_j) ds_j$$

where $f(s_j)$ is the pdf of the distribution of firms and $\rho_{ij}(s_j)$ is the probability that seller s_j serves any buyer in country i . Our empirical strategy uses the heterogeneity across sellers in their ability to reach foreign consumers, that is log-supermodular in firms' productivity and the level of search frictions.

In a model in which the only source of heterogeneity is consumers' preferences with respect to differentiated varieties, such heterogeneity does not exist and the moment used in the structural estimation becomes:

$$h_{ij}(M) = C_{B_i}^M \rho_{ij}^M (1 - \rho_{ij})^{B_i - M}$$

where

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

is homogenous across sellers and captures the likelihood that any seller from i is the preferred variety of any buyer from j .

A.7 A model of buyer acquisition under monopolistic competition

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 z is the realization of a Poisson variable with parameter $T_j z^{-\theta}$. Given exogenous input costs w_j and iceberg costs d_{ij} the number of firms serving market i at a cost below p is itself a Poisson variable of parameter $\mu_{ij}(p) = T_j \left(\frac{d_{ij} w_j}{p}\right)^{-\theta}$.

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

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

$p_{ij}(z_{s_j})$ 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, that all display an iso-

elastic demand function as in our model, and ii) the fixed cost of exporting is increasing in the number of buyers served:

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

where F_{ij} is a positive parameter, $\frac{B_{ij}(z_{s_j})}{B_i}$ measures the share of the market that seller s_j chooses to serve 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:

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

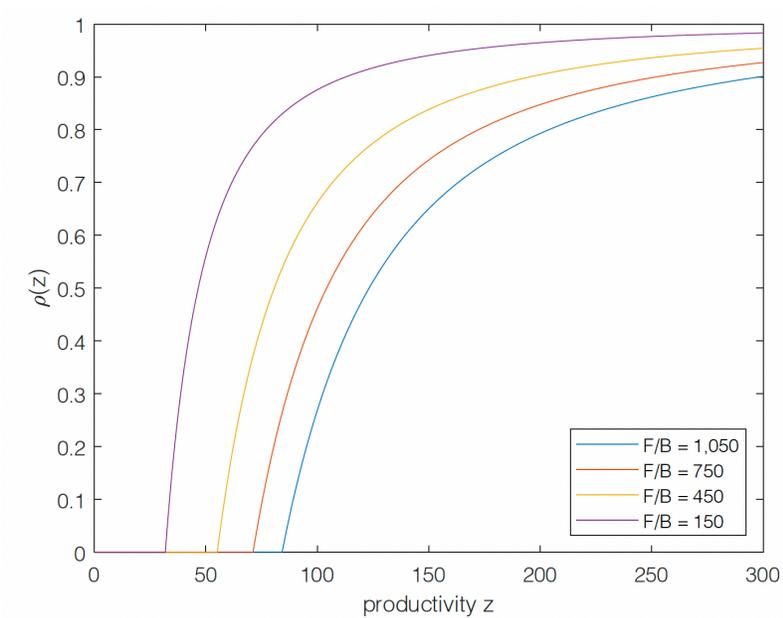
From this, it comes:

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

In this model, the buyer margin is decreasing in both iceberg and fixed export costs, especially at the bottom of the productivity distribution. This feature of the model is illustrated in Figure A.2, which reproduces Figure 2 in the context of the alternative model just discussed. As in the baseline model, the probability of serving a buyer is increasing in the seller's productivity. Reducing trade frictions, whether the fixed or the variable component of trade costs, however increases the export probability at every point of the productivity distribution. This is in contrast with our model which displays a non-linear impact of moving from a high to a low level of frictions.

