Frictions and adjustments in firm-to-firm trade*

François Fontaine† Julien Martin‡ Isabelle Mejean§

November 4, 2024

Abstract

We develop and structurally estimate a quantitative model of firm-to-firm trade in which the interaction of search frictions and comparative advantages shapes the structure of trade networks, generates heterogeneous prices and markups across sellers and within a seller across its downstream partners, and determines the magnitude of the international transmission of shocks. We recover the parameters of the model using a structural strategy that exploits firms' mobility along the network of French exporters, observed in firm-to-firm trade data. The estimated model is used to examine the incidence of domestic cost shocks on foreign partners, in the cross-section and over time.

1 Introduction

The incidence of cost shocks—arising from factors like exchange rates, productivity changes, or tariffs—is a central focus in international economics. Understanding the international transmission of these shocks implies digging into the determinants of prices within complex networks of firm-to-firm trade relationships. In recent years, significant progress has been made in describing and modeling the endogenous structure of these networks.[1] Most of this literature, however, imposes stylized pricing rules that fail to account for the large dispersion of prices and their dynamic over time, observed in firm-to-firm data.[2]

*The paper has benefited from insightful comments from participants at various seminars and conferences. We are grateful to Robin Lenoir and Pierre Coster for expert research assistance as well as Adrien Bilal, Jonathan Dingel and Daniel Xu for insightful discussions. Mejean gratefully acknowledges support from a public grant overseen by the French National Research Agency (ANR) as part of the "Investissements d'Avenir" program (Idex Grant Agreement No. ANR-11-IDEX-0003- 02/Labex ECODEC No. ANR-11-LABEX-0047 and Equipex reference: ANR-10-EQPX-17 – Centre d'accès sécurisé aux données – CASD) as well as the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 714597).

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

‡Université du Québec à Montréal and CEPR, Email address: martin.julien@uqam.ca

§Sciences Po, Department of Economics, CNRS, Paris, France. Email address: isabelle.mejean@sciencespo.fr.

[1] See among many others Bernard et al. (2019), Lim (2018), Huneeus (2018), Miyauchi (forthcoming), Bernard et al. (2022), Dhyne et al. (2023), Arkolakis et al. (2023), Demir et al. (2024), Eaton et al. (forthcoming).

[2] Fontaine et al. (2020) and Alviarez et al. (2021), and Burstein et al. (2024) dig into the dispersion of prices, in international and domestic firm-to-firm data respectively. They show that the dispersion of prices is not only

In this paper, we develop and structurally estimate a quantitative model of firm-to-firm trade in which the interaction of search frictions and comparative advantages (*i*) shapes the structure of trade networks, (*ii*) generates heterogeneous prices and markups across sellers and within a seller across its downstream partners, and (*iii*) determines the magnitude of the international transmission of shocks. We recover the parameters of the models using a structural strategy that exploits the mobility of firms along the network of French exporters, observed in exhaustive firm-to-firm trade data. The estimated model is used to examine the incidence of domestic cost shocks on foreign partners, in the cross-section and over time.

A first contribution to the literature is the introduction of price dynamics in a Ricardian model of firm-to-firm trade *à la* Eaton et al. (forthcoming). We model the equilibrium set of relationships linking final good producers and their input providers in an environment characterized by Ricardian comparative advantages and search frictions. Because of the frictions, final good producers meet with a random set of potential suppliers that offer perfect substitutes of the same input.[3] As a consequence, the degree of competition varies across buyers within the same input market. Under Bertrand competition, the markup set by a firm on its downstream partner depends on the second-best offer which the firm has met. These ingredients are thus sufficient to generate a distribution of markups that varies within and across firms.

Our approach differs from existing models that dig into the distribution of markups in imperfectly competitive firm-to-firm trade models (Alviarez et al., 2021, Dhyne et al., 2022). In these papers, markups are determined by the relative bargaining power of the firms involved into the transaction, a function of their size. In our model, seller and buyer characteristics affect prices, but the randomness associated with search frictions induces additional heterogeneity that cannot be explained by observed characteristics. This is consistent with Burstein et al. (2024) who document that a sizable share of the variance of markups cannot be explained by the observed size of firms involved into the relationship or other observable characteristics.

Our modeling strategy further sheds light on the dynamics of prices within and across relationships. Over time, buyers meet with new potential suppliers, which improves their outside option. Within existing relationships, this generates a downward trend in prices, driven by the supplier reducing its markup to retain its partner. Across relationships, new matches can generate endogenous switches, away from now less competitive input suppliers. Both predictions are borne out by the data. We show that prices tend to decrease with the age of the relationship.[4] Across relationships, we provide evidence that when foreign buyers switch

high across sellers within a market but also across the different downstream partners of the same firm. This suggests that markup rates are match-specific. Rich pricing patterns also show up in the sizable heterogeneity in the degree to which relative cost shocks are passed onto foreign consumers (Burstein and Gopinath, 2014). For instance, discussions on heterogeneous pass-through and incidence are key in the debate regarding the impact of tariffs imposed by the US and China (Flaaen et al., 2020, Fajgelbaum et al., 2020, Cavallo et al., 2021, Alviarez et al., 2021).

[3]The demand side of the model follows the assumptions in Antràs et al. (2017), who also describe the geography of import sourcing in a Ricardian setup in which varieties offered by domestic and foreign suppliers are perfect substitutes. They explain heterogeneous import sourcing patterns by the presence of fixed sourcing costs. Instead, sourcing patterns are shaped by search frictions in our model, as in Eaton et al. (forthcoming).

[4]This is true controlling for other sources of price adjustments, such as market-specific inflationary movements or firm-specific cost shocks. Similar patterns have been found in US data (Heise, 2024, Monarch and Schmidt-

across French suppliers, switches are directed towards more competitive input providers.

A second contribution is to introduce a new structural approach for estimating search and comparative advantage parameters in firm-to-firm trade data. This analysis leverages French customs data covering the universe of transactions involving French exporters and their European buyers from 2002 to 2006. The estimation uses maximum likelihood, a common tool to estimate search models of the labor market on duration data (Postel-Vinay and Robin, 2006). Identification is achieved using the mobility of firms along the supplier network, as measured by transaction and switch frequencies.[5] Intuitively, the rate at which importers switch from one French exporter to another, conditional on a transaction, is informative about the magnitude of the frictions faced by French firms in the corresponding foreign market. Instead, the censoring rate away from French firms, together with observed trade shares, are informative about the rate at which firms match with non-French firms.[6] However, the mapping between these empirical moments and the structural parameters is complex due to unobserved quality-adjusted cost differences between French exporters as well as unobserved switches toward non-French exporters. By working with unconditional hazard rates, as in Ridder and van den Berg (2003), we address these issues. Moreover, the use of unconditional switching frequencies implies that the estimation method remains robust across a range of alternative price-setting mechanisms. The simulated maximum likelihood estimator makes it possible to recover separately the parameters driving Ricardian comparative advantages and search frictions, at market the level.

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 23% of the observed variance in French firms' market shares, with the remaining 77% explained by Ricardian forces. Most of the cross-sectional variance is found across destinations, within a sector. Search frictions are closely associated with distance: The elasticity of relative meeting rates to distance is -0.77.[7]

The estimated model matches empirical transaction and switch frequencies almost perfectly. Whereas these two sets of moments are used for identification, we show that our model also reproduces the dynamics of prices within relationships observed in the data. After six months in a relationship, the price of French exports has decreased by 3% in our model on average, against 2.2% in the data. We also leverage the variability of the bilateral exchange rate between France and the UK to estimate the pass-through of exchange rate shocks into French export

Eisenlohr, 2023).

[5]From this point-of-view, we are not exposed to the identification issue underlined in Bernard and Zi (2022), who show that a large class of firm-to-firm trade models deliver observationally equivalent predictions regarding the cross-sectional structure of firm-to-firm data.

[6]Separation can also occur for other reasons, such as the firm's death. We account for this possibility using exogenous separation shocks which we calibrate using the long-run separation rate in the data.

[7]The systematic correlation between relative search frictions and distance from France provides interesting external validity to our estimates. It has long been recognized that the gravity structure of trade may in part reflect the impact of distance on the strength of information frictions (Rauch, 1999, Rauch and Trindade, 2002). Our results are consistent with this view, although the moment used to estimate frictions, the switching frequency of buyers in the firm-to-firm network, is not mechanically correlated with geography.

prices, and its sensitivity to search frictions. In the model as in the data, pass-through rates are found higher in markets in which the relative meeting rate of French firms is high. The level and sensitivity of pass-through rates to market frictions are quantitatively similar in the model and in the data.

Our third contribution is to explore how the sensitivity of prices to local competition within a firm's network affects the incidence of cost shocks. We study the adjustments of prices and trade relationships in the estimated model using a hypothetical 10% cost shock affecting all French firms.

Within surviving relationships, the average pass-through of the shock is equal to 38% on impact but varies across markets and over time. This average hides wide heterogeneity across relationships, within and across a market. Within a particular market, the pass-through varies between zero (50% of the relationships in the average market) and one (a third of all relationships). The reason is that the firm's incentive to pass the shock depends on the second-best option of its downstream partner. When the second best is non-French, and thus unaffected by the shock, the firm needs to adjust its markup one-to-one, to retain the buyer. If instead the second best option is also French, before and after the shock, the firm can pass the shock onto its downstream partner. Over time, the composition of input supplier networks changes, and the average pass-through decreases, to 29% after two years.[8] Of course, the relative prevalence of French firms in buyers' networks varies across markets, which explains that the pass-through rate increases in markets in which the relative meeting rate of French firms is high.

Beyond the pass-through of shocks within existing relationships, the overall incidence of the shock also depends on the price increase incurred by buyers switching away from French firms in the aftermath of the shock.[9] Accounting for both margins, we find that the overall incidence of the shock is equal to 35% on impact, and 26% after two years. The incidence is significantly higher, in particular in the short run, in markets where French firms face relatively low frictions compared to their competitors. In such markets, buyers' portfolios are more heavily weighted toward French sellers, limiting buyers' ability to mitigate the shock by switching away from French suppliers.

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 cost shocks. Key to our paper is the introduction of on-the-match search and second-best pricing, which induces dynamics *across* but also *within* firm-to-firm trade relationships.

Like several other recent contributions, we examine firm-to-firm trade in a model displaying search frictions.[10] Several papers introduce endogenous search efforts (Bernard et al., 2019,

[8] Existing evidence on the dynamics of pass-through rates rather points to the pass-through increasing over time, which the literature explains by the existence of nominal rigidities (Burstein and Gopinath, 2014). In our model, the sluggishness in price adjustments is due to search frictions. As buyers' networks expand over time, they can only gain market power to negotiate the pass-through down.

[9] By definition, the price increase associated with a switch is lower than the adjustment the firm would have incured when sticking with its existing partner. Switching is thus a way to reduce the incidence of the shock.

[10] Firm-to-firm trade has also been analyzed in monopolistic competition settings without search or matching

Demir et al., 2024, Eaton et al., 2022, Lu et al., 2024) and learning about product appeal (Eaton et al., 2021) to microfound the matching of firms in international markets. In Eaton et al. (forthcoming) and Miyauchi (forthcoming), search is random but there is congestion in the matching technology. Like in Lenoir et al. (2022), our model admits a simpler search structure in which the matching process is governed by pure random search. The benefit is that we can be more flexible on the heterogeneity of search parameters across countries and sectors, and estimate bilateral search parameters for 26 sectors and 14 destinations. Compared to these papers, our model also proposes a richer view of firm pricing strategies. A closely related paper is Grossman et al. (2024). The authors study theoretically the impact of tariffs in global supply chains in a general equilibrium model with endogenous search effort. Like in our model, the adjustment following an unanticipated price shock is governed by bilateral price adjustments and switches.

The second-best pricing at the core of our model generates rich patterns of price adjustments, within and across relationships. Dhyne et al. (2022) and Alviarez et al. (2021) also study the determinants of prices in (static) models of firm-to-firm relationships, both theoretically and empirically. These models impose an exogenous structure of oligopolistic competition across input suppliers. We instead assume perfect substitutability between input suppliers and endogeneize the strength of competition thanks to the combination of random search and Bertrand competition. This framework allows us to explore dynamic adjustments to cost shocks, at the intensive and the extensive margins.[11]

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, Drozd and Nosal, 2012, de Blas and Russ, 2015). Like Bernard et al. (2003) 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 a firm, across its relationships, as well as within firm-to-firm relationships. The variety of price adjustments predicted by the model in the aftermath of a shock talks to 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). It also relates to stylized facts on firm-to-firm pricing uncovered in Fontaine et al. (2020), Alviarez et al. (2021), and Burstein et al. (2024). The decline in prices within trade relationships is consistent with results in Heise (2024) based on US data. In our model this trend is driven by a drop in markups when buyers receive new offers. Heise (2024) instead assumes increasing relationship capital that lowers marginal production costs.

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 inspired by Ridder and van den Berg (2003) and, more generally, reminiscent of the estimation of the job-to-job transition parameters in job search

frictions (Bernard et al., 2018, Carballo et al., 2018, Lim, 2018). See Bernard and Moxnes (2018) for a review.

[11] In 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.

models (see Postel-Vinay and Robin (2006) for a survey). Like Ridder and van den Berg (2003), the identification of frictions exploits the frequency of switches. The key underlying assumption is that observed switches induce greater intertemporal profits for the buyer that switches. As the structural analysis does not use information on trade prices and quantities, the estimation remains robust under price-setting mechanisms beyond Bertrand competition. The parametric identification strategy allows us to estimate frictions that are market (i.e. sector$\times$country) specific. Alternatively, Miyauchi (forthcoming) 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 a number of facts regarding firm-to-firm trade relationships. Section 3 sets up the model, while Section 4 discusses the estimation and results. In Section 5 we study the estimated model's implications for the incidence of relative cost shocks. Finally, Section 6 concludes.

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, 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 input suppliers s in the model developed later and the European importers will be the empirical counterpart of the buyers b. The transaction is also characterized by a product category j at the 8-digit level of the combined nomenclature and the date t of the transaction at the monthly level.[12] 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.

The analysis covers the 2002-2006 period, which does not incur any substantial change in the product nomenclature, nor 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 4-digit 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

[12] Firms exporting less than 150K euros toward EU countries can fill a simplified export form that does not collect information on the exported product. As we use the product dimension in our analysis, we exclude the corresponding trade flows, which account for 1% of French exports to EU countries. We checked that surviving probabilities are not significantly different in the corresponding subset and our estimation sample.

specific product j and various French sellers s, over time.[13] The model explains the decision to purchase goods from a given seller in the context of frictional good markets whereby importers are willing to purchase a particular input j, 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. Our dataset covers 27 million transactions, 40,000 French exporters and 744,000 European importers. In the rest of the analysis, we 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 and an average of five transactions per relationship. 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 Four 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., 2018, Bernard et al., 2018, Lenoir et al., 2022, Eaton et al., forthcoming). The literature documents 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 these statistics in Appendix A.2 and compare the degree of connectedness in a monthly cross-section and over time. We also discuss how statistics vary within and across products. In this section, we focus on four additional stylized facts regarding the *dynamics* of firm-to-firm relationships, and its consequences for the dynamics of prices and quantities.

2.2.1 Hazard rates of firm-to-firm relationships

We first delve into the dynamics of firm-to-firm relationships. Our data indeed displays significant heterogeneity in the duration of relationships, with 40% of relationships not surviving the first transaction, whereas a significant number of relationships are long-lasting. We now use the dynamics within a firm-to-firm relationship to measure how separation rates evolve over time. Figure 1 illustrates the evolution of hazard rates, as defined by the probability that a relationship ends after x months, conditional on the relationship having survived up to that point.[14] Hazard rates are computed over the whole sample ("All" line) as well as in markets

[13] We cannot take the opposite perspective and focus on French importers and their foreign suppliers as the firm-to-firm dimension is not available in *import* data.

[14] Left-censoring is controlled for using data prior to 2002 to recover the full length of a relationship. In figure D14, we reproduce the graph using an alternative measure of tenure, the number of cumulated transactions in the relationship.

Figure 1: Hazard rate, over time

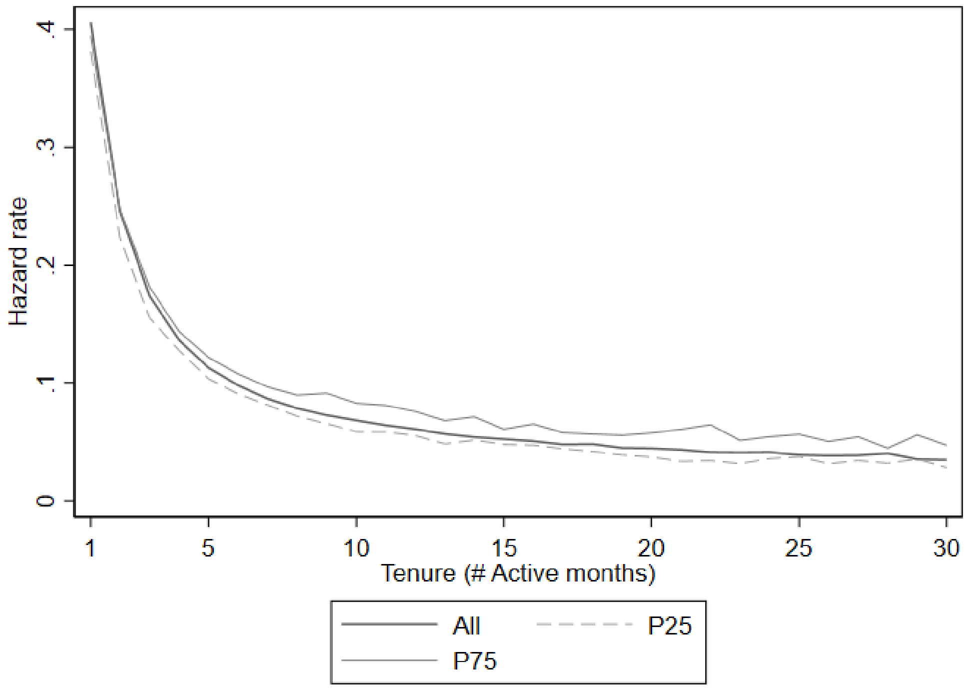

Notes: The hazard rate is defined as the probability of the relationship ending, conditional on tenure in the relationship and is calculated as the ratio of the density to the survival rate at tenure k. The figure is recovered from the 2002-2006 sample using the cumulated number of periods in the relationship as measure of tenure. The "All" line is computed from the whole dataset. The "P25" and "P75" lines correspond to the markets at the first and third quartile of the distribution of hazard rates at a tenure of 30 months.

lying at the first and third quartiles of the distribution of hazard rates after 30 months ("P25" and "P75" lines). The declining pattern implies that the survival rate increases over the course of a relationship. The dynamic is especially strong during the first year, whereas the probability of the relationship ending stabilizes after 1.5 to 2 years, at around 4%. The heterogeneity across markets is limited over long horizons, but the hazard rate after the first transaction and the speed at which it declines strongly vary across markets. We illustrate this heterogeneity using the wearing apparel sector as a case study in Figure D2 in the Appendix. This leads us to the first of our stylized facts.

