Relationship Stickiness: Measurement and Applications to International Economics

Preliminary and incomplete, please do not circulate.

Julien Martin* Isabelle Mejean† Mathieu Parenti‡

December 9, 2018

Abstract

This paper proposes a measure of stickiness in business relationships based on the duration of firm-to-firm trade. The measure is backed by a model in which a higher degree of stickiness translates into a lower probability of switching and longer firm-to-firm trade relationships. The model is brought to firm-to-firm export data which allow us to measure the duration of individual relationships and estimate the relationship stickiness of more than 4,000 HS6 products. Relationship stickiness is shown to correlate with various dimensions of trade flows such as Nunn (2007) measure of relationship specificity or the prevalence of intra-firm trade. Finally, we show that stickiness in firm-to-firm relationships is a significant driver of the response of trade to macroeconomic uncertainty.

1 Introduction

For a firm, changing suppliers can induce a substantial cost. While switching costs are limited for products anonymously purchased on spot markets, their size becomes non-negligible in presence of relationship-specific investments, of customization costs, or in markets displaying informational and/or contractual frictions. As a consequence, firm-to-firm relationships often display a strong degree of stickiness.[1] As shown by the (international economic) literature,

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

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

‡ULB and CEPR, Email address: mparenti@ulb.ac.be.

[1] In the literature, several terms are used to characterize sticky trade relationships, notably investment specificity (Feenstra and Hanson, 2005), relationship specificity (Nunn, 2007), locked-in effects (Antràs and Chor, 2013) or input specificity (Barrot and Sauvagnat, 2016). Throughout the paper, we will refer to this as relationship stickiness.

stickiness matters for several reasons. For instance, the interplay of relation specificity with the legal environment shapes the specialization of countries (Levchenko, 2007; Nunn, 2007). The degree of relation specificity also governs the decision to integrate suppliers at home or abroad (Acemoglu et al., 2009; Antràs and Chor, 2013). The purpose and design of commercial policy (shall) depend on the specificity of business relationships (Antras and Staiger, 2012). Last, input specificity is key for the propagation of shocks across regions and countries (Barrot and Sauvagnat, 2016).

While its role is increasingly important in modern economies featuring complex global value chains, measuring the extent of such "relationship stickiness" is a difficult task. Two measures of product specificity (Rauch, 1999; Nunn, 2007) have been widely used in the literature.[2] Rauch's measure is based on hand classification of product categories across three groups: differentiated products, products traded in organized markets, and products with posted prices. Differentiated product markets, he argues, are those that should be the most frictional. Nunn's measure uses Rauch's classification to assess the specificity of inputs entering production processes, a good being called more "specific" when it is produced with more differentiated inputs. While such classifications have proved useful, they focus on one source of switching costs only, namely the extent to which the product market is organized. In this paper, we instead construct a measure of relationship stickiness which is obtained from ex-post measures of the average length of firm-to-firm relationships in various product markets. This, we argue, allows to capture the impact of a wider set of product-market characteristics that contribute to the stickiness.

Our measure of business relationships' stickiness is calculated using data on the *duration* of firm-to-firm trade. The measure is backed by a simple theoretical framework. In a dynamic search model, importers receive offers randomly and decide to switch to a new supplier whenever its offer is sufficiently good in comparison with the price offered by the buyer's existing partner. In this environment, we show that larger switching costs and/or more frictions contribute to lengthening existing firm-to-firm relationships. We discuss how the model can be brought to the data under some parametric assumptions regarding the underlying distribution of offers. This exploits the French firm-to-firm export data which panel dimension allows to follow importers over time and compute the duration of each of their relationship with a French firm. The high degree of granularity in the data is exploited to control for heterogeneity in the quality of a match and, conditional on this, recover a product-level measure of

[2] Other related measures have been developed. This includes the share of wholesalers importing a product (Bernard et al., 2010), the suppliers' R&D expenses, the number of patents issued (Barrot and Sauvagnat, 2016), or the distance to final demand (Antras et al., 2012).

relationship specificity (RS) for more than 4,000 HS6 products. This measure is shown to be stable over time and across country samples. This is consistent with the view that relationship stickiness is in part the consequence of structural product-specific characteristics.

We compare our estimates with several measures of product attributes commonly used in the literature. Our index is positively correlated with both Nunn's and Rauch's measures, although the correlation is not huge due to our measure displaying more heterogeneity within an industry. As expected, more complex and more differentiated products are found to display more stickiness. This is also the case of products located upstream in the global value chain. However, our measure of stickiness is not a simple combination of existing measures. Taken together, common product attributes explain 10% of the dispersion in relationship stickiness across products. Our interpretation is that the remaining 90% capture other attributes of product markets that contribute to making firm-to-firm relationships stickier.

To further assess the relevance of our measure, we propose several exercises. First, we replicate Nunn (2007) estimation with more disaggregated data and confirm that countries with better institutions tend to specialize in the production of highly-specific products. Second, we show that our measure of relationship stickiness explains as much as 10% of the dispersion in the US share of intra-firm trade across NAICS industries. This is consistent with the view that stickiness is in part the consequence of contractual frictions along the value chain, which itself increases the propensity of producers to integrate their suppliers (Antràs and Chor, 2013). Finally, we show how the negative impact of distance is magnified in product markets displaying high relationship stickiness. A possible interpretation is that this explains by information frictions being large in those markets, which on the one hand increases the cost of switching to a new supplier and on the other hand induces the geographic concentration of trade (Rauch, 1999). An alternative interpretation is that stickier relationships are associated to higher monitoring costs which increase with distance (Head and Ries, 2008).