In this model, the number of sellers choosing to serve exactly $M \leq B_i$ buyers is

Figure A.2: *Probability of serving a buyer as a function of the seller's productivity in the model of buyer acquisition*



Notes: This figure illustrates how the probability of serving a buyer varies with the seller's productivity, for four different values of the fixed cost per buyer of serving market i .

equal to the number of sellers which productivity satisfies the following conditions:

$$\begin{aligned} \underline{z}(M) &\leq z < \underline{z}(M+1) \\ \text{where } \underline{z}(M) &\equiv \left(\frac{B_i - M}{B_i}\right)^{\frac{-1}{\lambda(\sigma-1)}} A_{ij} \\ \text{and } A_{ij} &\equiv \frac{\sigma}{\sigma-1} w_j d_{ij} \left[\sigma \frac{F_{ij}}{B_i \bar{X}_i}\right]^{\frac{1}{\sigma-1}} \end{aligned}$$

The moment used to identify search frictions would thus capture the variance of the following ratios:

$$\begin{aligned} \frac{h_{ij}(M)}{h_{ij}(1)} &= \frac{\underline{z}(M)^{-\theta} - \underline{z}(M+1)^{-\theta}}{\underline{z}(1)^{-\theta} - \underline{z}(2)^{-\theta}} \\ &= \frac{(B_i - M)^{\frac{\theta}{\lambda(\sigma-1)}} - (B_i - M - 1)^{\frac{\theta}{\lambda(\sigma-1)}}}{(B_i - 1)^{\frac{\theta}{\lambda(\sigma-1)}} - (B_i - 2)^{\frac{\theta}{\lambda(\sigma-1)}}} \end{aligned}$$

A.8 Two-sided heterogeneity

The model derived in the previous section can further be extended to handle two-sided heterogeneity, as in [Bernard et al. \(2018b\)](#). Suppose that the supply-side of the previous model is left unchanged but buyers in each destination are now heterogeneous in terms of their average demand:

$$c_{b_i}(p_{b_i}, \bar{X}_{b_i}) = p_{b_i}^{-\sigma} \bar{X}_{b_i}$$

In [Bernard et al. \(2018b\)](#), the heterogeneity comes from buyers combining inputs into a final good sold to final consumers. In their setting, the heterogeneity in demand ultimately comes from a random productivity component that affects the demand addressed to buyers, and in turn their network of suppliers. For the purpose of the appendix, it will be sufficient to assume that buyers are born with a random demand level, drawn from a Pareto distribution with shape Γ . With a discrete number of such buyers, the number of buyers that draw a \bar{X}_{b_i} above X is distributed Poisson of parameter $B_i \left(\frac{X}{\bar{X}_L}\right)^{-\Gamma}$.

Following [Bernard et al. \(2018b\)](#), sellers are assumed to decide whether to serve a buyer or not, given a fixed cost per buyer f_{ij} . Under these assumptions, a seller with productivity z chooses to serve all buyers which demand is sufficiently high to cover

the fixed cost per buyer. At the margin:

$$\begin{aligned} \pi_{ij}(z, \bar{X}(z)) &= 0 \\ \Leftrightarrow \bar{X}(z) &= \underbrace{f_{ij}\sigma \left(\frac{\sigma}{\sigma-1} d_{ij} w_j \right)^{\sigma-1}}_{F_{ij}} z^{1-\sigma} \end{aligned}$$

As discussed in [Bernard et al. \(2018b\)](#), the model thus displays negative assortative matching: High productivity sellers can afford serving relatively small buyers. As a consequence, high-productivity sellers are also those that serve more buyers, consistent with the data. The assortative matching also implies that the relative share of sellers at different points of the distribution of outdegrees reflects the shape of the Pareto distributions that parametrize the heterogeneity of sellers and buyers.

In this setting, the unconditional probability that a particular buyer is served by a seller of productivity z is the probability that this buyer's demand parameter is above the seller's threshold:

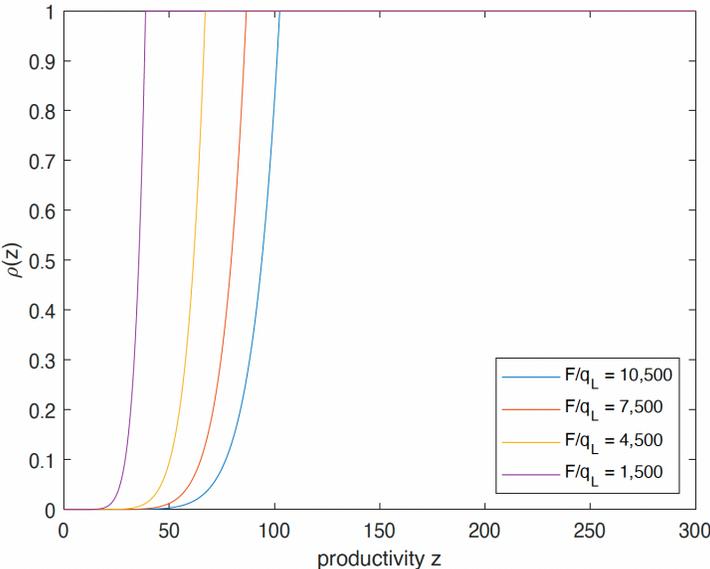
$$\begin{aligned} \rho_{ij}(z) &= Pr \left[\bar{X}_{b_i} \geq \bar{X}(z) \right] \\ &= Pr \left[\bar{X}_{b_i} \geq F_{ij} z^{1-\sigma} \right] \\ &= \left(\frac{F_{ij} z^{1-\sigma}}{X_L} \right)^{-\Gamma} \end{aligned}$$

The probability is depicted in [Figure A.3](#), which reproduces [Figure 2](#) in the context of the model just discussed. Again, the probability of serving a buyer is increasing in the seller's productivity, consistent with the data. Here as well, and contrary to our model, reducing trade frictions, whether the fixed or the variable component of trade costs, increases the export probability whatever the firm's productivity.

Finally, the number of sellers choosing to serve exactly $M \leq B_i$ buyers is equal to the number of sellers which productivity satisfies the following conditions:

$$\begin{aligned} \underline{z}(M) &\leq z < \underline{z}(M+1) \\ \text{where } \underline{z}(M) &\equiv \left(\frac{M}{B_i} \right)^{\frac{1}{\Gamma(\sigma-1)}} A_{ij} \\ \text{and } A_{ij} &\equiv \frac{\sigma}{\sigma-1} w_j d_{ij} \left[\sigma \frac{f_{ij}}{X_L} \right]^{\frac{1}{\sigma-1}} \end{aligned}$$

Figure A.3: *Probability of serving a buyer as a function of the seller's productivity under two-sided heterogeneity*



Notes: This figure illustrates how the probability of serving a buyer varies with the seller's productivity, for four different values of F_{ij}/q_L .

The moment used to identify search frictions would thus capture the variance of the following ratios:

$$\begin{aligned}\frac{h_{ij}(M)}{h_{ij}(1)} &= \frac{\underline{z}(M)^{-\theta} - \underline{z}(M+1)^{-\theta}}{\underline{z}(1)^{-\theta} - \underline{z}(2)^{-\theta}} \\ &= \frac{M^{\frac{-\theta}{\Gamma(\sigma-1)}} - (M+1)^{\frac{-\theta}{\Gamma(\sigma-1)}}}{1 - 2^{\frac{-\theta}{\Gamma(\sigma-1)}}}\end{aligned}$$

A.9 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

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

where $\rho_{ij}(\underline{z})$ 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:

$$h_{ij}(M) = \frac{\pi_{ij}}{\lambda_{ij}} C_{B_i}^M (B(\lambda_{ij}, M, B_i - M + 1) - B(\rho_{ij}(\underline{z}), M, B_i - M + 1)),$$

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

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

Then, the regularized incomplete beta function is

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

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

$$h_{ij}(M) = \frac{\pi_{ij}}{\lambda_{ij}} \frac{1}{M} \left(I_{\lambda_{ij}}(M, B_i - M + 1) - I_{\rho_{ij}(\underline{z})}(M, B_i - M + 1) \right).$$

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

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

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:

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

that is, the proportion of firms from j having exactly M buyers in destination i .³⁶ 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 $\left[\mathbf{1}\{B_{ij}(z_{s_j}) = M\}\right]_{s_j \in S_j}$, 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 $\sigma_{ij}^2(M)$. This is true for all $M \in [0, B_i]$.³⁷ The central limit theorem implies

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

where

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

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

We then consider the function

³⁶Here 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.