Fact 1 *The probability of the relationship ending declines over the length of a firm-to-firm relationship, before stabilizing after 18 months, on average.*

This pattern is consistent with evidence recovered from US import data, as discussed in Eaton et al. (2021) and Monarch and Schmidt-Eisenlohr (2023).

2.2.2 Importers' mobility across suppliers

We now turn to the importer's side of the graph and study the mobility of importers, along the supplier network. As discussed in Appendix A.2, this side of the network displays less cross-

Figure 2: Mobility of importers, over time

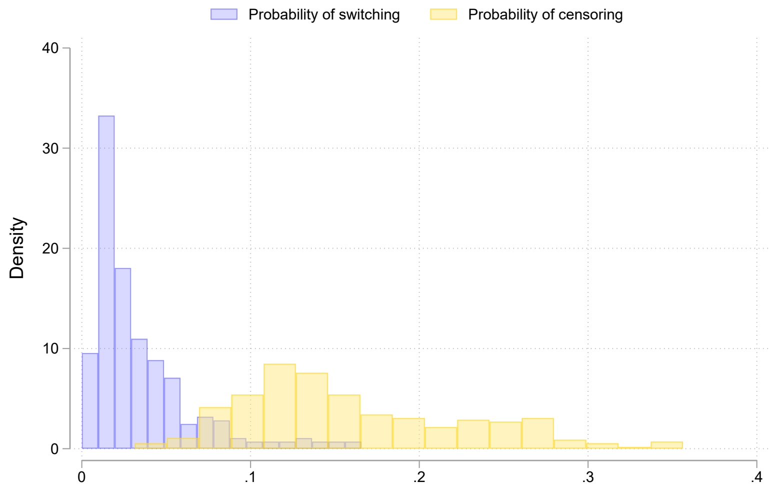

Note: The figure shows the distribution across country×sectors of the censoring and switch probabilities. The probabilities use as reference the population of importers purchasing products from France in January 2002. The probability of a censoring is defined as the share of firms that are never seen 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.

sectional heterogeneity than the exporter's side. More than 95% of importers interact with a single French firm over a particular month and product (see also Sugita et al., 2023, based on US data). However, the distribution of importers' in-degrees is 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. More specifically, around 25% of importers source multiple inputs from France, at a point in time. The degree of connectedness is stronger and more heterogeneous when we cumulate a firm's partners across periods. This fact is particularly interesting as it reveals a form of mobility, with importers switching across suppliers, over time.

The mobility is further illustrated in Figure 2. Here, we focus on a cross-section of importers that we observe purchasing a particular product from France in January 2002. We then follow these importers over the next two years 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) the next transaction takes place with the same partner, i.e. the relationship continues. Figure 2 shows the distribution of the estimated probabilities of a censoring and a switch, both computed at the country-sector level. 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, at 82% on average. This is consistent with Figure 1, which already showed that 60% of relationships involve at least two transactions, and thus at least one repeated transaction. Second, we observe that between 0 and 35% of importers (and 16% on average) are never seen interacting again with a French firm, which our model will explain by a combination of exogenous separations and the importer's decision to switch to a non-French supplier.[15] Finally, we also see a significant mass of switches, i.e. importers terminating a relationship before engaging into a new relationship with another French producer of the same product.[16] In the overall sample, the probability of a switch is low, at 5%. However, it varies substantially across markets, reaching 17% in the population of Belgian importers of beverages.

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.*

2.2.3 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 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:

$$\ln p_{bjst} = FE_{bjs} + FE_{ijt} + \sum_{k=2}^{K} \alpha_k \mathbb{1}(Tenure_{bjst} = k) + \varepsilon_{bjst} \tag{1}$$

where p_{bjst} is the unit value set on the transaction involving exporter s, product j, importer b that occurs at time t. The presence of relationship-specific fixed effects (FE_{bjs}) implies that the identification of other coefficients is within a firm-to-firm relationship. The α_k coefficients measure how the average price evolves when tenure into the relationship increases. The baseline regression also controls for country×product×period fixed effects (FE_{ijt}) to account for destination-specific inflation trends. We also tested models that control for seller×period or seller×product×period fixed effects to account for heterogeneous cost shocks. Additionally, to address composition effects, we estimate the α_k coefficients within bins based on relationship duration, ranging from the shortest relationships lasting two months to those exceeding one year.

Results are reported in Figure 3. They show a negative trend in prices, that is observed within short as well as within long relationships. The downard trend is consistent with recent evidence by Monarch and Schmidt-Eisenlohr (2023) and Heise (2024) showing a decline in the

[15] 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 switch to a (non-French) supplier not covered by our data. In the structural analysis, we will distinguish between these possibilities using two empirical moments, the censoring rate in Figure 2 and the long-run hazard rate in Figure 1.

[16] Lu et al. (2024), Monarch (2022), Sugita et al. (2023) also discuss the tendency of importers to switch across suppliers, over time, using different datasets, and slightly different definitions of a switch.

Figure 3: Price dynamics, within a firm-to-firm relationship

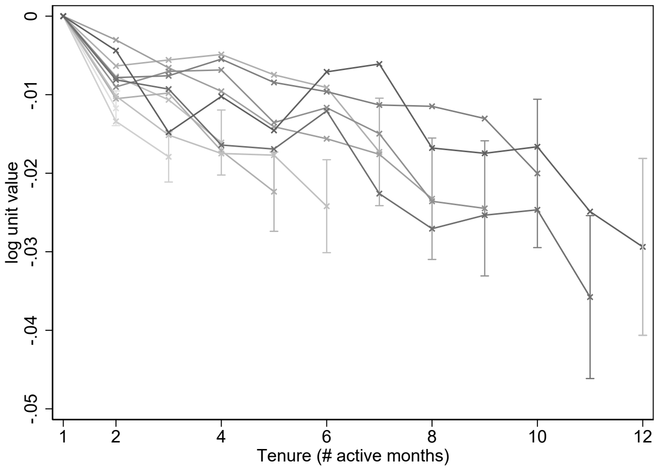

Note: This figure shows the evolution of prices within a firm-to-firm relationship. Coefficients are recovered from equation (1), allowing for different price trends across shorter and longer relationships. The figure reports the estimates and 95% confidence intervals at the final transaction within each bin of relationships. Tenure is measured by the number of periods since the beginning of the relationship. The lighter line corresponds to all realtionships of length two. The darker line is estimated on all relationships that last a year or more.

price paid by US importers as relationships age.[17] After six months, prices are on average 2.2% lower than in the initial transaction. After a year, they are almost 3% lower.

Results presented in the Appendix confirm the robustness of this finding. Figure D3 shows that the same pattern holds with a different measure of tenure, based on the number of shipments rather than the duration into the relationship. In Figure D4, we show that the relationship survives if we control for seller-specific cost shocks using seller×product×period fixed effects. The price decline within relationship is somewhat dampened but still significant, at 2% after a year in the relationship. In Figure D5, we study the dynamics of exported quantities. One may be concerned that the price decline is driven by increased shipments when the relationship ages, which would reveal a form of non-linear pricing. However, the quantity of exported goods tends to decline over time for short relationships while the pattern is inversed U-shaped for long relationships.

This leads us to our third stylized fact.

Fact 3 *Within a firm-to-firm relationship, prices tend to decrease over time. On average, the price decline reaches 2.2% after 6 months and 3% after a year.*

[17] Fitzgerald et al. (2023) document that the average prices charged by Irish exporters do not move after entry. Compared to us, their analysis is at the firm×destination×product×year level and the correlation between the average price, and tenure is identified across destination countries. Our fact instead focuses on prices within a firm-to-firm relationship, with identification achieved over time. In Figure D6 in the Appendix, we reproduce their empirical specification using our data, aggregated at the destination×country level and confirm the flat pattern that they document. If we adjust their specification to exploit the dynamics within firm-product-destination triplets, we recover a declining trend, that confirms results in Figure 3.

Figure 4: Price changes, conditional on a switch

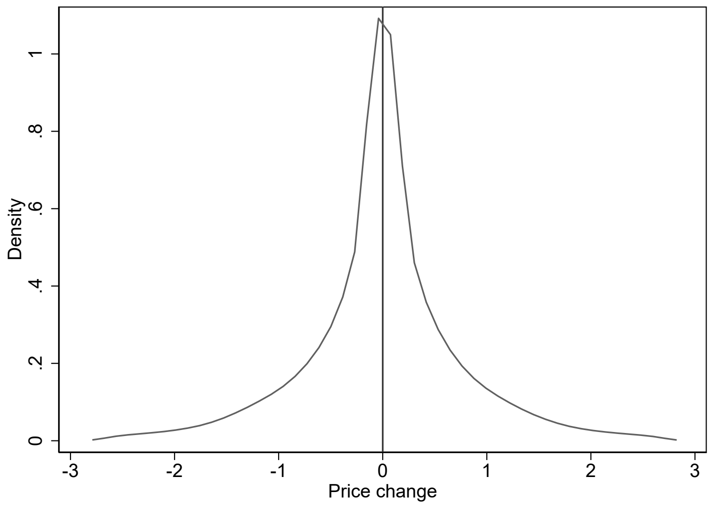

Note: This figure shows the kernel density of price changes, conditional on a switch. Price changes are computed in log differences, between the last transaction within a seller-buyer-product pair and the next transaction involving the same buyer×product but a different seller.

2.2.4 Price change conditional on a switch

Whereas the price decline within a firm-to-firm relationship is statistically significant, the behavior of prices following a switch is far less clear. This is illustrated in Figure 4 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. The density is centered around zero with a mean at .006 and a median at 0. Switches do not systematically involve a price saving for the importer.

In Figure 5, we explore a more comprehensive measure of the 'competitiveness' of the new input supplier. Switches might indeed be triggered by a combination of pure price motives and quality upgrading, thus explaining why the direction of the price adjustment conditional on a switch is not systematic (Monarch, 2022). As we lack a direct measure of a firm's position within the quality-adjusted price distribution, we use predictions from a large class of firm-to-firm trade models to estimate it. As discussed in Bernard and Zi (2022), many models of the matching of firms in firm-to-firm markets predict that a seller's degree increases with its competitiveness. This is true in our model as well. The expected number of buyers in an input supplier's portfolio is scaled by its competitiveness, as defined by a low quality-adjusted cost.

Based on these insights, we construct an index of sellers' positions within the distribution, using the following equation:

$$\ln n_{sjit} = FE_{sji} + FE_t + \sum_{k=2}^{K} \alpha_k \mathbb{1}(Experience_{sjit} = k) + \varepsilon_{bjst} \tag{2}$$

Figure 5: Evolution of sellers' attributes, within a buyer's sequence of French partners

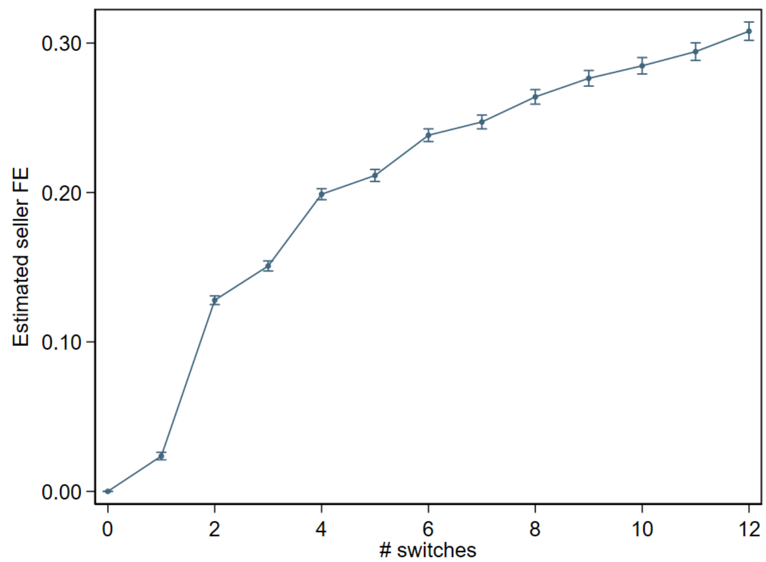

Note: This figure shows the evolution of the seller's cost-adjusted quality, across switches. Switches are defined as transitions from one French seller to another, within a 8-digit product market. Cost-adjusted qualities are estimated at the seller-product level as the fixed effect of a regression describing the dynamic accumulation of buyers, within a firm. The estimated equation controls for unobserved heterogeneity using a buyer×product fixed effect.

where n_{sjit} is the number of buyers from country i served by seller s with product j at time t. FE_{sji} and FE_t are individual and time fixed effects, respectively. $\mathbb{1}(Experience_{sjit} = k)$ is a dummy variable equal to one if the seller's experience in country i is equal to k. The α_k coefficients thus capture returns to experience and are depicted in Figure D7. We use the estimated FE_{sji} as a proxy for the seller's cost competitiveness in market ji. More efficient sellers tend to attract more buyers than other sellers with similar experience in the same market.

In a second step, we dig into the direction of buyers' switches, along the network of French input suppliers. More specifically, we estimate the following equation

$$\widehat{FE}_{b(\mathbf{s})ji} = FE_{bj} + \sum_{l=2}^{K} \alpha_l \mathbb{1}(Partner_{bjs} = l) + \varepsilon_{bji} \tag{3}$$

where $\widehat{FE}_{b(s)sji}$ is the "competitiveness" of buyer b's seller recovered from the previous equation. FE_{bj} is a buyer-product fixed effect and $\mathbb{1}(Partner_{bjs} = l)$ is a dummy that is equal to one if seller s is the lth partner of buyer b in a ranking recovered from the sequence of its French partners. In this equation α_l measures how $\widehat{FE}_{b(s)ji}$ changes when a buyer switches from its l-1th to its lth French supplier.

The pattern in Figure 5 shows a clear upward trend, indicating that supplier switches among French firms enable foreign buyers to move up the cost-adjusted quality ladder. Over time, foreign importers switch towards 'better' French suppliers, that accumulate more partners in that destination. This trend is consistent with a model in which switches are endogenous and systematically directed towards better French firms.[18] This leads us to the fourth and final stylized fact.

Fact 4 *Buyers switch towards sellers with lower quality-adjusted costs, but the observed price paid following a switch may increase or decrease.*

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

3 A Search Model of Firm-to-Firm Trade

3.1 The demand for intermediate goods

The demand side of input markets follows Antràs et al. (2017). In each country, final good producers use a CES bundle of imperfectly substitutable inputs, sourced domestically and abroad. Within an input type, input suppliers offer perfectly substitutable varieties at heterogeneous quality-adjusted costs. In the following, we denote p_j and q_j the price and quality of input j. The introduction of quality is motivated from fact 5: firms do not necessarily switch to cheaper suppliers but they switch to more competitive suppliers, as defined by a low quality-adjusted price.

The producer chooses the quantity x_j of each intermediate input to minimize variable costs, given output x_b, which we take as exogenous.[19] In equilibrium, the demand addressed to a given input provider is a function of its quality-adjusted price and variables specific to the final good supplier, namely its marginal cost and aggregate output:

$$p_j x_j = x_b \left(\frac{p_j}{q_j}\right)^{1-\eta} mc_b^{\eta} \tag{4}$$

where η is the elasticity of substitution between inputs and mc_b denotes the marginal cost of final good producer b:

$$mc_b = \left(\int_{j\in\Omega} \left(\frac{p_j}{q_j}\right)^{1-\eta} dj\right)^{\frac{1}{1-\eta}} \tag{5}$$

Given perfect substitutability between input suppliers of a given type, the final good producer chooses, for each input, the seller offering the lowest quality-adjusted price, p_j/q_j.

[18] An alternative interpretation of firm-to-firm switches is that different French exporters sell different varieties of the same input, including within the 8-digit product markets. If product differentiation was the sole motive for switches, we would not expect to see the systematic pattern in Figure 5. Instead, foreign buyers would switch back-and-forth across suppliers of differentiated varieties, of similar attributes.

[19] As in Antràs et al. (2017), the model is solved in partial equilibrium. The predictions do not vary if we introduce a monopolistically competitive market for final goods and endogeneize production quantities (Appendix C.1). The model could also handle shocks to the demand of final good producers. As long as the technology features constant returns to scale, these shocks would not affect equilibrium input prices. The only important simplifying assumption is that input search is independent across product markets.

3.2 Heterogeneous input suppliers

We now describe the production of inputs. As final good producers, input suppliers can be located in any country. For the sake of clarity and due to the limitation of our data, we focus here on suppliers located in France (country F) while accounting for the competition exerted by non-French competitors ($\overline{F}$). We also abstract from the j subscripts but all parameters introduced in this and the next sections are input-specific. As long as search takes place across separate (input-specific) markets, an assumption that we maintain throughout, we can estimate the model market by market and thus allow all parameters to vary across input types.

Input suppliers 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 country-specific cost of the input bundle. In the rest of the analysis, we do not seek to separate productivity and quality and thus characterize producers by their quality-adjusted productivity $z = e \times q$. The quality-adjusted unit cost of serving market i for a French firm of quality-adjusted productivity z reads:

$$c_{iF}(z) = \frac{\nu_F d_{iF}}{z}$$

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. Following Eaton and Kortum (2002), quality-adjusted cost is the single source of ex-ante heterogeneity across input suppliers in the model. In equilibrium, it explains heterogeneity in the structure of exporters' portfolios uses to proxy sellers' competitiveness in Section 2.2.4. It is a firm attribute, that we take as constant.[20]

When a buyer and a seller meet, the quality-adjusted productivity is drawn in a sampling distribution $F(z) = 1 - (z/\underline{z})^{-\theta}$, with support on $[\underline{z}, +\infty]$. As in Eaton and Kortum (2002), the probability that the cost of a French firm serving market i is below c reads

$$F_{iF}(c) = 1 - F\left(\nu_F d_{iF}/c\right) \tag{6}$$

Symmetrically, we define the probability that a non-French seller serves market i at a cost below c as $F_{i\bar{F}}(c) = 1 - F\left(\nu_{\bar{F}} d_{i\bar{F}}/c\right)$.