The last section of the paper develops an original empirical analysis of the impact of policy uncertainty on the creation of new trade relationships. Following the literature, notably Bloom (2009), it is now widely recognized that uncertainty on the economic environment can be an important threat to economic growth as it reduces firms' incentive to invest. We provide evidence consistent with this view using the formation of new firm-to-firm relationships as an outcome variable and the measure of policy uncertainty proposed by Baker et al. (2016). Namely, the number of new trade relationships formed between French exporters and importers from four destination countries within a month is found significantly lower in periods of high policy uncertainty. In the meantime, the number of relationships terminating during these periods is also lowered. More interestingly, we show that the impact of uncertainty is especially pronounced in product markets that feature a high degree of stickiness. This is consistent with the view that there is a sunk cost associated with switching from a supplier to another, which increases the cost of a switch in period of high uncertainty. Stickiness in firm-to-firm relationships is a significant driver of the response of the economy to policy uncertainty.

Related literature. Beyond the literature on stickiness in firm-to-firm relationships, which was already mentioned, our paper is related to the empirical literature on switching costs in international good markets. The closest paper is Monarch (2014) who structurally estimates the cost of switching across Chinese suppliers for US importers. The author finds that halving switching costs would reduce the US-China import price index by 15%. Because of computational issues, he focuses on 50 exported products. We develop a lighter procedure to estimate these costs which allows us to recover them for a wide range of products. Furthermore, we work with highly disaggregated seller-buyer relationships observed over various destinations. This allows us to purge our measure from country-specific costs and obtain a measure of relationship stickiness at the fine product level.

The paper further contributes to the large literature on the duration of trade relationships. The literature has mainly focused on the impact of size, distance, and product differentiation on the duration of trade relationships (eg. Besedes and Prusa, 2006). While most of the papers in the literature look at the duration of trade flows at the product-level, Schmidt-Eisenlohr and Monarch (2015) show that the survival probability of seller-buyer relationships increases with their size and age using matched US importer-exporter data. Instead we use the duration of seller-buyer relationships in international markets to back out a measure of relationship stickiness, controlling for individual characteristics.

The rest of the paper is organized as follows. Section 2 describes the data, their structure and provides stylized facts on the duration of firm-to-firm relationships. Then, Section 3 presents a parsimonious search model that allows deriving the estimated equation used to back out our measure of relationship stickiness. Section 4 discusses the recovered distribution of relationship stickiness indices and its stability. Section 5 is devoted to the application on policy uncertainty and the net creation of firm-to-firm relationships. Finally, Section 6 concludes.

2 Data

2.1 Dataset

Our measure of relationship stickiness takes advantage of a panel of firm-to-firm data provided to us by the French customs and described in Bergounhon et al. (2018). The data cover each single export transaction between French firms and their individual partners in the European Union. Importantly, the data allow to identify and follow over time both firms involved in the transaction, the exporting French firm and its client. Each transaction is also characterized by a product category (at the 8-digit level of the European combined nomenclature), a date (month and year) and the value of the shipment (in euros). The full analysis covers the period from 1996 to 2017 and is divided into two sub-periods, 1996 to 2006 and 2007 to 2017. Benchmark results are obtained from the first sub-period and we test their stability using the second one.[3] As the empirical analysis extensively uses the time dimension within a product, we adopt the procedure described in Behrens et al. (2018) to harmonize the definition of NC8 products over time.

For each product category, we observe all firm-to-firm relationships involving a French exporter and a European buyer, over time.[4] However, we do not observe another interesting part of the network, namely transactions between foreign importers and their non-French suppliers. Over 1996-2006, we observe as many as 121 millions firm-to-firm relationships. Table 1 provides descriptive statistics on the dimensionality of the data, in the overall European Union as well as in France's main destination countries. We observe almost 127,000 different French exporters over the period, that interact with two millions foreign importers. Many of these firms sell/purchase several products so that the dimensionality increases by an order of magnitude once products are controlled for (columns (4) and (5)). Finally, we observe a total of 23 millions firm-to-firm relationships, that thus interact over 5 transactions, on average.

2.2 Structure of the firm-to-firm data

While the dimensionality of the data is relatively large, not all transactions are used in the estimation sample. In particular, the analysis neglects all buyers that we observe over a single transaction throughout the period under study. These represent as many as 44% of

[3]Note that we also have access to data between 1993 and 1995, which we use to control for left-censoring in our measures of the duration of rirm-to-firm relationships.

[4]While the dataset is exhaustive, exports from the smallest French exporters cannot be exploited as these firms are allowed to fill a simplified form that does not specify the product category. In 2007, the simplified regime concerned 16,588 exporters (out of 68,386) accounting for 2% of transactions and less than .5% of the value of French exports.