³⁷Independence 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 theorem 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.

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

where $Var(\cdot)$ 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 $g(\cdot)$ is asymptotically normal:

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

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

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

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

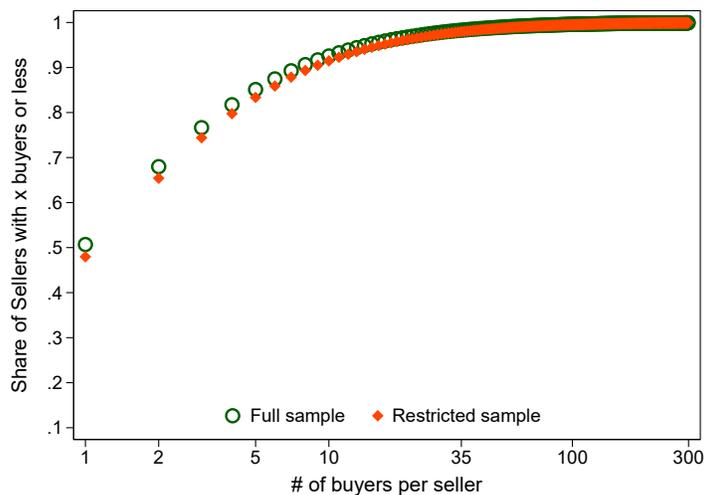
In practice, our estimation is implemented in two steps. First, we use an estimation of the $\Omega(\hat{\theta}_{ij})$ weight matrix using our observations $\nabla g(\hat{\theta}_{ij})$ and $\widehat{\Sigma}_{ij}$. Second, with the $\hat{\lambda}_{ij}$ estimated in the first step, we re-run our estimation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

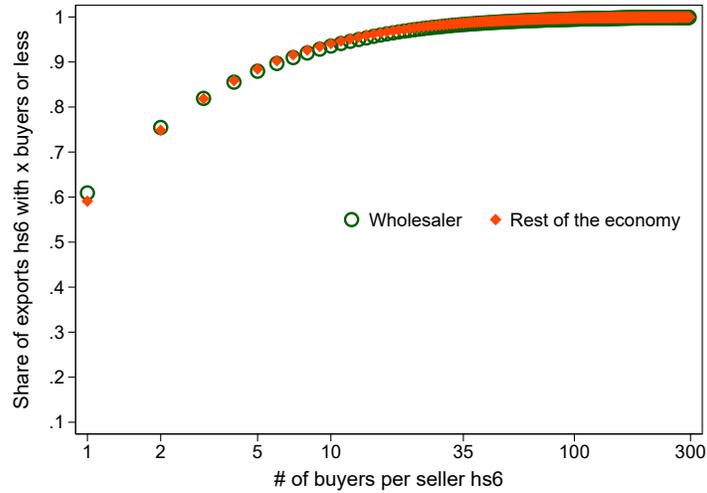
Figure A.4: *Number of buyers per seller, full and restricted sample*



Notes: This figure compares the number of buyers per seller, in the whole dataset and in the estimation dataset, restricted to the 70% of exporters that declare the product category of their exports (“Restricted sample”).