Until now, our environment closely follows Antràs et al. (2017). They next solve for the geographic composition of a firm's input portfolio under the assumption that there is a fixed cost of sourcing from any country. We depart from this assumption and instead follow Eaton et al. (forthcoming), in assuming that the geography of input portfolios is shaped by bilateral search frictions and comparative advantages.

[20]In Section D7, we formally derive the dynamics of buyers' acquisition as a function of the firm's quality-adjusted cost.

3.3 Matching and pricing on the intermediate good market

We now introduce the matching process, which we assume is governed by random search. A final good producer located in country i meets with French sellers at rate γ_{iF} and with sellers from other countries at a rate of $\gamma_{i\overline{F}}$ which is the sum of the non-French meeting rates.[21] Under the CES assumption, any match is potentially profitable but we put a limit to suppliers' market power by assuming that the final good producer can produce the input in-house, at a cost which constitutes its outside option before it starts accumulating knowledge about potential suppliers.[22] We solve for the steady state distribution of matches under the assumption of a stationnary mass of final good producers in each market. A buyer exits the market at exogenous rate μ and is replaced by a new buyer. The new buyer starts unmatched, meets input suppliers and maintains links.

Sellers Bertrand compete to supply inputs to final good producers. We rule out collusion between suppliers. We further assume that buyers always have the option to recall one of their previous suppliers and that there is no commitment beyond the current transaction. Both are important assumptions that simplify the price setting. The problem reduces to a static comparison of price-to-quality ratios, without intertemporal consideration.

Take a final good producer who has met with n potential suppliers over a given input. We index these sellers by their quality-adjusted serving cost : $c_1 \leq c_2 \leq ... \leq c_n$. The best supplier c_1 is able to set the price p such that the buyer is just indifferent between first best options. When the second-best cost is high, the firm may however prefer to set the monopoly price. In equilibrium, we thus have:

$$p\left(q_1, c_2\right) = Min\left\{c_2 q_1; \frac{\eta}{\eta - 1} c_1 q_1\right\}$$

The equilibrium price rule is thus similar as in Bernard et al. (2003). The difference with their setting is that the identity of the first- and second-best suppliers are random variables in presence of search frictions. This is an important property that delivers rich predictions regarding the distribution of markups. As the degree of competition within each final good producer's network is endogenous, markups not only depend on the seller's and the buyer's types but also the structure of the buyer's network, at a point in time.

The randomness induced by search frictions also implies that prices are renegotiated over time when 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, the new seller may have a quality-adjusted cost above $p\left(q_1, c_2\right)/q_1$, in which case nothing happens: the next transaction will be with the incumbent supplier at the same price. Instead,

[21] Search efforts are thus exogenous in our framework. Endogenous search efforts would create within heterogeneity in the matching rates but wouldn't change the direction of the job-to-job transitions. Assuming exogenous search effort simplifies the estimation, reduces the need for parametric assumptions and is unlikely to change the average meeting rate estimated on the data. See Demir et al. (2024) or Eaton et al. (2022) for firm-to-firm trade models with endogenous search efforts.

[22] In practice, we will calibrate the outside option to the highest realization of quality-adjusted costs in our simulations. The assumption is irrelevant for identification and only matters when we quantify the incidence of cost shocks in the estimated model.

if the new cost is below c_1, the next transaction (if there is no new match before) will be with the new seller at price

$$p\left(q', c_1\right) = Min\left\{c_1 q'; \frac{\eta}{\eta - 1} c' q'\right\}$$

Interestingly, that price can be above or below the previous price depending on the quality of the new seller's product. Consistent with fact 5, endogenous switches are associated with price adjustments that can be positive or negative, although the switch is always directed towards lower quality-adjusted cost firms. This unambiguously drives the buyer's quality-adjusted price down.

Lastly, the new seller may have a quality-adjusted cost in between c_1 and $p\left(q_1, c_2\right)/q_1$. 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:

$$p\left(q_1, c'\right) = q_1 c'$$

Hence, since $c' < Min\left\{c_2, \frac{\eta}{\eta-1} c_1\right\}$, 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 3.

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 average market, the firm has already reduced the price of her input by more than 30% after 6 months of activity.[23] Of course, the speed of these adjustments slows down over time, as it becomes increasingly difficult to meet with suppliers that compete with the firm's best match. A higher overall meeting rate helps firms gain bargaining power more rapidly.

3.4 Distributions and shares[24]

The distribution of suppliers. A buyer in country i meets suppliers at a rate $\gamma_i = \gamma_{iF} + \gamma_{i\overline{F}}$. A the same time, she exits the market at an exogenous rate μ and is replaced by a buyer unmatched with any supplier. At the steady state, there is thus a share u_i of unmatched buyers, which we derive using the flow equilibrium condition $B_i u_i \gamma_i = B_i (1 - u_i)\mu$,

[23]It should be noted that the magnitude of the price adjustment recovered in Figure 6 is not directly comparable to the empirical counterpart in Figure 3. The reason is that the data average price dynamics across 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. We will discuss the model's fit with respect to the price dynamics in Figure 3 in Section 4.3.

[24]Throughout 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.3, we derive analytical results regarding the dynamics of buyer acquisition. Since they are not directly used in the structural estimation, these derivations are left for the Appendix.

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

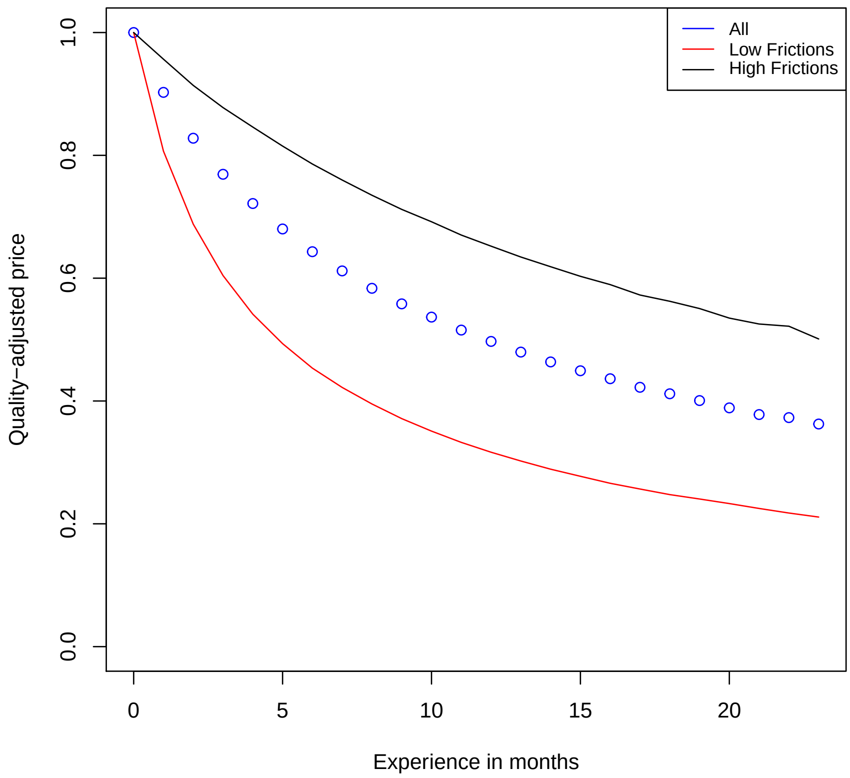

Notes: The figure shows the dynamic of average prices in the model's steady state, for average meeting rates (blue circles) and when meeting rates are either high (red line) or low (black line). The model is calibrated using parameters estimated in section 4 and $\eta = 3$. A buyer's experience is measured from its first match with a supplier located in any country.

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

When a buyer is matched with a new supplier, the associated quality-adjusted serving cost is drawn in $F_i(c)$. This distribution is a mixture of the country-specific distributions:

$$F_i(c) = \frac{\gamma_{iF}}{\gamma_i} F_{iF}(c) + \frac{\gamma_{i\overline{F}}}{\gamma_i} F_{i\overline{F}}(c) \tag{7}$$

for $c \in \left]0, max\left(\frac{\nu_F d_{iF}}{\underline{z}}, \frac{\nu_{\overline{F}} d_{i\overline{F}}}{\underline{z}}\right)\right]$ (see Appendix C.2 for details). It is useful to note that we can express the distribution of French suppliers as a shift of the distribution of the non-French suppliers:

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

where $\tau_{iF} \equiv \left(\frac{\nu_F d_{iF}}{\nu_{\overline{F}} d_{i\overline{F}}}\right)$ measures the comparative advantage of foreign suppliers over French firms.

At the steady state, it is straightforward to derive how buyers in country i are distributed across suppliers, along the cost-c dimension. It is a useful relationship as its shows how the matches are affected by frictions and how a shock to the meeting distribution translates into changes in the buyer distribution. Denote $L_i(c)$ the cdf of the buyer distribution and $\ell_i(c)$ its density. As long as buyers always choose the lowest cost supplier that they have met, $\ell_i(c)$ satisfies

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

with $\overline{L}_i(c) \equiv 1 - L_i(c)$. 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 ($\gamma_i \overline{L}_i(c) f_i(c)$). Simplifying and integrating by parts, we obtain

$$L_i(c) = \frac{\mu + \gamma_i}{\mu + \gamma_i F_i(c)} F_i(c) \tag{8}$$

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

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

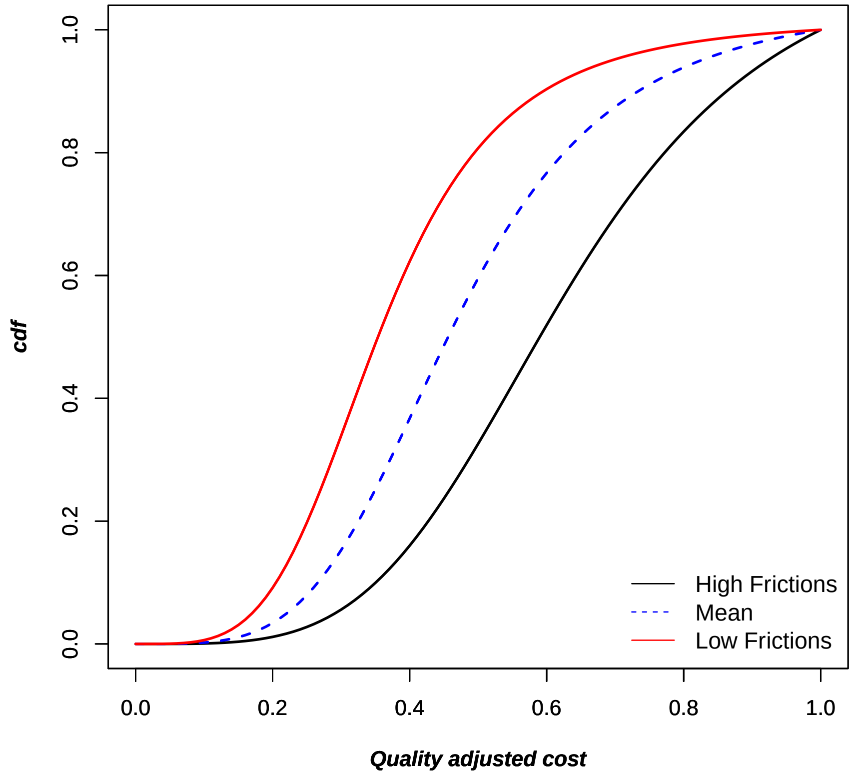

Notes: The figure reports the cumulated distribution of input costs paid by buyers in a particular market, at the average market in terms of meeting rates (blue dots) and when meeting rates are either high (red line) or low (black line). The calibration uses parameters estimated in section 4.

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 different 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.

Consider now the distribution of matches with French suppliers. The share of buyers matched with French sellers, π_{iF}, is an important object that depends both on the meeting rates and on the Ricardian advantage of French firms. Its expression plays a crucial role in the model's estimation and allows us to infer several implications from the estimated results. To derive this, we leverage the equilibrium condition, where the inflows and outflows are balanced, such that the density of buyers matched with a French seller at cost c, denoted by $\ell_{iF}(c)$, satisfies:

$$\underbrace{(1-u_i)\pi_{iF}\ell_{iF}(c)\,(\mu+\gamma_i F_i(c))}_{\text{outflows}} = \underbrace{u_i\gamma_{iF}f_{iF}(c)+(1-u_i)\overline{L}_i(c)\gamma_{iF}f_{iF}(c)}_{\text{inflows}} \tag{9}$$

This equation, together with the definition of $u_i = \mu/(\mu+\gamma_i)$, is used to solve for the equilibrium share of buyers matched with French firms. When μ is sufficiently close to zero, we have:[25]

$$\pi_{iF} = \frac{\gamma_{iF}}{\gamma_{iF} + \gamma_{i\overline{F}}\tau_{iF}^{\theta}} \tag{10}$$

Finally, we show that

$$\ell_{iF}(c) = \ell_i(c) \tag{11}$$

which means that buyers are identically distributed in terms of serving costs whatever the origin of their suppliers. As discussed in Appendix C.2, the invariance of serving costs across origin countries, conditional on a match, also implies that the expression for π_{iF} in equation (10) defines the share of country i's absorption that is sourced from France, if $\underline{z}$ 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.[26]

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 $\gamma_{i\overline{F}}$ 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 $\gamma_{i\overline{F}}$ and a decrease in τ_{iF} improve the likelihood that a French supplier serves market i. In comparison with the frictionless world in Eaton and Kortum (2002), heterogeneity in bilateral search frictions distort trade in favor of relatively low search / high meeting rates countries (Lenoir et al., 2022).

4 Structural estimation

4.1 Identification strategy

We estimate our model using simulated maximum likelihood, together with data on switch frequencies at the buyer level. At a very high level, our strategy can be understood as finding the structural parameters that best match the empirical switch frequencies in Figure 2. It is very similar to the estimation of search frictions using duration data in models of the labor market (e.g., Jolivet et al., 2006).[27] However, it has two distinctive features in comparison to a textbook

[25] Analytical details together with the formulas recovered in the general case when μ can take any value are provided in Appendix C.2. The interpretation of equation (10) discussed below goes through in the general case.

[26] Using this result, we can use equation (10) to define $\tau_{iF}^{-\theta}$ as a function of the meeting rates and the observed shares

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

This will be used in the estimation.

[27] In principle, our model can 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.

search model estimation. First, we don't want to put too much weight on our price bargaining assumption. While Bertrand competition is convenient to recover analytical solutions for the distribution of costs and the geography of trade, identification of search frictions relies on a looser assumption, namely that buyers switch if this improves their intertemporal profit. As such, our estimation of search frictions is robust to alternative price setting mechanisms.[28] Second, we want to avoid using data on either prices, quantities, or production costs, which are likely to be more noisy than the record of switch events. As in Ridder and van den Berg (2003), we use the fact that once we integrate over the distribution of matches, the probabilities to switch solely depend on the structural parameters, namely the matching rates, and the parameters shaping the distributions of serving costs. The fact that we do not rely on prices will be used later to assess the performance of the model with respect to untargeted moments.

The parameters to be estimated are the matching rates γ_{iF} and $\gamma_{i\overline{F}}$, the parameters shaping the distributions of serving costs, θ and τ_{iF}, and the rate at which a buyer exits the market, μ. One of the difficulties of our data is that switches to non-French sellers are not observed. If a switch is permanent this is materialized by the observation being censored : we do not observe any event until the end of the observation period. There are thus two ways for an observation to te censored, a switch outside France and an exit from the market. This means that μ is not separately identified from $\gamma_{i\overline{F}}$. For that reason, we calibrate μ. Our calibration strategy relies on the fact that the censoring events happen at joint rate $(\mu + \gamma_{i\overline{F}} F_{i\overline{F}}(c))$. Moreover, the longer a buyer stays with the same seller, the lower the serving cost : in the limit, as the duration of a relationship increases, we can infer that the serving cost is very low and $F_{iF}(c) \approx F_{i\overline{F}}(c) \approx 0$. In that case, the censoring rate converges to μ. Hence μ can be calibrated to fit the long-run empirical hazard rate of the buyer-seller relationships. Formally, with $H(c) = (\mu + \gamma_{iF} F_{iF}(c) + \gamma_{i\bar{F}} F_{i\bar{F}}(c))$, the hazard rate can be defined as

$$\mathbb{E}\left[\frac{H(c)e^{H(c)t}}{e^{H(c)t}}\right] = \mathbb{E}\left[H(c)\right] \xrightarrow[t\to\infty]{} \mu$$

since the probability to receive better offers goes to 0 in expectation, conditional on survival. In practice, we fit the average asymptotic hazard rate across markets, which implies $\mu = 0.04$ (see Figures 1 and D2).