Table 1: ***Summary statistics on the structure of the dataset***

	# transac.	# sellers	# buyers	# sellers ×products	# buyers ×products	# buyer×seller ×products
	(1)	(2)	(3)	(4)	(5)	(6)
EU27	121,526,703	126,996	2,057,919	1,582,000	16,619,222	23,386,550
Belgium	23,500,617	84,636	251,294	711,020	2,786,005	4,187,021
Germany	23,201,027	71,051	427,656	583,639	2,934593	4,141,553
Italy	15,256,499	61,362	346,477	457,109	2,490,812	3,355,340
Spain	15,168,339	61,268	323,535	475,219	2,301,084	3,030,330
UK	12,353,052	56,279	200,869	402,491	1,502,953	2,259,541
Netherlands	7,388,882	50,860	124,535	310,222	899,134	1,327,122
Portugal	5,731,671	37,886	86,531	271,456	869,896	1,207,535
Luxembourg	3,615,290	35,635	30,584	219,743	556,545	767,670
Austria	2,758,348	25,846	47,473	145,426	467,014	602,812
Greece	2,401,544	22,736	41,636	157,892	426,317	580,281

Notes: This table is based on French customs firm-to-firm data for 1995-2007. The first line corresponds to all countries while the rest of the table gives statistics for individual countries, limited to the 10 more popular destinations (in terms of the number of transactions). Column (1) reports the number of transactions, a transaction being defined as a trade flow in a given month and year, involving a particular seller-buyer pair, for a given product. Column (2) is the number of exporters observed over the period. Column (3) is the number of importers. Column (4) is the number of seller-product pairs. Column (5) is the number of buyer-product pairs. Column (6) is the number of seller-buyer-product triplets observed over time, also called "relationships" in the rest of the paper.

the importers in the raw data (see the top panel of Figure 1) and **X**% of the value of trade. The reason why we abstract from these one-time transactions is that our empirical strategy is based on the duration of a relationship which is not defined when the importer is observed just once. One may argue that neglecting these transactions bias our measure of stickiness up since these firms may have switched to a non-French partner immediately after their first transaction. But censoring at the beginning and end of the period is also likely to generate a lot of these observations.[5] Our strategy thus consists in ignoring these transactions to estimate the degree of stickiness at the product-level, and use them ex-post to check whether they indeed concentrate in those sectors that we estimate display less stickiness.

The remaining 56% of foreign buyers also display a lot of heterogeneity regarding the "intensity" of their relationships with French exporters. This heterogeneity is illustrated in the top panel of Figure 1. 16% of foreign buyers are observed over two transactions, with **X%** of these firms interacting twice with the same French firm and the rest switching between the first and the second transactions. At the other side of the spectrum, more than 25% of foreign buyers are observed over more than 10 different transactions, eventually involving several partners. These firms are good candidates to observe the duration of their relationships with French firms.

The bottom panel of Figure 1 shows the distribution in the number of French partners, individual buyers interact with, over the whole timespan of their presence in the dataset. Overall, 67% of buyers have a single partner in France while less than 7% have three partners or more (see the circles line). Of course, interacting with a single partner in France is more likely to happen for firms that are involved in a small number of transactions. The other three distributions thus show the number of partners by buyers, for importers that are involved in at least 5, 10 or 50 transactions. Even within the subset of importers that we observe over as many as 50 transactions, we do observe a substantial proportion of "faithful" buyers that always interact with the same exporter again and again. Such behavior is consistent with the idea that some firm-to-firm relationships in international markets are especially sticky. The question that the empirical analysis addressees is whether this is more likely to be the case in some sectors than others.

We terminate this raw description of the data structure with a last stylized fact later used to justify the theoretical framework under study. Namely, we will now argue that the network under study displays many-to-one matching, once the product dimension is controlled for. At a point in time (defined by a particular month in a particular year), we observe most buyers

[5]Indeed, we do observe that the time distribution of these one-time transactions is not homogenous across years with more than one over eight one-time transaction observed in 2007, the last year.

Figure 1: ***Distribution of the number of transactions and the number of French partners, per buyer***