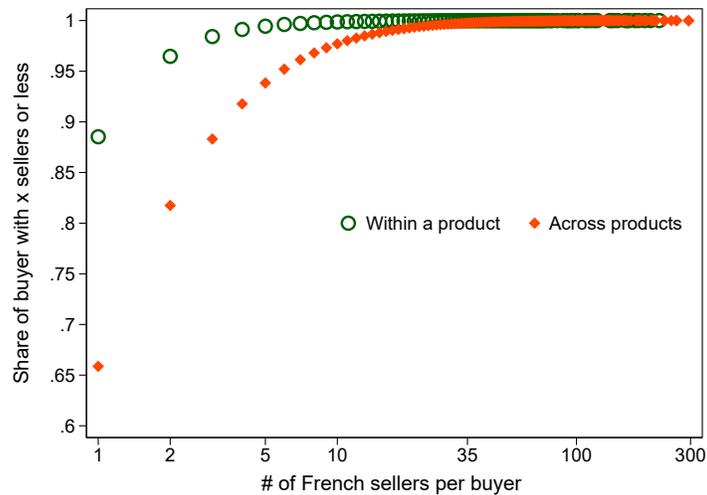
$$\begin{aligned}
 \frac{\partial g(\theta_{ij}(\lambda_{ij}))}{\partial \lambda_{ij}} &= \frac{2}{3}(m_1 - \bar{m}) \frac{\partial \theta_{ij}(\lambda_{ij}, 2)/\theta_{ij}(\lambda_{ij}, 1)}{\partial \lambda_{ij}} \\
 &+ \frac{2}{3}(m_2 - \bar{m}) \sum_{M=3}^6 \frac{\partial \theta_{ij}(\lambda_{ij},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

Figure A.5: *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.6: *Number of sellers per buyer, Within and across products*



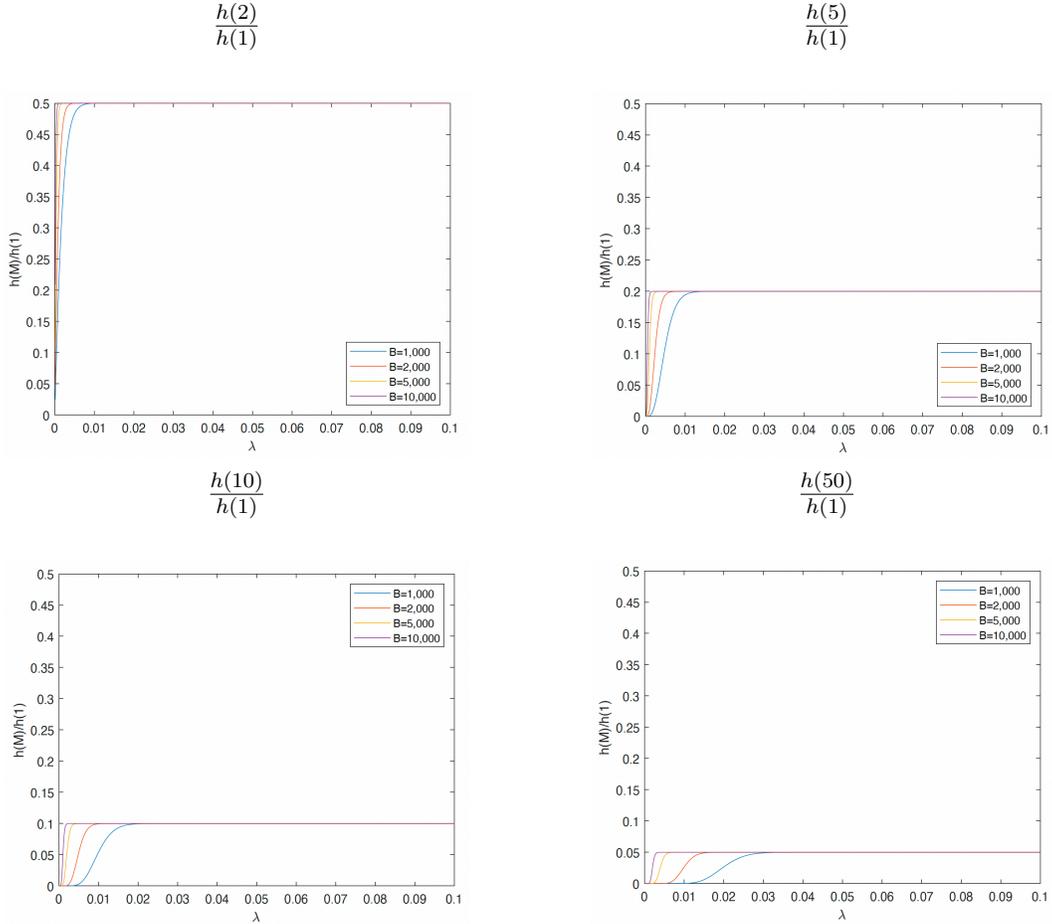
Notes: This figure compares the number of sellers per buyer, computed within a product and across products.

