

Frictions and adjustments in firm-to-firm trade

François Fontaine* Julien Martin† Isabelle Mejean‡

January 4, 2023

Abstract

We study bilateral trade adjustments in a dynamic Ricardian model of trade with search frictions. The model generates an endogenous network of firm-to-firm trade relationships that displays price bargaining within and across firm-to-firm relationships. Following a foreign shock, firms sourcing inputs from abroad have three options: absorb the shock, renegotiate with their current supplier or switch to a supplier in another country. The relative importance of these adjustment margins depends on the interplay between Ricardian comparative advantages, search frictions and firms' individual characteristics, including the history of the relationship. We exploit French firm-to-firm trade data to estimate the model structurally and quantify the relative importance of these adjustment margins in 26 sectors and 14 EU countries.

1 Introduction

How do search frictions shape the adjustment of trade to relative price shocks? We answer this question in a dynamic Ricardian model of firm-to-firm trade in which markups are endogenous and the matching of firms is governed by random search. In our model, shocks to relative prices induce a reshuffling of firm-to-firm trade relationships. Firms sourcing inputs from a country whose production costs increase have three options: fully absorb the shock, bargain with their current supplier, or switch to a supplier in another country. The relative importance of these adjustment margins depends on the interplay between Ricardian comparative advantages, search frictions, and firms' individual characteristics. We exploit panel data on the universe of transactions involving French exporters and their European partners to estimate search frictions across markets. We then quantify the role of frictions in shaping the adjustment of trade flows and trade prices.

We study an environment in which firms face frictions when taking decisions on the sourcing of their inputs. Our model borrows ingredients from the labor literature, most notably models

*Paris School of Economics and Université Paris Panthéon Sorbonne, Email address: francois.fontaine@univ-paris1.fr.

†Université du Québec à Montréal and CEPR

‡Sciences Po and CEPR, Email address: isabelle.mejean@sciencespo.fr.

of on-the-match search as in [Postel-Vinay and Robin \(2002\)](#), to enrich a dynamic Ricardian model of trade with endogenous markups. To produce, firms purchase an endogenous set of inputs at the lowest possible quality-adjusted price. Following [Eaton and Kortum \(2002\)](#), intermediate producers are heterogeneous in terms of their cost of serving a given destination, the interaction of technology and geography shaping comparative advantages. We add random search to the Ricardian structure ([Lenoir et al., 2022](#), [Eaton et al., 2022b](#)).¹ Because input markets display bilateral search frictions, each producer chooses from a discrete number of quality-adjusted prices, which expands randomly when new input suppliers are met. Random search introduces a rich structure of competition in which input suppliers compete to retain the buyers they have met by adjusting their price dynamically.

In equilibrium, more competitive suppliers serve a broader set of buyers, on average. This property of the model helps to reproduce the heterogeneity across firms in the buyer’s margin that has already been documented in the literature on firm-to-firm trade ([Bernard et al., 2018](#), [Carballo et al., 2013](#), [Lenoir et al., 2022](#), [Eaton et al., 2022b](#)). In addition to the cross-sectional pattern, the model reproduces three novel empirical facts regarding the dynamics of firm-to-firm relationships.² First, the relationship between a firm’s accumulation of foreign partners and its experience in the destination country is increasing and concave. Upon entry into a foreign market, exporters continue to randomly match with foreign buyers, which drives the growth of their portfolio. Because downstream firms endogenously decide who to source their inputs from, each supplier converges toward a finite portfolio whose size is a function of its productivity. Second, the hazard rate of a relationship is shown to be decreasing over the length of the relationship. After six months, the probability of the relationship ending is already a quarter of what it was when the relationship started. Our model reproduces the stylized fact through endogenous selection. In the early stages of a relationship, the probability that the buyer meets a lower quality-adjusted cost supplier is high. Over time, only the best relationships survive and the probability of a separation decreases. Third, within a firm-to-firm relationship, we find evidence of prices being renegotiated downwards: After a year, prices are on average 1.2% lower than in the initial transaction. In the model, suppliers compress their markups over time as downstream firms exploit new matches to bargain with their current supplier or switch to a more competitive one.

Both the steady-state structure of trade networks and their dynamic adjustment to relative price shocks depend on the interaction of search frictions and Ricardian comparative advantages. We propose a structural approach to estimating these parameters in firm-to-firm trade data. The estimation relies on unconditional inference and maximum likelihood, as in [Ridder and van den Berg \(2003\)](#). Identification is achieved using transaction and switch frequencies that we observe in the data for all European importers interacting with French exporters at a point in time. Intuitively, the rate at which importers switch from one French exporter to another,

¹In this context, search frictions can be interpreted as a parsimonious way of capturing the various contractual and informational frictions affecting the establishment of production linkages that as previously been discussed in the literature (see, e.g., [Antras and Helpman, 2004](#), [Grossman and Helpman, 2021](#)).

²These dynamics patterns are recovered from an exhaustive dataset of firm-to-firm transactions involving French exporters and their European partners, between 2002 and 2006. The dataset is described in details in [Bergounhon et al. \(2018\)](#). A cross-section of these data has been exploited recently in [Lenoir et al. \(2022\)](#) and [Eaton et al. \(2022b\)](#).

conditional on a transaction, is informative about the magnitude of the frictions. However, the mapping between the empirical moments and the structural parameters is complex due to unobserved quality-adjusted cost differences between French exporters and unobserved switches toward non-French exporters. Working with unconditional hazard rates in the model and the data solves the first issue. We further exploit the model’s structure to take into account the endogenous censoring in the data and match it with observed trade shares. In the end, the simulated maximum likelihood estimator makes it possible to recover the relative size of search frictions faced by French exporters in each of their (European) export markets, together with estimates for the overall meeting rate and Ricardian advantage of French firms there.

We estimate the model’s structural parameters for 14 EU countries and 26 different sectors. Results reveal a substantial degree of heterogeneity in the magnitude of Ricardian comparative advantages and the level of search frictions, across countries and sectors. Over the 330 country-sector pairs that constitute our sample of estimated parameters, we find that search frictions explain as much as 40% of the observed variance in French firms’ market shares. However, the importance of Ricardian comparative advantages increases when we focus on the heterogeneity across sectors within a destination, or across destinations within a sector. Point estimates suggest that relative search frictions faced by French firms in European markets are lower in neighboring countries and in sectors that constitute the core of France’s international competitiveness, such as chemical products, motor vehicles, electrical products or beverages.

Armed with these estimates, we can then study dynamic adjustments to relative price shocks using a counterfactual. Namely, we simulate a uniform, permanent 10% increase in the cost of French products and study its impact on export prices and product-level trade.³ We first study how the magnitude of bilateral trade adjustments varies along the distribution of estimated search frictions. In the aggregate, the 10% shock to French relative prices is associated with a -32% drop in the value of bilateral trade, an elasticity of 3.2.⁴ Across products, the elasticity of trade systematically decreases with relative meeting rates, i.e. trade adjustments are muted in markets in which French firms face relatively low matching frictions. Relative meeting rates indeed shape the geographic structure of individual firms’ network of suppliers, which is a key determinant of firms’ ability to move quickly away from French firms when the shock hits. We confirm this prediction of the model using estimated product- and country-specific trade elasticities.

Beyond its tractable aggregate properties, our model also displays rich heterogeneity across firms in terms of the magnitude and mechanism of the adjustment. On impact, 20% of trade

³The nature of the simulated shock implies that the pass-through rates and trade elasticities discussed therein should be compared with empirical moments estimated from permanent shocks such as changes in tariffs. Adjustments to changes in temporary shocks, such as exchange rates, are notoriously smaller (see, e.g., [Fontagné et al., 2018](#), for a discussion of the international elasticity puzzle). In our simulations, we can further control that the shock is a pure shock to relative prices. Empirically, these shocks are difficult to identify in the French context as French exporters are exposed to the same tariffs and exchange rates than the majority of their (European) competitors.

⁴This value is in the ballpark of estimates of export elasticities to permanent relative price shocks recovered from firm-level data (see, e.g. [Fontagné et al. \(2018\)](#), [Fitzgerald and Haller \(2018\)](#)). The empirical literature also points to some heterogeneity across countries and sectors, see [Caliendo and Parro \(2015\)](#), [Imbs and Mejean \(2015\)](#), and [Boehm et al. \(2022\)](#), among others. In our model, such heterogeneity is driven by the joint effect of heterogeneous comparative advantages and relative meeting rates.

relationships involving a French supplier are disrupted as a consequence of importers switching to non-French alternative suppliers. Within the remaining set of relationships, the pass-through of the shock on foreign buyers is limited, as most French exporters are forced to absorb the shock into their margin to remain competitive. As a consequence, the median markup rate among French suppliers decreases from 1.28 to 1.19. On average, these adjustments lead to a 20% pass-through of the shock onto foreign buyers.⁵ Over time, individual buyers gain market power through further matches. In the aggregate, however, the shock is inflationary as it permanently reduces the strength of competition in input markets. The model also delivers rich heterogeneity across sector×countries and across firms, within a market. Simulated pass-through rates are found to be lower in product markets in which French market shares are low, as a consequence of high bilateral search frictions or a weak comparative advantage. In markets with a high level of (multilateral) search frictions, French exporters instead face less competition and the incidence falls more heavily on the buyer side. Our model also rationalizes heterogeneous pass-through rates across sellers and buyers. Everything else equal, more experienced buyers have already accumulated a larger network, on average, which increases their bargaining power and helps them contain the price increase. On the other side of the market, high-productivity sellers are more likely to survive the shock and maintain relatively high markup, consistent with evidence in [Berman et al. \(2012\)](#).

Related literature: Our paper contributes to the fast-growing literature on firm-to-firm trade in international markets, and the more established literature on firm-level trade and price adjustments following relative price shocks. In contrast with these literatures, the main innovation that we borrow from the labor literature is the introduction of price renegotiations. This ingredient allows us to study price dynamics *across* but also *within* firm-to-firm relationships.

Like several other recent contributions, we examine firm-to-firm trade in a model displaying search frictions ([Miyachi, 2019](#), [Chor and Ma, 2020](#), [Demir et al., 2021](#), [Eaton et al., 2022b](#), [Lenoir et al., 2022](#)).⁶ [Eaton et al. \(2022a\)](#) model supplier-retailer relationships in presence of endogenous search efforts and estimate the model in the context of the US apparel industry. [Eaton et al. \(2021\)](#) model firm-to-firm dynamics in presence of search frictions when importers learn about product appeal, and estimate their model using Colombian export data. Our model admits a simpler search structure in which heterogeneity is limited to the cross-country dimension. The benefit is that we can estimate bilateral search parameters for 26 sectors and 14 destinations. Compared to these papers, our model proposes a richer view of firm pricing strategies, which we can confront to the data as in [Fontaine et al. \(2020\)](#).

In our model, firms retain their current partners by adjusting their markup down. In this respect, our work is related to [Kikkawa et al. \(2019\)](#) and [Alviarez et al. \(2021\)](#) who also study pricing in a (static) model of firm-to-firm relationships, both theoretically and empirically.

⁵As a matter of comparison, [Fontagné et al. \(2018\)](#) estimate that French exporters absorb about a third of shocks to EU-wide tariffs, thus passing a larger share of the shock onto foreign consumers. Their shock is, however, less biased against French firms than ours as tariff adjustments are equally felt by French firms and all of their European competitors.

⁶Firm-to-firm trade has also been analyzed in monopolistic competition settings without search or matching frictions ([Bernard et al., 2018](#), [Carballo et al., 2018](#), [Lim, 2018](#)). See [Bernard and Moxnes \(2018\)](#) for a review.

Non-constant markups in these papers are driven by oligopolistic competition. We propose a complementary perspective. Introducing random search and firm dynamics in a Ricardian setting generates rich adjustment patterns of markups, in the steady state and in the aftermath of a shock.⁷ More broadly we contribute to the theoretical literature on prices and variable markups in international trade (see, among many others [Bernard et al., 2003](#), [Atkeson and Burstein, 2008](#), [de Blas and Russ, 2015](#)). Like [Atkeson and Burstein \(2008\)](#) and [de Blas and Russ \(2015\)](#) we examine markups in a Ricardian context with Bertrand competition. The presence of search frictions allows us to discuss the determinants of markups *within* firm-to-firm relationships, and generate new insights on the dynamics of markups, in the steady state and following a shock. The variety of price adjustments predicted by the model in the aftermath of a shock is consistent with the vast empirical literature on heterogeneous cost pass-through ([Berman et al., 2012](#), [Fitzgerald and Haller, 2014](#), [Amiti et al., 2014](#), [Garetto, 2016](#), [De Loecker et al., 2016](#), [Auer et al., 2021](#))

We also contribute to literature on firm dynamics in international markets reviewed in [Alessandria et al. \(2021a\)](#). In this literature, individual dynamics are often linked to firm heterogeneity and sunk costs of entry in international markets (see, e.g., [Roberts and Tybout, 1997](#), [Das et al., 2007](#), [Alessandria and Choi, 2007](#), [Impullitti et al., 2013](#), [Alessandria and Choi, 2014](#)). A strand of the literature also emphasizes customer accumulation as a source of firm dynamics ([Arkolakis, 2010](#), [Drozd and Nosal, 2012](#), [Gourio and Rudanko, 2014](#), [Fitzgerald et al., 2017](#), [Piveteau, 2020](#)).⁸ Here as well, increasing fixed costs associated with serving a larger customer base are used to explain the heterogeneity observed in the data. Our model does not display any sunk cost. Like [Eaton et al. \(2022a\)](#) and [Eaton et al. \(2021\)](#), we instead emphasize the role of search frictions for firm-to-firm trade dynamics.

Finally, our model borrows from the labor literature, most notably models of on-the-job search and wage renegotiation as in [Postel-Vinay and Robin \(2002\)](#), [Cahuc et al. \(2006\)](#), [Bagger et al. \(2014\)](#). Our estimation strategy is very similar to [Ridder and van den Berg \(2003\)](#) and, more generally, reminiscent of the estimation of the job-to-job transition parameter in job search models. Like [Ridder and van den Berg \(2003\)](#), the identification of frictions exploits the frequency of switches and buyer acquisition. The parametric identification strategy allows us to estimate frictions that are market (i.e. sector \times country) specific. Alternatively, [Miyachi \(2019\)](#) proposes to estimate frictions non-parametrically using exogenous separations attributable to firms' death. Although interesting, such a strategy exploits relatively rare events which makes it difficult to recover market-specific measures of frictions.

The rest of the paper is organized as follows. Section 2 presents the data and documents new facts on firm-to-firm trade dynamics. Section 3 sets up the model, while Section 4 discusses the estimation. In Section 5.2 we study the model's performances in reproducing the observed dynamics of trade following relative price shocks. Finally, Section 6 concludes.

⁷In this regard, our paper relates to [Alessandria \(2004\)](#) in which deviations from the law of one price in international markets emerge from the presence of search frictions.

⁸Several papers also relate individual export decisions to learning ([Nguyen, 2012](#), [Berman et al., 2019](#), [Arkolakis et al., 2018](#)).

2 Data and stylized facts

2.1 Data

Throughout the paper, we use data provided by the French customs, covering every export transaction involving a French firm and one of its partners in the European Union. Importantly, these data identify the French exporter by its siren number *and* the European importer, identified by an anonymized VAT number that includes a code for the buyer’s country of origin i . French exporters’ siren numbers are used to identify the sellers s in the model developed in next section and the European importers will be the empirical counterpart of the buyers b . The transaction is also characterized by a product category p at the 8-digit level of the combined nomenclature and the date t of the transaction identified by a particular month within a year. Whereas the dataset is exhaustive, the customs form is simplified below a threshold defined over annual exports in the European Union. Unfortunately, the product category, that is heavily used in the estimation is missing for firms below the threshold. We thus chose the estimation period so as to minimize the associated attrition.⁹ Finally, we observe the value of the transaction and the quantity exported, which we use to compute the unit value of each transaction, our proxy for prices.¹⁰

The analysis covers the 2002-2006 period, which does not incur any substantial change in the combined nomenclature, nor in the declaration rules for exports. The sample is further restricted to the 14 historical members of the European Union. Product codes affected by yearly changes in the combined nomenclature are harmonized over time using the algorithm proposed by [Behrens et al. \(2019\)](#). As the raw data goes back to 1995, we can use the pre-sample period to control for censoring. In this case, the matching of firm-to-firm relationships in- and out-of-sample is based on hs4 products whose definition is invariant over time.

In the rest of the analysis, the focus is on European importers, and their interactions with French sellers. We follow the history of transactions involving a particular buyer b for a specific product p and various French sellers s , over time. The model explains the decision to purchase goods from a particular seller in the context of frictional good markets whereby importers are willing to purchase a particular input p , to French or other producers. In this model, input purchases cannot be intermediated through wholesalers. We thus exclude from the analysis all transactions that involve a wholesaler, whether on the export or the import side. [Appendix A](#) explains how we identify wholesalers and how they contribute to aggregate trade.

[Table A1](#) in the Appendix shows statistics on the dimensionality of the estimation sample. After dropping intermediaries, the data covers a total of 27 million transactions that involve almost 40,000 French exporters and 744,000 European importers. In the rest of the analysis, we

⁹Before 2011, the declaration threshold was set at 150,000 euros. Since 2011, it has more than doubled, at 460,000 euros.

¹⁰For the vast majority of products, the quantity is declared in kilograms. For some 8-digit products, the customs authorities ask firms to declare the quantity of exports in some identified physical units (e.g. liters for wine, number of units for living animals, number of carats for diamonds, etc), sometimes complemented with the weight of the merchandise. We use the physical quantity whenever available and the weight of the merchandises elsewhere. Using different units across products is innocuous, as our analysis always controls for product fixed effects.

define a “relationship” as the set of transactions involving a particular pair of firms interacting over a specific product. There are 5.6 million such relationships in the estimation sample that we thus see interacting over five transactions, on average. Of course, the intensity of these relationships is strongly heterogeneous, with a number of relationships being short-lived whereas others induce a large number of subsequent transactions. Our analysis mostly exploits this heterogeneity to discuss the dynamics of firm-to-firm trade relationships.

2.2 Five stylized facts on firm-to-firm relationships

Previous papers using similar data on firm-to-firm relationships have extensively discussed the cross-sectional properties of the network of sellers and buyers in international markets (Carballo et al., 2013, Bernard et al., 2018, Lenoir et al., 2022, Eaton et al., 2022b). The literature shows a strong degree of heterogeneity in firms’ in- and out-degrees, as measured by the number of partners an exporter sells to as well as the number of exporters an importer is connected to. We reproduce some of these results in Appendix A.2 and show that the degree of connectedness varies depending on whether we focus on a single cross-section or if we cumulate relationships over time, and on whether we condition or not on the product being traded. In this section, we depart from the literature and focus more specifically on the panel dimension of the data. We recover five novel stylized facts on the dynamics of firm-to-firm networks, and its consequences for the dynamics of prices and quantities.

2.2.1 Exporters’ accumulation of buyers

We start from the exporters’ side of the network and study how firms accumulate buyers over time. As already discussed in the literature and confirmed in Appendix A.2, French exporters display strongly heterogeneous connectedness to foreign markets, as measured by the number of foreign partners they serve. On average, large exporters serve more buyers in more countries.¹¹ A still open question is the extent to which this cross-sectional heterogeneity relates to firms’ dynamic ability to expand through the buyer margin. This question echoes the literature on export dynamics, which studies the evolution of a firm’s exports, posterior to entry into a destination (Fitzgerald et al., 2016, for instance). In comparison with this literature, we are able to further delve into the structure of a firm’s export portfolio, over time.

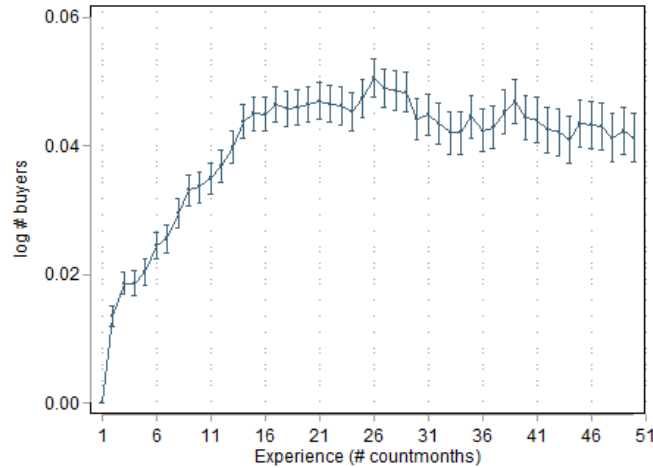
To this end, we estimate the following equation:

$$\ln n_{spit} = FE_{spi} + FE_t + \sum_{k=2}^K \alpha_k \mathbb{1}(Experience_{spit} = k) + \varepsilon_{bpst} \quad (1)$$

where n_{spit} is the number of buyers from country i served by seller s with product p at time t . FE_{spi} and FE_t are individual and time fixed effects, respectively. $\mathbb{1}(Experience_{spit} = k)$ is a

¹¹The heterogeneity in exporters’ ability to serve a large number of foreign partners is explained in the model by the interaction of exporters’ productivity heterogeneity and the history of their matches with foreign firms. From that point of view, we follow the recent literature on matching in international goods markets (Lenoir et al., 2022, Eaton et al., 2022b). The same data pattern can also be rationalized using models of two-way heterogeneity as discussed in Carballo et al. (2013) and Bernard et al. (2018).

Figure 1: Acquisition of buyers, over time



Note: The figure shows the evolution of a seller’s stock of buyers, over time, recovered from equation (1). The figure reports the estimates and their 95% confidence intervals. Experience is measured by the number of periods since first entry.