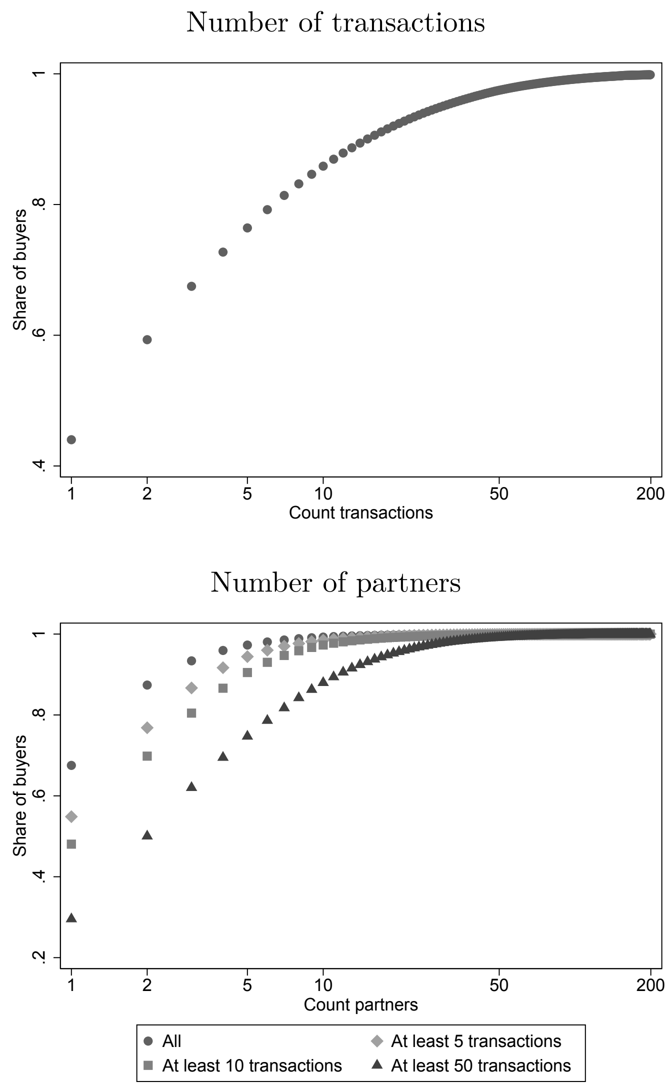

Notes: Cumulated distributions of the number of transactions (top panel) and the number of partners, per foreign buyer (×product) (bottom panel). A transaction is the purchase of a particular good, to a given seller, in a given month. A partner is a French exporting firm. The number of partners is calculated over the sub-sample of importers that are involved in at least two transactions over the period of analysis ("All") and at least 5, 10 and 50 transactions.

Figure 2: ***Distribution of the number of partners, per buyer/seller and date (month×year)***

Number of partners per buyer

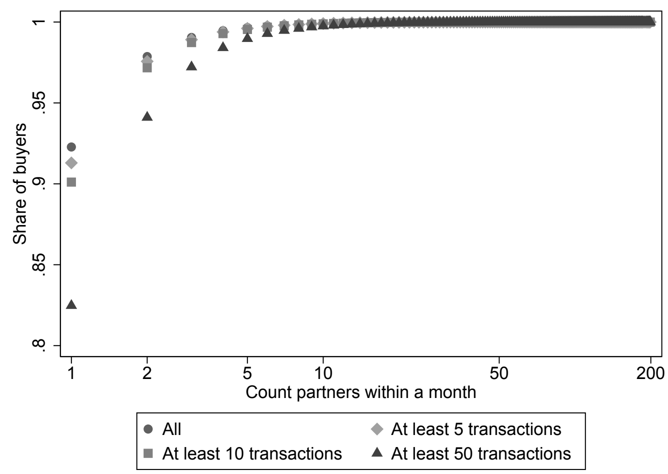

Number of partners per seller

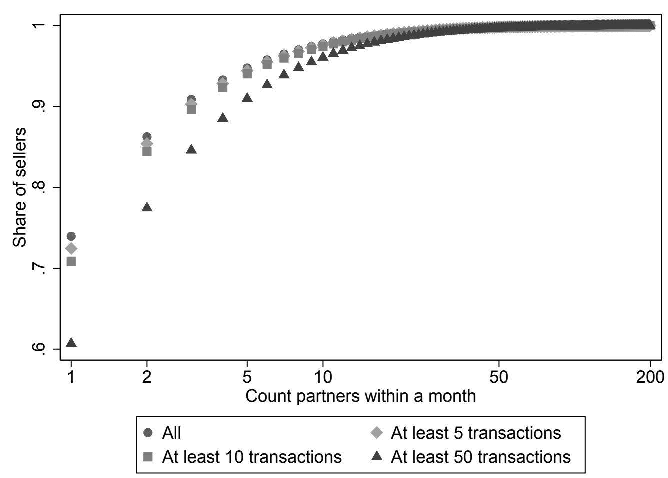

Notes: Cumulated distributions of the number of partners a French exporter interacts with in a given country (bottom panel) and the number of partners a foreign buyer (×product) interacts with within a particular month (top panel). The number of partners is calculated over the sub-sample of importers (resp. exporters) that are involved in at least two transactions over the period of analysis ("All") and at least 5, 10 and 50 transactions.

purchasing a particular product from a single seller while instead sellers simultaneously serve several importers (even within a country). This is illustrated in Figure 2 which shows the distribution in the number of sellers interacting with a given importer during a particular month (top panel) and the distribution in the number of partners from the same country a French exporter is interacting with (bottom panel). More than 90% of importers have only one French supplier for a given product within a given month. Even when we concentrate on importers that we observe over many (i.e. at least 50) transactions, this proportion is high, above 80%. Instead, 26% of French exporters sell the same product within the same month to several partners located in the same country, the proportion increasing to 53% when we pool partners located in different countries.[6]

2.3 Duration of firm-to-firm relationships

Using the time-series of each buyer's interactions with French firms, it is now possible to construct the main statistic at the root of the estimation, namely the duration of a buyer's interactions with French firms.[7] While the statistics is a priori straightforward to compute, a number of practical difficulties arise which we now describe.

In the model, the duration of a relationship is written as a function of the probability of a switch, i.e. of an importer leaving its current partner to start interacting with a new one. In the data, the two objects do not map exactly because of the heterogeneity in the frequency of trade transactions. This is illustrated in Table 2 which first three lines compare statistics on i) the mean duration of a buyer's relationships with French suppliers, ii) the inverse of the probability of this buyer's switching to a new supplier and iii) the inverse of the probability of switching, conditional on trade. Would buyers purchase French products at regular intervals, e.g. every month, the three statistics would deliver the exact same information. As shown in the fourth line of Table 2, the frequency of transactions is neither close to one, nor homogenous across buyers. On average, the probability of a transaction occuring in a given month is equal to .358, which corresponds to a transaction every 2.8

[6] This is in contrast with Bernard et al. (2017) who use qualitatively similar data and find that the matching between exporters and importers display many-to-many relationships. Beyond their data covering a different country, a possible reason for such discrepancy is that they do not condition on a particular product while we do. Indeed, we do see in our data that buyers often interact with several French exporters in a given month, although to purchase different products. Once we condition on a given product, purchasing from multiple French exporters becomes very rare.