Table A1: *Product- and firm-level gravity equations: Robustness to the proxy for information frictions*

	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	-0.472*** (0.065)	-0.247*** (0.030)	-0.134*** (0.023)	-0.091** (0.044)	-0.080* (0.048)	-0.115*** (0.025)	0.036 (0.041)
log Import Demand	0.864*** (0.015)	0.261*** (0.006)	0.155*** (0.005)	0.447*** (0.010)	0.463*** (0.010)	0.199*** (0.007)	0.264*** (0.009)
log GDP per Capita	-0.197*** (0.039)	-0.0843*** (0.019)	-0.009 (0.013)	-0.104*** (0.025)	-0.153*** (0.030)	-0.106*** (0.018)	-0.047** (0.020)
lof French migrants	0.363*** (0.020)	0.217*** (0.009)	0.096*** (0.007)	0.050*** (0.012)	0.210*** (0.014)	0.114*** (0.008)	0.096*** (0.012)
Observations	66,335	66,335	66,335	66,335	633,136	633,136	633,136
R-squared	0.640	0.785	0.427	0.578	0.687	0.432	0.716
Fixed effects	Product	Product	Product	Product	Firm	Firm	Firm
log Distance	-0.730*** (0.060)	-0.351*** (0.028)	-0.170*** (0.020)	-0.208*** (0.040)	-0.171*** (0.044)	-0.133*** (0.023)	-0.039 (0.039)
log Import Demand	0.880*** (0.015)	0.277*** (0.006)	0.164*** (0.005)	0.439*** (0.009)	0.480*** (0.011)	0.213*** (0.008)	0.267*** (0.009)
log GDP per Capita	0.106*** (0.037)	0.091*** (0.018)	0.068*** (0.011)	-0.052** (0.023)	-0.028 (0.026)	-0.040** (0.016)	0.012 (0.018)
log Social Connectedness	0.342*** (0.025)	0.244*** (0.012)	0.116*** (0.008)	-0.018 (0.015)	0.230*** (0.015)	0.146*** (0.010)	0.084*** (0.012)
Observations	66,335	66,335	66,335	66,335	633,136	633,136	633,136
R-squared	0.635	0.780	0.426	0.578	0.686	0.431	0.716
Fixed effects	Product	Product	Product	Product	Firm	Firm	Firm