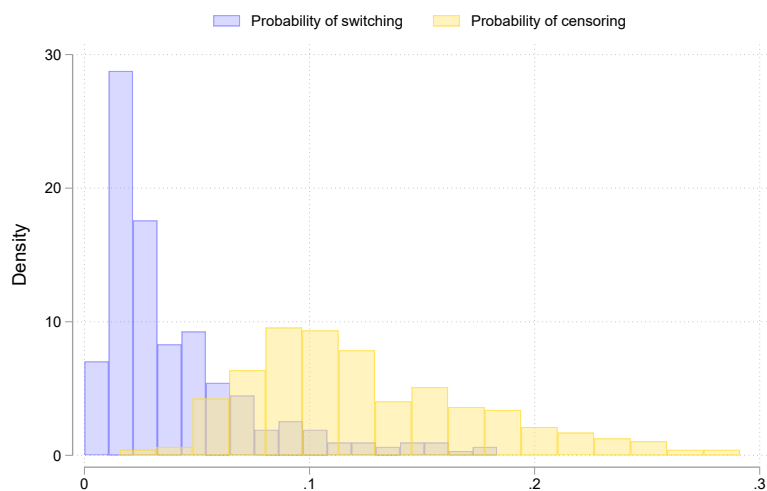
dummy variable equal to one if the seller’s experience in country i is equal to k . We tested with two alternative definitions of a seller’s experience, either the number of periods or the number of export transactions since first entry into the destination.¹² This specification makes it possible to visualize the mean growth of sellers’ stock of buyers, over time, controlling for unobserved heterogeneity across sellers and periods. The α_k coefficient can be interpreted as the (log of the) number of clients a seller with experience k has on average at time t in destination i , normalized by the number of its clients at entry.

Results are shown in Figure 1 and reveal a clear positive correlation between a firm’s experience in a destination and the number of clients it serves there. The relationship is concave meaning that the accumulation of buyers is especially strong in the early stages of the firm’s export experience. After 6 months in the market, the number of clients served has increased by about 2.3%. After 2 years, the number of clients served is on average 4.5% larger than at entry. Whereas these numbers may seem small in light of the average number of partners served by a firm, it should be noted that the regression is affected by composition effects. Not all firms remain active over 2 consecutive years and those that do tend to be the largest ones. Since their number of clients at entry is already larger than the average, the 4.5% growth is significant for these firms. This leads us to the first of our stylized facts.

Fact 1 *Posterior to entry into a destination, a seller’s portfolio of clients tends to grow over time, at a decreasing rate.*

¹²We control for censoring by using information prior to the estimation period to recover the full history of a firm’s experience in a destination. As a consequence, exporters do not necessarily enter the estimation sample used to recover the coefficients of equation (1) with an experience of one. Figure 1 defines experience in terms of the number of periods since first entry into the destination and Figure D1 reproduces the same picture using the cumulated number of transactions.

Figure 2: Mobility of importers, over time



Note: The figure shows the distribution across country×sectors of the censoring and switch probabilities. The probabilities use as a reference the population of importers purchasing products from France in January 2002. The probability of a censoring is defined as the share of firms that are never seen interacting with a French firm again during the next 12 months. The probability of a switch is defined as the proportion of firms that are observed purchasing the product from a different firm in their next transaction with French firms. The population of firms that neither exit the sample nor switch is the share of firms that interact with the same firm at least once during the next 12 months.

2.2.2 Importers’ connectedness to France

We now turn to the other side of the graph, and focus on the population of foreign importers that interact with French exporters. As discussed in Appendix A.2, this side of the network displays much less cross-sectional heterogeneity as the vast majority of importers (more than 95%) interact with a single French firm over a particular month and product. The distribution of importers’ indegrees is, however, shifted down when cumulating all the products purchased by the same firm. The downward shift implies that importers combine purchases over multiple inputs sourced from different suppliers in France or elsewhere. More specifically, around 25% of importers source multiple inputs from France, at a point in time. The degree of connectedness is also shown to be stronger and more heterogeneous when we cumulate all the individual partners that interact with the same firm over a particular product, over time. This fact is particularly interesting as it reveals a form of mobility for importers that switch across suppliers, over time.

This mobility is further illustrated in Figure 2, which shows the distribution of switching and censoring probabilities, among the population of firms observed interacting with a French firm at a point in time. To construct these probabilities, we first select a cross-section of importers that we observe purchasing a product from a French firm in January 2002. We then follow these importers over the next year and observe whether i) they are never seen interacting with a French firm again (a “censoring”), ii) their next transaction with a French firm is with a different supplier (a “switch”) or iii) whether the next transaction takes place with the same partner, i.e. they continue the relationship. Figure 2 shows the probability of a censoring and

the probability of a switch, computed for each sector and country in our data. First of all, we observe that these probabilities are rather small, which implies that the probability of the transaction being followed by a second transaction with the same firm is high - see Section 2.2.3 for more details. Second, we observe that between 0 and 30% of importers (and 13% on average) are never seen interacting again with a French firm, which our model will explain by the combination of exogenous separations and the importer's decision to switch to a non-French supplier. Finally, we also see a significant mass of switches. In the overall sample, the probability of a switch is equal to 5% but reaches 18% in the population of Belgian importers of forestry products.

This leads us to the second of our motivating stylized facts:

Fact 2 *Whereas importers interact with a single supplier at a point in time, they maintain repeated relationships with their suppliers and switch from one French supplier to another, at a 5% rate on average.*

It should be noted that it is difficult to interpret censoring probabilities as they result from two possible events: The importer may stop purchasing the product or she may switch to a (non-French) supplier not covered by our data. In the empirical analysis, we will distinguish between these possibilities using a simulated method of moments.

2.2.3 Dynamics of firm-to-firm relationships

We next delve deeper into the dynamics of firm-to-firm relationships by studying the likelihood that a particular relationship ends. Our data indeed displays significant heterogeneity in the duration of relationships, with 60% of relationships not surviving the first transaction, whereas a significant number are long-lasting. We now use the dynamics within a firm-to-firm relationship to measure how the probability of a relationship ending evolves over its course.

More specifically, Figure 3 illustrates the evolution of hazard rates over the history of a relationship. The hazard rate is defined as the probability that a relationship ends after x months, conditional on the relationship having survived up to that point.¹³ The declining pattern is consistent with the survival rate increasing over the course of a relationship. The dynamic is especially strong during the first year of the relationship, whereas the probability of the relationship ending stabilizes after 1.5 to 2 years, at around 4%. This leads us to the third of our stylized facts.

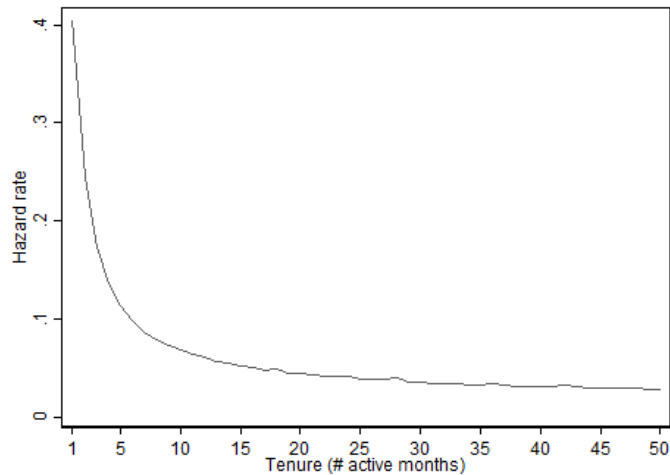
Fact 3 *The probability of the relationship ending declines over the length of a firm-to-firm relationship, before stabilizing after 18 months, on average.*

2.2.4 Price dynamics within and across relationships

The last set of stylized facts concerns the dynamics of unit values, over time. Since we observe the unit value at the level of each firm-to-firm transaction, it is possible to measure the extent

¹³Here as well, left-censoring is controlled for using data prior to 2002 to recover the full length of a relationship. In figure D2, we reproduce the graph using the number of cumulated transactions in the relationship to measure tenure.

Figure 3: Hazard rate, over time



Notes: The hazard rate is defined as the probability of the relationship ending, conditional on tenure in the relationship and is calculated as the ratio of the density to the survival rate at tenure k . The figure is recovered from the 2002-2006 sample using the cumulated number of periods in the relationship as measure of tenure.

to which export prices change over the course of a relationship, and in case of a switch. To this end, we first estimate the following equation:

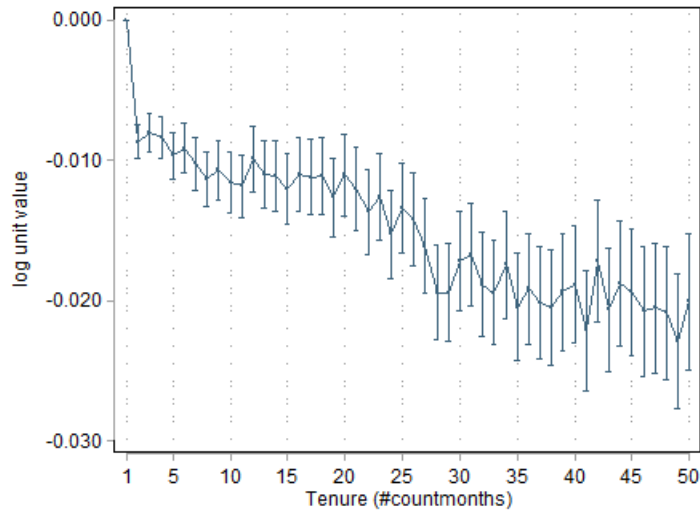
$$\ln p_{bpst} = FE_{bps} + FE_{ipt} + \sum_{k=2}^K \alpha_k \mathbb{1}(Tenure_{bpst} = k) + \varepsilon_{bpst} \quad (2)$$

where p_{bpst} is the unit value set on the transaction involving exporter s , product p , importer b that occurs at time t . The presence of relationship-specific fixed effects (FE_{bps}) implies that the identification of other coefficients is within a firm-to-firm relationship. The baseline regression also controls for country \times product \times period fixed effects (FE_{ipt}) to account for destination-specific inflation trends. The coefficients of interest are the α_k coefficients that measure the average price change after a tenure of k .

Results are reported in Figure 4. They show a negative trend in prices, at least over the first 2 years. The relationship becomes fuzzier over long tenures due to the small number of long-lasting relationships. But the price decline seems to persist. Note that the rate at which prices decline is moderate. After a year, prices are on average 1.2% lower than in the initial transaction. After 3 years, they are 2% lower. Finally, we show in Figure D4 that the same pattern is observed if we estimate the relationships separately on short and long tenures. On the other hand, the dynamics in the value of exports does not show any clear pattern over time. Whereas the value of the transaction seems to increase consistently between the first and second transactions, the dynamic after the second transaction is either relatively stable or decreasing.¹⁴ This leads us to our fourth stylized fact.

¹⁴The jump observed between the first and the second transaction may be indicative of a form of partial month effect or a learning phenomenon whereby importers first test the seller over relatively small quantities, before establishing a more stable relationship over larger quantities. Given that the dynamic does not persist

Figure 4: Price dynamics, within a firm-to-firm relationship



Note: This figure shows the evolution of prices within a firm-to-firm relationship. Coefficients are recovered from equation (2). The figure reports the estimates and their 95% confidence intervals. Tenure is measured by the number of periods since the beginning of the relationship.

Fact 4 *Within a firm-to-firm relationship, prices tend to decrease over time.*

While the price decline within a firm-to-firm relationship is statistically significant, the behaviour of prices following a switch is far less clear. This is illustrated in Figure 5 which shows the kernel density of price changes, conditional on a switch. Namely, we compute the price growth between the last transaction within a firm-to-firm relationship and the next transaction involving the same buyer and product but a different seller.¹⁵ The density is centered around zero with a mean at .006 and a median at 0. This means that when a firm switches to a new partner, it is equally likely to incur a drop as it is to incur a rise in the unit value it pays for the same good. This leads us to the fifth and final stylized fact.

Fact 5 *After a switch, the unit value paid by the importer is equally likely to increase or decrease.*

In the next section, we develop a dynamic search model that reproduces the stylized facts just described.

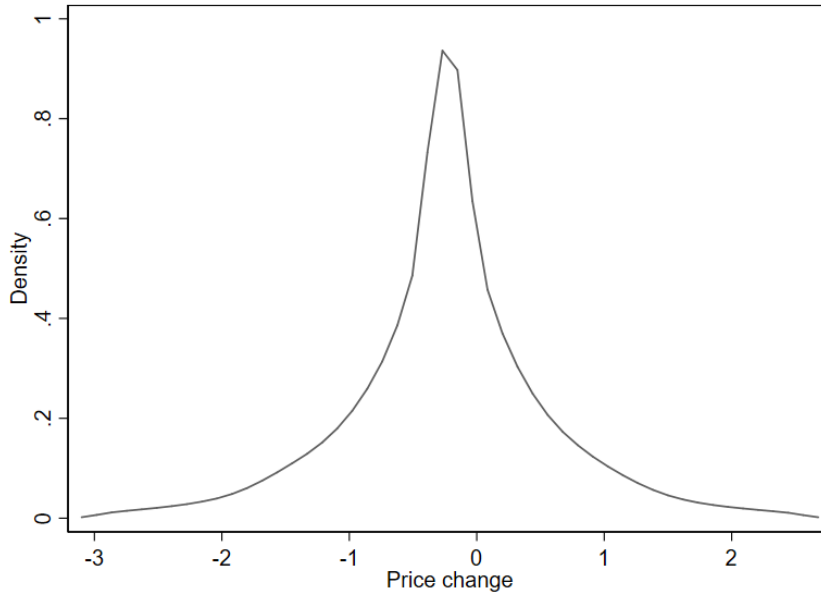
3 A Search Model of Firm-to-Firm Trade

Our model pictures an environment where heterogeneous sellers are matched randomly on the international markets with buyers of intermediate inputs. The sellers compete to retain the buyers by changing prices over time. Moreover, intermediate inputs are of different quality and competition hinges on the price-to-quality ratio rather than the input price only.

after the second transaction, we do not seek to model it afterwards.

¹⁵Here a switch denotes a situation in which we observe the same importer interacting with different French exporters. In the model, a switch will also designate a situation in which an importer terminates a relationship to start purchasing its input from a firm located in a different country.

Figure 5: Price changes, conditional on a switch



Note: This figure shows the kernel density of price changes, conditional on a switch. Price changes are computed in log differences, between the last transaction within a seller-buyer-product pair and the next transaction involving the same buyer×product but a different seller.

3.1 The demand for intermediate goods

In each country i , final good producers produce using inputs bought from intermediate good producers. Their production function is assumed CES and inputs are vertically differentiated. As we will see in the next subsection, the quality chosen by the final good producer for each input and the number of inputs depend on the producer’s network. We denote M_b the (endogenous) number of inputs used by producer b , q_j the quality of input j and p_j the price of that input. The producer chooses the quantity x_j of each intermediate input to minimize its costs, given the quantity x_b to be produced, which is taken to be exogenous.¹⁶

$$\left\{ \begin{array}{l} \min_{x_j} \sum_{j=1}^{M_b} p_j x_j \\ \text{s.t.} \\ x_b = \left(\sum_{j=1}^{M_b} (q_j x_j)^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}} \end{array} \right.$$

The solution of this program gives the demand addressed to each input provider, as a

¹⁶We discuss in Appendix C how the partial equilibrium model of trade in intermediates can be integrated into a general equilibrium structure whereby the demand expressed by final good producers is determined by their relative price, in a continuum of horizontally differentiated varieties.

function of its quality-adjusted price and the firm's network of input suppliers:

$$p_j x_j = x_b \left(\frac{p_j}{q_j} \right)^{1-\eta} \left(\sum_{j'=1}^{M_b} \left(\frac{p_{j'}}{q_{j'}} \right)^{1-\eta} \right)^{\frac{\eta}{1-\eta}} \quad (3)$$

and the final good producer's marginal cost of production:

$$mc_b = \left(\sum_{j=1}^{M_b} \left(\frac{p_j}{q_j} \right)^{1-\eta} \right)^{\frac{1}{1-\eta}} \quad (4)$$

From this, it becomes clear that it is optimal for the final good producer to choose, for each intermediate input, the seller offering the lowest price adjusted for quality, p_j/q_j , to minimize the marginal cost of production. Conditional on a chosen seller, the value of the transaction depends on the demand shifter that is specific to the buyer, on the price-to-quality ratio offered by the seller, and on the buyer's marginal cost of production. x_b and mc_b are buyer-specific and vary over time. As explained below, p_j/q_j is seller-specific and varies over time, both within a seller-buyer match and when the buyer switches to a new input provider.

3.2 A model of heterogeneous input suppliers

To understand the dynamics of input purchases, we need to describe both the matching process between the final good producers and the intermediate good producers, and how the prices are set. Buyers can be located in different countries and so can the sellers. For the sake of clarity and due to the limitation of our data, we focus here on sellers located in France (country F) and it is useful to keep in mind that we only observe buyers in the EU. Finally, we assume that search occurs (simultaneously) on as many separate markets as there are input types. What follows describes the output of the search and matching process for a given type j . For the sake of brevity, we no longer specify the type of inputs although all parameters in this subsection need to be understood as being input-specific. All coefficients that are indexed by F are further assumed to be heterogeneous across producing countries.

Intermediate good producers produce with a constant-return-to-scale technology and face iceberg transportation costs. They differ in terms of their productivity e , the quality q of their input and the cost of the input bundle, which is assumed to be country-specific. The unit cost of serving market i for a French firm of productivity e reads:

$$\frac{\nu_F d_{iF}}{e}$$

where d_{iF} is the bilateral iceberg cost, and ν_F is a unit cost shifter that determines how France compares with respect to other countries in terms of producing costs. As we will see when we discuss the matching of buyers and sellers, quality-adjusted productivities and quality-adjusted serving costs are the key variables. We thus define the quality-adjusted serving cost $c_{iF}(e, q)$ as,

$$c_{iF}(e, q) = \frac{\nu_F d_{iF}}{e q} \equiv \frac{\nu_F d_{iF}}{z} = c_{iF}(z)$$

As in [Lenoir et al. \(2022\)](#) and [Eaton et al. \(2022b\)](#), quality-adjusted cost is the single source of ex-ante heterogeneity across firms in the model, which explains the ex-post distribution in exporters' portfolios of buyers discussed in [Section 2.2.1](#).

When a buyer and a seller meet, the quality-adjusted productivity, $z = e \times q$, is drawn in a sampling distribution $F(z)$, with support on $[\underline{z}, +\infty]$. We follow [Eaton and Kortum \(2002\)](#) in assuming that the measure of firms in France with efficiency above z reads

$$\mu_F^Z(z) = T_F z^{-\theta}$$

which implies that $F(z) = 1 - (z/\underline{z})^{-\theta}$ and that the total measure of sellers in France (noted as S_F) is $T_F \underline{z}^{-\theta}$. Hence, whenever a French seller and a buyer from i are matched, the probability that the serving cost is below c reads

$$F_{iF}(c) = 1 - F(\nu_F d_{iF}/c) = \bar{F}(\nu_F d_{iF}/c) \quad (5)$$

Symmetrically, we define the probability that a buyer from i meets with a seller from another country ζ that offers a cost below c as $F_{i\zeta}(c) = \bar{F}(\nu_\zeta d_{i\zeta}/c)$.

3.3 Matching and pricing on the intermediate good market

Buyer-seller matching. In our framework, the buyers purchase intermediate goods of different qualities to produce the final goods. Under the CES assumptions, any buyer-seller match is potentially profitable. There are B_i buyers in country i . A buyer exits the market at exogenous rate μ and is replaced by a new buyer. New buyers start unmatched but, over time, they come into contact with multiple sellers and maintain links. A buyer meets with French sellers at rate γ_{iF} and with sellers from other countries at a rate of $\gamma_{i\bar{F}}$ which is the sum of the non-French meeting rates.

Borrowing from the model of the labor market in [Postel-Vinay and Robin \(2002\)](#), we assume that sellers Bertrand compete to supply goods to buyers and that there is no collusion between suppliers. We further assume that buyers always have the option to recall one of their previous sellers and that there is no commitment beyond the current transaction. Both are important assumptions that simplify the price setting.¹⁷ Assuming a buyer with n potential sellers, we index these sellers by their quality-adjusted serving cost : $c_1 \leq c_2 \leq \dots \leq c_n$.

Price dynamics. Among all the possible sellers known by the buyer for a given input, the best supplier is the one able to serve with the lowest quality-adjusted price that minimizes the buyer's marginal cost (4). Consider the two sellers with the lowest quality-adjusted serving cost $c_1 \leq c_2$, whose respective qualities are denoted q_1 and q_2 . The best supplier is able to set

¹⁷Namely, the wage equation in [Postel-Vinay and Robin \(2002\)](#) is affected by intertemporal considerations that are absent from our setting. The reason is that workers cannot recall previous employers whereas buyers can recall previously met input suppliers.

the price p such that the buyer is indifferent between buying the good from her or from the other seller when that seller offers the best quality-adjusted price c_2 :

$$p(q_1, c_2) = q_1 c_2$$

The price-setting mechanism thus pushes the quality-adjusted price (p/q) to be equal to the quality-adjusted cost of the second-best supplier (c_2).

Prices can be renegotiated over time as buyers meet with new, potentially more productive, sellers. Consider that the buyer matches with a new seller with quality-adjusted serving cost c' . We can distinguish three cases. First, if the new cost is below c_1 , the next transaction (if there is no new match before) will be with this new seller at price

$$p(q', c_1) = q' c_1$$

Interestingly, that price can be above or below the previous price if the quality of the good offered by the new seller is high enough. This is consistent with evidence in Figure 5: Conditional price changes are equally likely to be positive or negative. However, the quality-adjusted price p/q is always lower.¹⁸ Second, the new seller may have a quality-adjusted cost above the second best, $c' \geq c_2$ and nothing happens: the next transaction will be with the incumbent supplier at the same price.