[7] Note that all the statistics are treated for left censoring. Namely, we use information on trade between 1993 and 1996 to differentiate a relationship that starts in 1996 or later from a relationship that pre-existed the estimation period. The same strategy is applied to control for right censoring.

months. 25% of buyers however purchase French products more than once every two months while in the first quartile of the distribution, firms purchase products less than once every 8.3 months. Because of heterogeneous frequencies, the three available measures of duration are not equivalent. In general, one can show that the mean duration is in between the two measures calculated from switching probabilities. In the data, the three statistics are correlated at more than 80% meaning that heterogeneity in the frequency of transactions does not completely distort the distribution of trade durations, across buyers and products.

Table 2: ***Descriptive statistics on the duration of firm-to-firm relationships***

	Mean	Median	P25	P75
Mean duration	24	13	3	34
$1/\mathbb{P}(\text{switch})$	9	19	8	42
$1/\mathbb{P}(\text{switch}\vert\text{Trade})$	3	3	2	7
Frequency of transactions	0.358	0.246	0.120	0.500
Proba Recall	0.054	0.000	0.000	0.000

This table gives statistics on the distribution of durations, frequencies and recall probabilities, across importers connected to French firms. The first three lines correspond to alternative measures of the mean duration of a buyer's relationships with French firms, namely the mean number of months between the first transaction and the first transaction with a different supplier ("Mean duration"), the inverse of the switching probability recovered as the number of switching episodes divided by the total number of months the buyer is present in the data ("$1/\mathbb{P}(\text{switch})$") and the inverse of the switching probability conditional on a transaction, computed as the number of switching episodes over the total number of transactions ("$1/\mathbb{P}(\text{switch}|\text{Trade})$"). "Frequency of transactions" is the probability of observing a transaction within a month, computed as the number of transactions divided by the total number of months the buyer is present in the data. Finally, "Proba Recall" is the probability that a buyer switches to a French exporter that it had already interacted with in the past and is computed as the number of recalls of an already known supplier divided by the number of switching episodes. Statistics calculated on the dataset covering the 1996-2006 period.

Second, in the model and in statistics in Table 2, the duration of a relationship is calculated between the first time a seller and buyer interact and the first time the same buyer interacts with a different French exporter. In case the buyer interacts again with its first supplier later in time, a new relationship is supposed to be created, and its duration calculated independently of the first relationship these two firms have shared. Abstracting for the whole history of the buyer's interactions with French sellers and focusing on the probability to switch to a new supplier just met greatly simplifies the analysis. Moreover, the probability of a "recall", i.e. of a buyer switching back to a supplier it knows from before, is very small in the data (last

Table 3: ***Duration, switching probabilities, and the size of trade flows***

	(1)	(2)	(3)	(4)	(5)	(6)
				Log of		
	duration	duration	duration	1/P(switch)	1/P(switch |Trade)	duration
Mean exports (log)	0.050*** (0.000)	0.053*** (0.000)	0.059*** (0.000)	0.125*** (0.000)	0.047*** (0.000)	0.091*** (0.000)
Observations	9,797,609	9,797,609	9,797,609	9,797,609	9,797,609	25,132,896
(Within) R^2	0.005	0.005	0.005	0.046	0.004	0.014
Fixed effects	Country	Product	Product× country	Product× country	Product× country	Product× buyer

This table correlates the duration of a relationship with a measure of the size of this transaction. Statistics calculated on the dataset covering the 1996-2006 period. **Table needs to be updated**

line in Table 2).

While the statistics in Table 2 are calculated for each importer, averaging across all relationships it has with a given French supplier, our analysis further exploits the granularity of the data to control for individual determinants of the duration of firm-to-firm relationships. In particular, results in Table 3 show that the duration of trade relationships is positively correlated with the size of the transaction, which we use as proxy for the quality of the match between the buyer and its supplier. This is true both across buyers within a product and within a buyer, across the different suppliers it meets throughout its interactions with French firms. This correlation is fully taken into account in our empirical framework which recovers a measure of the mean duration of trade relationships, *conditional on the quality of a match.*

3 Revealed relationship stickiness

3.1 Theoretical framework

Previous section shows how firm-to-firm trade data can be used to measure and document heterogeneity in the duration of business relationships across firms and products. We now build a theoretical framework that relates these durations to characteristics of the firms and the products they sell. In particular, we will assume that products systematically vary in terms of their degree of business stickiness. Such cross-sectional heterogeneity might be explained by the products sold being more or less substitutable, by the size of relationship-specific investments varying across sectors or other sector-specific characteristics. The micro-foundations at the root of such stickiness are not explicitly modeled. Instead, the model features a parameter γ of stickiness, that is proportional to the cost of switching from an input provider to the other.