Conditional on survival, the censoring and switching rates illustrated in Figure 2 can be used to pin down the matching rates γ_{iF} and $\gamma_{i\overline{F}}$. However, identification is not straightforward, for three reasons. First, the switching 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 outside France. Third, we only observe transitions that give rise to at least one transaction. Considering the first problem,

[28] Alternative price solutions have been used in the related literature. Grossman et al. (2024) and Eaton et al. (2021) use different forms of Nash bargaining, accounting for the presence of multiple sellers. In these settings the share of the surplus that the supplier gets is exogenous. Under our Bertrand competition assumption, the best seller extracts all the surplus by matching the buyer's second best option. While the prices are different, our setting shares with theirs the prediction that buyers switch towards sellers that generate higher surpluses which is our key identification assumption.

our method uses the fact that unconditional hazard rates, that is the hazard rates taken in expectation given the match distribution ($\ell_{iF}(c)$ in the model), only depend on the structural parameters. This is a surprising property, already used in Ridder and van den Berg (2003) or Jolivet et al. (2006), that relaxes the need to estimate the distribution parameters, apart from $\tau_{iF}^{-\theta}$. As an example, consider again $H(c)$, the hazard rate. Using equations (10) and (11):

$$\int_{c_{inf}}^{c^{sup}} H(c)dL_{iF}(c) = \frac{\gamma_{iF}\tau_{iF}^{-\theta} + \gamma_{i\bar{F}}}{\gamma_{iF} + \gamma_{i\bar{F}}} \int_0^1 \frac{\mu(\mu + \gamma_{iF} + \gamma_{i\bar{F}})}{\mu + \gamma_{iF}\tau_{iF}^{-\theta}x + \gamma_{i\bar{F}}x} dx \tag{12}$$

This demonstrates that the unconditional hazard rate solely depends on the matching rate parameters γ_{iF} and $\gamma_{i\bar{F}}$, the exit rate μ, and $\tau_{iF}^{-\theta}$, the parameter driving comparative advantages. Note that $\tau_{iF}^{-\theta}$ can be replaced by a function of the matching rates and the observed French share, using equation (10).[29] Given the definition of the share and the calibrated value for μ, $\gamma_{i\bar{F}}$ and γ_{iF} are then identified using the observed switches between French sellers and the censoring rates, reported in Figure 2.

Our estimation strategy has one important property : it is agnostic about the exact determinant of the switching decision. For example, if the prices were Nash-bargained, the buyers would still move towards sellers with lower quality-adjusted serving costs and equation (8), that defines the equilibrium distribution of buyers across sellers, would remain unchanged and so does the frequencies of the switches. Combined with unconditional inference, this is even stronger. Assume $c = \nu(\Pi)$, a monotonically decreasing function of the buyer intertemporal value in a match. Under that assumption $F(c)$ is the distribution of match value offered by the seller to the buyer when they meet and $L_{iF}(c)$ is the distribution of values for buyers among sellers, conditional on buyers switching if an offer increases their intertemporal value. Under this scenario, both (8) and thus (12) remain unchanged and our empirical strategy still identifies meeting rates, independently of price strategies.

The second 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 analytically the related densities. The solution is to rely on a simulated maximum likelihood 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 estimation procedure is implemented in practice.

Finally, we have to deal with the fact that any observation implies a transaction. A new match or a switch are only observed at the time of a transaction. In practice, this means that we need to record simulated data in accordance with this limitation and that we have to model how transactions happen. Following the way we model other events, we assume that transactions are exponentially distributed. To gain a bit of flexibility, while keeping the model simple, we consider a mixture model with two types of buyers, one buying more frequently

[29] In 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.

than the other. Formally, a buyer i can be of type 1 with probability p and a transaction rate t^1_{iF}, or of type 2 with probability $1-p$ and transaction rate t^2_{iF}. Since the estimation is performed at the level of a specific product and destination, these probabilities vary in these dimensions.[30] The key assumption is that the transaction rate is uncorrelated with the match quality. This is a natural assumption since our model is silent about the determinant of the transaction frequency. Moreover, under this assumption, the transaction rates are identified separately from the match distributions, using transaction frequencies.

4.2 Estimation results

The model is estimated for 330 sector×country pairs for which we have at least 100 observations. For each pair, we recover an estimate of the meeting rate of French and non-French firms (γ_{iF} and $\gamma_{i\overline{F}}$), from which we can compute the search advantage of French firms ($\gamma_{iF}/\gamma_{i\overline{F}}$) and the overall meeting rate, a proxy for the degree of frictions ($\gamma_{iF}+\gamma_{i\overline{F}}$). Conditional on the search parameters, we recover Ricardian comparative advantages ($\tau_{iF}^{-\theta}$) from the observed trade shares. Finally, we also estimate the average frequency of transactions ($pt^1_{iF}+(1-p)t^2_{iF}$). As 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. In doing this, we keep the granularity of the data but 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.

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, for the transaction rates and for French comparative advantages are all 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. Figures D9, D10 and D11 provide additional statistics on the mean value of the estimated parameters, by country and sector.

In Table 1, we focus on the model's predictions regarding the geography of trade. One of the model's predictions, as summarized in equation (10), is that bilateral trade flows are shaped by the combination of Ricardian comparative advantages and relative meeting rates. French firms are expected to capture a larger market share in countries where they have a comparative advantage (high $\tau_{iF}^{-\theta}$) and where their relative meeting rate with foreign buyers is higher (high $\gamma_{iF}/\gamma_{i\overline{F}}$). Columns (1) and (2) in Table 1 dig into this prediction, quantifying the relative contributions of each determinant. The model attributes 23.5% of the cross-sectional variance in French firms' penetration of European markets to differences in relative meeting rates, while the remaining 76.5% is explained by comparative advantages. Search frictions

[30] In theory, we could introduce more than two types. In practice, we manage to fit the transaction frequencies almost perfectly with two types (Figure 8).

Table 1: Contribution of estimated parameters to the geography of trade

	Dep. Var				
	$\ln \frac{\gamma^j_{iF}}{\gamma^j_{i\bar{F}}}$ (1)	$\ln \left(\tau^j_{iF\bar{F}}\right)^{-\theta}$ (2)	$\ln \frac{\gamma^j_{iF}}{\gamma^j_{i\bar{F}}}$ (3)	$\ln \frac{\gamma^j_{iF}}{\gamma^j_{i\bar{F}}}$ (4)	$\ln \frac{\gamma^j_{iF}}{\gamma^j_{i\bar{F}}}$ (5)
$\ln \frac{\pi_{iF}}{1-\pi_{iF}}$	0.235[a]	0.765[a]	0.078	0.193[a]	
	(.061)	(.061)	(.059)	(.068)	
ln distance					-0.765[a]
					(.103)
Obs.	330	330	330	330	330
Adjusted R^2	.040	.321	.205	.160	.270
Country FE	No	No	Yes	No	No
Product FE	No	No	No	Yes	Yes

Note: The table correlates estimated search and Ricardian comparative advantages to trade share ratios. In the model, we have

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

and thus the variance decomposition in columns (1) and (2) is thus exact. In column (3) (resp. (4)), we control for the unobserved heterogeneity of relative search frictions across countries (resp. sectors). Column (5) correlates estimated relative search frictions with distance from France, controlling for unobserved heterogeneity across sectors.

thus explain a sizable share of bilateral trade, although less than in Eaton et al. (forthcoming) who estimate search parameters at the aggregate level and use a different (cross-sectional) identification strategy. Columns (3) and (4) dig into the contribution of relative meeting rates and show that most of the explanatory power is across countries, within a sector (column (4)). Instead, when we control for country fixed effects, relative meeting rates play a negligible role in explaining the geography of trade.

Finally, column (5) shows that distance is an important correlate of relative meeting rates: the relative meeting rate of French firms is significantly larger in nearby than in distant countries. In quantitative terms, the -.765 elasticity implies that distance explains a .219 gap between the average relative meeting rate faced by French firms in Belgium and Finland, a third of the actual estimated difference. This result provides interesting external validity to our empirical strategy. It has long been recognized that the gravity structure of trade may in part reflect the impact of distance on the strength of information frictions (Rauch, 1999, Rauch and Trindade, 2002). Our results are consistent with this view, although the moment used to estimate frictions, the switching frequency of buyers in the firm-to-firm network, is not mechanically correlated with geography.

4.3 Model Fit

Targeted moments: Figure 8 illustrates how our model performs in matching the observed distribution of transaction and switch frequencies.[31] We compare the distribution in our data

[31]Figures D12 and D13 in the Appendix provide more detailed results by country.

with the distribution obtained in simulated data using market-specific estimated parameters. The top panel illustrates the match vis-à-vis transaction frequencies, i.e. the outcome variable is the probability that we observe a given number of transactions in the next 24 months for each of the buyers in the data. The match between the data and the simulated model is almost perfect. This is not surprising: transaction frequencies are high in the data thus offering a useful source of identification. The bottom panel illustrates how we perform vis-à-vis switching frequencies (from France to France). The model performs well on the probability of observing at least one switch. The corresponding probability is 5.7% in the data and 6.9% in the simulated model. Conditional on at least one switch, our model slightly overestimates the probability that we observe several switches during the two-year observation period, while under-estimating the probability of a single switch (2.5% in the data and 1.8% in the model). Overall, the estimated model replicates the data well.

Out-of-sample. Whereas the model has rich predictions regarding the evolution of prices, prices are not used for identification. Comparing the dynamics of prices in the model and in the data thus offers an informative out-of-sample assessment of the model's fit. To this aim, we provide two validation exercises.

In the first exercise, we compare the dynamics of prices within relationships in the simulated model and in the data. As shown in Section 2.2, prices have a tendency to decline over time, within a relationship. The model replicates the declining pattern (Figure 9). At steady state, the downward trend of prices matches empirical patterns estimated on relatively short-lived (i.e. less than 6 months) relationships. After 4 months, the model thus predicts a 1.9% decline in prices, which is exactly what is estimated in the data in the subsect of relationships that last exactly four months. After 6 months, the model predicts -3% which is in the confidence interval of estimates recovered from the data. Instead, the model overestimates markup adjustments within longer relationships. After 12 months, the model thus predicts a 6% drop in prices on average, compared to a 2% to 4% drop in the data. A possible source of discrepancy is that the model ignores any nominal rigidity, that could undermine buyers' ability to benefit from declining markup rates. Since these rigidities are probably more prevalent within long-term contracts, they may dampen the declining trend associated with search frictions in our model.

In the second exercise, we compare how prices adjust to shocks depending on the level of search frictions, in the model and in the data. In the data, we examine price adjustments to nominal exchange rate shocks, using French exports to the UK as a case study. The data and econometric model are explained in Appendix A.3. The results of this estimation are reported in the first two columns of Table 2. In the first column, we see that the average pass-through of the pound/euro exchange rate into French export prices is equal to 89%. The elasticity is in line with estimates in Berman et al. (2012) who find an average exchange rate pass-through of 92% in French firm-level data. In the second column, the exchange rate is interacted with estimated relative meeting rates. The interaction is positive and significant, suggesting that French exporters are able to pass a larger share of exchange rate movements in markets characterized by relatively low search frictions. Quantitatively, the pass-through

Figure 8: Model fit on targeted moments

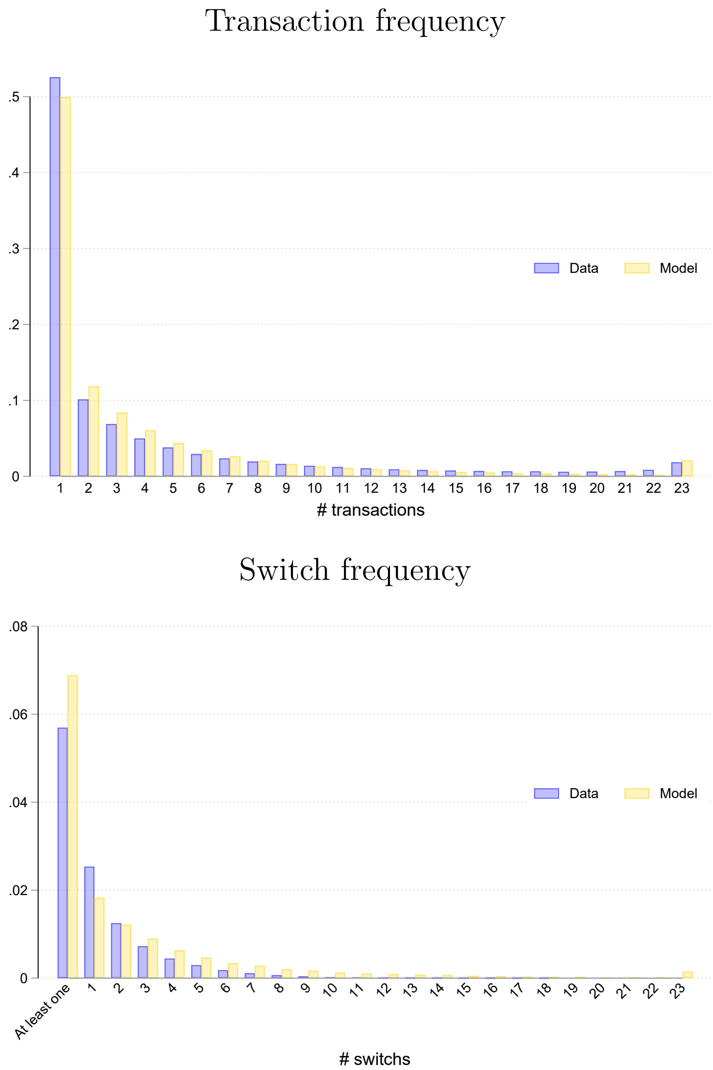

Note: The figures show moments in the data (blue bars) and estimated model (yellow bars). The top panel describes the number of transactions per importer, over the two-year observation period. The second panel shows the probability of at least one switch (first bars) and the probability of exactly one to 23 switches.

Figure 9: Model fit: price dynamics within relationships

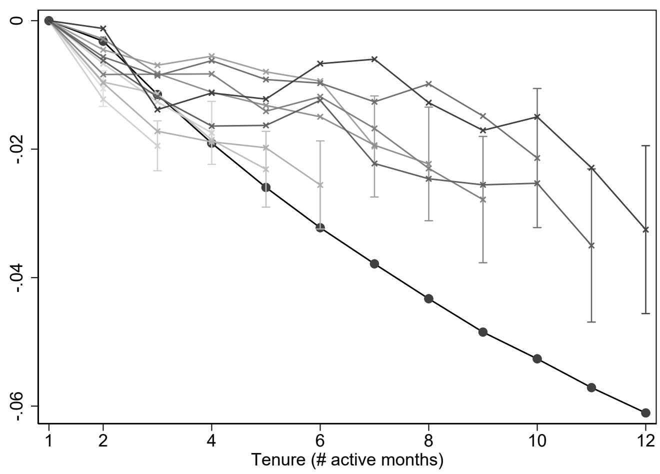

Note: This figure shows the evolution of prices within a firm-to-firm relationship in our data and in the calibrated model. Coefficients obtained from the data are recovered from equation (1) estimated by tenure.[32] Tenure is measured by the number of periods since the beginning of the relationship. A similar equation is estimated using simulated data. The figure reports the estimates and their 95% confidence intervals.

increases from 83 to 92% when moving from the first to the last decile of the distribution.

We then compare these estimates to predictions recovered from the simulated data. To this aim, we simulate a 10% shock in the relative competitiveness of British firms relative to their foreign competitors (including French sellers).[33] We then estimate the pass-through and its heterogeneity across markets in these simulated data. The results are reported in columns (3) and (4) of Table 2. The average pass-through is 75%, in the same order of magnitude as the one estimated in the data. In column (4), the interaction of the shock with the relative meeting rate faced by French exporters is positive and significant. The model thus reproduces the heterogeneity in pass-through rates along the distribution of relative meeting rates. Going from the first to the top decile in terms of meeting rates increases pass-through by 7pp, from 72% to 79%, to be compared to the 9pp increase estimated in the data.

Overall, these results confirm that the dynamics of prices, which the structural estimation does not use as an identifying moment, is well-reproduced by the estimated model, both qualitatively and quantitatively.

5 Search frictions and the incidence of cost shocks

We use the estimated model to quantify the incidence of relative cost shocks by comparing adjustments in firm-to-firm relationships in a counterfactual scenario, where French exporters

[33] As shocks to the Euro/Pound exchange rate affect all European exporters the same way, we simulate a cost shock affecting British sellers, with a depreciation of the pound being equivalent in the simulated data to an improvement in the cost-competitiveness of British sellers.

Table 2: Pass-through rates, in the model and the data

	Dependent variable: log p			
	Data		Simulated model	
	(1)	(2)	(3)	(4)
log cost shock	0.893***	0.830***	0.745***	0.704***
	(0.014)	(0.030)	(0.001)	(0.003)
- × Relative meeting rate		0.164**		0.080***
		(0.064)		(0.006)
FE	sji	sji	sji	sji
Obs.	1,527,909	1,527,909	178,688	178,688

Notes: Robust standard errors in parenthesis. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$. Price adjustments are computed on impact in simulated data and on a quarterly basis in actual data. The cost shock in actual data is measured using the euro/gbp exchange rate.

face a unilateral 10% increase in production costs, relative to the steady state. The incidence reflects both the pass-through within continuing relationships and the cost shifts for buyers who switch away from French suppliers following the shock. We examine, in sequence, the pass-through, the prevalence of switching, and the total incidence of the shock.

Pass-through. Within continuing relationship, we can distinguish three degrees of pass-through in response to the 10% increase in the cost of French sellers: (i) nil if the second-best supplier is not French and its cost is below the firm's monopoly price, (ii) complete if the second-best option is also French or the firm prices at its monopoly markup, before and after the shock, or (iii) incomplete if the shock either changes the identity of the second-best supplier, from French to non-French, or forces the firm to switch from the monopoly price to the quality-adjusted cost of the non-French second-best supplier. Figure 10 shows the frequency of complete, incomplete, and zero pass-through rates, on impact and at various time horizons.

In our quantification, zero pass-through is the most likely outcome at the firm-level in the short-run (49% of surviving relationships), but 37% of surviving relationships display full pass-through. Averaging across relationships gives a pass-through rate of 38%. Over time, firms are increasingly likely to compete with non-French suppliers and full pass-through becomes less frequent. For this reason, the average pass-through rate, conditional on the relationship surviving, tends to decrease, from 38% on impact to 29% after two years. As discussed bellow, the pass-through and the incidence varies across markets depending on the level of search frictions.