Lastly, the new seller may have a quality-adjusted cost in between the best supplier and the former second best: $c_1 \leq c' < c_2$. In that case, the next transaction will be with the same supplier but at a lower price because the incumbent supplier has to match the utility level that the new supplier could provide:

$$p(q_1, c') = q_1 c'$$

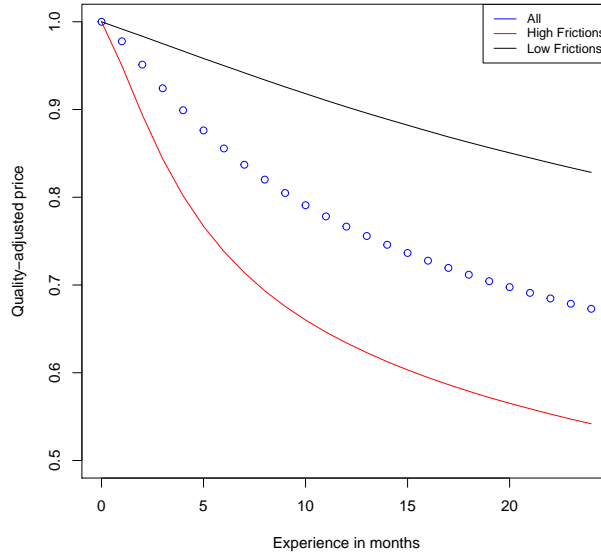
Hence, since $c' < c_2$, the new price offered by supplier 1 will be lower. Our model thus predicts that *within* a buyer-seller relationship, the price tends to decrease, as confirmed by evidence in Figure 4.

The magnitude of these price adjustments is illustrated in Figure 6, for an average level of meeting rates (blue circles) and when meeting rates are either low or high (black and red lines, resp.). The price dynamic is computed taking as reference a buyer entering the market at time 0 with the first supplier she has met and gaining bargaining power over time, through new matches. Upon entry, firms rapidly gain market power, which allows them to negotiate lower prices. In the mean market, the firm has already reduced the price of her input by more than 10% after 6 months of activity.¹⁹ Of course, the speed of these adjustments slows down

¹⁸Although testing this prediction of the model is tricky as quality-adjusted prices are not observed, Figure D5 in the Appendix presents suggestive evidence that is consistent with the prediction. Namely, we proxy a firm's position in the quality-adjusted price distribution by the individual fixed effect recovered from the estimation of equation (1). In theory, differences in the mean ability of input suppliers to accumulate a large number of buyers in a destination reflect their competitiveness there. In a second stage, we regress this proxy on the buyer-specific fixed effect and a dummy for the rank of the seller in the sequence of the buyer's French partners. As expected, we observe buyers climbing the distribution of sellers' attributes, over time.

¹⁹It should be noted that the magnitude of the price adjustment recovered in Figure 6 is not directly comparable to the empirical counterpart in Figure 4. The reason is that the data average price dynamics across

Figure 6: Mean dynamics of prices, as a function of the buyer’s experience



Notes: The figure shows dynamics of average prices in the model’s steady state, for average meeting rates (blue circles) and when meeting rates are either high (red line) or low (black line). The model is calibrated using parameters estimated in section 4. A buyer’s experience is measured from its first match with a supplier located in any country.

over time, as it becomes increasingly difficult to meet with suppliers that compete with the firm’s best match. As illustrated by the black and red lines of Figure 6, the dynamic of prices post-entry is also sensitive to the level of frictions. A higher overall meeting rate helps firms gain bargaining power more rapidly.

3.4 Distributions and shares

The distribution of suppliers. ²⁰ The overall meeting rate for a buyer, noted as γ_i , is the sum of the rates at which she meets French and non-French suppliers, i.e. $\gamma_i = \gamma_{iF} + \gamma_{i\bar{F}}$. When a buyer is new to a given input market, she meets her first supplier at that rate and exits the market at an exogenous rate μ . To ensure steady state, we assume that when a buyer exits, she is replaced by an unmatched buyer. Hence, the share u_i of buyers that are unmatched satisfies $B_i u_i \gamma_i = B_i (1 - u_i) \mu$ and thus

relationships involving buyers of heterogeneous experience, who are matched with a French supplier. Instead, Figure 6 controls for any source of composition effects by focusing on buyers with the shortest experience, who have met with a single supplier. As shown in Figure 6, the rate of downward price adjustments decreases with experience. As a consequence, the magnitude of price adjustments in Figure 6 constitutes an upper bound for any empirical price adjustments recovered from a set of firm-to-firm relationships.

²⁰Throughout this section, we take the point of view of the buyers and derive analytical results regarding the distribution of the suppliers they meet. The model also has predictions for the other side of the market. In Appendix C.2, we derive analytical results regarding the dynamics of buyer acquisition. These predictions help understand how the model reproduces the dynamics observed in the data and summarized in Facts 1. Since they are not directly used in the structural estimation, these derivations are left for the Appendix.

$$u_i = \frac{\mu}{\gamma_i + \mu}$$

After the first match, buyers keep on searching for new suppliers. The overall quality-adjusted serving cost distribution is a mixture of the country-specific ones, noted as $F_i(c)$, with

$$F_i(c) = \frac{\gamma_{iF}}{\gamma_i} F_{iF}(c) + \frac{\gamma_{i\bar{F}}}{\gamma_i} F_{i\bar{F}}(c) \quad (6)$$

for $c \in \left] 0, \max\left(\frac{\nu_F d_{iF}}{\underline{z}}, \frac{\nu_{\bar{F}} d_{i\bar{F}}}{\underline{z}}\right)\right]$ (see Appendix C.1 for details).

Before looking at how buyers are distributed among sellers, it is useful to note the relationship between the distributions of French and non-French suppliers:

$$F_{iF}(c) \equiv \left(\frac{\nu_F d_{iF}}{c \underline{z}}\right)^{-\theta} = \tau_{iF}^{-\theta} F_{i\bar{F}}(c)$$

where $\tau_{iF} \equiv \left(\frac{\nu_F d_{iF}}{\nu_{\bar{F}} d_{i\bar{F}}}\right)$ measures the comparative advantage of foreign suppliers over French firms. The distribution of French costs is a translation of the distribution for the other countries with $\tau_{iF}^{-\theta}$ being the cost shifter.²¹

Armed with this model, it now becomes possible to derive the distribution of costs faced by final producers in country i . We denote $L_i(c)$ its cumulated distribution function and $\ell_i(c)$ its probability density function. As long as buyers always choose the lowest cost supplier that they have met, $\ell_i(c)$ satisfies

$$\underbrace{B_i(1 - u_i)\ell_i(c) (\mu + \gamma_i F_i(c))}_{\text{outflows}} = \underbrace{B_i(1 - u_i)\gamma_i \bar{L}_i(c) f_i(c) + B u_i \gamma_i f_i(c)}_{\text{inflows}}$$

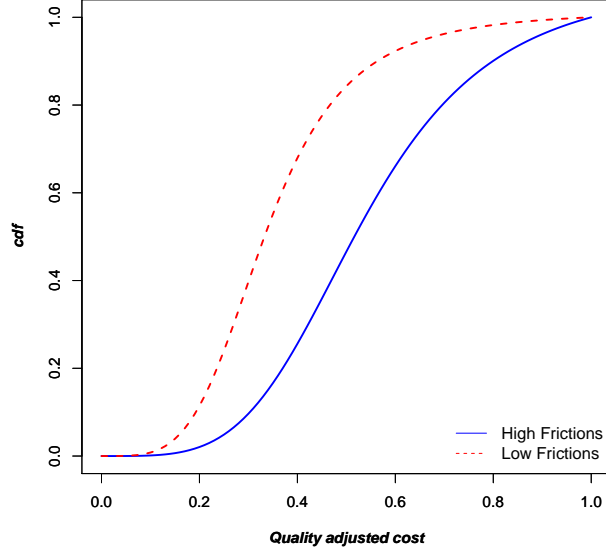
with $\bar{L}_i(c) \equiv 1 - L_i(c)$. The outflows are the sum of buyers exiting the market (μ) and buyers switching when they meet with a lower quality-adjusted cost supplier ($\gamma_i F_i(c)$). The inflows correspond to unmatched buyers meeting a cost- c supplier ($\gamma_i f_i(c)$) and buyers previously matched with sellers of serving cost higher than c ($\gamma_i \bar{L}_i(c) f_i(c)$). Simplifying and integrating by parts, we obtain

$$L_i(c) = \frac{\mu + \gamma_i}{\mu + \gamma_i F_i(c)} F_i(c) \quad (7)$$

The distribution of buyers among sellers hinges on the distribution of matches, $F_i(c)$, but also on the meeting frictions that slow down reallocations, hence the efficiency of the market. The

²¹It should be noted that the maximum serving cost for France is $\nu_F d_{iF}/\underline{z}$ and $\nu_{\bar{F}} d_{i\bar{F}}/\underline{z}$ for the other countries.

Figure 7: Cumulated distribution of the costs paid by individual buyers as a function of meeting frictions



Notes: The figure reports the cumulated distribution of input costs paid by buyers in a particular market, displaying either high or low overall frictions. The calibration uses parameters estimated in section 4.

relationship between search frictions and the efficiency of the matching process is illustrated in Figure 7 which shows the distribution of input costs, conditional on matches, for two values of meeting probabilities. Decreasing the meeting rate pushes the whole distribution of input costs to the right, i.e. lower meeting rates increase the mean cost of inputs for buyers.

The distribution of buyers among French suppliers. Consider the shares of buyers matched with a French seller, π_{iF} . Again, in equilibrium, flows in and out are balanced such that the density of buyers matched with a French-seller at cost c , noted as $\ell_{iF}(c)$, satisfies

$$\underbrace{(1 - u_i)\pi_{iF}\ell_{iF}(c)}_{\text{outflows}} (\mu + \gamma_i F_i(c)) = \underbrace{u_i \gamma_{iF} f_{iF}(c) + (1 - u_i) \bar{L}_i(c) \gamma_{iF} f_{iF}(c)}_{\text{inflows}} \quad (8)$$

As discussed in Appendix C.1, this equation, together with the definition of $u_i = \mu/(\mu + \gamma_i)$, can be used to solve for the equilibrium share of buyers matched with French firms. When μ is sufficiently close to zero, we have:²²

$$\pi_{iF} = \frac{\gamma_{iF}}{\gamma_{iF} + \gamma_{iF} \bar{\tau}_{iF}^\theta} \quad (9)$$

Finally, we show that

$$\ell_{iF}(c) = \ell_i(c) \quad (10)$$

²²Analytical details together with the formulas recovered in the general case when μ can take any value are provided in Appendix C.1. The interpretation of equation (9) discussed below goes through in the general case.

which means that buyers are identically distributed in terms of serving costs whatever the origin of their suppliers. As discussed in Appendix C.1, the invariance of serving costs across origin countries, conditional on a match, also implies that the expression for π_{iF} in equation (9) defines the share of country i 's absorption that is sourced from France, if z is small enough. As in Eaton and Kortum (2002), the geography of bilateral trade flows is entirely summarized by the probability that a buyer in i ends up purchasing inputs from France.²³

Two forces shape the geography of bilateral trade in our setting, the strength of matching frictions and Ricardian comparative advantages. The ratio of γ_{iF} over $\gamma_{i\bar{F}}$ indicates how easy it is for a buyer to meet a French supplier in comparison with non-French suppliers. $\tau_{iF}^{-\theta}$ instead reflects French suppliers' competitiveness, conditional on a match. Both an increase in γ_{iF} over $\gamma_{i\bar{F}}$ and a decrease in τ_{iF} improve the likelihood that a French supplier serves market i . As already discussed in Lenoir et al. (2022), introducing search frictions in a Ricardian model of trade helps refine our understanding of the geography of trade. As in Eaton and Kortum (2002), the interaction of technology and geography reflected in τ_{iF} shapes the Ricardian advantage of French firms in market i . In comparison with the frictionless world in Eaton and Kortum (2002), heterogeneity in bilateral search frictions distorts trade in favor of relatively low search / high meeting rates countries.

4 Structural estimation

4.1 Identification strategy

We estimate our model using an unconditional inference method and maximum likelihood, and our strategy uses transaction and switch frequencies to recover the value of the structural parameters. In the following section, we present that method and discuss identification. As identification does not use information contained in actual trade transactions (i.e. prices and quantities), we will later evaluate the performance of the model by considering untargeted moments computed from price and quantity data.

Before jumping into the estimation procedure, it is important to note that our framework models the determinant of matches and transitions while, in the data, any observation implies a transaction. A new match or a switch is only observed at the time of a transaction. For that reason, we need to make assumptions about the transaction process. As for any event in the model, transactions are assumed to be exponentially distributed. We note t_{iF} , the rate at which a buyer from country i makes a transaction with a French seller. We allow for a discrete heterogeneity in the transaction rates : a buyer can be of type 1 with probability p (transaction rate noted t_{iF}^1), or type 2 with probability $1 - p$ (transaction rate noted t_{iF}^2). Since the estimation is performed at the level of a specific product and destination, these probabilities

²³Using this result, we can use equation (9) to define $\tau_{iF}^{-\theta}$ as a function of the meeting rates and the *observed* shares

$$\tau_{iF}^{-\theta} = \frac{\pi_{iF}}{(1 - \pi_{iF})} \frac{\gamma_{i\bar{F}}}{\gamma_{iF}} \quad (11)$$

This will be used in the estimation.

are allowed to vary in these dimensions. While we allow for heterogeneity, a strong assumption is that the transaction rate is uncorrelated with the match quality.

The structural parameters of our model are the matching rates γ_{iF} and $\gamma_{i\bar{F}}$, the parameters shaping the distributions of serving costs, θ and τ_{iF} , the transaction rates t_{iF}^1 and t_{iF}^2 , and the rate at which a buyer exits the market, μ . Our dataset records transactions but also switches when buyers are observed making a transaction with a new supplier. Note that a switch can be intermediated first by i) an unobserved switch to a foreign seller before the buyer is matched again with a (more efficient) French supplier, or ii) a switch to a French seller that does not give rise to a transaction. Considering which moments of the data contribute to the identification of these parameters, the transaction rates t_{iF}^1 and t_{iF}^2 are identified by the frequencies of the transactions. Our assumption that they are independent of the match quality obviously eases identification. In the same way, the frequency at which we observe switches between French sellers helps to pin down γ_{iF} .²⁴ However, this is not as straightforward as for the transaction rates, for two reasons. First, the switch rate between two French sellers depends on the cost at which the buyer is currently served, for which we have no direct information. Second, the switch can be indirect if the buyer first made an unobserved transition.

With respect to the first problem, our method uses the fact that *unconditional* hazard rates only depend on the structural parameters. As an example, consider the overall hazard rate for a buyer matched with a French seller of quality c . It adds the exit rate and the matching rate and reads

$$H(c) = \mu + \gamma_{iF}F_{iF}(c) + \gamma_{i\bar{F}}F_{i\bar{F}}(c) \quad (12)$$

As noted, we do not observe c in the data but our model gives us the distribution of buyers among French sellers $\ell_{iF}(c)$. Using (9) and (10), we compute the unconditional hazard rate

$$\begin{aligned} \int_{c_{inf}}^{c^{sup}} H(c)dL_{iF}(c) &= \frac{\gamma_{iF}\tau_{iF}^{-\theta} + \gamma_{i\bar{F}}}{\gamma_{iF} + \gamma_{i\bar{F}}} \int_{c_{inf}}^{c^{sup}} \frac{\mu(\mu + \gamma_{iF} + \gamma_{i\bar{F}})dF_{i\bar{F}}(c)}{\mu + \gamma_{iF}\tau_{iF}^{-\theta}F_{i\bar{F}}(c) + \gamma_{i\bar{F}}F_{i\bar{F}}(c)} \\ &= \frac{\gamma_{iF}\tau_{iF}^{-\theta} + \gamma_{i\bar{F}}}{\gamma_{iF} + \gamma_{i\bar{F}}} \int_0^1 \frac{\mu(\mu + \gamma_{iF} + \gamma_{i\bar{F}})}{\mu + \gamma_{iF}\tau_{iF}^{-\theta}x + \gamma_{i\bar{F}}x} dx \end{aligned} \quad (13)$$

Equation (13) shows that *unconditional* hazard rates, i.e. the expected hazard rates over the observed buyer population, solely depend on the matching rate parameters γ_{iF} and $\gamma_{i\bar{F}}$, the exit rate μ , and τ_{iF} , the relative cost advantage of France vs the rest of the world. This reasoning is true for the densities/hazard rates of any type of events described by our model since they are all a combination of the matches/transaction rates and serving-cost distribution, integrated by $\ell_{iF}(c)$ (or $\ell_{i\bar{F}}(c)$).

Given that we already have to estimate γ_{iF} and $\gamma_{i\bar{F}}$, it is unclear how τ_{iF} , θ and μ would be

²⁴In principle, our model could be estimated as a duration model or using frequencies/number of switches and transactions. We chose the later for the sake of practicality but durations and frequencies are the two faces of the same coin: since the events are exponentially distributed, the underlying Poisson process also describes the distribution of the number of events within a certain time frame.

identified separately. However, τ_{iF} or even $\tau_{iF}^{-\theta}$ does not need to be estimated. Indeed, $\tau_{iF}^{-\theta}$ can be replaced by a function of the matching rates and the observed French share, using equation (9).²⁵ Finally, exits are used to pin down $\gamma_{i\bar{F}}$. Note that exits are never directly observed: they have to be essentially inferred using censoring rates. For that reason, we calibrate the exit rate μ to 0.01. μ being given and the transaction rates identified, $\gamma_{i\bar{F}}$ is identified by the share of observations that disappear after one transaction while γ_{iF} is identified using the switching rate.

The remaining difficulty is that observed switches can be intermediated by a number of unobserved switches and especially by switches toward foreign sellers. For this reason, it is impossible to derive the related densities that would be necessary for a regular maximum likelihood procedure. The solution is to rely on a *simulated* estimation method: for given values of the structural parameters we simulate our model and compute the required frequencies. We then choose the estimated parameters that best reproduce the empirical frequencies. Section B in the Appendix details how the simulated likelihood procedure is implemented in practice.

4.2 Estimation results

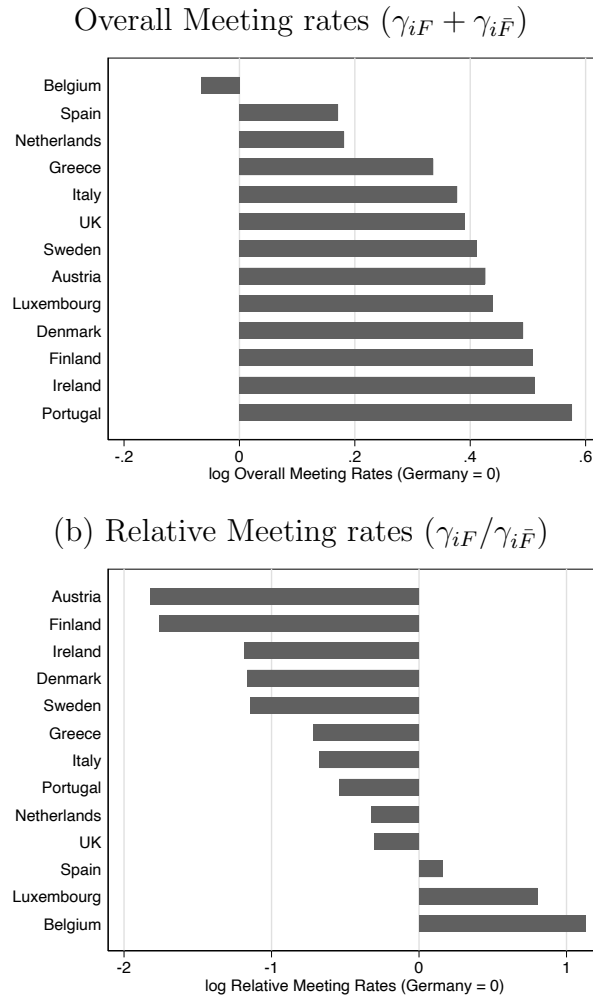
The model is estimated for 330 sector \times country pairs for which the number of observations and switches is sufficient for identification. Because the estimation strategy is demanding, we chose to pool observations observed at the level of a (product-specific) transaction across products within broader CPA2/NACE sectors and constrain the model's parameters to equality across products within a sector. Whereas this implies losing granularity on the estimated parameters, we gain significantly in terms of the precision of estimates. Figures D6 and D7 in the Appendix display how we match the distributions of transactions and the distributions of switches : we compare the distribution in our data with the distribution we obtain in simulated data where we simulate each market at the estimated values. We have an almost perfect match over the distribution of transactions. This is not surprising : transaction frequencies are high in the data thus offering a useful source of identification and, even if the transaction rates are not formally separately identified, they only depend on the other parameters through censoring. If we consider the distribution of the number of switches, we match the data reasonably well, although the distribution on simulated data is more skewed than the data counterpart.

Figure D8 in the Appendix displays the distributions of the estimated parameters, by country. The figures show the large dispersion in estimates: the distributions for the meeting rates γ_{iF} and $\gamma_{i\bar{F}}$, for the transaction rates $p t_{iF}^1 + (1 - p)t_{iF}^2$ and for French comparative advantages $\tau_{iF}^{-\theta}$ are skewed. French firms tend to be at a disadvantage in foreign markets, from the point of view of their Ricardian comparative advantage and their relative meeting rates, i.e. $\tau_{iF}^{-\theta}$ and $\gamma_{iF}/\gamma_{i\bar{F}}$ tend to cluster below 1. This is not surprising given that the parameters systematically compare French exporters to all possible competitors located in any other country, including the destination country itself.

To gain more interesting insights into how estimated coefficients vary across countries and

²⁵In practice, we do not use this approximated formula, but instead use the formulas provided in the Appendix, accounting for the possibility of different supports for the two F -distributions.