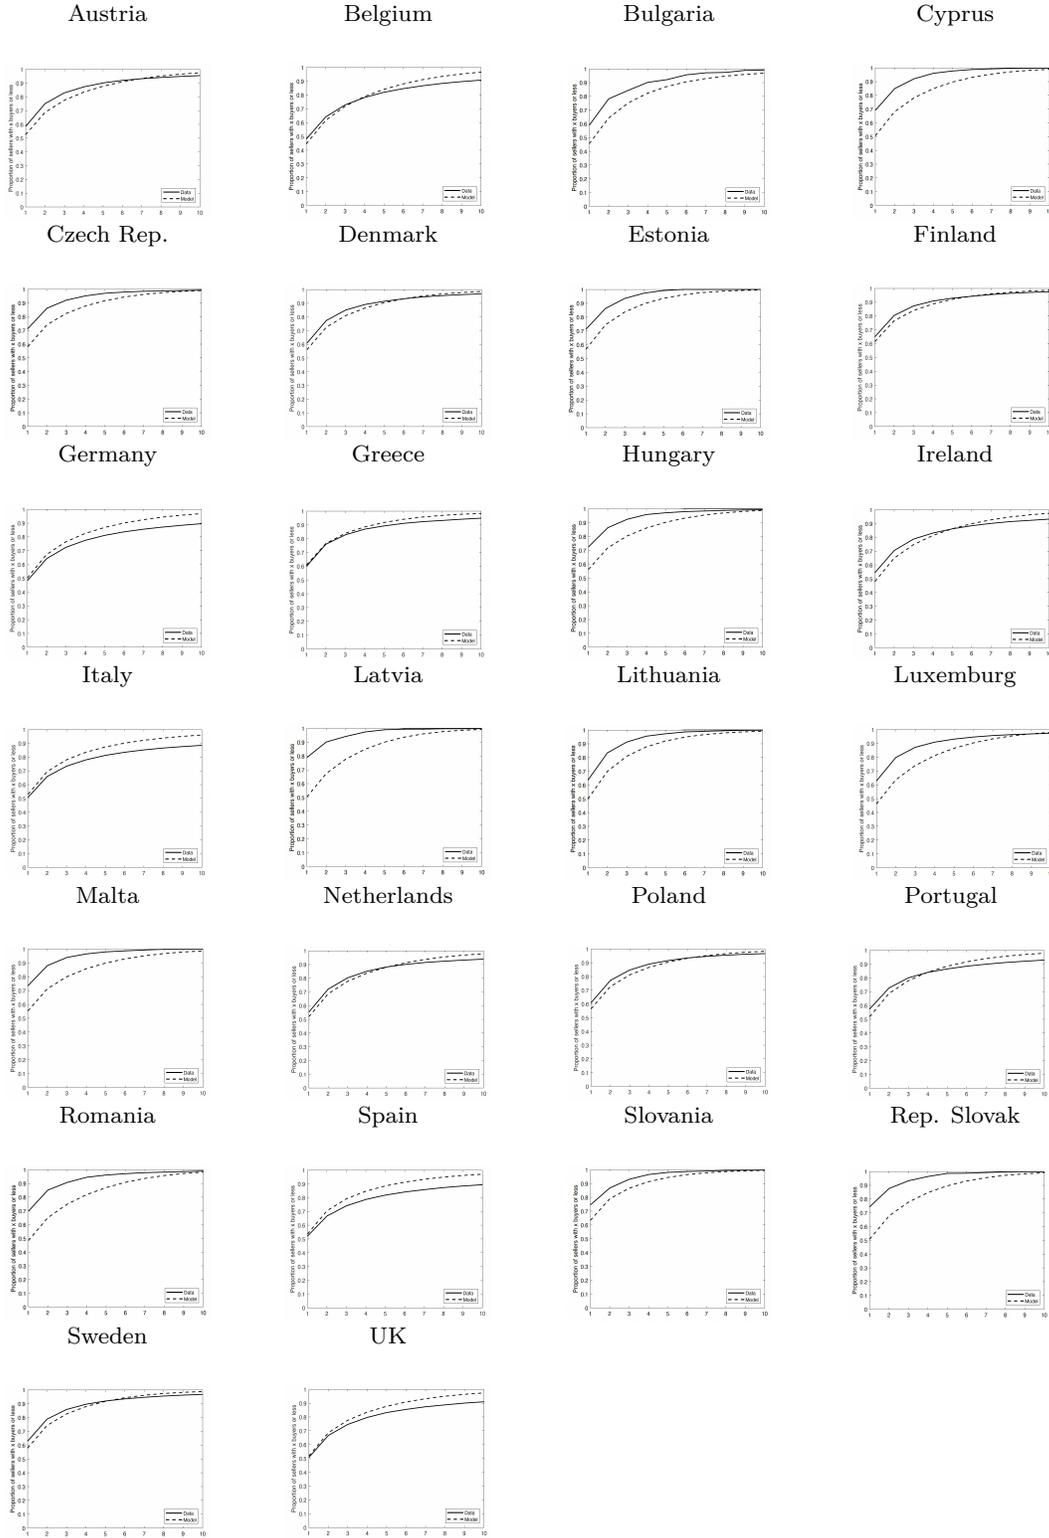
Notes: Standard errors, clustered in the country×HS2 chapter dimension, are in parentheses, with ***, ** and * respectively denoting significance at the 1%, 5% and 10% levels. “log Distance” is the log of the weighted distance between France and the destination. “log Import demand” is the log of the 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. “log French migrants” is the log of the number of French migrants per 1000 of inhabitants in the destination country. log Social connectedness is the log of the social connectedness between France the destination as measured by [Bailey et al. \(2020\)](#) using anonymized Facebook data. 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 the log of firm-level bilateral exports as left-hand-side variable, 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)).