Suppose that an importer is willing to buy a certain product. Every period, it receives an offer $\tilde{p}$ from a new input supplier with a probability λ and ultimately decides whether to stick to its existing partner or to switch and benefit from this offer. Suppose that $\tilde{p}$ is the (quality-ajusted) price at which the new input supplier is willing to sell the product. It is the realization of a random variable P drawn into a cumulated distribution function $H_P(p) = \mathbb{P}(P \leq p)$. Conditional upon its current deal p, a firm may decide to switch suppliers as soon as it receives a deal that is not only better but also covers its switching cost. Namely, the firm decides to switch whenever $\tilde{p} < \frac{p}{\gamma}$ where $\gamma > 1$. This occurs with a probability $\lambda H_P(p/\gamma)$.[8]

Under these conditions, the expected length of a buyer-seller relationship, conditional on its price is given by

$$\mathbb{E}[\mathcal{T}|p] = \sum_{k=1}^{+\infty} k(1 - \lambda H_P(p/\gamma))^{k-1} \lambda H_P(p/\gamma) = \frac{1}{\lambda H_P(p/\gamma)} \tag{1}$$

The duration of a relationship is just the inverse of the probability of switching. It is a function of the firm's existing deal p, the sectoral degree of business stickiness as measured by γ and the frequency of offers λ which reflects the extent of frictions in that market. Everything else equal, a firm which has met a more competitive supplier is more likely to interact with it over a long relationship. But conditional on a price, larger switching costs and less frequent offers are also expected to lengthen firm-to-firm relationships. These product characteristics are what we want our measure of relationship stickiness to capture. We now explain how to estimate it using observed durations in the data.

Estimating equation (1) raises several difficulties. First, one needs to make a parametric assumption regarding the underlying distribution of prices. Second, quality-adjusted prices are not directly observable and their proxy - even abstracting from quality differences - typically noisy. To circumvent these issues, we make two additional assumptions. Let us first suppose that the demand of imports is iso-elastic with σ the price elasticity of demand. Under this assumption, equation (1) can be written conditional on the size of the transaction r instead

[8]This implicitly assumes p to be determined prior to the arrival of a new offer, i.e. we do not let the firm and its supplier re-negotiate over the price when a better offer arrives. We discuss later how such renegotiation "on-the-match" can affect the duration of firm-to-firm relationships.

of the (unobserved) price offered by the supplier:

$$\mathbb{E}[\mathcal{T}|r] = \frac{1}{\lambda\left(1 - H_R(r\gamma^{\sigma-1})\right)}$$

where $H_R(r) = \mathbb{P}(R \leq r)$ is the cumulative distribution of sales.

Second, suppose that the distribution of prices is inverse-Pareto so that transactions are distributed Pareto:

$$H_R(r) = 1 - \left(\frac{r}{R_{min}}\right)^{-\frac{k}{\sigma-1}}$$

with R_{min} the scale parameter and k the shape of the distribution of quality-adjusted prices. Together, these assumptions imply that the expected duration of a relationship, conditional on the size of the transaction, is log-linear in the size of the transaction:

$$\mathbb{E}[\mathcal{T}|r] = \eta\left(\frac{r}{R_{min}}\right)^{\frac{k}{\sigma-1}} \tag{2}$$

where $\eta \equiv \frac{\gamma^k}{\lambda}$ is our measure of relationship stickiness. As explained in the introduction, it is a function of several parameters of the model, that capture various forces that tend to lengthen firm-to-firm relationships, conditional on a match. From that point-of-view, it cannot be interpreted structurally. In particular, larger switching costs and less frequent offers both increase the expected duration of a match. The sensitivity of durations to switching costs is itself a function of k as more dispersion in the underlying price distribution means that a buyer is more likely to meet with a seller which price is sufficiently lower than the price it currently pays to cover the switching cost.

Section 3.2 explains how this equation can be confronted to the data to recover an estimate of η, our measure of relationship stickiness. Before this, let us discuss the generality of the above prediction. Equation (2) relies on three important assumptions. First, firms are assumed to post their price without renegotiation; second, the demand is supposed iso-elastic and third the distribution of sales is parametrized to a Pareto function. While the second and third assumptions facilitate the estimation of Equation (1), it is legitimate to ask whether the theoretical framework would still be insightful would firms be able to overbid competitors' offers. Our argument is that this would not turn out our main results regarding the length of firm-to-firm relationships, while potentially deriving richer predictions regarding the dynamics of prices. Assume for instance that the buyer triggers a Bertrand-type competition between its existing supplier and any newly met seller.[9] Either its current supplier matches its competitor's offer by decreasing its mark-up, either the buyer starts a new relationship. The switch happens whenever the newly met seller has a marginal cost that is sufficiently

[9]See Postel-Vinay and Robin (2002) for an application of this assumption to on-the-job search models.

below that of the buyer's current partner, namely when $c' < c/\gamma$ (where c' and c respectively denote the quality-adjusted marginal cost of delivering the product, of the newly met and the existing supplier, respectively). Under this scenario, the switching probability is the same as before and Equation (1) holds while prices are expected to increase over time within a buyer-seller pair.