Figure 8: Dispersion in estimated search parameters, across countries

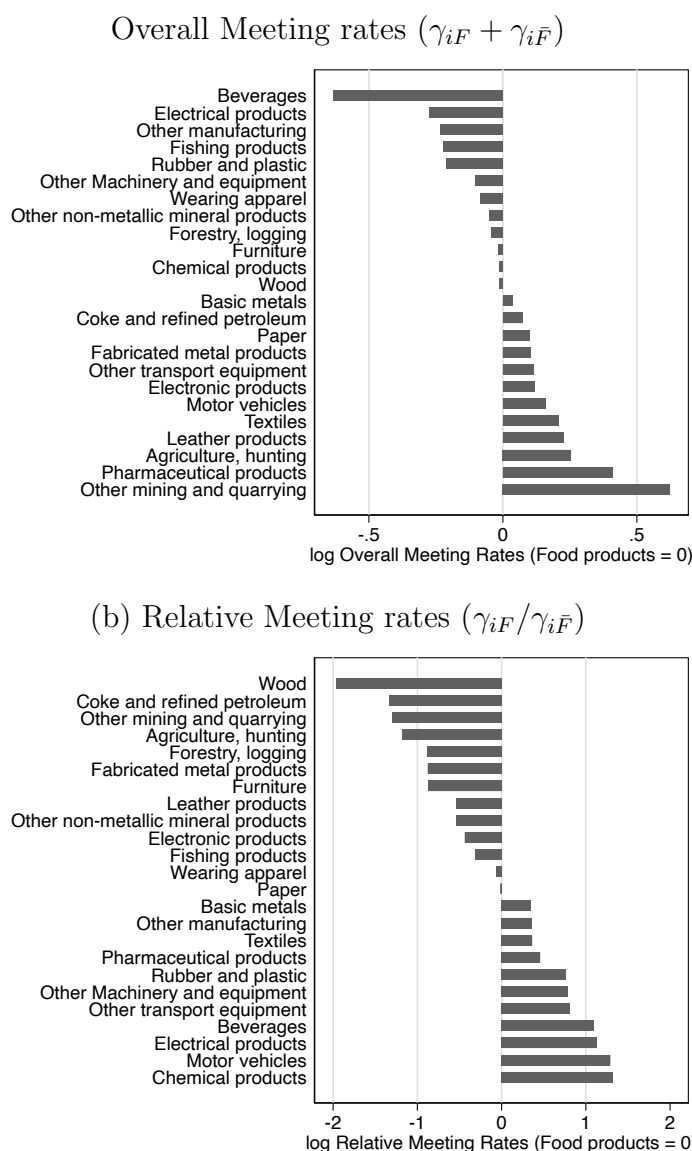


Note: The figure shows the mean value of estimated parameters, by country. All values are normalized by the mean estimate for Germany.

sectors, Figure 8 and 9 show the mean value of estimated parameters, by country and sector, respectively. In both cases, the variability in the dimension that is ignored is controlled for using a fixed effect in a two-way fixed effect decomposition of estimated parameters. The dispersions are shown for estimated overall meeting rates (top panels) and the relative meeting rate of French firms (bottom panels), with the former being indicative of the overall magnitude of frictions whereas the later captures the comparative advantage of French firms in terms of meeting rates. Figure D9 shows similar histograms for estimated Ricardian advantages.

Consider first the dispersion across countries, shown in figure 8. The decomposition reveals sizable heterogeneity in the overall degree of frictions, with the overall meeting rate being 60% larger in Portugal than in Belgium, on average. However, the relative advantage of French firms in terms of meeting rates is hardly correlated with the overall level of frictions. Here, geographical and cultural proximity seems to help, since the destinations in which French firms have relatively high meeting rates are all neighboring countries, namely Belgium, followed by Luxembourg, Spain, Germany and the UK. This result is consistent with the trade literature

Figure 9: Dispersion in estimated search parameters, across sectors



Note: The figure shows the mean value of estimated parameters, by sector. All values are normalized by the mean estimate of the Food products industry.

that attributes part of the gravity structure of trade to the impact of information frictions and their correlation with distance (Rauch, 1999, 2001).

The dispersion in average meeting rates across sectors is even stronger, with a 120% gap between the least frictional sector (“Other Mining and quarrying”) and the most frictional one (“Beverages”). In relative terms, French firms are sometimes at an advantage in more frictional markets such as beverages or electrical products. But the two sectors in which relative meeting rates are the highest are chemical products and motor vehicles, two sectors that are at the core of France’s international competitiveness. This result confirms the role of search frictions in shaping comparative advantages, besides Ricardian forces. Finally, although these average statistics are useful, it is important to point out that our estimates display significant heterogeneity across products within a country and across countries within a product. These

Table 1: Ricardian versus frictional determinants of trade

	Dep. Var			
	$\ln \frac{\gamma_{iF}}{\gamma_{i\bar{F}}}$ (1)	$\ln \tau_{iF\bar{F}}^{-\theta}$ (2)	$\ln \frac{\gamma_{iF}}{\gamma_{i\bar{F}}}$ (3)	$\ln \frac{\gamma_{iF}}{\gamma_{i\bar{F}}}$ (4)
$\ln \frac{\pi_{iF}}{1-\pi_{iF}}$	0.405 ^a (0.075)	0.595 ^a (0.075)	0.286 ^a (0.074)	0.254 ^a (0.082)
Obs.	330	330	330	330
Adjusted R^2	.083	.151	.209	.217
Country FE	No	No	Yes	No
Product FE	No	No	No	Yes

Note: The RHS is based on predicted trade shares using:

$$\frac{\pi_{iF}}{1-\pi_{iF}} = \frac{\gamma_{iF}}{\gamma_{i\bar{F}}} \tau_{iF\bar{F}}^{-\theta}$$

The variance decomposition in columns (1) and (2) is thus exact.

dimensions of heterogeneity explain 86 and 54% of the variance of overall and relative meeting rates, respectively.

5 Impact of search frictions on trade adjustments

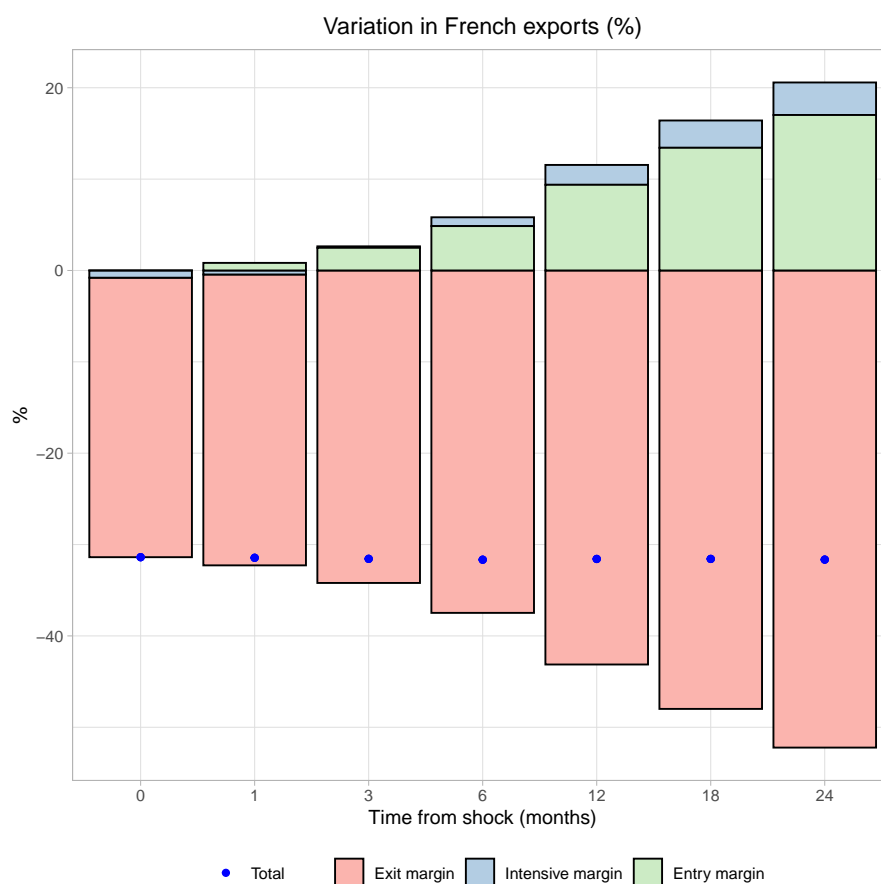
Armed with our estimates, we can now examine how search frictions shape trade adjustments in our model. To this end, we rely on a counterfactual. We simulate a uniform 10% permanent shock to French relative prices and study the adjustment of trade, at product level and on individual firms. We first discuss how product-level trade adjusts to the shock, before delving into heterogeneity across firms.

5.1 Search frictions and product-level trade elasticities

As explained in Section 3.4, our model delivers a geography of product-level trade that is shaped by the interaction of Ricardian comparative advantages and search frictions. Since the empirical strategy allows us to identify both parameters, we can first quantify the relative contribution of these forces. Results are summarized in Table 1. Consider first columns (1) and (2), which give the unconditional variance decomposition of trade shares in terms of search frictions and comparative advantages. Overall, relative meeting rates explain 40% of the variance in French firms' foreign market shares, the rest being attributable to Ricardian comparative advantages. The contribution of search frictions is reduced when we focus on the variance of trade shares across products within a country or across countries within a product, as in columns (3) and (4). Still, the explanatory power of search frictions remains sizable.

What does the prominent role of search frictions imply for the elasticity of trade to shocks? To examine this question, we rely on our counterfactual. Figure 10 illustrates the aggregate impact of the shock. The elasticity of trade with respect to the 10% cost shock is high, at 3.2.

Figure 10: Impact of a 10% relative price shock on aggregate trade



Notes: The figure shows how aggregate trade adjusts to a 10% relative price shock. The blue dots represent the overall trade adjustment. The pink bars quantify the drop in trade attributable to disrupted firm-to-firm relationships (“Exit margin”). The green bars measure the volume of trade attributable to trade relationships that are formed posterior to the shock and involve a French firm (“Entry margin”). Finally, the blue bars capture the intensive margin, i.e. the change in the volume of trade within existing relationships (“Intensive margin”).

In our simulations, this value is reached on impact because most buyers in the steady state of our model are hit by a shock at a time when they have already accumulated enough matches so that they can immediately adjust.²⁶ Initially, most of the adjustment takes place at the extensive margin, i.e. buyers switch to non-French suppliers. As explained in the next section, these switches concern 20% of the population of buyers on impact, a proportion that increases over time to reach more than 60% after two years. Of course, buyers keep on meeting with (French and non-French) suppliers and the additional disruptions are somewhat compensated by new firm-to-firm relationships, although at a permanently lower rate than before the shock (the green bars in Figure 10). Finally, the shock also has consequences at the intensive margin (the blue bars in Figure 10). On impact, part of the shock is passed on prices and the volume of trade diminishes as a consequence. In the longer-run, price renegotiations within firm-to-firm relationships more than compensate for the initial decline.

Beyond these aggregate effects, the model predicts some heterogeneity in the strength of adjustments across markets due to differences in the magnitude of relative search frictions. Simple comparative statics convey insightful intuitions on how the *elasticity* of trade varies with search frictions. Using the approximation of trade shares in equation (9), we have:

$$\varepsilon_{iF} = \frac{\theta \pi_{iF}}{\tau_{iF}^{-\theta} \gamma_{iF} / \gamma_{i\bar{F}}} = \frac{\theta}{1 + \frac{\gamma_{iF}}{\gamma_{i\bar{F}}} \tau_{iF}^{-\theta}} \quad (14)$$

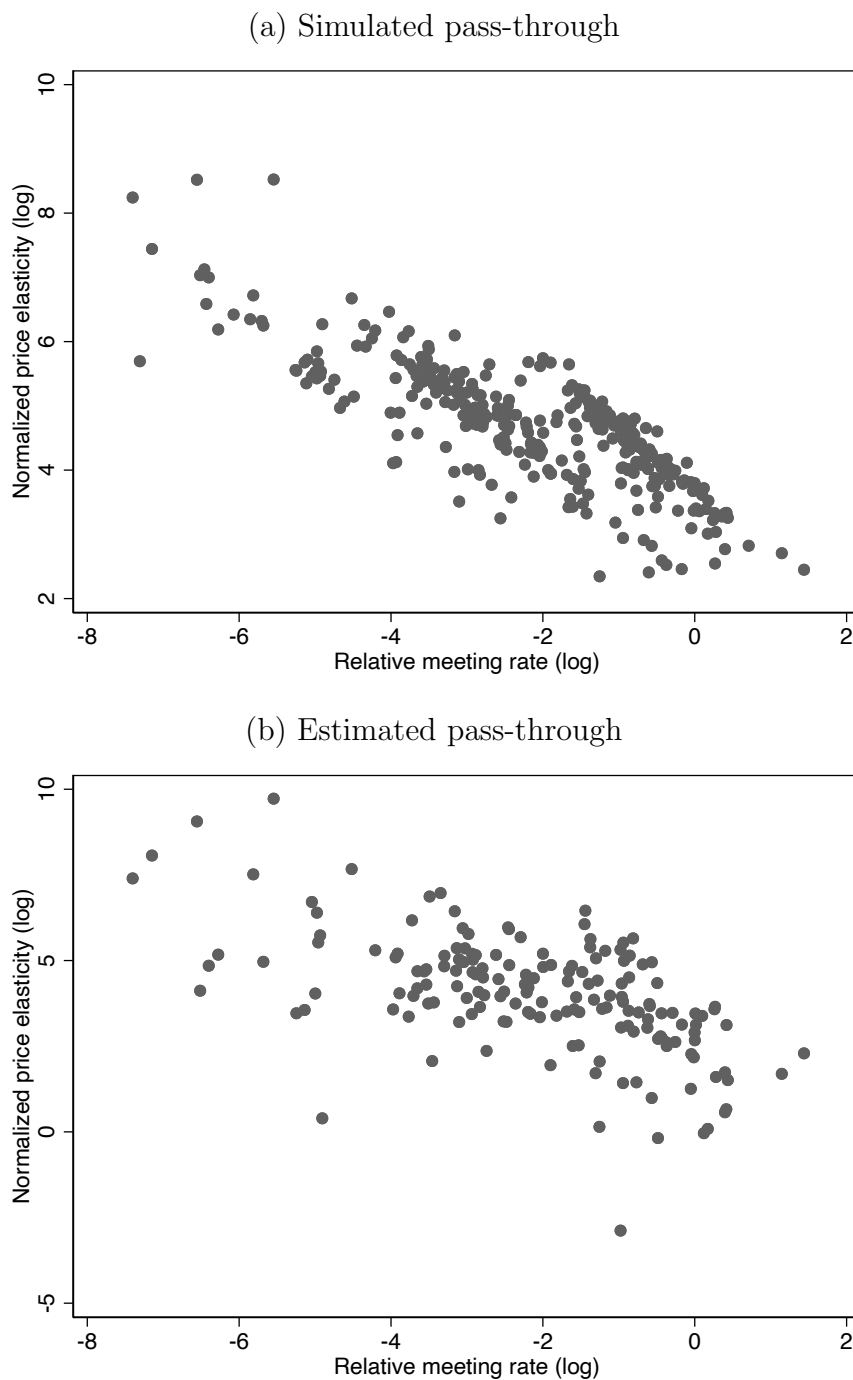
where $\varepsilon_{iF} \equiv -\frac{d \ln \pi_{iF}}{d \ln \tau_{iF}}$ denotes the elasticity of market shares with respect to a shock on France's relative costs. For low-enough French market shares, it is also the elasticity of bilateral trade to the shock. In the long-run, the model thus converges toward a log-linear relationship linking the price elasticity of trade, normalized by French firms' trade share ($\varepsilon_{iF} / \pi_{iF}$) and the comparative advantage of French firms in terms of costs ($\tau_{iF}^{-\theta}$) and meeting rates ($\gamma_{iF} / \gamma_{i\bar{F}}$). As in [Eaton and Kortum \(2002\)](#), the elasticity of trade with respect to relative price shocks is muted in markets in which French firms have a high comparative advantage, where the relative price shock needs to be large in order to induce buyers to switch and trade to adjust at the extensive margin.²⁷ Conditional on Ricardian comparative advantages, the model predicts that the elasticity of trade is also lower in markets in which the relative meeting rate of French firms is larger. A high meeting rate for France indeed implies that the direct competitors of French firms in foreign markets are more likely to be French. As a consequence, the shock does not deteriorate French firms' competitiveness as much as it would in a market in which their competitors are mostly non-French.

The interaction is illustrated in the top panel of Figure 11. Based on the simulated data recovered from the counterfactual exercise, we compute the trade elasticity, for each country and sector pair, which we normalize by trade shares. The normalized elasticities are then correlated

²⁶A key assumption of our model for the fast adjustment recovered here is the fact that buyers keep in memory their past matches and do not commit with their current partner beyond the next transaction. When hit by the shock, buyers can thus immediately use the bargaining power accumulated through past matches.

²⁷Our empirical analysis is performed on highly disaggregated data. In our context, comparative advantages must thus be understood as a mix of industry-specific comparative advantages, destination-specific trade costs, and *granular* comparative advantages based on firms' individual quality-adjusted costs ([Gaubert and Itskhoki, 2021](#)).

Figure 11: Distribution of elasticities, along the distribution of relative meeting rates



Notes: The figure in panel (a) shows a scatter plot of simulated elasticities, normalized by the initial market share of French firms, against relative frictions, both in logs. In the model, the relationship is linear as shown in equation (14). The figure in panel (b) shows a scatter plot of normalized trade elasticities against estimated relative meeting rates, both in logs. The normalized trade elasticities are defined as the ratio of the estimated long-run elasticity of trade over French firms' market share. The figure is restricted to product*country pairs in which the estimated long-run elasticity is negative.

with estimated meeting rates. As expected, the correlation is significantly negative, at 83%, i.e. the response of trade to the cost shock is particularly strong in product markets in which French firms tend to compete with non-French competitors, due to a relative disadvantage in meeting probabilities. The relationship between trade elasticities and estimated meeting rates is confirmed in the bottom panel of Figure 11, using external measures of trade elasticities. Here, we estimate the long-run elasticity of trade to real exchange rate shocks for each sector and country using an error correction model inspired from [Alessandria et al. \(2021b\)](#).²⁸ Here as well, the correlation with estimated relative meeting rates is negative and significant, at -61%.

5.2 Prices, markups, and trade adjustments

A novel feature of our model lies in its rich predictions regarding the incidence of relative price shocks, along the distribution of buyers and sellers. Depending on the impact of the shock on the strength of competition in each firm’s supplier network, a shock can go from having zero consequences on the buyer’s input prices to being fully passed onto the buyer. Such richness is useful in as much as it can help rationalize the strong degree of heterogeneity in pass-through rates across firms and markets, as has been emphasized in the previous literature ([Burstein and Gopinath, 2014](#), [Berman et al., 2012](#), [Cavallo et al., 2021](#), [Amiti et al., 2019](#)). To examine the quantitative importance of these different forms of adjustment and their timing, we again rely on the counterfactual shock to French sellers’ relative prices. We trace out the average effect of this shock and explore its heterogeneity taking as given the comparative advantage of French firms and the level of search frictions across (country×sector) markets.

Figure 12 summarizes the main drivers of the adjustment, on impact (red bars) and 24 months after the shock (blue bars), focusing on trade relationships that are directly hit by the shock, i.e. those involving a French supplier. On impact, the shock mostly bears on the supplier side. Nearly 70% of the relationships entailing a French supplier display zero pass-through, which means that the shock is entirely absorbed into the supplier’s markup. The prevalence of zero pass-through explains the left shift in the distribution of French sellers’ markups, illustrated in Figure 13. On impact, the median markup of French firms drops from 1.28 to 1.19, a 7% drop.²⁹ Such adjustments occur when the French firm’s competitor is not French, but the firm has a sufficiently low quality-adjusted cost relative to her competitor to absorb the shock. In 25% of cases, the seller cannot absorb the shock, but the buyer has a competitive (non-French) alternative and thus switches to a non-French buyer. In that case, the

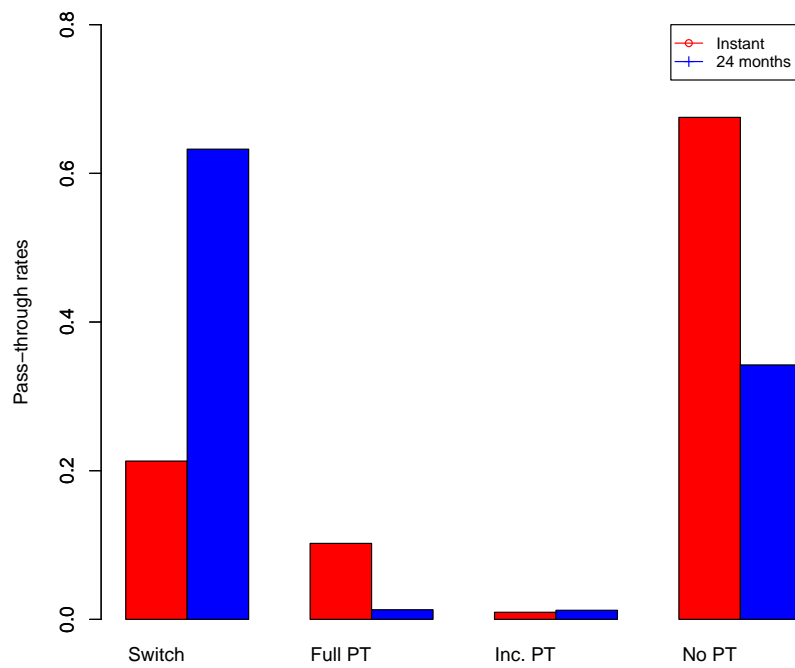
²⁸Namely, we first aggregate the trade data at the monthly and product-country level. The panel is merged with country-specific real exchange rate series recovered using Eurostat data on nominal exchange rate and PPI for France and each of its European partners. Finally, we estimate the following model for each sector×country pair:

$$d \ln X_{pct} = \beta^{SR} d \ln RER_{ct} + \beta^2 \ln RER_{ct-1} + \beta^3 \ln X_{pct-1} + FE_{pt} + \varepsilon_{pct}$$

where X_{pct} is the value of exports of product p to country c at time t , RER_{ct} is the real exchange rate between France and country c , defined such that an increase in RER denotes a real appreciation for French firms, and FE_{pt} is the product×period fixed effect. In this equation, β^{SR} estimates the short-run elasticity of trade to RER shocks while the long-run elasticity is defined as $\beta^{LR} = -\beta^2/\beta^3$.