This shows how a given shock leads to highly heterogeneous pass-through across firm-to-firm relationships. The average pass-through then depends on the prevalence of full, nil, and incomplete pass-through. The average pass-through we estimate is in the same ballpark as Burstein and Gopinath (2014) who estimate a 20% pass-through of exchange rate shocks into US import prices or Cavallo et al. (2021) who estimate a 33% pass-through of foreign retaliatory tariffs on US export prices during the Trump trade war. However, reducing the level of frictions or changing the nature of the shock generates higher degress of pass-through.[34]

[34] We find a much higher pass-through rate when estimating the impact of bilateral exchange rate movements in Table 2. The reason is that simulations in this section are for a shock to French suppliers relative to all their competitors whereas the change in exchange rate is not specific to French suppliers and thus only affect their

Figure 10: Frequency of pass-through rates, conditional on survival

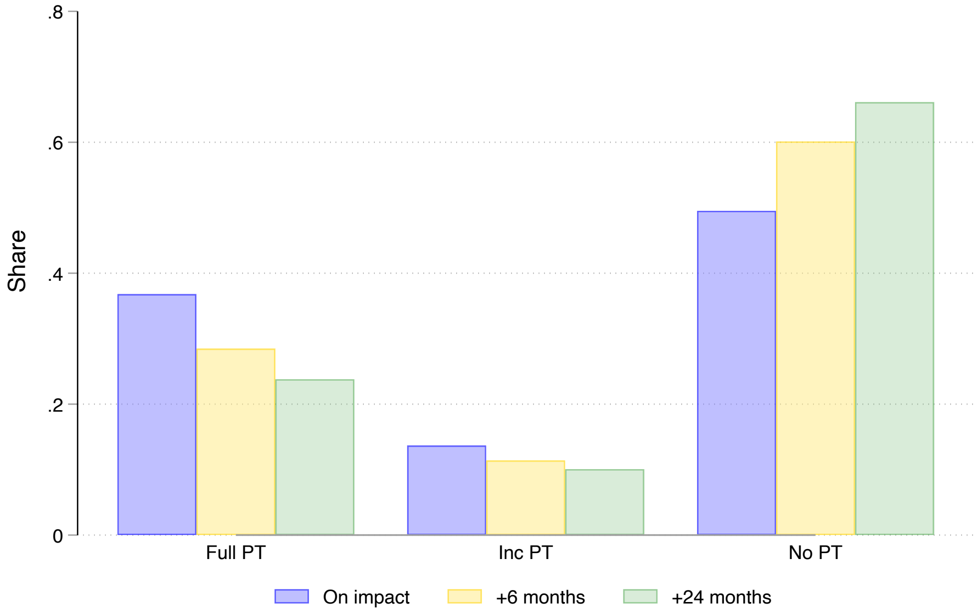

Note: The figure shows the frequency of complete, incomplete and zero pass-through rates, among the subsample of surviving firm-to-firm relationships. The frequencies are computed at the time of the shock ("On Impact") as well as six, and twenty four months after the shock.

Switches. Importers also adjust to input cost shocks by changing the sourcing of their inputs (see, e.g., Gopinath and Neiman, 2014, Devereux et al., forthcoming, Flaaen et al., 2020). In our model, changes along the extensive margin occur when the shock modifies the ranking of suppliers in the buyer's network. This drives instantaneous switches, away from France. As shown in Figure D14 in the appendix, the 10% cost induces a separation rate of 27% on average. These extensive margin adjustments reduce buyers' exposure to the cost shock. However, they are still inflationary and as such contribute to the overall incidence of the shock. Among the buyers that switch instantaneously, the price increase averages 2.6%, a 26% incidence rate.

As the relative competitiveness of French firms is permanently lowered, the rate at which buyers switch away from French firms also increases in the long run. After two years, the frequency of switches directly attributable to the shock is still substantial, at 8%. The frequency of switches varies along the distribution of relative meeting rates as buyers are more likely to know non-French competitive suppliers in markets in which the relative meeting rate of French firms is low.

Overall incidence. Combining these elements, we can quantify the incidence of a 10% shock and dissect the channels of the adjustment. We decompose the average evolution of prices, among the population of French firms' downstream partners as follows:

competitiveness vis-à-vis British producers. An almost full pass-through of Trump tariffs into US import prices from China as for instance been found by Fajgelbaum et al. (2020) or Cavallo et al. (2021).

$$d \ln \tilde{P}_t^{shock} \equiv \frac{1}{N} \left(\sum_b d \ln P_{bt}^{shock} - d \ln P_{bt}^{no\ shock} \right) = w_t^S \sum_{b \in S_t} d \ln \tilde{P}_{bt}^{shock} + (1 - w_t^S) \sum_{b \notin S_t} d \ln \tilde{P}_{bt}^{shock}$$

where $d \ln P_{bt}^{shock}$ denotes the growth of the price paid by buyer b, between period 0 (i.e. just before the shock) and period t. $d \ln P_{bt}^{no\ shock}$ is the baseline scenario calculated from a panel simulated on the same population of firms, using the same structural parameters, but without a shock. N is the number of buyers matched with French firms at the time of the shock. $d \ln \tilde{P}_t^{shock}$ is the average price variation attributable to the shock, in the population of buyers matched with French firms at the time of the shock. The weight w_t^S denotes the share of buyers that have switched away from French firms at date t, and $d \ln \tilde{P}_{bt}^{shock}$ denotes the evolution of prices for buyer b, normalized by the no shock counterfactual. The incidence is the increase in the cost of buyers trading with French sellers on impact.

The results of this decomposition are reported in Figure 11. For the average market, the price paid by foreign buyers increases by 3.5% on impact, a 35% incidence rate. About 75% of the effect is driven by pass-through within surviving relationships. The rest is explained by switches away from French suppliers. The overall incidence rate decreases over time because the prevalence of full pass-through drops and switches are directed to increasingly good suppliers. But after two years, the average price paid by buyers exposed to French firms at the time of the shock is still 1.4% higher than in the no-shock counterfactual.

We also see than the incidence varies across markets depending on the level of search frictions. On impact, moving from the first to the fourth quartile of the distribution of relative meeting rates (from high to low relative frictions) increases the average incidence by 65% from 24% to 39%. The reason is that the portfolio of potential suppliers of buyers located in markets in which French suppliers enjoy relatively low frictions is biased towards other French firms. As a consequence, French suppliers can transmit a higher fraction of the shock while retaining their buyers. In markets in which buyers are less likely to meet French suppliers, incidence is lower on impact since buyers can switch or threaten to switch to non-French suppliers. The difference, however, attenuates over time.

This exercise shows that switches are quantitatively important and deserve attention when looking at the incidence of a cost shock. Furthermore, the incidence of a shock leads to substantial variation across markets depending on the level of frictions.

6 Conclusion

We build and estimate a Ricardian model of firm-to-firm trade in frictional input markets. The model delivers empirically-consistent predictions on both the cross-section and the evolution of firm-to-firm trade networks. Assuming Bertrand competition within each buyer's random set of potential suppliers, we recover rich predictions on the distribution of markups across sellers and their downstream partners, and over time. The model's analytical tractability allows us to

Figure 11: Incidence of the shock, over time

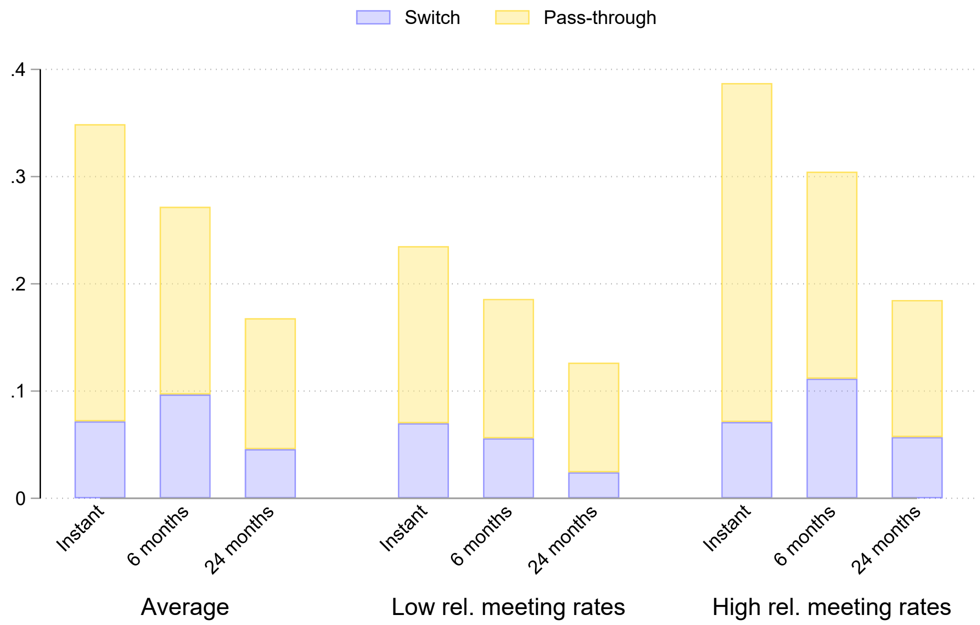

Note: The figure shows the elasticity of foreign buyers' prices to the shock on their supplier's relative cost. Incidence is measured on the whole population of buyers that are matched with a French firm at the time of the shock. The "Instant" effect is calculated at the moment of the shock, and is entirely driven by adjustments over the buyer's existing network. The "6 months" and "24 months" bars are instead calculated 6 and 24 periods after the shock, respectively. All results are in relative terms with the steady state dynamic of prices in the absence of a shock. The first three bars correspond to the average market. The next six bars compare incidence rates calculated at the first and third quartiles of the distribution of estimated relative meeting rates. "Switch" is the growth of prices calculated on the subset of firms that have switched away from French firms. "Pass-through" is the average evolution of prices among surviving relationships.

estimate the parameters driving comparative advantages and search frictions, separately for 330 (country×sector) markets. Identification of search parameters is achieved using the mobility of firms along the distribution of potential suppliers of the same input, which we observe in the time-series of firm-to-firm trade data.

The model implies novel predictions regarding the pass-through of cost shocks into import prices. First, the degree of pass-through varies between zero and one in equilibrium, and depends on the strength of competition characterizing each firm-to-firm relationship. When the shock shifts the price of a firm, relative to direct competitors in contact with the downstream partner, a low pass-through rate is needed for the firm to retain its buyer. Instead, shocks that affect a large share of buyers' network display high pass-through rates. For this reason, the average pass-through rate varies across markets. The pass-through of French sellers is higher where France faces lower search frictions, a prediction that we confirm in the data. Third, the average pass-through rate decreases over time, as new matches increase the buyer's outside option. The incidence of the shock on foreign buyers is further dampened by endogenous switches away from sellers which cost has increased. Overall, the estimated model implies an average incidence rate of unilateral shocks to French production costs of 35% on impact and 26% after two years.

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 or change their prices immediately after the shock, which would accentuate the persistence of the 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.

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.

Antràs, Pol, Teresa C. Fort, and Felix Tintelnot, "The Margins of Global Sourcing: Theory and Evidence from US Firms," *American Economic Review*, September 2017, *107* (9), 2514–2564.

Arkolakis, Costas, Federico Huneeus, and Yuhei Miyauchi, "Spatial Production Networks," Working Paper 30954, National Bureau of Economic Research February 2023.

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 Fontaine, 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.

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.

Bernard, Andrew B. and Andreas Moxnes, "Networks and Trade," *Annual Review of Economics*, August 2018, *10* (1), 65–85.

_ and Yuan Zi, "Sparse Production Networks," NBER Working Papers 30496, National Bureau of Economic Research, Inc September 2022.

_ , Andreas Moxnes, and Karen Helene Ulltveit-Moe, "Two-Sided Heterogeneity and Trade," *The Review of Economics and Statistics*, July 2018, *100* (3), 424–439.

_ , _ , and Yukiko U. Saito, "Production Networks, Geography, and Firm Performance," *Journal of Political Economy*, 2019, *127* (2), 639–688.

—, Emmanuel Dhyne, Glenn Magerman, Kalina Manova, and Andreas Moxnes, "The Origins of Firm Heterogeneity: A Production Network Approach," *Journal of Political Economy*, 2022, *130* (7), 1765–1804.

—, 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.

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.

—, Javier Cravino, and Marco Rojas, "Input price dispersion across buyers and misallocation," Working Papers Central Bank of Chile 1006, Central Bank of Chile January 2024.

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.

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.

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.

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 Yi Xu, and Kelly Kaili Yang, "O-Ring Production Networks," *Journal of Political Economy*, 2024, *132* (1), 200–247.

Devereux, Michael B, Wei Dong, and Ben Tomlin, "Trade Flows and Exchange Rates: Importers, Exporters and Products," *Journal of International Economics*, forthcoming.

Dhyne, Emmanuel, Ayumu Ken Kikkawa, and Glenn Magerman, "Imperfect Competition in Firm-to-Firm Trade," *Journal of the European Economic Association*, 03 2022, *20* (5), 1933–1970.

—, —, Xianglong Kong, Magne Mogstad, and Felix Tintelnot, "Endogenous production networks with fixed costs," *Journal of International Economics*, 2023, *145*, 103841.

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,” *Econometrica*, forthcoming.

Fajgelbaum, Pablo D, Pinelopi K Goldberg, Patrick J Kennedy, and Amit K Khandelwal, “The Return to Protectionism,” *The Quarterly Journal of Economics*, 2020, *135* (1), 1–55.

Fitzgerald, Doireann and Stefanie Haller, “Pricing-to-Market: Evidence From Plant-Level Prices,” *Review of Economic Studies*, 2014, *81* (2), 761–786.

—, —, and Yaniv Yedid-Levi, “How Exporters Grow,” *The Review of Economic Studies*, 07 2023, p. rdad070.

Flaaen, Aaron, Ali Hortaçsu, and Felix Tintelnot, “The Production Relocation and Price Effects of US Trade Policy: The Case of Washing Machines,” *American Economic Review*, July 2020, *110* (7), 2103–27.

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.

Gopinath, Gita and Brent Neiman, “Trade Adjustment and Productivity in Large Crises,” *American Economic Review*, March 2014, *104* (3), 793–831.

Gourieroux, Christian, Alain Monfort, and Eric Renault, “Indirect Inference,” *Journal of Applied Econometrics*, 1993, *8*, 85–118.

Grossman, Gene M., Elhanan Helpman, and Stephen J. Redding, “When Tariffs Disrupt Global Supply Chains,” *American Economic Review*, April 2024, *114* (4), 988–1029.

Heise, Sebastian, “Firm-to-Firm Relationships and the Pass-Through of Shocks: Theory and Evidence,” *The Review of Economics and Statistics*, 09 2024, pp. 1–45.

Huneeus, Federico, "Production network dynamics and the propagation of shocks," *Graduate thesis, Princeton University, Princeton, NJ*, 2018, *52.*

Jolivet, Gregory, Fabien Postel-Vinay, and Jean-Marc Robin, "The empirical content of the job search model: Labor mobility and wage distributions in Europe and the US," *European Economic Review*, May 2006, *50* (4), 877–907.

Lenoir, Clemence, Julien Martin, and Isabelle Mejean, "Search Frictions in International Goods Markets," *Journal of the European Economic Association*, 2022, *21*, 326–366.

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.

Lu, Dan, Asier Mariscal, and Luis-Fernando Mejía, "How firms accumulate inputs: Evidence from import switching," *Journal of International Economics*, 2024, *148*, 103847.

Miyauchi, Yuhei, "Matching and Agglomeration: Theory and Evidence from Japanese Firm-to-Firm Trade," *Econometrica*, forthcoming.

Monarch, Ryan, "It's Not You, It's Me: Prices, Quality, and Switching in U.S.-China Trade Relationships," *The Review of Economics and Statistics*, December 2022, *104* (5), 909–928.

_ and Tim Schmidt-Eisenlohr, "Longevity and the value of trade relationships," *Journal of International Economics*, 2023, *145* (C).

Postel-Vinay, Fabien and Jean-Marc Robin, "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity," *Econometrica*, November 2002, *70* (6), 2295–2350.

_ and _ , "Microeconometric search-matching models and matched employer-employee data," SciencePo Working papers Main hal-03587666, HAL 2006.

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

Rauch, James E and Vitor Trindade, "Ethnic Chinese networks in international trade," *Review of Economics and Statistics*, 2002, *84* (1), 116–130.

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.

Sugita, Yoichi, Kensuke Teshima, and Enrique Seira, "Assortative Matching of Exporters and Importers," *The Review of Economics and Statistics*, 11 2023, *105* (6), 1544–1561.

For Online Publication

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 newtork 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

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.

within a particular month and over a particular product.[35] This number decreases somewhat, to 75%, when we do not condition over a particular product (upper-right panel), which means 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.

[35] As 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.

Figure A1: Cumulated distribution of European importers' indegrees

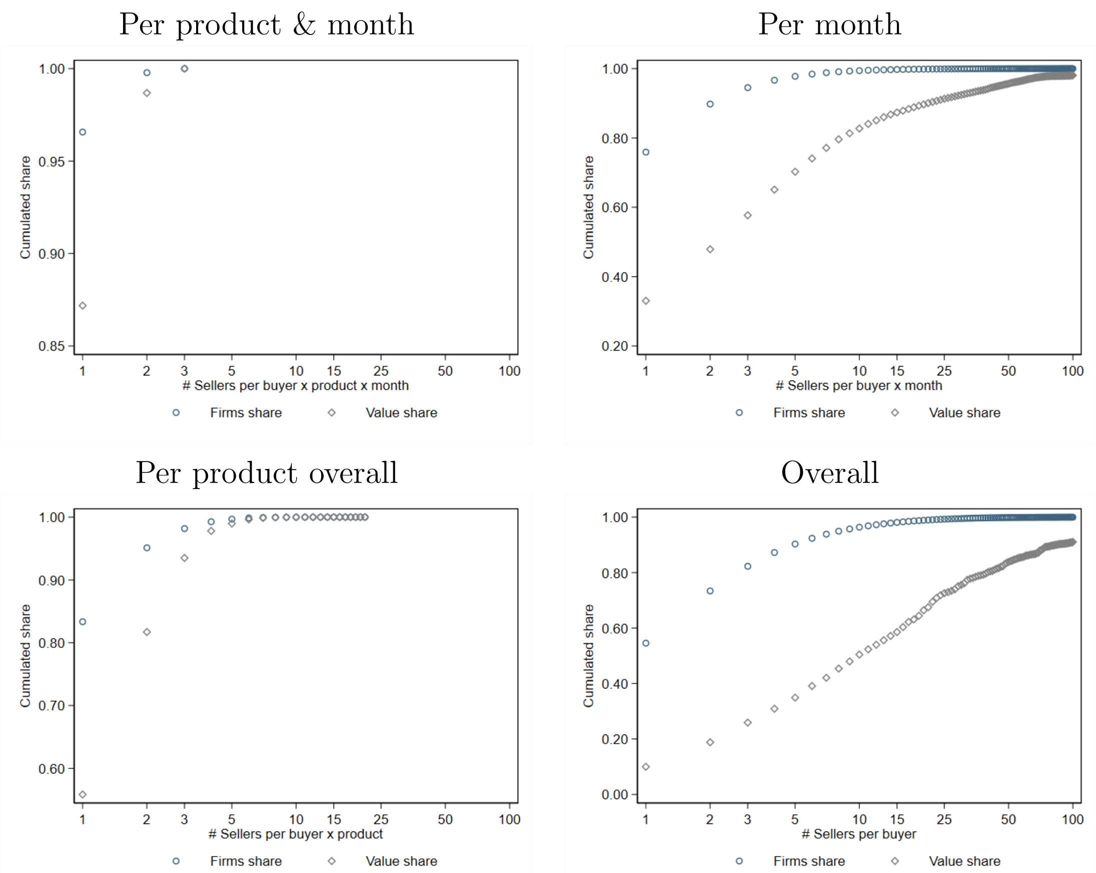

The figure illustrates the heterogeneity across European importers' in their "indegrees", i.e. their number of partners in France. The x-axis corresponds to a number of partners and the y-axis is the cumulated share of firms (blue circles) with x sellers or less, and their cumulated contribution to aggregate 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.