3.2 Estimation strategy

Equation (2) is about the expected duration of a relationship, conditional on its size. To make this operational, we work with the expected duration of relationships, conditional on their size interval ($\mathbb{E}[\mathcal{T}|r \in Interval_d]$ where $Interval_d$ is a particular quantile of the size distribution). The empirical counterpart of this object is simply the mean duration of trade relationships that fall in a given size interval (Dur_{dpc}). We show in the appendix that under the conditions described in section 3.1, we can recover the relationship stickiness (RS) of products by estimating the following equation:

$$log(Dur_{dpc}) = FE_p + FE_c + \alpha A_d + \epsilon_{dkc} \tag{3}$$

where Dur_{dpc} is the mean duration of trade relationships in size-bin d, for product p in destination c, FE_c is a country fixed effect, FE_p is a product fixed effect, A_d is a term which is proportional to the size interval, and ϵ_{dkc} is the error term. Importantly, in the case of a Pareto distribution, A_d is independent of all the parameters that enter η. This is the reason why the product fixed effect can be used to recover information on the heterogeneity of η across product categories.

To compute the mean duration (Dur_{dpc}) we proceed as follows: i) we compute the size of a relationship as the average size of its transaction in constant euro, ii) we assign all the trade relationships to a size-decile (specific to a destination and a product category), iii) we take the average duration within each bin.[10]

We expect α to be positive as larger transactions are expected to last longer on average. The country fixed effects are aimed at capturing difference in the length of transaction that are country specific. The product fixed effects are our estimates of relationship stickiness ($\ln \eta_k$). Because our estimates of specificity are product fixed effects, they cannot be interpreted in level. However they give the ranking and the relative stickiness of the products.

[10] As explained in the appendix we truncate the first and tenth deciles to be fit with the theoretical framework.

Table 4: ***Correlation with other measures***

Measure	Corr(η,.)	OLS η
$\mathbf{1}_{differentiated}$ (Rauch)	.06***	.01
Share of not homogen. products (Nunn)	.07**	.09
Upstreamness (Antras et al.)	.14***	.16***
Elasticity of subs. (Imbs & Mejean)	-.11***	-.30***
Product complexity (Haussman & Hidalgo)	.25***	.10***
Observations		3,877
R^2	-	.11

4 Results

4.1 Correlation with other product-level attributes

Measures of differentiation and complexity. Table 4 presents a basic comparison of relationship stickiness with other measures used in the literature to characterize product categories. The first column reports the pairwise correlation of these indicators. The last column reports the coefficients of a regression of our RS measure on all the other characteristics.

The degree of product stickiness is positively correlated with other measures of product specificity such as the one developed by Rauch (2002) and Nunn (2007). Our measure also correlates with the level of differentiation of products as measured by Imbs and Mejean (2015). Interestingly enough, we find a positive correlation between the level of upstreamness of a product and its the degre of stickiness. This suggests that products far from the final demand entail more buyer-specific investment, more elaborated contracts or more customization than products dedicated to final consumption. Last, we find a positive correlation between the degree of stickiness of a product and its level of complexity as measured by Hausmann and Hidalgo (2014). In the second column, we regress our estimates of relationship stickiness on the different measures we just discussed. When taken in conjunction with other product attributes, the correlation between our estimates and the two main measures of specificity is no more significant while the other measures remain highly significant.

In Figure 3, we explore the correlation between our RS measure and the measure developed by (Nunn, 2007). We see the weak correlation between the two measures. The vertical lines of dots are the consequence of the different levels of aggregation of the two measures. In particular, we see that within a NAICS industry (the level at which Nunn's measure is computed), there is a strong level of dispersion in the degree of stickiness across product

Figure 3: ***Share of non-homogeneous products vs RS measure***

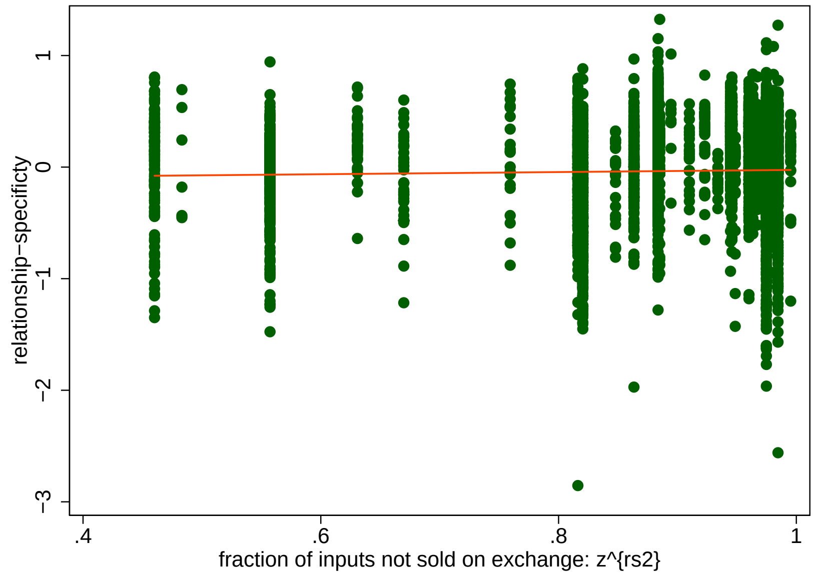

categories. Zooming on specific industries, this high level of dispersion seems to make sense. For instance, in the car industry, some parts are very specific to model of cars while other parts are not. (Nunn, 2007) measure assigns a single level of specificity to this industry. Furthermore in this industry the final products (the cars) need not be customized or do not entail specific contracts. Zooming on the chemical industry leads to similar observations. In this industry, some products are standardized while other are highly customized depending on the customer and they require specific formulation (the specialty chemicals).