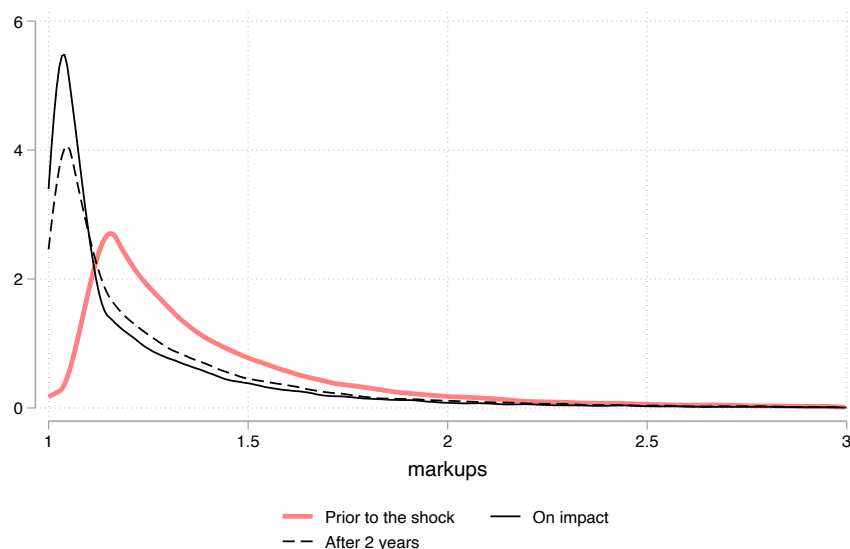
²⁹The drop in markups following a positive relative price shock is consistent with evidence recovered from trade liberalization episodes, e.g. [De Loecker et al. \(2016\)](#) who also point to important heterogeneity along the distribution of firms.

Figure 12: Heterogeneous pass-through, posterior to the shock



Notes: The figure shows the result of simulations of a 10% relative price shock. The simulations assume: $\theta=4$ and use the estimates of search frictions and comparative advantages described above.

Figure 13: Markups of French sellers



Notes: The figure shows the distribution of markups among transactions involving a French seller before the shock, on impact, and 2 years after the shock.

price of the buyer increases by an average 4%. The incidence thus partly falls onto the buyer, although French sellers are most severely hit as they lose clients on which they were charging strictly positive markups prior to the shock. Finally, about 10% of the buyers involved in relationships with French firms are stuck with their current partner *and* bear the cost of the adjustment, through input price increases. When it happens, the pass-through is almost always complete.³⁰ Complete pass-through happens in our model when the buyer's second best option is a French supplier, before and after the shock. This outcome is more likely when France's Ricardian advantage is high and bilateral search frictions are moderate.

The incidence of the shock thus varies depending on the buyer's bargaining power at the time of the shock. Buyers whose network is mostly composed of French suppliers are left badly equipped to negotiate, which translates into a larger pass-through. The comparison of the red and blue bars in Figure 12 shows that the prevalence of full pass-through decreases over time, whereas the probability of a switch increases from 25% to more than 60%. These dynamics are the consequences of the buyer continuously meeting with new suppliers, which improves her bargaining power.

Figure 14 summarizes the evolution of prices for buyers active with a French supplier at the time of the shock. In the top panel, the pass-through is computed for bins of buyers grouped based on the duration of their relationship post shock. The red line shows the evolution of prices for buyers that switch immediately to a non-French supplier. On average, the switch is associated with a 4.5% price increase. The initial price increase is somewhat compensated over

³⁰Incomplete pass-through occurs when the French supplier remains the buyer's best option but the shock affects the buyer's second best option, which switches from French to non-French. In such situations, the buyer can negotiate the price downward and the pass-through is incomplete. In practice, incomplete pass-through is a rare outcome in our simulations.

time, through downward price renegotiation and further switches. On average, however, it takes 22 months for switchers to recover the pre-shock price levels. In the same figure, the different shades of gray summarize the evolution of prices within relationships that are disrupted 1 to 23 months after the shock. On impact, the size of the pass-through varies between 10 and 50%, on average, and is negatively correlated with the subsequent length of the relationship, i.e. firms that pass a larger share of the shock onto buyers are also more likely to be replaced by non-French suppliers in the near future. The downward trend in prices from period 1 is again triggered by subsequent matches that improve the buyer’s bargaining power.

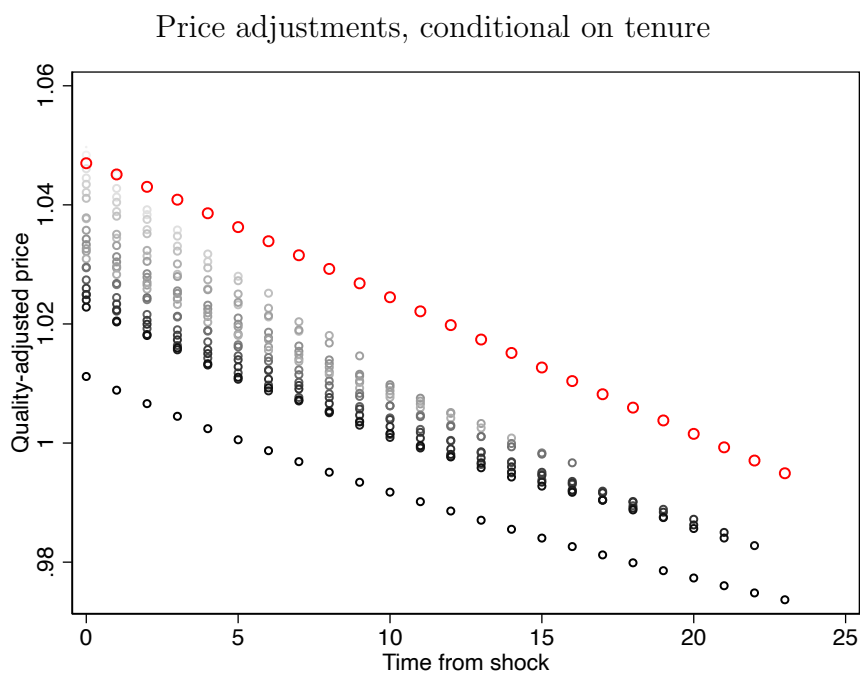
The bottom panel in Figure 14 summarizes these heterogeneous price dynamics using the average evolution of prices paid by buyers.³¹ On impact, there is a 19% pass-through rate that continues to increase, reaching 26% after 2 years. As a matter of comparison, Fontagné et al. (2018) estimate that French exporters absorb about one third of tariff shocks. Whereas this share may seem low in comparison with our simulations, we have to keep in mind that the nature of the shock is different as well. The shock that we simulate is specific to French exporters while tariff shocks affect all European producers. Given that the competitiveness shock that we simulate is larger in scale, it is not surprising that its incidence on French firms is larger as well.³² Interestingly, the incidence of the shock is also felt by buyers that are not directly exposed on impact, i.e. those that are matched with non-French sellers (red line in the bottom panel of Figure 14). The reason is that the shock permanently shifts the distribution of sellers’ costs to the right, thus reducing the amount of competition in input markets. As a consequence, the average price paid by buyers to non-French sellers increases by .4% on impact, and remains .3% above its pre-shock level in the long run. Overall, we observe that the bulk of the adjustment takes place within a 6-month window after the shock. Such relatively fast adjustment is consistent with observations following large and sudden shocks such as the 2015 appreciation of the Swiss Franc studied in Auer et al. (2021).

We conclude this section by investigating the heterogeneity in the magnitude of the adjustments just described, across sector×country pairs, and across firms, within a sector×country. Results are summarized in Table 2. The impact of the shock on buyers’ prices is larger in markets displaying high *relative* meeting rates, in which the probability of the buyer being able to switch is low (columns (1)-(2)) while price adjustments, conditional on no switch, are large on average (columns (5)-(6)). Instead, high *overall* meeting rates increase the likelihood that the buyer’s network at the time of the shock is wide enough so that the buyer can reduce her exposure to the shock, by switching or negotiating a price cut. The heterogeneity induced by sector×country differences in the magnitude of frictions translates into heterogeneous adjust-

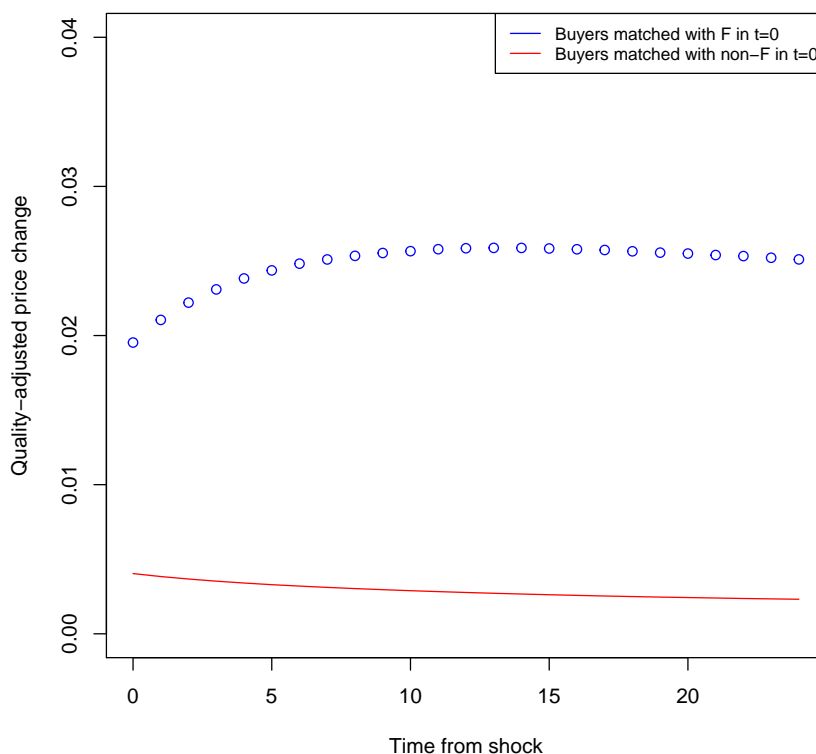
³¹Here, the dynamic of prices is normalized by the evolution of prices in the absence of the shock so that the average price increase can be interpreted as being *caused* by the relative price shock.

³²Cavallo et al. (2021) find that, during the US-China trade war, Chinese exporters absorbed less than 2% of US tariffs hikes, whereas US exporters absorbed more than 30% of Chinese tariff hikes. In the context of our model, these different levels of pass-through can be explained by the comparative advantage of China in sectors targeted by the US: using trade data, we indeed found that the average market share of Chinese exporters in the US in sectors targeted by the tariffs is high, at 20%, whereas the market share of US firms in sectors targeted by the Chinese authorities is around 10%. Differences in the bargaining power of targeted exporters may thus explain asymmetric pass-through rates. Likewise, the 20% average pass-through rate measured for French firms reflects the small market share of France in most foreign markets, which reduces their bargaining power.

Figure 14: Price adjustments, posterior to the shock



Aggregate normalized price adjustments



Notes: The figure shows the average dynamics of prices posterior to the shock. The top panel shows the dynamics of prices conditional on the length of the post-shock tenure. The red line corresponds to buyers that immediately switch to a non-French supplier. The gray lines are for buyers that do not immediately switch. The darker the gray, the shorter the post-shock tenure of the relationship. The bottom panel summarizes how these rich price dynamics affect the average price paid by buyers initially matched with French firms (blue circles) and compare it with the average price paid by firms matched with non-French partners. In both cases, these price adjustments are normalized by the corresponding dynamics in the absence of a shock.

ments across countries, as illustrated in Figure D10 in the Appendix. On impact, the incidence of the shock on foreign buyers varies, from a 15% average pass-through in Austria, to more than 25% in Belgium. The model also displays heterogeneity across firms within a sector×country. The switching probability is positively associated with the buyer’s experience at the time of the shock, while it is lower if the buyer’s current match displays a relatively low cost (column (3)-(4)). The buyer’s past experience also affects her bargaining power conditional on no switch. A longer experience is associated with a wider network and thus a better option value, on average, which explains the lower price growth (column (7)-(8)).³³ The heterogeneity across buyers and suppliers is systematically more pronounced in markets displaying low relative meeting rates and high overall meeting rates.

Table 2: Determinants of instantaneous extensive and intensive adjustments to the relative price shock

	Dependent variable							
	Switching probability				Price adjustment			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Relative meeting rate	-0.065*** (0.000)	-0.068*** (0.000)			0.023*** (0.000)	0.021*** (0.000)		
Overall meeting rate	0.047*** (0.001)	0.052*** (0.001)			-0.007*** (0.000)	-0.006*** (0.000)		
Buyer’s experience		0.077*** (0.001)	0.146*** (0.001)	0.149*** (0.001)		-0.017*** (0.000)	-0.012*** (0.000)	-0.012*** (0.000)
× $\gamma_F/\gamma_{\bar{F}}$				-0.019*** (0.001)				-0.013*** (0.000)
× $(\gamma_F + \gamma_{\bar{F}})$				0.039*** (0.001)				0.012*** (0.000)
log quality-adjusted cost		0.032*** (0.001)	0.209*** (0.001)	0.196*** (0.002)				
× $\gamma_F/\gamma_{\bar{F}}$				-0.115*** (0.001)				
× $(\gamma_F + \gamma_{\bar{F}})$				0.249*** (0.003)				
Country*Product FE	No	No	Yes	Yes	No	No	Yes	Yes
Obs.	1,385,095	1,385,095	1,385,095	1,385,095	781,402	781,402	781,402	781,402
R-squared	.018	.029	.080	.107	.143	.198	.321	.348
R-squared (within)			.039	.067			.039	.077

Notes: Robust standard errors in parenthesis. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

All in all, the results thus display a high level of heterogeneity, across markets and across firms within a market. The rich heterogeneity is recovered from the introduction of bargaining in a dynamic model of trade in frictional markets. In this set-up, the structure of a firm’s network at the time of the shock determines its ability to bargain, now and in the future.

6 Conclusion

We build and estimate a Ricardian model of firm-to-firm trade and on-the-match search. The model delivers empirically-consistent predictions on both the cross-section and the evolution of

³³In our model, the price adjustment conditional on no switch does not directly depend on the quality-adjusted cost of the buyer’s supplier. Instead, the price adjustment solely reflects changes in the buyer’s second-best options.

sellers' portfolio of buyers, as well as the dynamics of markups and prices, within and across firm-to-firm relationships. We estimate the model's structural parameters for 14 EU countries and 26 different sectors. In the context of our model, heterogeneity in search frictions faced by French firms explains 40% of the observed variance in France' market shares.

In a counterfactual analysis, we show the importance of search frictions for the adjustment of trade to relative price shocks. The elasticity of trade is systematically larger in markets displaying relatively high bilateral frictions, both in the model and in the data. Search frictions further induce rich price and markup adjustments. After a 10% cost shock hitting French firms, more than 65% of buyers negotiate so that sellers compress their markups and absorb the bulk of the shock. However, not all sellers can absorb the shock through their margin. About 20% of the buyers thus switch away from French suppliers to limit the price increase. These extensive adjustments drive bilateral trade down. Over time, buyers continue to meet with new sellers which allows them to gain bargaining power and renegotiate their prices down, within or across relationships. All in all, the relative price shock induces an aggregate 19% pass-through on impact, which increases to 26% after 2 years. The analysis shows that the level of pass-through and thus the incidence of the shock varies across markets with the level of search frictions. Buyers bear a large fraction of the price increase in markets where French firms face lower search frictions relative to their competitors.

The tractability of the model makes it possible to identify its parameters in a transparent way. Despite its simplicity, the model also produces rich patterns of price and markup adjustments, that are consistent with empirical evidence. The downside of the simplicity is naturally that we miss some potentially interesting aspects of firms' price setting. In particular, our model and identification strategy crucially rely on the assumption that firms do not commit to a price beyond the next transaction. In future research, it would be interesting to extend the model to contractual rigidities. Intuitively, contractual rigidities will add an intertemporal element to price-setting behaviors. Firms engaged in long-run contracts will not be able to switch immediately after the shock, which should enrich the model with a more interesting trade dynamic. A natural consequence that directly follows the logic of the paper is that these contractual rigidities will reduce buyers' bargaining power when confronted with a price shock. An open question is what moment in the data would potentially help identify such contractual rigidities.

References

- Alessandria, George, “International Deviations From The Law Of One Price: The Role Of Search Frictions And Market Share,” *International Economic Review*, November 2004, 45 (4), 1263–1291.
- Alessandria, George A., Costas Arkolakis, and Kim J. Ruhl, “Firm Dynamics and Trade,” *Annual Review of Economics*, 2021, 13, 253–280.
- , Shafaat Y. Khan, Armen Khederlarian, Kim J. Ruhl, and Joseph B. Steinberg, “Trade-Policy Dynamics: Evidence from 60 Years of U.S.-China Trade,” NBER Working Papers 29122, National Bureau of Economic Research, Inc August 2021.
- Alessandria, George and Horag Choi, “Do Sunk Costs of Exporting Matter for Net Export Dynamics?,” *The Quarterly Journal of Economics*, 2007, 122 (1), 289–336.
- and —, “Establishment heterogeneity, exporter dynamics, and the effects of trade liberalization,” *Journal of International Economics*, 2014, 94 (2), 207–223.
- Alviarez, Vanessa, Michele Fioretti, Ayumu Ken Kikkawa, and Monica Morlacco, “Two-Sided Market Power In Firm-to-Firm Trade,” 2021.
- Amiti, Mary, Oleg Itskhoki, and Jozef Konings, “Importers, Exporters, and Exchange Rate Disconnect,” *American Economic Review*, July 2014, 104 (7), 1942–78.
- , —, and —, “International Shocks, Variable Markups, and Domestic Prices,” *Review of Economic Studies*, 2019, 86 (6), 2356–2402.
- Antras, Pol and Elhanan Helpman, “Global Sourcing,” *Journal of Political Economy*, June 2004, 112 (3), 552–580.
- Arkolakis, Costas, “Market Penetration Costs and the New Consumers Margin in International Trade,” *Journal of Political Economy*, 2010, 118 (6), 1151 – 1199.
- , Theodore Papageorgiou, and Olga Timoshenko, “Firm Learning and Growth,” *Review of Economic Dynamics*, January 2018, 27, 146–168.
- Atkeson, Andrew and Ariel Burstein, “Pricing-to-Market, Trade Costs, and International Relative Prices,” *American Economic Review*, December 2008, 98 (5), 1998–2031.
- Auer, Raphael, Ariel Burstein, and Sarah M. Lein, “Exchange Rates and Prices: Evidence from the 2015 Swiss Franc Appreciation,” *American Economic Review*, February 2021, 111 (2), 652–86.
- Bagger, Jesper, Francois, Fabien Postel-Vinay, and Jean-Marc Robin, “Tenure, Experience, Human Capital, and Wages: A Tractable Equilibrium Search Model of Wage Dynamics,” *American Economic Review*, June 2014, 104 (6), 1551–1596.

- Behrens, Kristian, Brahim Boualam, and Julien Martin, “Are Clusters Resilient? Evidence from Canadian Textile Industries,” *Journal of Economic Geography*, 2019.
- Bergounhon, Flora, Clémence Lenoir, and Isabelle Mejean, “A guideline to French firm-level trade data,” 2018.
- Berman, Nicolas, Philippe Martin, and Thierry Mayer, “How do Different Exporters React to Exchange Rate Changes?,” *The Quarterly Journal of Economics*, 2012, *127* (1), 437–492.
- , Vincent Rebeyrol, and Vincent Vicard, “Demand Learning and Firm Dynamics: Evidence from Exporters,” *The Review of Economics and Statistics*, March 2019, *101* (1), 91–106.
- Bernard, Andrew B. and Andreas Moxnes, “Networks and Trade,” *Annual Review of Economics*, August 2018, *10* (1), 65–85.
- , —, and Karen Helene Ulltveit-Moe, “Two-Sided Heterogeneity and Trade,” *The Review of Economics and Statistics*, July 2018, *100* (3), 424–439.
- , Jonathan Eaton, J. Bradford Jensen, and Samuel Kortum, “Plants and Productivity in International Trade,” *American Economic Review*, 2003, *93* (4), 1268–1290.
- , Marco Grazzi, and Chiara Tomasi, “Intermediaries in International Trade: Products and Destinations,” *The Review of Economics and Statistics*, October 2015, *97* (4), 916–920.
- Boehm, Christoph E., Andrei A. Levchenko, and Nitya Pandalai-Nayar, “The Long and Short (Run) of Trade Elasticities,” *the American Economic Review*, 2022.
- Burstein, Ariel and Gita Gopinath, “International Prices and Exchange Rates,” in G. Gopinath, . Helpman, and K. Rogoff, eds., *Handbook of International Economics*, Vol. 4 of *Handbook of International Economics*, Elsevier, 2014, chapter 0, pp. 391–451.
- Cahuc, Pierre, Fabien Postel-Vinay, and Jean-Marc Robin, “Wage Bargaining with On-the-Job Search: Theory and Evidence,” *Econometrica*, March 2006, *74* (2), 323–364.
- Caliendo, Lorenzo and Fernando Parro, “Estimates of the Trade and Welfare Effects of NAFTA,” *Review of Economic Studies*, 2015, *82* (1), 1–44.
- Carballo, Jeronimo, Gianmarco I.P. Ottaviano, and Christian Volpe Martincus, “The buyer margins of firms’ exports,” *Journal of International Economics*, 2018, *112* (C), 33–49.
- , Gianmarco Ottaviano, and Christian Volpe Martincus, “The Buyer Margins of Firms’ Exports,” CEPR Discussion Papers 9584, C.E.P.R. Discussion Papers August 2013.
- Cavallo, Alberto, Gita Gopinath, Brent Neiman, and Jenny Tang, “Tariff Pass-Through at the Border and at the Store: Evidence from US Trade Policy,” *American Economic Review: Insights*, March 2021, *3* (1), 19–34.
- Chor, Davin and Lin Ma, “Contracting Frictions in Global Sourcing: Implications for Welfare,” 2020.