Figure A.7: *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).

Figure A.8: *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 $h_i^k(M)$, at the country and product level, before aggregating across products using information on the relative number of producers of each good in France.

Table A2: *French sellers and EU buyers, 2007*

	<i>Number of</i>			<i>Number of</i>		
	Exporters (1)	Importers (2)	Pairs (3)	Exporter-HS6 (4)	Importer-HS6 (5)	Triplets (6)
Overall	44,280	572,585	1,260,237	184,108	2,388,274	2,879,221
Austria	8,206	14,023	28,119	21,367	52,911	61,556
Belgium	29,486	71,283	214,106	97,265	378,962	482,607
Bulgaria	2,294	2,287	3,657	5,738	6,872	7,617
Cyprus	2,361	1,628	3,736	7,257	8,332	10,036
Czech Republic	6,848	6,117	13,198	16,533	21,449	25,157
Denmark	8,359	8,834	20,853	21,093	37,371	46,565
Estonia	1,803	1,235	2,495	5,231	5,471	6,353
Finland	5,257	5,167	11,594	13,696	21,926	26,046
Germany	24,650	117,937	236,559	73,660	391,170	462,749
Greece	7,793	11,260	25,414	26,036	55,548	68,510
Hungary	5,376	4,439	9,556	12,899	16,278	18,642
Ireland	6,355	6,669	16,266	17,911	38,043	49,179
Italy	20,129	95,915	183,356	63,390	375,947	439,230
Latvia	2,063	1,355	2,948	5,897	6,058	7,430
Lithuania	2,914	1,854	4,699	7,224	7,294	9,879
Luxembourg	10,730	7,646	28,555	31,313	54,766	70,079
Malta	1,783	931	2,554	4,696	4,698	5,773
Netherlands	16,444	33,642	69,845	43,475	131,264	157,822
Poland	9,734	12,858	30,230	24,673	43,477	52,630
Portugal	11,648	19,676	42,923	35,034	95,352	113,494
Romania	5,038	4,856	9,502	12,487	16,411	18,397
Slovakia	3,271	2,306	5,002	7,335	8,072	9,396
Slovenia	2,840	2,226	4,386	7,510	8,628	9,755
Spain	21,638	77,596	159,645	70,379	359,774	419,999
Sweden	7,683	10,203	20,409	20,209	39,375	45,560
UK	18,898	50,642	110,630	55,228	202,825	254,760

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 A3: *Number of buyers per seller across destination countries*

	Mean	Median	p75	Sh. with 1 buyer
	(1)	(2)	(3)	(4)
Austria	2.9	1	2	63%
Belgium	5.0	2	4	50%
Bulgaria	1.3	1	1	81%
Cyprus	1.4	1	1	80%
Czech Republic	1.5	1	1	76%
Denmark	2.2	1	2	66%
Estonia	1.2	1	1	85%
Finland	1.9	1	2	70%
Germany	6.3	2	4	50%
Greece	2.6	1	2	64%
Hungary	1.4	1	1	78%
Ireland	2.7	1	2	63%
Italy	6.9	1	3	53%
Latvia	1.3	1	1	84%
Lithuania	1.4	1	1	79%
Luxembourg	2.2	1	2	66%
Malta	1.2	1	1	86%
Netherlands	3.6	1	2	59%
Poland	2.1	1	2	67%
Portugal	3.2	1	2	62%
Romania	1.5	1	1	77%
Slovenia	1.3	1	1	82%
Slovakia	1.3	1	1	84%
Spain	6.0	1	3	54%
Sweden	2.3	1	2	67%
United Kingdom	4.6	1	3	53%
Across countries	15.6	3	10	32%

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.