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' quality-adjusted productivity 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 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.

Figure A2: Cumulated distribution of French exporters' outdegrees

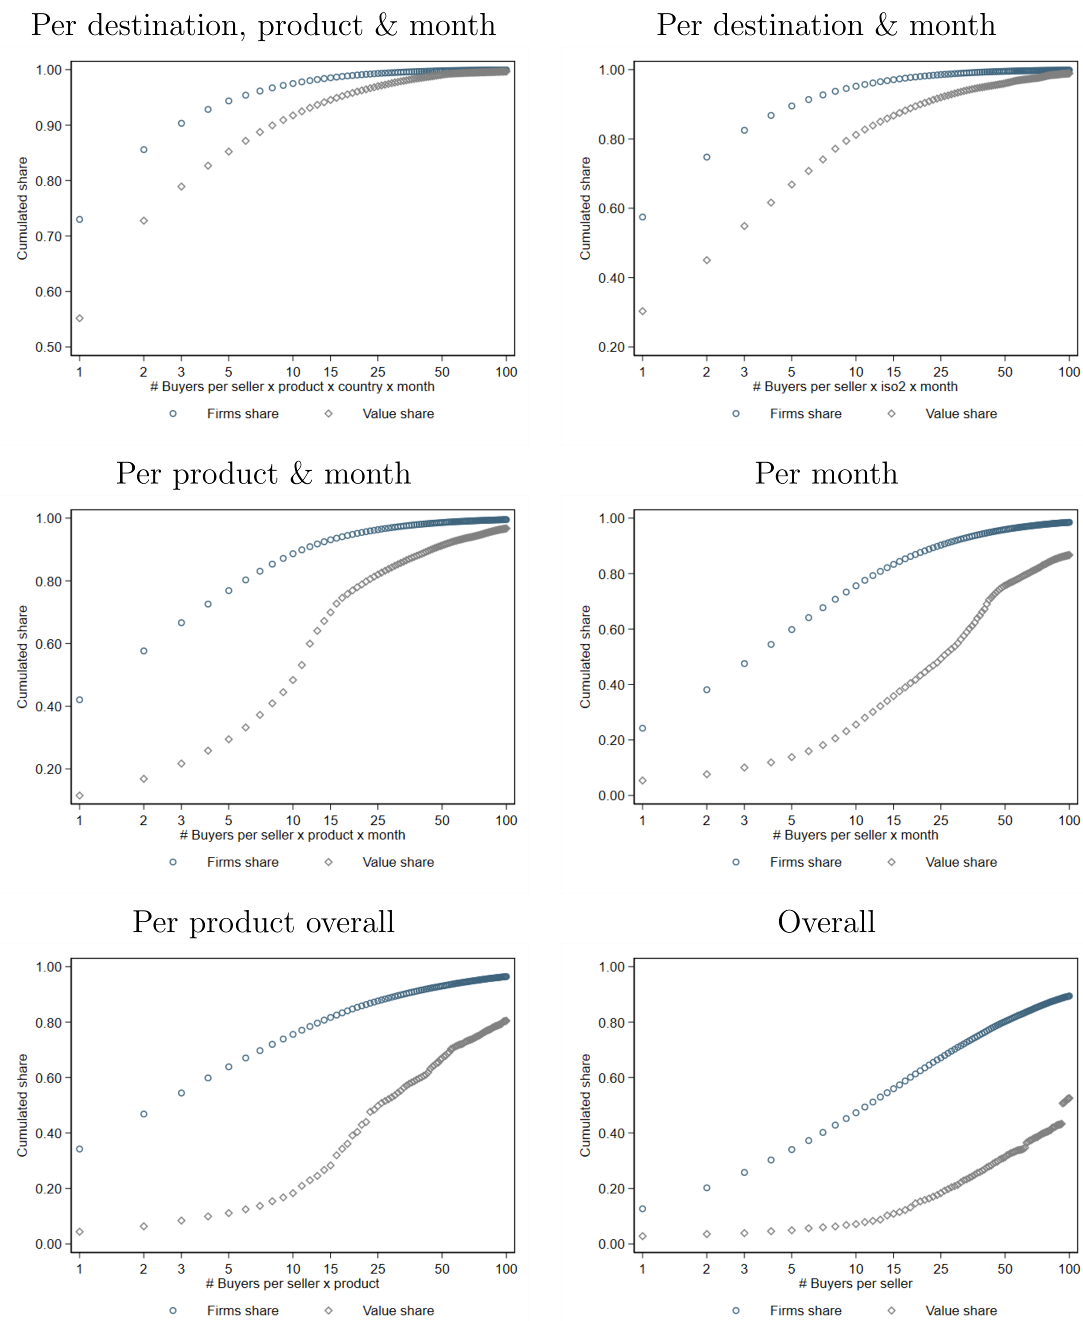

The figure illustrates the heterogeneity across French exporters' in their "outdegrees", i.e. their number of partners within the European Union. The x-axis corresponds to a number of partners and the y-axis is the cumulated share of firms (blue circles) with x buyers or less, and their cumulated contribution 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.

A.3 Estimation of pass-through rates

We seek to compare how pass-through varies across markets with different levels of frictions in the data and in our simulated model. In the data, we consider the pass-through of GBP/EUR nominal exchange rate into French export prices to the UK.

French Customs data. Information on prices comes from the same French Customs data, as in the rest of the paper. We however use a longer panel, which we aggregate at the quarterly level.Our proxy for prices is thus the quarterly unit value of all transactions involving a given French seller × EU buyer × 8-digit product. Product nomenclatures are harmonized over time using the concordance algorithm developed in Behrens et al. (2019).

Nominal exchange rate. We examine the impact of nominal exchange rate shocks on export prices. We use the quarterly series of the UK pound against the euro provided by the ECB (https://data.ecb.europa.eu/data/datasets/EXR/EXR.Q.GBP.EUR.SP00.A).

Estimation. We follow the trade literature (eg. Berman et al., 2012) by regressing the log of unit values on the log of the exchange rate. Unit values are converted into pounds so that we can interpret estimated coefficients in terms of pass-through rates. We further augment the equation with an interaction between the ER and the market-specific meeting rate estimated in Section 4. All specifications include firm×product×destination fixed effects so that identification is achieved using the variability over time. More specifically, we estimate:

$$\log(uv_{sbjit}) = \alpha \log(ER_t) + \beta \log(ER_t) \times Meeting_{ji} + FE_{sji} + \varepsilon_{sbjit} \tag{A1}$$

where s, b, j, i, t denote seller, buyer, product, destination, and year, respectively and $Meeting_{ji}$ is the relative meeting rates of French sellers in market ji.

Simulation. There is no perfect mapping between shocks in the model and nominal exchange rate shocks. In the model, shocks hit either French or non-French firms. A nominal exchange rate shock may instead capture a change in the value of the pound against all other currencies. The shock thus not only hits the French exporters but also their non-British competitors.[36] More generally, the shock hits all exporters exporting in the same currency as French firms. To see why the source of the shock matters in our model, consider the following example. A French exporter interacts with a British firm. The downstream firm's outside option is a German supplier. If the French firm is hit by a French-specific shock, its optimal markup needs to adjust to retain its partner. If instead the shock affects all European exporters as in the case when the pound depreciates, the French firm that competes with a German exporter can pass the shock onto its British partner.

[36] Here we abstract from the question of the invoicing currency, which would add a layer of complexity to the reasoning. The French customs do not collect the invoicing currency on export flows.

To build a shock that mimics a nominal shock on the euro/pound exchange rate (hitting French firms and some non-French sellers) we proceed as follows:

- simulate a 10% shock on French exporters
- in case the outside option is not French, assume this supplier is also hit by the shock with a certain probability
- whether the competitor is hit or not by the shock is determined as follows: make a draw from a standard uniform distribution, and assume the competitor is from the UK if the draw if lower than the share of UK firms in (non-French) consumption in the UK.

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.

The empirical moments include the switch and transaction frequencies, in each market. To construct these moments, we proceed as follows. The starting point is the estimation sample, which includes all firm-to-firm transactions involving French exporters and European importers over 2002-2006. The estimation strategy takes the point of view of individual buyers and we first reshape the dataset accordingly. More specifically, we keep all the importer×product pairs observed matched with a French firm at least once between 2002 and 2003. Starting from the first transaction in this time window, we retain all the subsequent transactions involving these buyers and French firms, over a maximum of two years.[37] For each buyer, we then compute the number of transactions, defined by the number of months in which a transaction with a French firm is recorded, and the number of switches, as defined by the number of times the buyer is involved into a transaction with a seller she was not interacting with the period before.[38] This leaves us with an estimation sample that contains, for each country×sector market, the number of active buyers, the number of their transactions and the number of their switches.

We then simulate comparable statistics using our model. 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 the probabilities needed to compute the likelihood, on simulated data. Namely

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

[37]Imposing a constant 2-year window permits to control for right-censoring. In the structural estimation, every buyer is observed for exactly two years.

[38]While the definition of a switch is straightforward when a buyer interacts with a single seller in each period, it is trickier for the minority of buyers that perform more than one transaction within the same month. In such case, a switch is recorded if at least one of the buyer's sellers was not interacting with the buyer at the previous period. We also tried with a stricter definition in which a switch requires that none of the sellers observed matched with a buyer at time t was in the supplier set of the firm at the previous period of activity.

where $n \in \{1, ..., 5\}$ ($n = 5$ means at least 5 transactions) and $s \in \{0, 1, 2\}$ ($s = 2$ means at least 2 switches). Since we reduce the state space, and thus approximate the likelihood, our estimation method falls within the class of indirect inference methods (Gourieroux et al., 1993).

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

i. For each dataset (country×sector), 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 December 2003 in the data. Then, we follow that buyer for 24 months recording any subsequent transaction or any switch towards French sellers.

iii. Using simulated data, we compute the frequencies $P(\text{ transactions } = n \cap \text{ switches} = \text{s})$, $\forall(n, s)$. 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 demand for final goods

The model in Section 3 takes the demand addressed to final good producers as given. In this section, we incoporate the partial equilibrium model into a more general structure in which the demand addressed to final good producers emanates from final consumers.

There is a representative consumer in the country, which preferences are CES across a continuum of differentiated varieties:

$$C = \left[\int_0^1 c(\omega)^{\frac{\sigma-1}{\sigma}} d\omega \right]^{\frac{\sigma}{\sigma-1}}$$

where σ is the elasticity of substitution across varieties and we have normalized the mass of differentiated varieties.

The representative consumer maximizes consumption under her budget constraint, given revenues recovered from her inelastic labor supply and residual profits $R = wL + \Pi$. In equilibrium, the demand addressed to each final good producer is:

$$c(\omega) = \left(\frac{p(\omega)}{P}\right)^{-\sigma} \frac{R}{P}$$

with $P = \left[\int_0^1 p(\omega)^{1-\sigma} d\omega\right]^{\frac{1}{1-\sigma}}$ the ideal price index.

In equilibrium, each variety is produced by a single final good producer. The quantity produced by a buyer b is then:

$$x_b = \left(\frac{p_b}{P}\right)^{-\sigma} \frac{R}{P}$$

where p_b is the price set by buyer b on the representative consumer, which maximizes profits:

$$\begin{cases} & \pi_b = (p_b - mc_b)x_b \\ s.t. & x_b = \left(\frac{p_b}{P}\right)^{-\sigma} \frac{R}{P} \end{cases}$$

In equilibrium, the price is defined as:

$$p_b = \frac{\sigma}{\sigma - 1} mc_b$$

with mc_b the marginal cost, defined in the main text.

C.2 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

$$F_i(c) = \frac{\gamma_{iF}}{\gamma_i} F_{iF}(c) + \frac{\gamma_{i\bar{F}}}{\gamma_i} F_{i\bar{F}}(c) \tag{A2}$$

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

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

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 simplying one gets

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

Then integrating by part

$$L_i(c) = \frac{\mu + \gamma_i}{\mu + \gamma_i F_i(c)} F_i(c) \tag{A3}$$

The distribution of buyers among French suppliers. Consider the share 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

$$\underbrace{(1-u_i)\pi_{iF}\ell_{iF}(c)\left(\mu + \gamma_i F_i(c)\right)}_{\text{outflows}} = \underbrace{u_i \gamma_{iF} f_{iF}(c) + (1-u_i)\bar{L}_i(c)\gamma_{iF} f_{iF}(c)}_{\text{inflows}} \tag{A4}$$

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

$$\pi_{iF}\ell_{iF}(c) = \frac{\gamma_{iF}}{\gamma_i} \frac{\mu(\mu+\gamma_i)}{\left(\mu + \gamma_i F_i(c)\right)^2} f_{iF}(c) \tag{A5}$$

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

$$(1-\pi_{iF})\ell_{i\bar{F}}(c) = \frac{\gamma_{i\bar{F}}}{\gamma_i} \frac{\mu(\mu+\gamma_i)}{\left(\mu + \gamma_i F_i(c)\right)^2} f_{i\bar{F}}(c) \tag{A6}$$

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

$$\begin{aligned}
&(1-u_i)\pi_{iF}\left(\mu + \gamma_{i\bar{F}} \int_0^{c_F^{max}} F_{i\bar{F}}(c)\ell_{iF}(c)dc\right) = u_i \gamma_{iF} \\
&+ (1-u_i)(1-\pi_{iF})\gamma_{iF}\left(\bar{L}_{i\bar{F}}(c^{max}) + \int_0^{c^{max}} F_{iF}(c)\ell_{i\bar{F}}(c)dc\right)
\end{aligned} \tag{A7}$$

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 (A5) and (A6),

$$(1-\pi_{iF})\tau_{iF}^{-\theta}\gamma_{iF}\ell_{i\bar{F}}(c) = \pi_{iF}\gamma_{i\bar{F}}\ell_{iF}(c) \tag{A8}$$

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

$$L_{i\bar{F}}(c_F^{max}) = \frac{\pi_{iF}}{(1-\pi_{iF})}\frac{\gamma_{i\bar{F}}}{\gamma_{iF}}\tau_{iF}^{\theta} \tag{A9}$$

which can be substituted in (A7) to obtain

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

The integrals cancel out and, after simplification, one gets

$$\pi_{iF} = \frac{\gamma_{iF}}{\gamma_{iF}+\gamma_{i\bar{F}}}\frac{\mu+\gamma_{iF}+\gamma_{i\bar{F}}}{\mu+\gamma_{iF}+\gamma_{i\bar{F}}\tau_{iF}^{\theta}} \tag{A10}$$

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

$$\begin{aligned}(1-u_i)\pi_{iF}\left(\mu + \gamma_{i\bar{F}}\bar{L}_{iF}(c_{\bar{F}}^{max}) + \gamma_{i\bar{F}}\int_0^{c_{\bar{F}}^{max}} F_{i\bar{F}}(c)\ell_{iF}(c)dc\right) &= u_i\gamma_{iF}\\ &+ (1-u_i)(1-\pi_{iF})\gamma_{iF}\int_0^{c_{\bar{F}}^{max}} F_{iF}(c)\ell_{i\bar{F}}(c)dc\end{aligned} \tag{A11}$$

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

$$(1-u_i)\pi_{iF}\left(\mu + \gamma_{i\bar{F}}\bar{L}_{iF}(c_{\bar{F}}^{max})\right) = u_i\gamma_{iF} \tag{A12}$$

We get an expression for $\bar{L}_{iF}(c_{\bar{F}}^{max})$ by integrating (A8) up to $c_{\bar{F}}^{max}$,

$$L_{iF}(c_{\bar{F}}^{max}) = \frac{1-\pi_{iF}}{\pi_{iF}}\frac{\gamma_{iF}}{\gamma_{i\bar{F}}}\tau_{iF}^{-\theta} \tag{A13}$$

Finally, using that expression, we obtain

$$\pi_{iF} = \frac{\gamma_{iF}}{\gamma_{iF}+\gamma_{i\bar{F}}}\frac{\mu+(\gamma_{iF}+\gamma_{i\bar{F}})\tau_{iF}^{-\theta}}{\mu+\gamma_{iF}\tau_{iF}^{-\theta}+\gamma_{i\bar{F}}} \tag{A14}$$

The share of French sellers when $\mu \approx 0$. Interestingly, when μ is close to zero, (A10) and (A14) are approximately equal

$$\pi_{iF} = \frac{\gamma_{iF}}{\gamma_{iF} + \gamma_{i\bar{F}}\tau_{iF}^{\theta}} \tag{A15}$$

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

$$p_j x_j = \alpha_b \left(\frac{p_j}{q_j}\right)^{1-\eta} \left(\int_{j' \in \Omega} \left(\frac{p_{j'}}{q_{j'}}\right)^{1-\eta} dj'\right)^{\frac{\eta}{1-\eta}} \tag{A16}$$

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

$$\mathbb{E}\left[\frac{p}{q}|F\right] = \int_0^c \int_0^{+\infty} \tilde{c}\ell(c)\ell_{iF}(\tilde{c}|c)d\tilde{c}dc \tag{A17}$$

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

$$(1-u_i)\pi_{iF}\bar{L}_{iF}(\tilde{c}|c)\left(\mu + \gamma_i F_i(\tilde{c})\right) = \left(u_i + (1-u_i)\bar{L}_i(\tilde{c})\right)\gamma_{iF} f_{iF}(c) \tag{A18}$$

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

$$(1-u_i)(1-\pi_{iF})\bar{L}_{i\bar{F}}(\tilde{c}|c)\left(\mu + \gamma_i F_i(\tilde{c})\right) = \left(u_i + (1-u_i)\bar{L}_i(\tilde{c})\right)\gamma_{i\bar{F}} f_{i\bar{F}}(c) \tag{A19}$$

Note that

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

Hence $L_{iF}(\tilde{c}|c) = L_{i\bar{F}}(\tilde{c}|c) = L_i(\tilde{c}|c)$ and $\mathbb{E}\left[p/q|F\right] = \mathbb{E}\left[p/q|\bar{F}\right]$. For that reason, the expected quantity does not depend on the country of origin of the supplier and the share of French products in import portfolios map with the share of French suppliers in the buyer's portfolio.