- Das, Sanghamitra, Mark J. Roberts, and James R. Tybout, “Market Entry Costs, Producer Heterogeneity, and Export Dynamics,” *Econometrica*, May 2007, 75 (3), 837–873.
- de Blas, Beatriz and Katheryn N. Russ, “Understanding Markups in the Open Economy,” *American Economic Journal: Macroeconomics*, April 2015, 7 (2), 157–180.
- Demir, Banu, Ana Cecília Fieler, Daniel Xu, and Kelly Kaili Yang, “O-Ring Production Networks,” NBER Working Papers 28433, National Bureau of Economic Research, Inc February 2021.
- Drozd, Lukasz A. and Jaromir B. Nosal, “Understanding International Prices: Customers as Capital,” *American Economic Review*, February 2012, 102 (1), 364–395.
- Eaton, Jonathan and Samuel Kortum, “Technology, Geography, and Trade,” *Econometrica*, 2002, 70 (5), 1741–1779.
- , David Jinkins, James R. Tybout, and Daniel Xu, “Two-Sided Search in International Markets,” NBER Working Papers 29684, National Bureau of Economic Research, Inc January 2022.
- , Marcela Eslava, David Jinkins, C. J. Krizan, and James Tybout, “A Search and Learning Model of Export Dynamics,” Working Papers 21-17, Center for Economic Studies, U.S. Census Bureau August 2021.
- , Samuel Kortum, and Francis Kramarz, “Firm-to-Firm Trade: Imports, Exports, and the Labor Market,” NBER Working Papers X, National bureau of economic researc 2022.
- Fitzgerald, Doireann and Stefanie Haller, “Pricing-to-Market: Evidence From Plant-Level Prices,” *Review of Economic Studies*, 2014, 81 (2), 761–786.
- and —, “Exporters and shocks,” *Journal of International Economics*, 2018, 113, 154–171.
- , —, and Yaniv Yedid-Levi, “How Exporters Grow,” NBER Working Papers 21935, National Bureau of Economic Research, Inc January 2016.
- , —, and —, “How Exporters Grow,” Working Papers 201702, School of Economics, University College Dublin January 2017.
- Fontagné, Lionel, Philippe Martin, and Gianluca Orefice, “The international elasticity puzzle is worse than you think,” *Journal of International Economics*, 2018, 115, 115–129.
- Fontaine, François, Julien Martin, and Isabelle Mejean, “Price discrimination within and across EMU markets: Evidence from French exporters,” *Journal of International Economics*, 2020, 124 (C).
- Garetto, Stefania, “Firms’ heterogeneity, incomplete information, and pass-through,” *Journal of International Economics*, 2016, 101 (C), 168–179.

- Gaubert, Cecile and Oleg Itskhoki, “Granular Comparative Advantage,” *Journal of Political Economy*, 2021, 129 (3), 871–939.
- Gourio, François and Leena Rudanko, “Customer Capital,” *Review of Economic Studies*, 2014, 81 (3), 1102–1136.
- Grossman, Gene and Elhanan Helpman, “When Tariffs Disrupt Global Supply Chains,” 2021.
- Imbs, Jean and Isabelle Mejean, “Elasticity Optimism,” *American Economic Journal: Macroeconomics*, July 2015, 7 (3), 43–83.
- Impullitti, Giammario, Alfonso A. Irarrazabal, and Luca David Opromolla, “A theory of entry into and exit from export markets,” *Journal of International Economics*, 2013, 90 (1), 75–90.
- Kikkawa, Ayumu Ken, Glenn Magerman, and Emmanuel Dhyne, “Imperfect competition in firm-to-firm trade,” Working Paper Research 363, National Bank of Belgium January 2019.
- Lenoir, Clemence, Julien Martin, and Isabelle Mejean, “Search Frictions in International Goods Markets,” *Journal of the European Economic Association*, 2022, *forthcoming*.
- Lim, Kevin, “Endogenous Production Networks and the Business Cycle,” 2018.
- Loecker, Jan De, Pinelopi K. Goldberg, Amit K. Khandelwal, and Nina Pavcnik, “Prices, Markups, and Trade Reform,” *Econometrica*, March 2016, 84, 445–510.
- Melitz, Marc J., “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, November 2003, 71 (6), 1695–1725.
- Miyauchi, Yuhei, “Matching and Agglomeration: Theory and Evidence from Japanese Firm-to-Firm Trade,” Boston University - Department of Economics - Working Papers Series WP2020-007, Boston University - Department of Economics December 2019.
- Nguyen, Daniel X., “Demand uncertainty: Exporting delays and exporting failures,” *Journal of International Economics*, 2012, 86 (2), 336–344.
- Piveteau, Paul, “An Empirical Dynamic Model of Trade with Consumer Accumulation,” *American Economic Journal: Microeconomics*, 2020, *forthcoming*.
- Postel-Vinay, Fabien and Jean-Marc Robin, “Equilibrium Wage Dispersion with Worker and Employer Heterogeneity,” *Econometrica*, November 2002, 70 (6), 2295–2350.
- Rauch, James E., “Networks versus markets in international trade,” *Journal of International Economics*, June 1999, 48 (1), 7–35.
- , “Business and Social Networks in International Trade,” *Journal of Economic Literature*, December 2001, 39 (4), 1177–1203.
- Ridder, Geert and Gerard J. van den Berg, “Measuring Labor Market Frictions: A Cross-Country Comparison,” *Journal of the European Economic Association*, March 2003, 1 (1), 224–244.

Roberts, Mark J and James R Tybout, "The Decision to Export in Colombia: An Empirical Model of Entry with Sunk Costs," *American Economic Review*, September 1997, 87 (4), 545-564.

A Data appendix

A.1 Construction of the estimated sample

Wholesalers: Whereas the raw data cover each single transaction involving French exporters and their partners in the European Union, the model tackles the choice of input suppliers in frictional markets in a context with no trade intermediation. Ideally, we would thus be willing to exclude from the sample all intermediated transactions. For French firms, we can use information from INSEE about the firm’s sector of activity and remove all firms that are either wholesalers or retailers. As noted by [Bernard et al. \(2015\)](#), intermediaries are important traders in international markets. In our data, French intermediaries represent 40% of the population of exporters and 15% of the total value of exports. Unfortunately, we do not know the activity of the importing firm. As a proxy for wholesaling activities, we thus measure the maximum number of French sellers a particular foreign firm interacts with for a given product, over a particular month. Our argument is that firms purchasing the same product to many different French exporters are more likely to intermediate trade than firms which purchase a particular good to a single French exporter. In our data, only 5% of importers ever purchase the same product from two different French exporters in a particular month but some importers simultaneously interact with more than 50 producers of specific accessories of motor vehicles or Bordeaux wine. Despite their small number, the combined share of overall trade intermediated by these multi-seller importers is high, at 23%, which again is consistent with evidence in [Bernard et al. \(2015\)](#). We thus exclude from the estimation sample the one percent of importers that display the maximum number of simultaneous suppliers within the same month. This excludes all firms that ever interacted with more than 3 French exporters in the same month. The remaining sample covers 75% of the total value of trade in the raw data and 4.7 million importer×product pairs.

The dimensionality of the recovered estimation sample is summarized in [Table A1](#).

A.2 Statistics on the connectivity of the graph

A now standard way of measuring the connectivity in such seller-buyer networks consists in measuring the in- and out-degrees at each node, i.e. the number of partners firms at each side of the network are connected to. [Figures A1](#) and [A2](#) illustrate the heterogeneity in this measure of connectivity, among European importers and French exporters, respectively.

Focusing first on importers, [Figure A1](#) illustrates the strong sparsity of this side of the network as the vast majority of European importers are connected with a single French exporter. As shown in the upper-left panel, more than 95% of the European importers that ever interact with a French exporter over the 2002-2006 period never interact with more than one firm within a particular month and over a particular product.³⁴ This number decreases somewhat, to 75%, when we do not condition over a particular product (upper-right panel), which means

³⁴As explained in [section A.1](#), this number is somewhat inflated artificially since we dropped firms purchasing the same product to many different exporters on the ground of the argument that these are more likely to be wholesalers. Remember however that the selection is based on the top 1% of firms with the highest indegree and does not change this figure much as a consequence.

Table A1: Dimensionality of the estimation sample

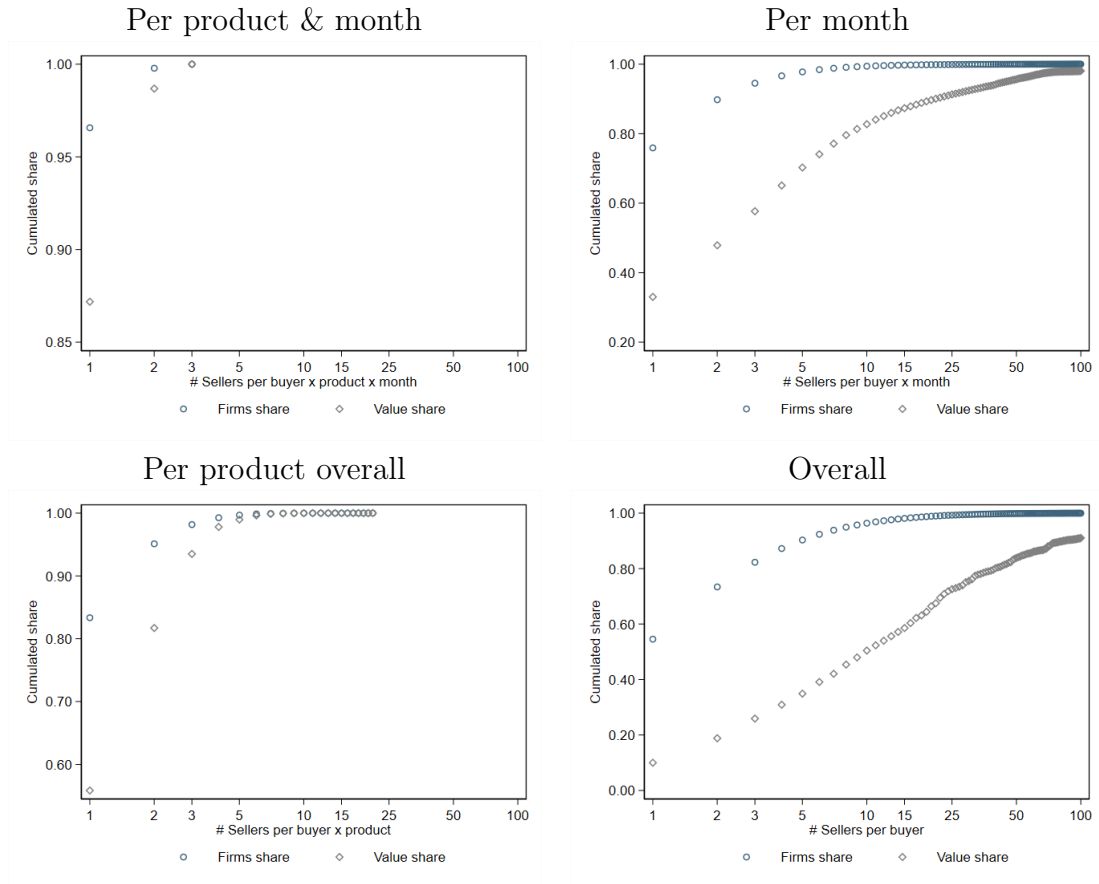
	Transactions (1)	Exporters s (2)	Importers $b(i)$ (3)	$sb(i)j$ Triplets (4)
All	27,442,785	39,751	744,118	5,646,587
Austria	787,990	9,669	20,765	157,550
Belgium	4,501,923	27,786	86,174	927,695
Denmark	577,165	9,478	14,326	116,695
Finland	357,670	6,261	7,718	69,181
Germany	5,731,010	24,683	181,630	1,122,918
Greece	634,143	8,415	14,950	136,556
Ireland	426,605	7,221	9,207	104,659
Italy	3,613,227	20,395	129,124	812,073
Luxembourg	479,248	10,922	8,047	97,417
Netherlands	1,869,157	17,344	46,071	375,632
Portugal	1,165,765	12,625	26,545	259,340
Spain	3,639,465	21,362	104,745	732,013
Sweden	637,453	8,975	15,298	121,086
United Kingdom	3,021,964	19,885	79,518	613,772

Notes: This table shows statistics on the dimensionality of the estimation dataset. The dataset covers the period from 2002 to 2006 and the EU15 countries. Trade intermediaries are neglected on the sellers' and buyers' sides, as explained in Section A.1.

that a non-negligible number of European importers interact with several French exporters simultaneously, to purchase different products. Whereas firms connected with multiple partners are relatively rare in the cross-section of the data, their number naturally increases when we cumulate their partners over time as in the bottom panels of Figure A1. Then, the share of firms that we never see interacting with two different exporters over the same product is reduced to 83%. This result is particularly important for the purpose of our exercise as the estimation exploits moments on firms that switch from one supplier to the other, after accumulating contacts over time. The shift of the distributions between the upper and the bottom panels of Figure A1 indicates that such switches are not uncommon.

Whereas importers interacting simultaneously with several exporters are rare, the reciprocal is not true, as illustrated in Figure A2. The upper left panel thus shows the cumulated distribution of sellers that interact with a given number of importers from a particular destination over a given month and for a particular product, as well as their contribution to aggregate trade. 70% of the sample is composed of French exporters that interact with a single firm in their typical destination at a point in time. When we cumulate across destinations as in the middle left panel, there are still 40% of exporters that serve a single importer in a single destination. These firms are however small, on average, and cumulate only 17% of the overall value of trade. At the other extreme of the distribution about 10% of exporters interact with more than 10 European importers but they cumulate almost 50% of French exports. The heterogeneity in exporters' ability to serve a large number of foreign partners is explained in the model by the interaction of exporters' productivity heterogeneity and the history of their matches with foreign firms. The deterministic dimension can explain why the distribution of these outdegrees is not fundamentally different when we cumulate French exporters' partners over time as in the bottom left panel. Whereas cumulating over time significantly shifted the distribution down

Figure A1: Cumulated distribution of European importers' indegrees



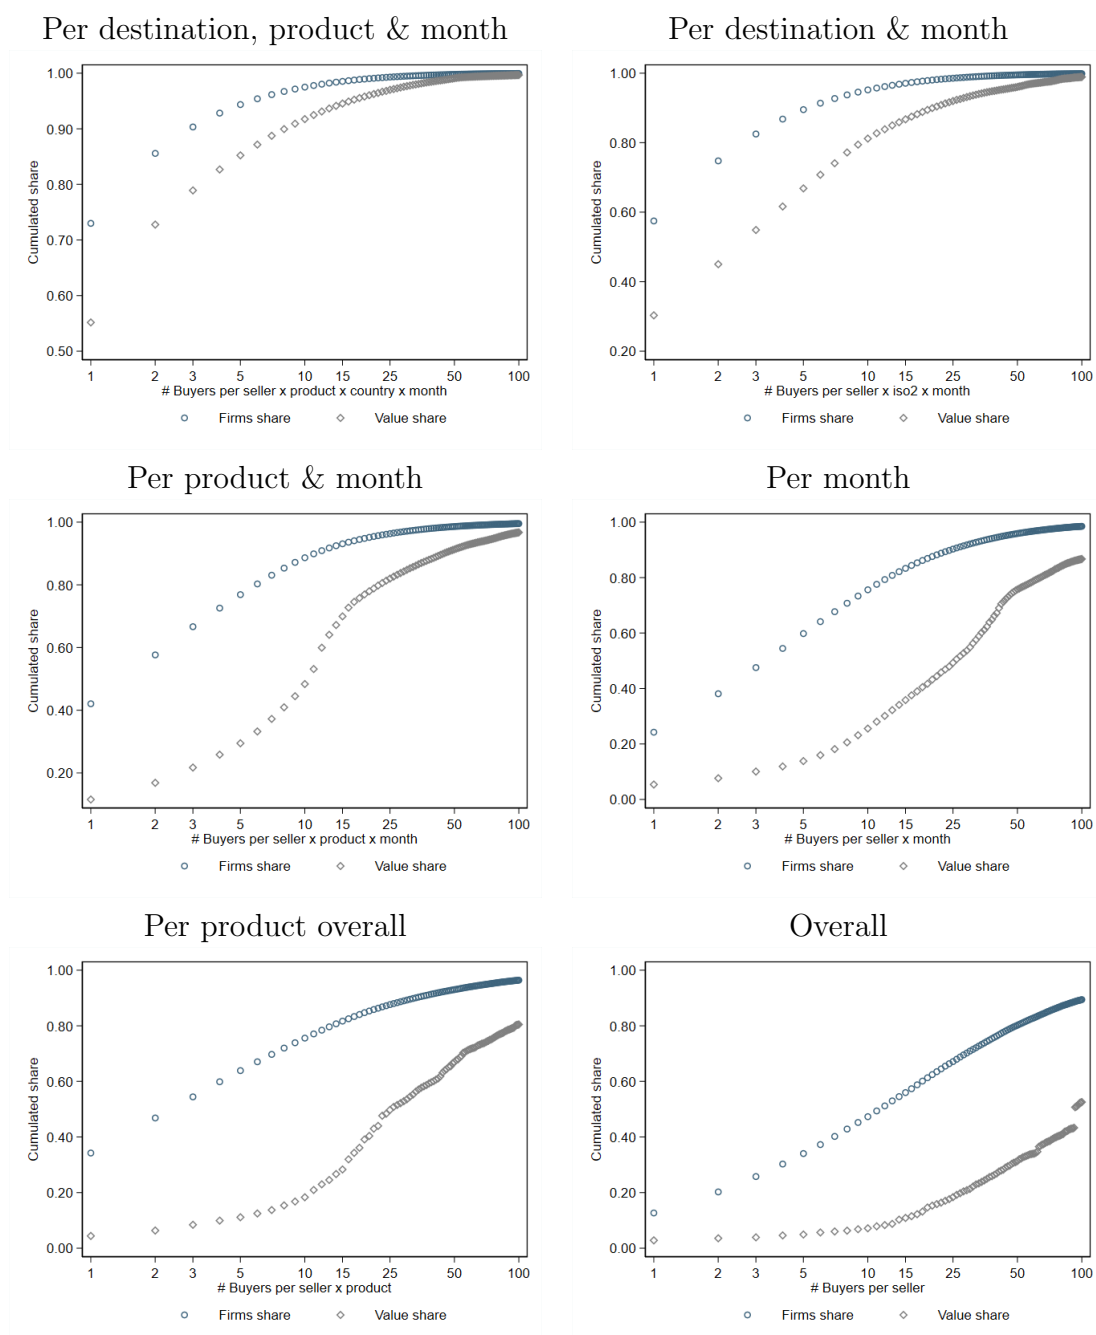
The figure illustrates the heterogeneity across European importers' in their “indegrees”, i.e. their number of partners in France. The x-axis corresponds to a number of partners and the y-axis is the cumulated share of firms (blue circles) with x sellers or less, and their cumulated contribution to aggregate trade (gray squares). The two upper panels measure a firm's indegree in the cross-section, i.e. within a particular month. The bottom panels instead cumulate partners over the whole period of activity of the firm. The left panels treat multi-product importers as independent units whereas the right panels cumulate partners over the firm's portfolio of imported products.

when we were focusing on importers in Figure A1, the same is not true when we take the point of view of exporters. Here as well, cumulating partners over the exporter's portfolio of products as we do in the right panels of Figure A2 shifts the distributions down. The reason is that the vast majority of exporters do not serve the same importers with their different products.

A.3 Trade elasticity estimates

Trade data The monthly series of bilateral trade at product-level are directly recovered from our main dataset. Namely, we aggregate all transactions at the level of a particular month, destination country, and for a particular month, over a period from January 1999 to December 2019. We obtain an unbalanced panel of 27 countries and 252 months, which we then merge with control variables recovered from external sources.

Figure A2: Cumulated distribution of French exporters' outdegrees



The figure illustrates the heterogeneity across French exporters' in their “outdegrees”, i.e. their number of partners within the European Union. The x-axis corresponds to a number of partners and the y-axis is the cumulated share of firms (blue circles) with x buyers or less, and their cumulated contribution to aggregate trade (grey squares). The two upper panels measure a firm’s outdegree in the cross-section, i.e. within a particular month. The bottom panels instead cumulate partners over the whole period of activity of the firm. The left panels treat multi-product firms as independent units whereas the right panels cumulate partners over the firm’s portfolio of products.

Nominal Exchange Rate The bilateral nominal exchange rates are sourced from Eurostat (series *ert_bil_eur_m* and *ert_bil_conv_m* for pre-Euro exchange rates among euro area countries) and correspond to the average monthly value.

Producer Price Index We obtain producer price indices (PPI) from Eurostat (series *sts_inpp-m*). We use the index that covers the largest set of sectors, namely industries B to E in the NACE classification.³⁵ Indices are seasonally adjusted, and normalized at 100 in January 2015.

Real Exchange Rate The real exchange rate is computed using the previous variables, as:

$$RER_{jt} = \frac{NER_{jt} \cdot PPI_{Ft}}{PPI_{jt}}$$

where $j = F$ stands for France.

Variations in the real exchange rate for euro area countries hence solely comes from variations in relative PPIs. For the non-EMU countries, nominal exchange rate movements add a lot more volatility. Summary statistics about the volatility of the real exchange rate in each country can be found in table A2, which shows the mean and median of the rolling 12-month standard deviation of the variable.

Table A2: 12 month volatility of the real exchange rate

Country	Mean	Median
Austria	0.0057	0.0049
Belgium	0.0138	0.0142
Germany	0.0065	0.0054
Denmark	0.1193	0.0867
Spain	0.0066	0.0070
Finland	0.0087	0.0071
United Kingdom	0.0191	0.0169
Greece	0.0151	0.0133
Ireland	0.0192	0.0164
Italy	0.0048	0.0047
Luxembourg	0.0155	0.0130
Netherlands	0.0146	0.0120
Portugal	0.0136	0.0093
Sweden	0.1567	0.1243

Import demand We obtain countries' import data from Eurostat (series *DS-1060915*). More precisely, we compute monthly total imports originating from countries outside the EU-28, for all countries in the sample.

³⁵B-Mining and quarrying, C-Manufacturing, D-Electricity, gas, steam and air conditioning supply, E-Water supply; sewerage, waste management and remediation activities

B Details on the estimation procedure

As explained in the main text, our estimation of the model's parameters uses a simulated likelihood approach. For given values of the structural parameters, we simulate our model, compute the needed frequencies, and compare them with their empirical counterpart. For that purpose, we need to limit the set of possible events. In the following, we record up to 5 transactions, remembering that you need at least one transaction to be part of the sample, and up to 2 switches. Then, we compute on simulated data the probabilities needed for the likelihood, namely

$$P(\text{ transactions } = n \cap \text{ switches } = s)$$

where $n \in \{1, \dots, 5\}$ ($n = 5$ means at least 5 transactions) and $s \in \{0, 1, 2\}$ ($s = 2$ means at least 2 switches).

Now we have defined the relevant set of events, the exact procedure to compute the likelihood given values of the structural parameters is as follow

- i. For each dataset (country×market), we simulate 100 times more buyers than there are in the dataset. If a buyer exits the market in our simulations (μ shock), it is replaced by an unmatched buyer. This ensures that the steady state assumption holds. Note that some of these simulated buyers will not be used to compute the frequencies. This is the case when they are never matched with a French seller.
- ii. We first simulate 2000 months of buyers' history to reach steady state. After this step, we sample according to the way the estimation sample is generated. Hence we simulate for 24 months and we record any buyer observed making a transaction with a French seller, as we record any buyer between January 2002 and January 2004 in the data. Then, we follow that buyer for 24 months recording any subsequent transaction or any switch. In the same way, the estimation sample follows any recorded buyer for up to 24 months, that is up to January 2006 at the latest.
- iii. Using simulated data, we compute the frequencies $P(\text{ transactions } = n \cap \text{ switches } = s)$, $\forall(n, s)$. We denote these frequencies $P^{sim}(n, s|\omega)$ where ω is the vector of parameters' value.
- iv. If there are J buyers in real data, indexed by j , the log-likelihood is as follow

$$\mathcal{L}(\omega) = \sum_{j=1}^J P^{sim}(n_j, s_j|\omega)$$

Finally, our estimates are obtained by maximizing the likelihood:

$$\hat{\omega} = \arg \max \mathcal{L}(\omega)$$

C The model: additional derivations

C.1 The distributions of buyers across sellers

The distribution of buyers among all suppliers. The overall quality-adjusted serving cost distribution is a mixture of the country-specific ones, noted $F_i(c)$, with

$$F_i(c) = \frac{\gamma_{iF}}{\gamma_i} F_{iF}(c) + \frac{\gamma_{i\bar{F}}}{\gamma_i} F_{i\bar{F}}(c) \quad (15)$$

for all $c \in]0, \max\left(\frac{v_F d_{iF}}{z}, \frac{v_{\bar{F}} d_{i\bar{F}}}{z}\right)]$. Notice that the two distributions (F and \bar{F}) are defined for all c with, for example, $F_{iF}(c) = 1$ and $f_{iF}(c) = 0$ for $c > v_F d_{iF}/z$. This ensures that $F_i(c)$ is continuous and properly defined on the whole support.

Let us now derive the distribution of buyers across sellers irrespective of the origin of the seller. That distribution is noted $L_i(c)$ and the corresponding flow equation is simply

$$\underbrace{B_i(1 - u_i)\ell_i(c) (\mu + \gamma_i F_i(c))}_{\text{outflows}} = \underbrace{B_i(1 - u_i)\gamma_i \bar{L}_i(c) f_i(c) + B u_i \gamma_i f_i(c)}_{\text{inflows}}$$

with $\bar{L}_i(c) \equiv 1 - L_i(c)$. The outflows are the sum of buyers exiting the market (μ) and buyers switching when they meet with a lower quality-adjusted cost supplier ($\gamma_i F_i(c)$). The inflows correspond to unmatched buyers meeting a cost- c supplier ($\gamma_i f_i(c)$) and buyers previously matched with sellers of serving cost higher than c ($\gamma_i \bar{L}_i(c) f_i(c)$). Using

$$u_i = \frac{\mu}{\gamma_i + \mu}$$

and simplifying one gets

$$\ell_i(c) (\mu + \gamma_i F_i(c)) = \gamma_i \bar{L}_i(c) f_i(c) + \mu f_i(c)$$

Then integrating by part

$$L_i(c) = \frac{\mu + \gamma_i}{\mu + \gamma_i F_i(c)} F_i(c) \quad (16)$$

The distribution of buyers among French suppliers. Consider the shares of buyers matched with a French seller, π_{iF} . Again, in equilibrium, flows in and out are balanced such that the density of buyers matched with a French-seller at cost c , noted $\ell_{iF}(c)$, satisfies

$$\underbrace{(1 - u_i)\pi_{iF}\ell_{iF}(c) (\mu + \gamma_i F_i(c))}_{\text{outflows}} = \underbrace{u_i \gamma_{iF} f_{iF}(c) + (1 - u_i)\bar{L}_i(c)\gamma_{iF} f_{iF}(c)}_{\text{inflows}} \quad (17)$$

Substituting $\bar{L}_i(c) = 1 - L_i(c)$ by its expression in equation and using $u_i = \mu/(\mu + \gamma_i)$, one gets

$$\pi_{iF} \ell_{iF}(c) = \frac{\gamma_{iF}}{\gamma_i} \frac{\mu(\mu + \gamma_i)}{(\mu + \gamma_i F_i(c))^2} f_{iF}(c) \quad (18)$$

and similarly if we consider the density of buyers matched with non-French sellers

$$(1 - \pi_{iF}) \ell_{i\bar{F}}(c) = \frac{\gamma_{i\bar{F}}}{\gamma_i} \frac{\mu(\mu + \gamma_i)}{(\mu + \gamma_i F_i(c))^2} f_{i\bar{F}}(c) \quad (19)$$

The trade shares when $v_F d_{iF} < v_{\bar{F}} d_{i\bar{F}}$. We denote $c_F^{max} = v_F d_{iF} / \underline{z}$ the highest serving costs among French suppliers. Consider matches with French sellers, the flows in and out satisfy

$$\begin{aligned} (1 - u_i) \pi_{iF} \left(\mu + \gamma_{i\bar{F}} \int_0^{c_F^{max}} F_{i\bar{F}}(c) \ell_{iF}(c) dc \right) &= u_i \gamma_{iF} \\ + (1 - u_i) (1 - \pi_{iF}) \gamma_{iF} \left(\bar{L}_{i\bar{F}}(c_F^{max}) + \int_0^{c_F^{max}} F_{iF}(c) \ell_{i\bar{F}}(c) dc \right) & \end{aligned} \quad (20)$$

Note that, whenever $F_{iF}(c)$ and $F_{i\bar{F}}(c)$ have common support (that is up to c_F^{max}), we have $f_{iF}(c)/f_{i\bar{F}}(c) = \tau_{iF}^{-\theta} \forall c$. Combining equations (18) and (19),

$$(1 - \pi_{iF}) \tau_{iF}^{-\theta} \gamma_{iF} \ell_{i\bar{F}}(c) = \pi_{iF} \gamma_{i\bar{F}} \ell_{iF}(c) \quad (21)$$

Integrating up to c_F^{max} and simplifying, one gets

$$L_{i\bar{F}}(c_F^{max}) = \frac{\pi_{iF}}{(1 - \pi_{iF})} \frac{\gamma_{i\bar{F}}}{\gamma_{iF}} \tau_{iF}^{-\theta} \quad (22)$$

which can be substituted in (20) to obtain

$$\begin{aligned} (1 - u_i) \pi_{iF} \left(\mu + \gamma_{i\bar{F}} \int_0^{c_F^{max}} F_{i\bar{F}}(c) \ell_{iF}(c) dc \right) &= u_i \gamma_{iF} + (1 - u_i) (\gamma_{iF} - \pi_{iF} (\gamma_{iF} + \gamma_{i\bar{F}} \tau_{iF}^{-\theta})) \\ + (1 - u_i) (1 - \pi_{iF}) \gamma_{iF} \int_0^{c_F^{max}} F_{iF}(c) \frac{\pi_{iF}}{1 - \pi_{iF}} \frac{\gamma_{i\bar{F}}}{\gamma_{iF}} \tau_{iF}^{-\theta} \ell_{iF}(c) dc & \end{aligned}$$

The integrals cancel out and, after simplification, one gets

$$\pi_{iF} = \frac{\gamma_{iF}}{\gamma_{iF} + \gamma_{i\bar{F}}} \frac{\mu + \gamma_{iF} + \gamma_{i\bar{F}}}{\mu + \gamma_{iF} + \gamma_{i\bar{F}} \tau_{iF}^{-\theta}} \quad (23)$$

The trade shares when $v_F d_{iF} > v_{\bar{F}} d_{i\bar{F}}$. We can derive in a similar manner the trade share, denoting $c_{\bar{F}}^{max} = v_{\bar{F}} d_{i\bar{F}} / \underline{z}$. We start with

$$\begin{aligned}
& (1 - u_i)\pi_{iF} \left(\mu + \gamma_{i\bar{F}} \bar{L}_{iF}(c_F^{max}) + \gamma_{i\bar{F}} \int_0^{c_F^{max}} F_{i\bar{F}}(c) \ell_{iF}(c) dc \right) = u_i \gamma_{iF} \\
& + (1 - u_i)(1 - \pi_{iF}) \gamma_{iF} \int_0^{c_F^{max}} F_{iF}(c) \ell_{i\bar{F}}(c) dc
\end{aligned} \tag{24}$$

Since $(1 - \pi_{iF})\gamma_{iF}\tau_{iF}^{-\theta}\ell_{i\bar{F}}(c) = \pi_{iF}\gamma_{i\bar{F}}\ell_{iF}(c)$, the integrals cancel out

$$(1 - u_i)\pi_{iF} \left(\mu + \gamma_{i\bar{F}} \bar{L}_{iF}(c_F^{max}) \right) = u_i \gamma_{iF} \tag{25}$$

We get an expression for $\bar{L}_{iF}(c_F^{max})$ by integrating (21) up to c_F^{max} ,

$$L_{iF}(c_F^{max}) = \frac{1 - \pi_{iF}}{\pi_{iF}} \frac{\gamma_{iF}}{\gamma_{i\bar{F}}} \tau_{iF}^{-\theta} \tag{26}$$

Finally, using that expression, we obtain

$$\pi_{iF} = \frac{\gamma_{iF}}{\gamma_{iF} + \gamma_{i\bar{F}}} \frac{\mu + (\gamma_{iF} + \gamma_{i\bar{F}})\tau_{iF}^{-\theta}}{\mu + \gamma_{iF}\tau_{iF}^{-\theta} + \gamma_{i\bar{F}}} \tag{27}$$

The share of French sellers when $\mu \approx 0$. Interestingly, when μ is close to zero, (23) and (27) are approximately equal

$$\pi_{iF} = \frac{\gamma_{iF}}{\gamma_{iF} + \gamma_{i\bar{F}}\tau_{iF}^{-\theta}} \tag{28}$$

The trade shares. π_{iF} is the share of French sellers among the providers but it is also the trade share when we assume $\underline{z} \rightarrow 0$ as in Eaton and Kortum (2002). To demonstrate the equivalence, first notice that the demand of input j by a buyer reads

$$p_j x_j = \alpha_b \left(\frac{p_j}{q_j} \right)^{1-\eta} \left(\sum_{s=1}^{M_b} \left(\frac{p_s}{q_s} \right)^{1-\eta} \right)^{\frac{\eta}{1-\eta}} \tag{29}$$

where, given our price setting mechanism, p_j/q_j is the quality adjusted cost to serve of the second best supplier. Consider a buyer whose best supplier is French, the expected price-to-quality is

$$\mathbb{E}[p/q|F] = \int_0^c \int_0^{+\infty} \tilde{c} \ell(c) \ell_{iF}(\tilde{c}|c) d\tilde{c} dc \tag{30}$$

where $\ell_{iF}(\tilde{c}|c)$ denotes the pdf of the price distribution *conditional* on being matched with a French supplier c ($L_{iF}(\tilde{c}|c)$ the cdf).

Working with the complementary cumulative distribution, $\bar{L}_{iF}(\tilde{c}|c)$ we have in steady state

$$(1 - u_i)\pi_{iF}\bar{L}_{iF}(\tilde{c}|c)(\mu + \gamma_i F_i(\tilde{c})) = (u_i + (1 - u_i)\bar{L}_i(\tilde{c}))\gamma_{iF}f_{iF}(c) \quad (31)$$

and the equivalent cdf conditional on being match with a non-French supplier, $\bar{L}_{i\bar{F}}(\tilde{c}|c)$,

$$(1 - u_i)(1 - \pi_{iF})\bar{L}_{i\bar{F}}(\tilde{c}|c)(\mu + \gamma_i F_i(\tilde{c})) = (u_i + (1 - u_i)\bar{L}_i(\tilde{c}))\gamma_{i\bar{F}}f_{i\bar{F}}(c) \quad (32)$$

Remark that

$$\begin{aligned} \frac{\gamma_{iF}f_{iF}(c)}{\pi_{iF}} &= \frac{\gamma_{iF}f_{iF}(c)(\gamma_{iF} + \gamma_{i\bar{F}}\tau_{iF}^\theta)}{\gamma_{iF}} = \frac{\gamma_{i\bar{F}}\tau_{iF}^{-\theta}f_{i\bar{F}}(c)(\gamma_{iF} + \gamma_{i\bar{F}}\tau_{iF}^\theta)}{\gamma_{i\bar{F}}} \\ &= \frac{\gamma_{i\bar{F}}f_{i\bar{F}}(c)}{1 - \pi_{iF}} \end{aligned}$$

Hence $L_{iF}(\tilde{c}|c) = L_{i\bar{F}}(\tilde{c}|c) = L_i(\tilde{c}|c)$ and $\mathbb{E}[p/q|F] = \mathbb{E}[p/q|\bar{F}]$. For that reason, the expected quantity does not depend on the country of origin of the supplier and the trade shares of France simply follows the share of French sellers in the buyers' portfolio.

C.2 Buyer acquisition on the seller's side

The stylized facts in Section 2.2 characterize three things, how prices change as buyers move between sellers, how they change as the buyer-seller relationship last, and how sellers acquire buyers over time. We now focus on the acquisition of buyers.

Over time, within a product category, French sellers meet with buyers from i at rate λ_{iF} and consistency implies: $\gamma_{iF}B_i = \lambda_{iF}S_F$. Consider now a French seller that can serve market i at quality-adjusted cost c . Its number of buyers, noted $n_i(c)$, evolves as new profitable links are formed and old links are dissolved when a buyer exit (rate μ) or meets a seller with a lower quality adjusted serving cost. $n_i(c)$ dynamics thus follows

$$\dot{n}_{it}(c) = -n_{it}(c)(\mu + \gamma_i F_i(c)) + \lambda_{iF}((1 - u_i)\bar{L}_i(c) + u_i)$$

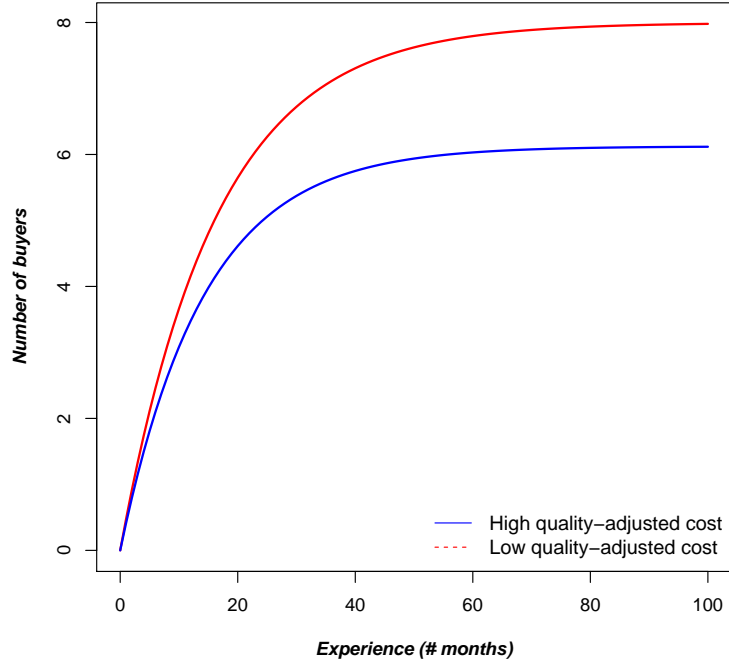
with $n_{i0} = 0$. Hence, given equation (7), the expected number of buyers in period t has a simple solution

$$n_{it}(c) = \frac{\lambda_{iF}\mu}{(\mu + \gamma_i F_i(c))^2} (1 - e^{-(\mu + \gamma_i F_i(c))t}) \quad (33)$$

The expected number of buyers is thus increasing over time as it converges toward a steady state³⁶

³⁶At the aggregate level, the number of buyers for each suppliers follows a stationary distribution with mean $n_i(c)$ and where the probability to have k buyers, noted $p(k, i, c)$, reads $p(k|i, c) = \frac{1}{k!}n_i(c)^k e^{-n_i(c)}$.

Figure A3: Expected number of buyers as a function of the seller’s experience, for two levels of quality-adjusted costs



Notes: The figure shows the simulated expected number of buyers in a seller’s portfolio as a function of its experience in the market. The blue (resp. red) line is for a high (resp. a low) quality-adjusted cost seller.

$$n_i(c) = \frac{\lambda_i F \mu}{(\mu + \gamma_i F_i(c))^2} \quad (34)$$

This implies that the number of buyers served is higher for sellers with lower serving-costs. Moreover, before reaching steady state, they grow at higher pace since they retain current buyers and find new buyers to serve more easily (see equation (C.2)). The relationship between a seller’s experience and the expected size of her portfolio of customers is illustrated in Figure A3 and can be compared with Figure 1 in the data. The heterogeneity in sellers’ number of buyers at the steady state helps explain the cross-sectional heterogeneity in sellers’ outdegrees discussed in Section A.2 and in the literature before us.

C.3 Plugging the model into a general equilibrium structure

The representative consumer There is one representative consumer in each country i . She consumes a bundle of differentiated goods and her utility function reads

$$U_i = \left(\sum_{b=1}^{B_i} x_b^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}$$

with $\sigma > 1$ the (constant) elasticity of substitution between varieties. For simplicity, final goods are assumed non-traded but it would be easy to plug trade in final goods in this set-up, following Melitz (2003). Noting R_i the consumer's income, which consists of labor income and residual profits, we obtain the demand for variety b

$$x_b = \frac{R_i}{P_i} \left(\frac{p_b}{P_i} \right)^{-\sigma} \quad (35)$$

where p_b stands for the price of variety b and P_i is the ideal price index

$$P_i = \left(\sum_{b=1}^{B_i} p_b^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$$

Remark that the representative consumer has no intertemporal choice. Within a certain time frame, she spends R_i to collect the different available varieties. The number of varieties depends on the number of differentiated good producers (our *buyers*) that are matched with at least one intermediate good producers (our *sellers*). The surplus generated by consumption depends on the level of final good prices, which itself is endogenous to the history of matches between final producers and their input suppliers.

The final good producer Given the level of demand x_b , the final good producer chooses the quantity of each input that minimizes costs

$$\left\{ \begin{array}{l} \min_{\{x_j\}} \sum_{j=1}^{M_b} p_j x_j \\ \text{s.t.} \\ x_b = \left(\sum_{j=1}^{M_{ib}} (q_j x_j)^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}} \end{array} \right. \quad (36)$$

where η is the elasticity of substitution between inputs, that can be below one in case of complementarities, and M_b is the (endogenous) number of inputs incorporated in production.

In optimum, the nominal demand for input j writes:

$$p_j x_j = x_b \left(\frac{p_j}{q_j} \right)^{1-\eta} \left(\sum_{j=1}^{M_b} \left(\frac{p_j}{q_j} \right)^{1-\eta} \right)^{\frac{\eta}{1-\eta}}$$

and the marginal cost:

$$mc_b = \left(\sum_{j=1}^{M_b} \left(\frac{p_j}{q_j} \right)^{1-\eta} \right)^{\frac{1}{1-\eta}}$$

Given the CES structure, it is optimal for the final good producer to choose the lowest quality-adjusted price for each input, among the set of input suppliers it has met. The love-for-variety property implies that M_b increases incrementally each type the final good producer meets with a supplier offering an input it does not already purchase.

The intermediate good producers Closing the model requires to further specify the production function of intermediate good producers. As explained in the text, intermediate good producers have a constant-return-to-scale technology that is linear in some input. Suppose that this input is labor, which is hired on the domestic market at a rate ν_F .