C.3 Buyer acquisition on the seller's side

Our model has predictions for the acquisition of buyers, conditional on an input suppplier's type. Although these predictions are not key for identification, we now show that the model qualitatively replicates the evidence.

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. The dynamic of $n_i(c)$ thus follows

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

with $n_{i0} = 0$. Hence, given equation (8), the expected number of buyers in period t has a simple solution

$$n_{it}(c) \quad = \tfrac{\lambda_{iF}\mu}{(\mu+\gamma_i F_i(c))^2}\left(1 - e^{-(\mu+\gamma_i F_i(c))t}\right) \tag{A20}$$

The expected number of buyers is thus increasing over time as it converges toward a steady state.[39]

$$n_i(c) = \frac{\lambda_{iF}\mu}{(\mu + \gamma_i F_i(c))^2} \tag{A21}$$

This implies that the number of buyers served is higher for sellers with lower quality-adjusted 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.3)). 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 D7 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.

[39] At 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

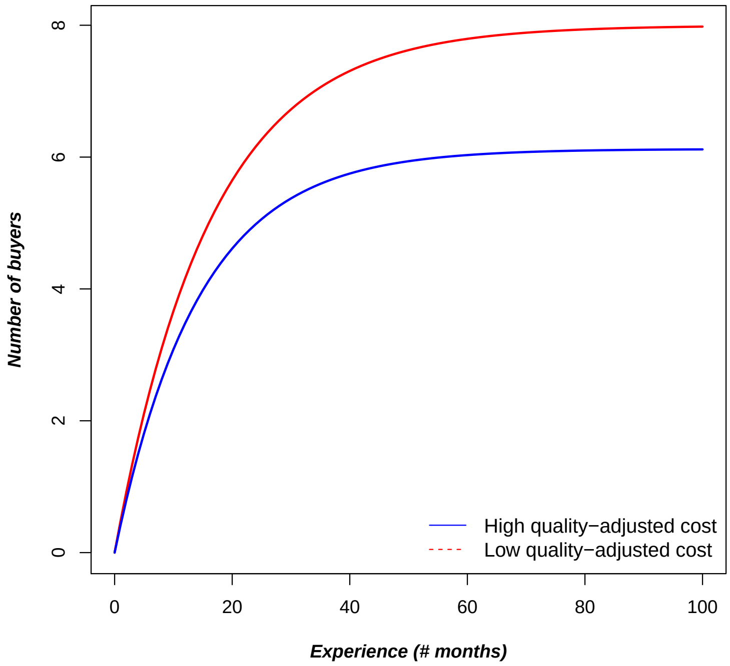

Notes: The figure shows the simulated expected number of buyers in a seller's portfolio as a function of its experience in the market. The blue (resp. red) line is for a high (resp. a low) quality-adjusted cost seller.

Figure D1: Hazard rate over time, Alternative tenure definition

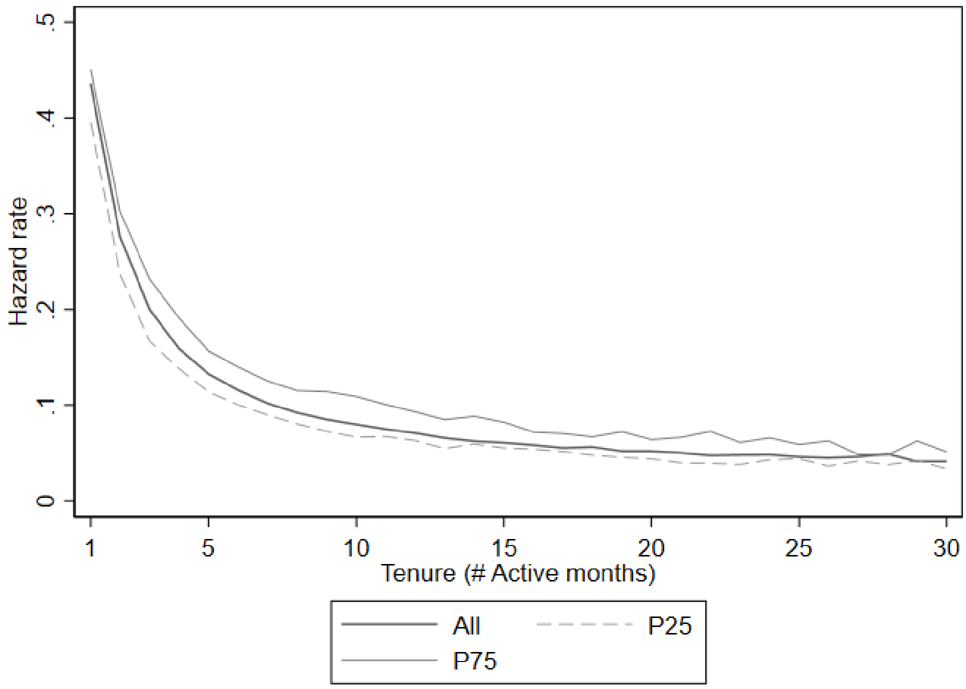

Notes: The hazard rate is defined as the probability of the relationship ending, conditional on tenure into the relationship and is calculated as the ratio of the density to the survival rate at tenure k. The figure is recovered from the 2002-2006 sample using the cumulated number of transactions since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

D Additional results

Figure D2: Hazard rate, over time, by country within the wearing apparel sector

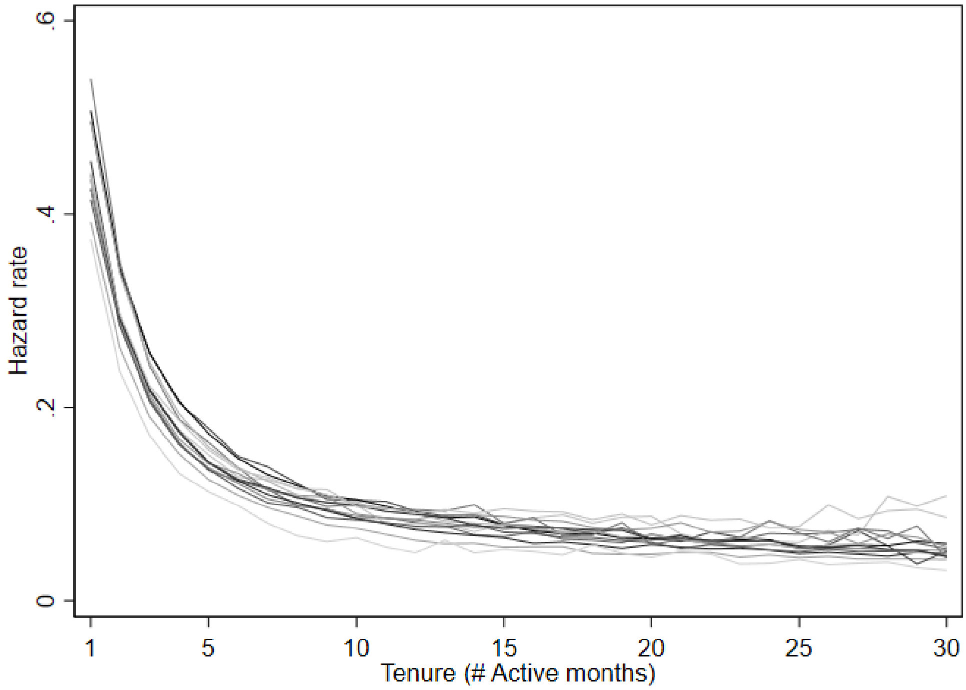

Notes: The hazard rate is defined as the probability of the relationship ending, conditional on tenure in the relationship and is calculated as the ratio of the density to the survival rate at tenure k. The figure is recovered from the 2002-2006 sample and the sector of wearing apparel. Tenure is defined based on the cumulated number of periods in the relationship as measure of tenure. Each line correspond to a destination country.

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

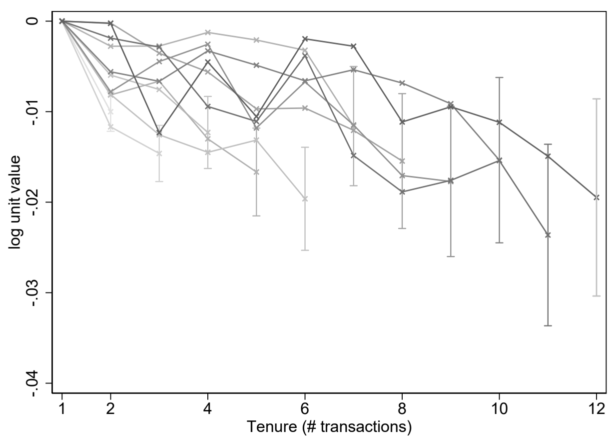

Note: These figures show the evolution of prices within a firm-to-firm relationship. Coefficients are recovered from equation (1), with tenure defined in terms of the number of transactions since the beginning of the relationship. The figure reports the estimates and their 95% confidence intervals.

Figure D4: Price dynamics within a firm-to-firm relationship, Controlling for seller-specific shocks

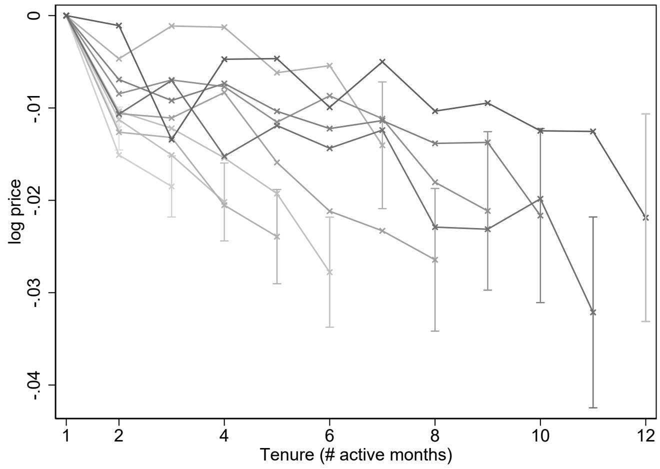

Note: This figure shows the evolution of exported prices within a firm-to-firm relationship. Coefficients are recovered from equation (1), using relationship-specific fixed effects as well as seller×product×period fixed effects, to control for firm-specific cost shocks. The relationship is estimated for increasingly long relationships from 2 months (lightest grey) to 12 months or more (darkest grey). 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: Quantity dynamics within a firm-to-firm relationship, by tenure

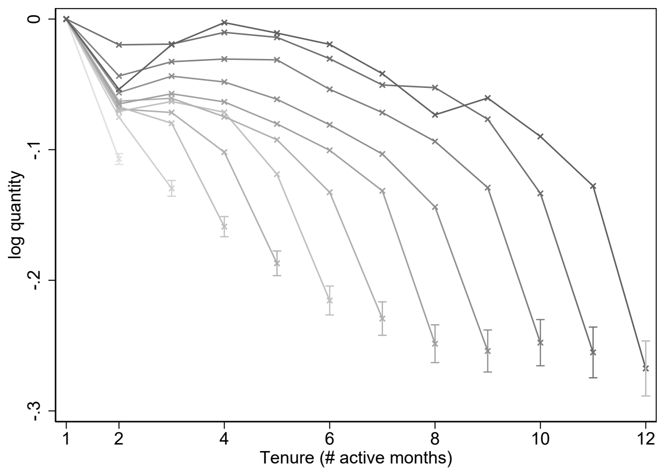

Note: This figure shows the evolution of exported quantities within a firm-to-firm relationship. Coefficients are recovered from equation (1), controlling for composition effects using overall tenure dummies from 2 months (lightest grey) to 12 months or more (darkest grey). The figure reports the estimates and their 95% confidence periods. Tenure is measured by the number of months since the beginning of the relationship.

Figure D6: Price dynamics: Comparison with Fitzgerald and Haller (2014)

Identification across destinations
(Fitzgerald and Haller, 2014)

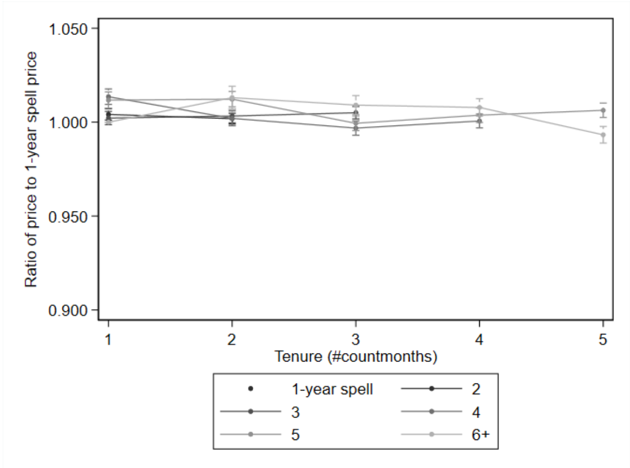

Identification over time (us)

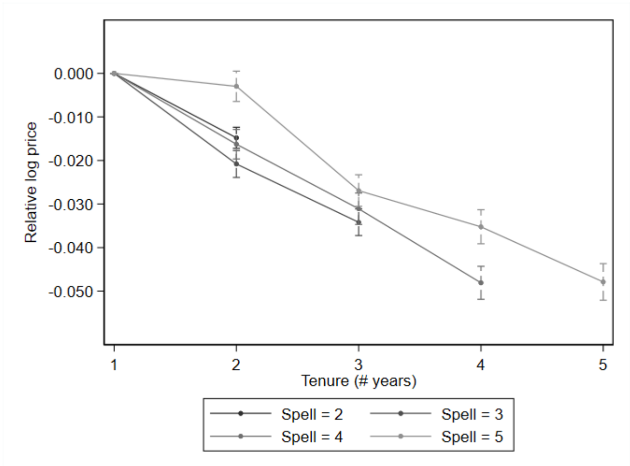

Notes: The figure compares the dynamics of prices as measured in data aggregated at the firm×country×year level, using either the specification in Fitzgerald and Haller (2014) or a specification that exploits the time dimension, within a relationship, as we do in Figure 4. Fitzgerald and Haller (2014) specification is as follows

$$\begin{aligned} \ln p_{ijst} &= FE_{sjt} + FE_{ijt} + \beta X_{ijst} + \sum_{d=2}^{T} \delta_d \mathbb{1}(Tenure_{ijst} = d) \\ &+ \sum_{k=2}^{6} \sum_{d=2}^{t} \gamma_{kd} \mathbb{1}(Spell_{ijs} = k) \mathbb{1}(Tenure_{ijst} = d) + \varepsilon_{ijst} \end{aligned}$$

where i, j, s, t respectively denote a destination, product, exporter and year. $\mathbb{1}(Tenure_{ijst} = d)$ is a dummy equal to one when the tenure into the relationship is d and $\mathbb{1}(Spell_{ijs} = k)$ is a dummy identifying relationships of overall spell k. X_{ijst} is a set of controls that contains dummies for left and right censoring. The top panel shows the estimated γ_{kd} coefficients. The specification in the bottom panel is inspired from the specification in equation (1), which uses the variation over time, within a relationship. Instead of having FE_{sjt} fixed effects, we thus use the equation above with FE_{isj} fixed effects. Again, the figure shows the estimated γ_{kd} coefficients.

Figure D7: Acquisition of buyers over time

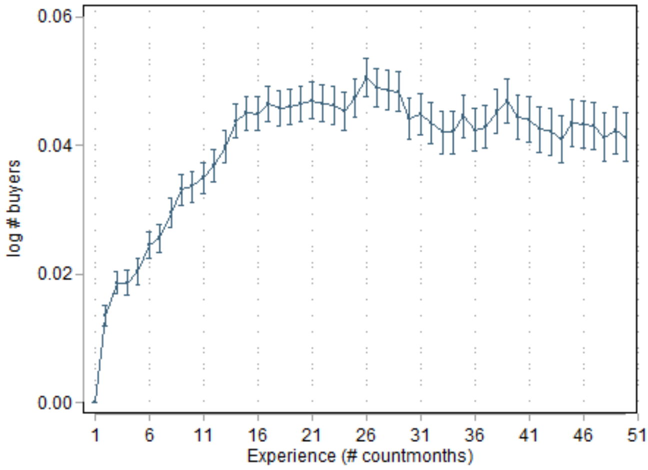

Note: The figure shows the evolution of a seller's stock of buyers, over time, recovered from equation (2).The figure reports the estimates and their 95% confidence intervals. Experience is measured by the number of months since first entry.