The model can be solved in general equilibrium using the labor market clearing condition:

$$L = \sum_{\tau=0}^{\infty} \int_{z=\underline{z}}^{\infty} l^d(z, \tau) dF(z)$$

where L is the aggregate supply of labor and $l^d(z, \tau)$ the demand of labor of a firm of quality-adjusted productivity z with experience τ . Given there is a measure of firms at each point of the joint productivity and experience distribution, we can use the expected demand of labor to substitute for $l^d(z, \tau)$:

$$l^d(z, \tau) = \sum_i n_{i\tau}(z) \mathbb{E} \left[x_b mc_b \left(\frac{p}{q} \right)^{-\eta} \mid z \right]$$

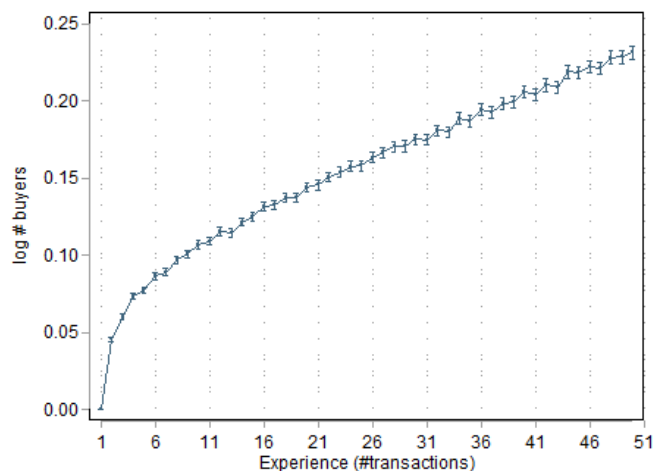
with $n_{i\tau}(z)$ defined in equation (33), x_b and mc_b which are functions of the expected distribution of quality-adjusted prices, across buyers in country i and the distribution of the quality-adjusted price derived in (30).

Solving the model in general equilibrium thus requires solving for the level of wage ν_F such that the labor market clears.³⁷

D Additional results

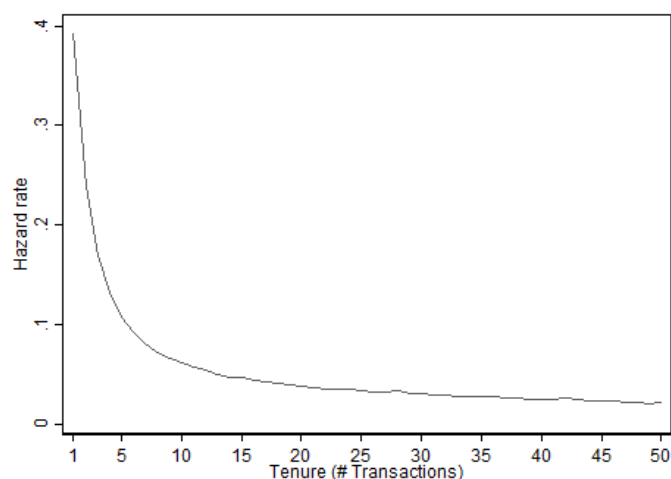
³⁷Here, we consider the case of a small open economy in which we solve for the wage rate in France, given foreign input prices.

Figure D1: Acquisition of buyers over time, Alternative definition of experience



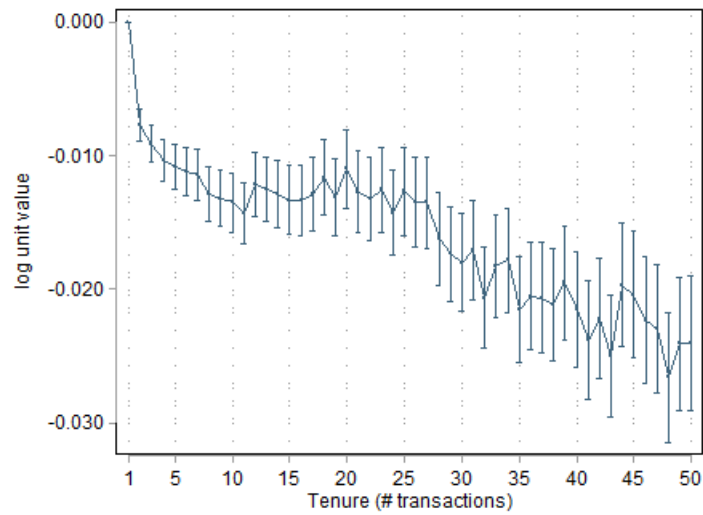
Note: The figure shows the evolution of a seller's stock of buyers, over time, recovered from equation (1). The figure reports the estimates and their 95% confidence intervals. Experience is measured by the cumulated number of transactions since first entry.

Figure D2: Hazard rate over time, Alternative tenure definition



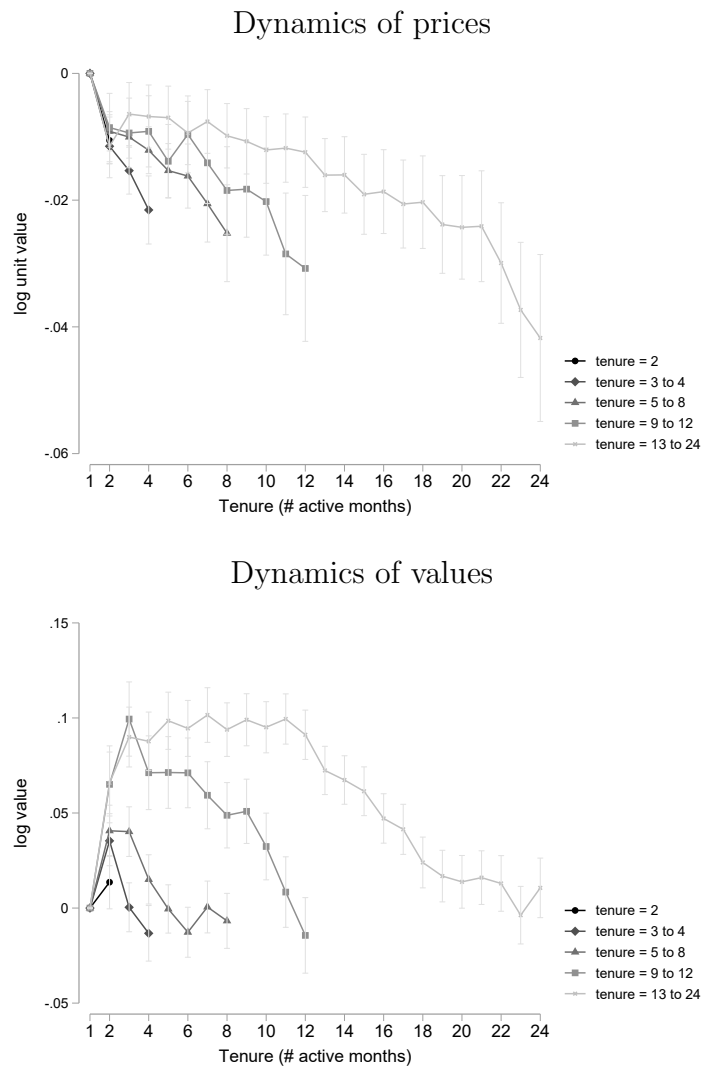
Notes: The hazard rate is defined as the probability of the relationship ending, conditional on tenure into the relationship and is calculated as the ratio of the density to the survival rate at tenure k . The figure is recovered from the 2002-2006 sample using the cumulated number of transactions since the start of the relationship as measure of tenure.

Figure D3: Price dynamics within a firm-to-firm relationship, Alternative tenure definition



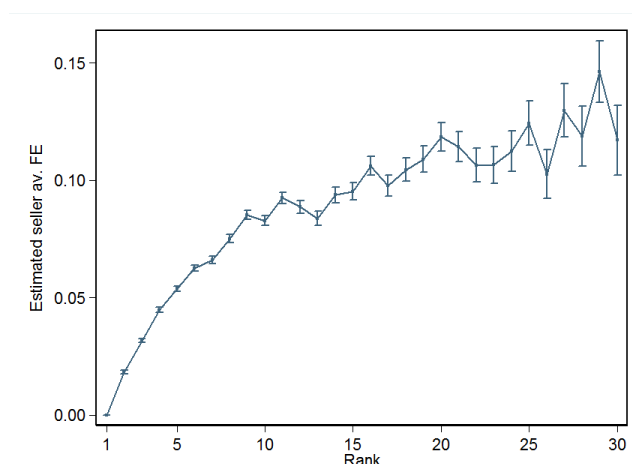
Note: This figure shows the evolution of prices within a firm-to-firm relationship. Coefficients are recovered from equation (2). The figure reports the estimates and their 95% confidence intervals. Tenure is measured by the cumulated number of transactions since the beginning of the relationship.

Figure D4: Price and quantity dynamics within a firm-to-firm relationship, by tenure



Note: This figure shows the evolution of prices and exported values within a firm-to-firm relationship. Coefficients are recovered from equation (2) which is estimated separately by subset of tenure lengths. The figure reports the estimates and their 95% confidence periods. Tenure is measured by the number of months since the beginning of the relationship.

Figure D5: Evolution of sellers' attributes, within a buyer's sequence of French partners

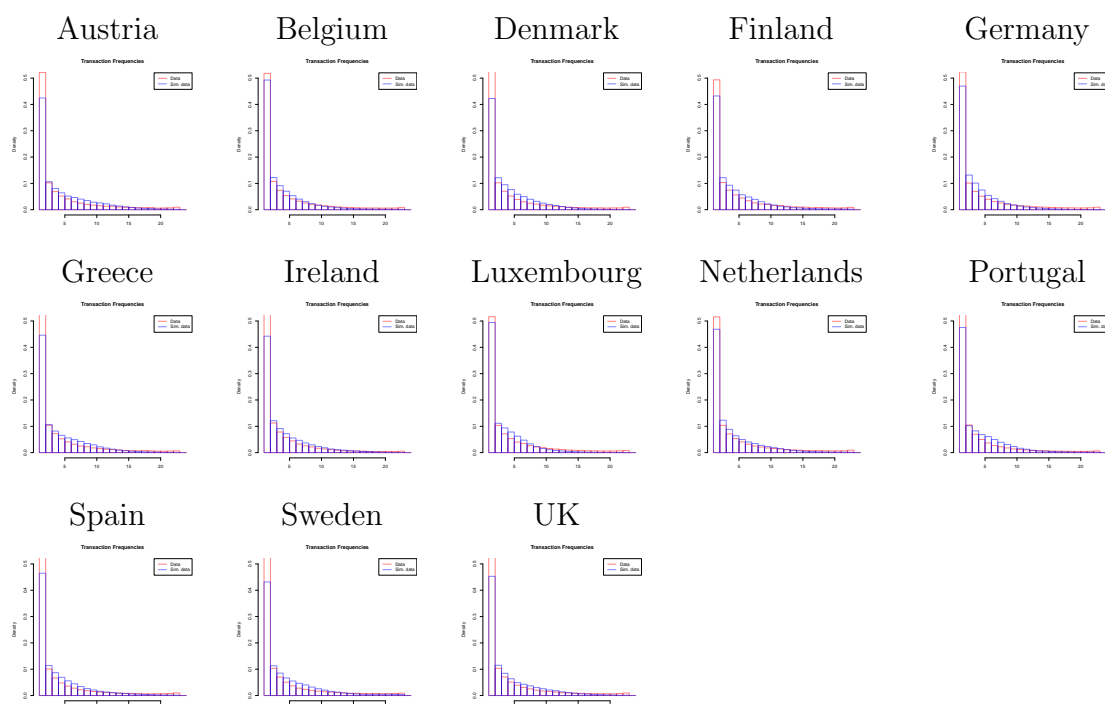


Note: This figure shows how buyers climb along the distribution of sellers' attributes, over time. It is recovered from the estimation of:

$$\hat{F}E_{spi} = FE_{bp} + \sum_{l=2}^K \alpha_l \mathbb{1}(Partner_{bps} = l) + \varepsilon_{bps}$$

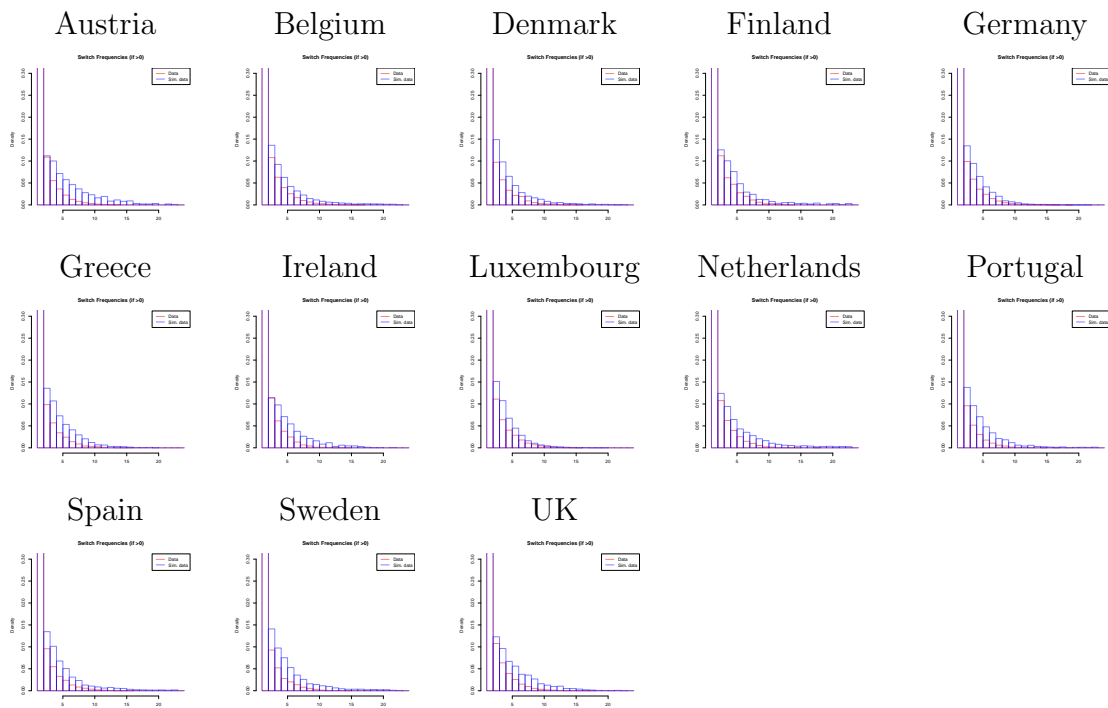
where $\hat{F}E_{spi}$ is the fixed effect recovered from the estimation of equation (1), that we interpret as a proxy for the seller's quality-adjusted price in market i , FE_{bp} is a buyer-product fixed effect and $\mathbb{1}(Partner_{bps} = l)$ is a dummy if seller s is the l th partner of buyer b when sellers are ranked sequentially based on their history of transactions with buyer b . In this equation α_l measures how $\hat{F}E_{spi}$ improves when a buyer switches from its $l-1$ th to its l th French supplier.

Figure D6: Goodness of fit: Transaction frequencies



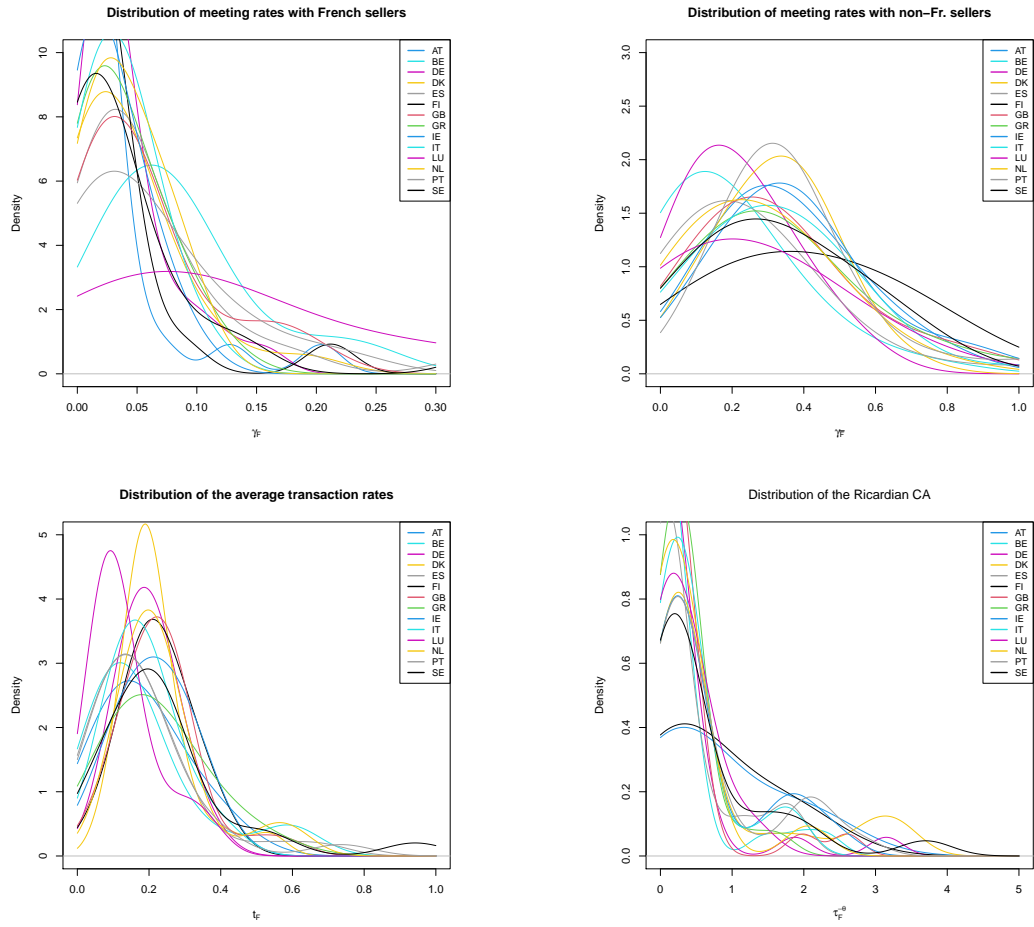
Note: Figure shows how we fit the transaction frequencies, by comparing the actual and simulated data.

Figure D7: Goodness of fit: Switch frequencies



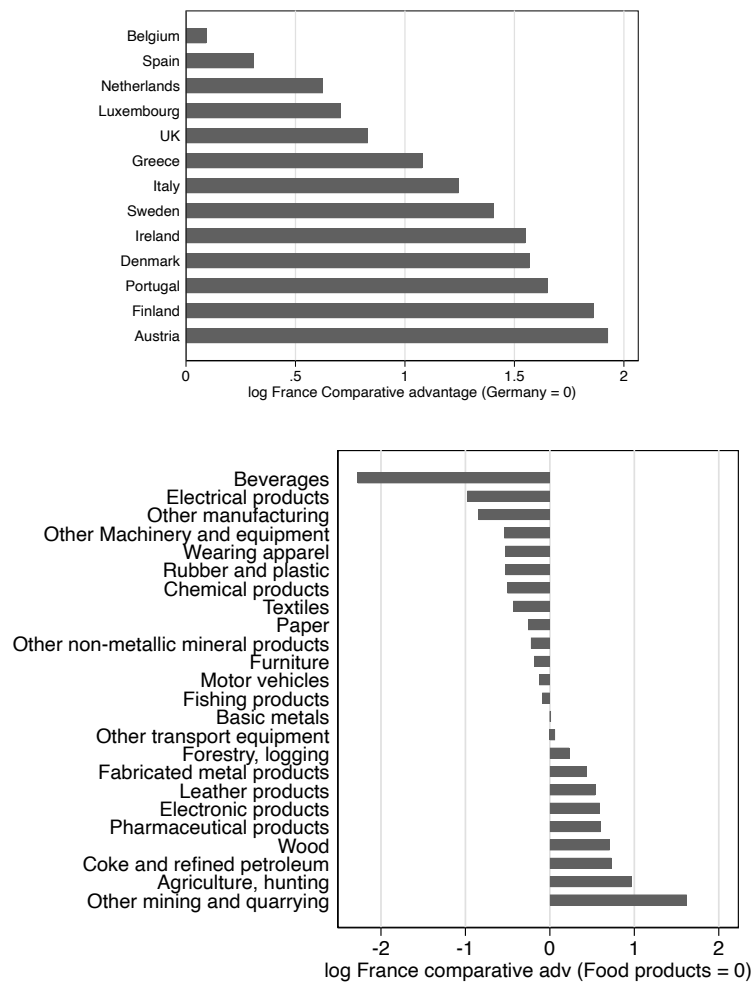
Note: Figure shows how we fit the switch frequencies, by comparing the actual and simulated data.

Figure D8: Densities for the estimated parameters



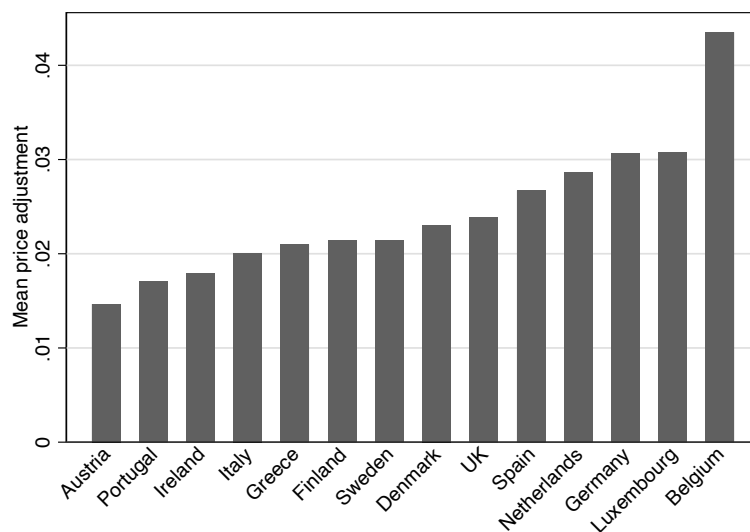
Note: The figure shows the country-specific distributions of estimated parameters. The figure is restricted to sector×country pairs for which we observe at least 100 buyers and 2 switches.

Figure D9: Dispersion in estimated Ricardian comparative advantages, across countries and sectors



Note: The figure shows the mean value of estimated Ricardian comparative advantages, by country and sector.

Figure D10: Average instantaneous price adjustment, by importing country



Note: The figure shows the mean adjustment of prices paid by buyers located in various countries, at the period of the 10% shock on French firms' production costs .