Figure D8: Densities for the estimated parameters

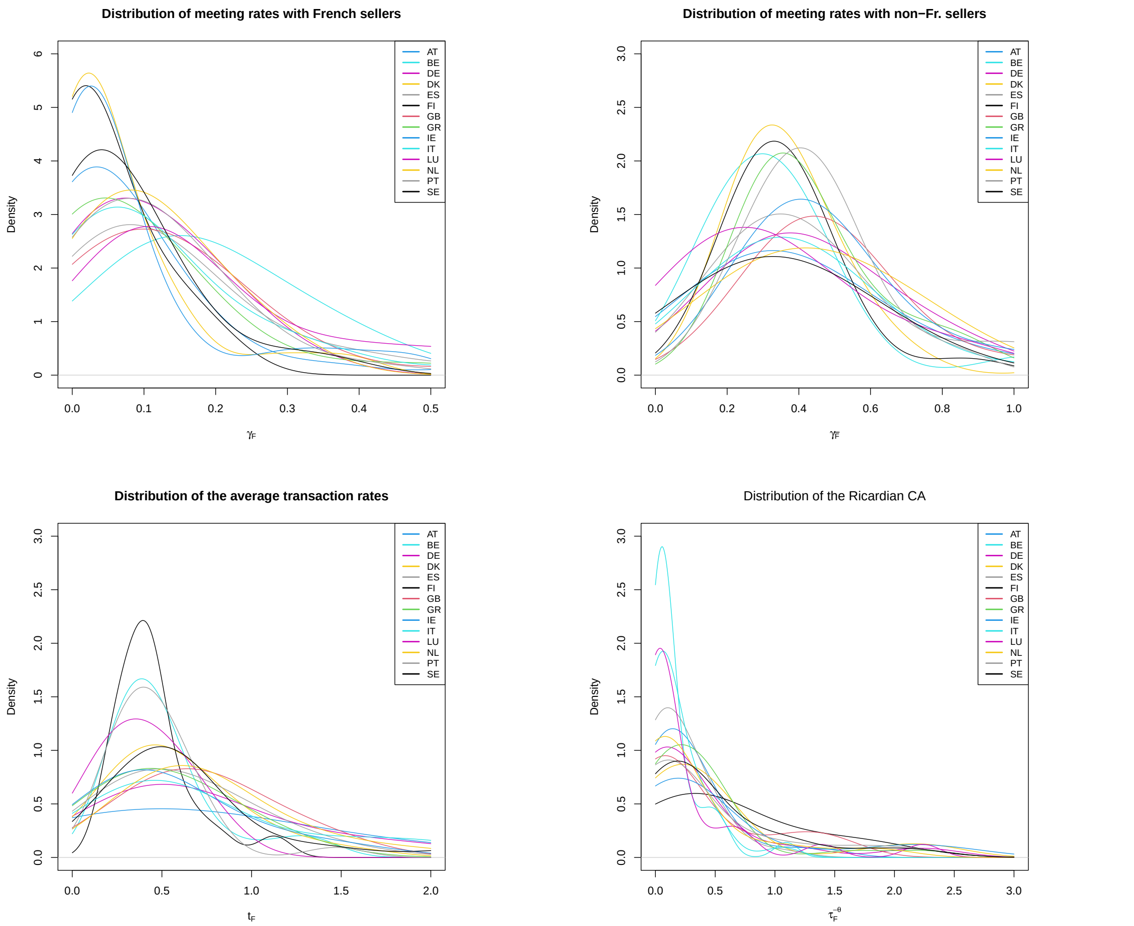

Note: The figure shows the country-specific distributions of estimated parameters. The figure is restricted to sector×country pairs for which we observe at least 100 buyers and 2 switches.

Figure D9: Dispersion in relative meeting rates ($\gamma_{iF}/\gamma_{i\bar{F}}$), across countries and sectors

(a) Across countries

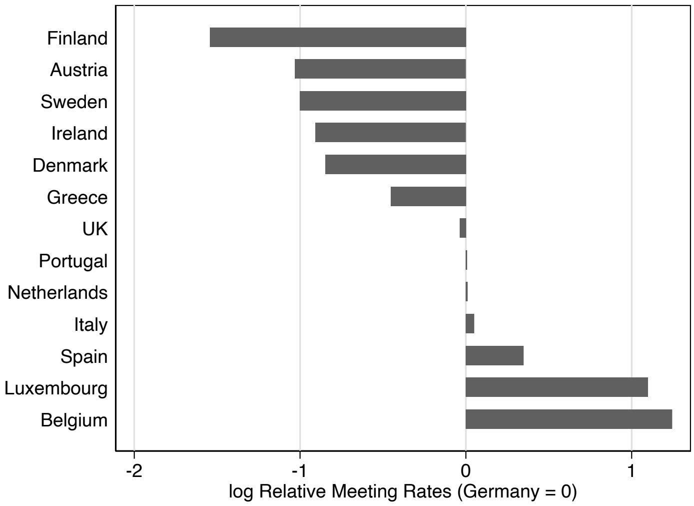

(b) Across sectors

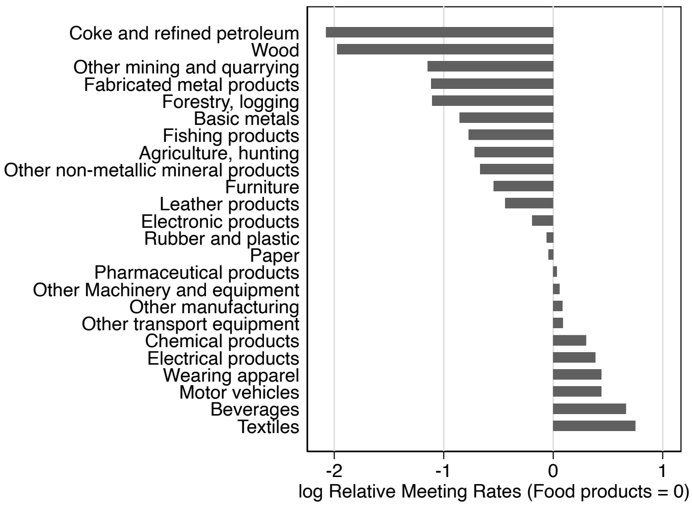

Note: The figure shows the mean value of relative meeting rates, by country and sector. All values are normalized by the mean estimate of Germany (top panel) and the Food industry (bottom panel).

Figure D10: Dispersion in overall meeting rates ($\gamma_{iF} + \gamma_{i\bar{F}}$), across countries and sectors

(a) Across countries

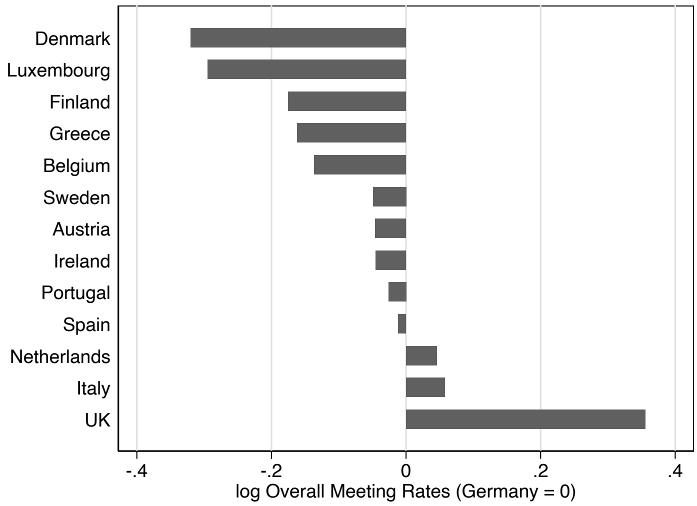

(b) Across sectors

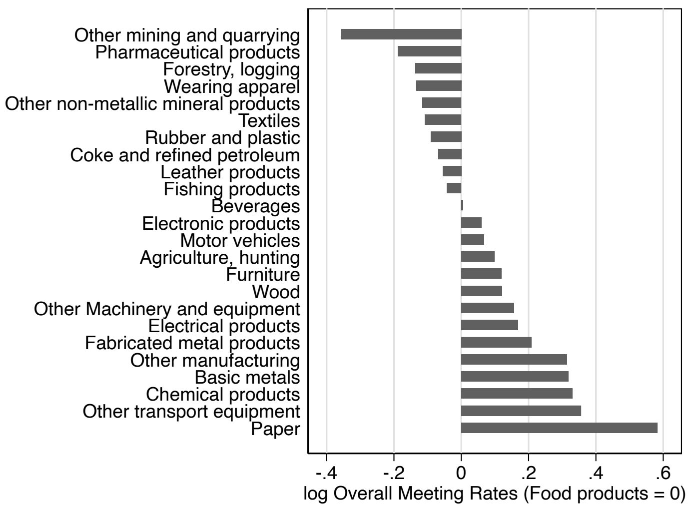

Note: The figure shows the mean value of overall meeting rates, by country and sector. All values are normalized by the mean estimate of Germany (top panel) and the Food industry (bottom panel).

Figure D11: Dispersion in estimated comparative advantages ($\tau_{iF}^{-\theta}$), across countries and across sectors

(a) Across countries

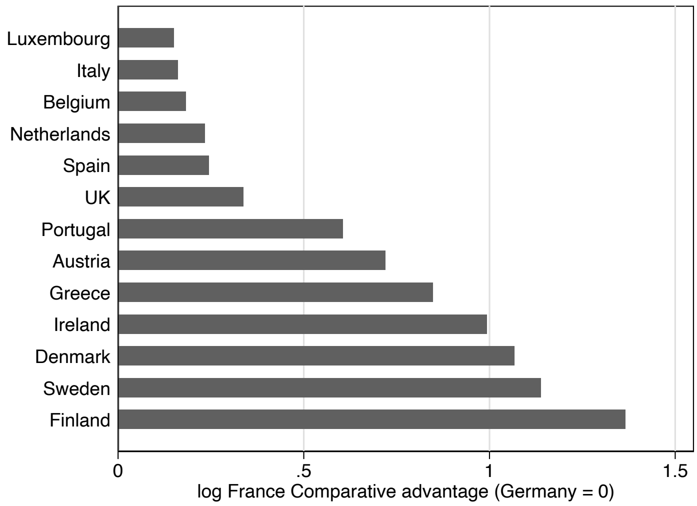

(b) Across sectors

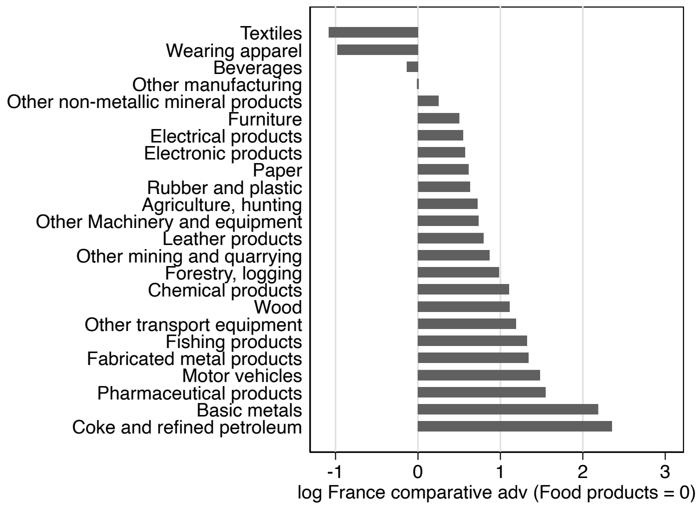

Note: The figure shows the mean value of comparative advantages, by country and sector. All values are normalized by the mean estimate of Germany (top panel) and the Food industry (bottom panel).

Figure D12: Goodness of fit: Transaction frequencies

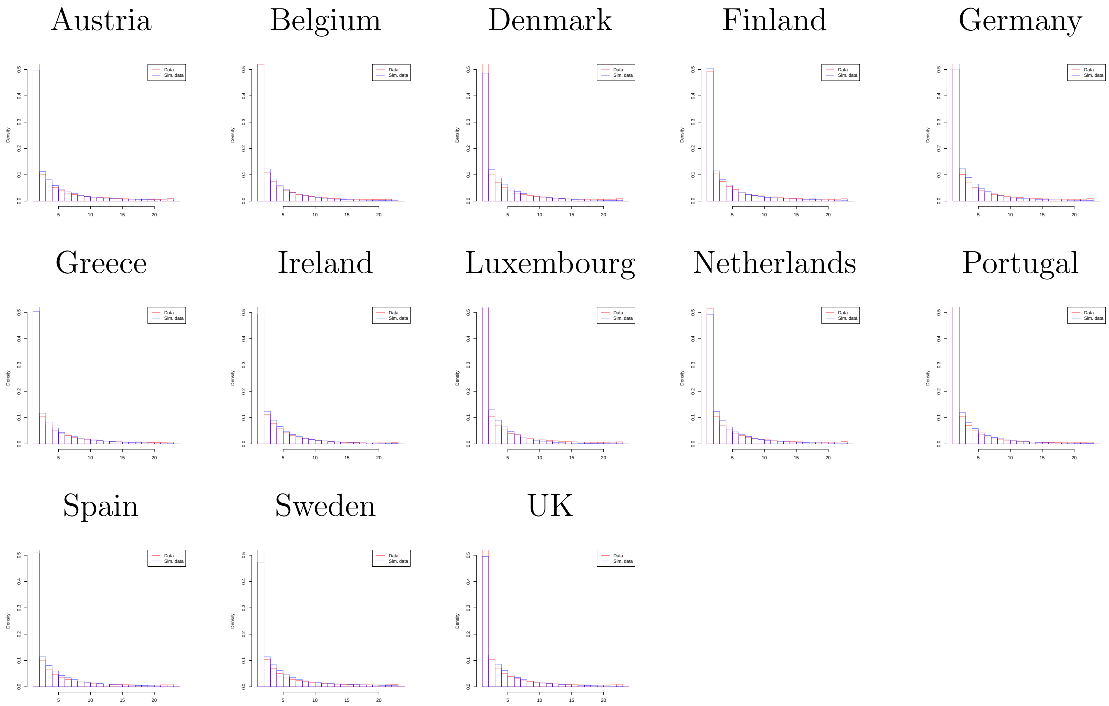

Note: Figure shows how we fit the transaction frequencies, by comparing the actual and simulated data.

Figure D13: Goodness of fit: Switch frequencies

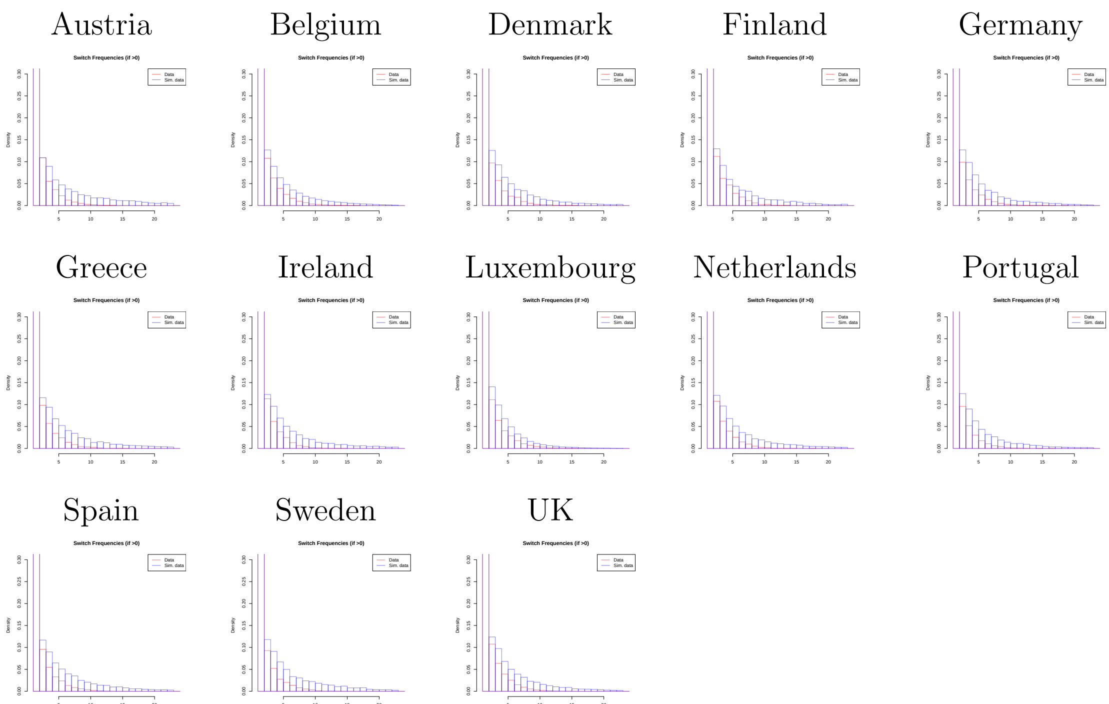

Note: Figure shows how we fit the switch frequencies, by comparing the actual and simulated data.

Figure D14: Switching rates, over time and across markets

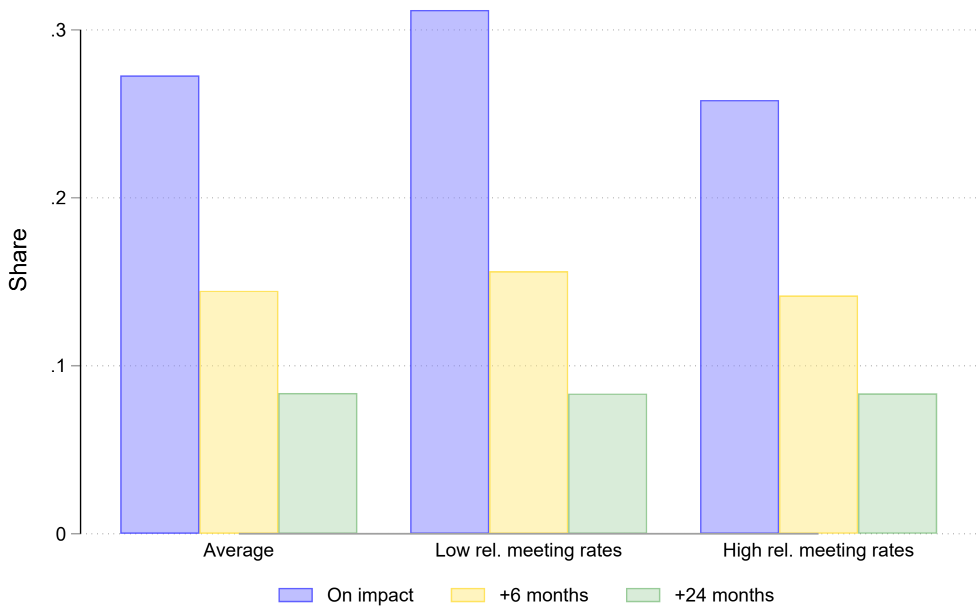

Notes: The graph shows the switching frequency following the shock and after 6, and 24 months, on average and in relatively low and high meeting rate markets, as defined by the first and fourth quartiles of the distribution of estimated relative meeting rates. Switching frequencies are normalized by the corresponding probabilities in the counterfactual without a shock.