In Figure 4, we present the distribution of our RS measure across broad economic categories (BEC). At first sight, the Figure mainly shows that there is substantial dispersion in the level of stickiness within BEC categories. However, if one looks more closely at the figure, interesting patterns emerge. More specifically, we find that final goods like cars or consumer goods display on average a lower level of stickiness while parts and accessories or food processed for the industry have a higher degree of stickiness.

Relationship stickiness and intra-firm trade. Many theories suggest that the strength of relationships between a firm and its supplier may lead to vertical integration. In Table 5, we examine this feature by looking at the correlation between the level of intrafirm trade in a product category and its degree of relationship stickiness. Of course we do not expect this correlation to be perfect as vertical integration further depend on country and firm

Figure 4: ***Relationship stickiness across Broad Economic Categories***

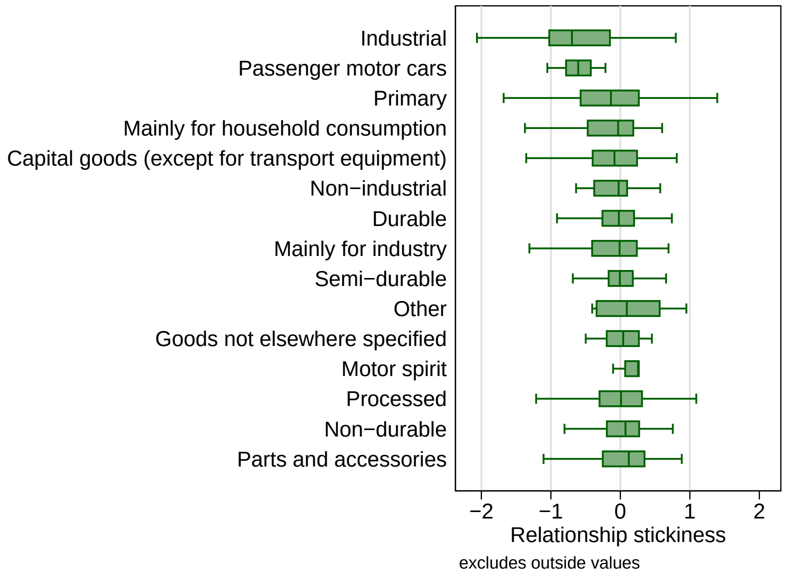

characteristics. It is worth noting that we explore this question using US data while our RS estimates are obtained from French data. We find a positive correlation between the level of relationship stickiness of a product and its share of intrafirm trade. Relationship stickiness explains (in terms of R^2) 6 to 8 percent of the dispersion in the share of intrafirm trade across product categories.

4.2 External validation.

One way to verify the relevance of our measure is to use it in situations where we have strong prior regarding its impact on economic outcomes. Nunn (2007) and Levchenko (2007) provide strong evidence that countries with good contract enforcement specialize in the production of goods for which relationship-specific investments are most important. We use this well established results to test the validity of our measure. We reproduce the same exercise as Nunn (2007) but working with more disaggregated data (6-digit level of the Harmonized Nomenclature) and using our own measure of relationship stickiness. The results are reported in Table 6. In every regression we further control for the relation-specificity measure developed by Nunn (2007). In the first three columns, we follow Nunn and explain the value of countries' exports at the product level. In columns (4) and (5), we deviate from Nunn (2007) and consider measures of specialization that allow us to account for product-country pairs with zero trade flows. We confirm Nunn (2007) findings that countries with good contract enforcement

Table 5: ***Share of intrafirm trade and relationship stickiness***

	(1)	(2)	(3)	(4)
	Share of intra-firm			
	exports		*imports*	
RS (η)	0.177***	0.180***	0.140***	0.138***
	(0.040)	(0.041)	(0.030)	(0.031)
Nunn' measure		0.406***		0.199***
		(0.063)		(0.046)
Upstreamness		0.060***		0.015
		(0.016)		(0.011)
Elasticity (σ)		0.002		-0.005
		(0.006)		(0.004)
Observations	378	378	378	378
R-squared	0.058	0.166	0.071	0.119

Robust standard errors in parentheses with *, **, *** denoting significance at the 10, 5 and 1% levels.

specialize in the production of more relationship specific goods. We view these findings as additional evidence that our measure capture meaningful product characteristics.

In Table 7, we interact distance with our measure of relationship stickiness in an otherwise standard gravity equation for bilateral trade. The HS6 product-level bilateral trade flows are from BACI for year 2005 (Gaulier and Zignago, 2010). Distance is the weighted distance between countries' main cities (Mayer and Zignago, 2011). We find that the negative impact of distance on trade flows is stronger for product categories that exhibit a higher degree of relationship stickiness. A possible interpretation is that this explains by information frictions being large in those markets, which on the one hand increases the cost of switching to a new supplier and on the other hand induces the geographic concentration of trade (Rauch, 1999). An alternative interpretation is that stickier relationships are associated to higher monitoring costs which increase with distance (Head and Ries, 2008).

4.3 Stability of our estimates

Cross-country stability. As a first assessment of the stability of our estimates, we estimate RS country by country. We focus on countries where France exports a large set of HS6 product categories, and where many exporters are active (to be able to compute the average duration by product-category X decile). The measure seems quite stable across countries. The pairwise

Table 6: ***Institutional comparative advantage***

	(1)	(2)	(3)	(4)	(5)
		log(exports)		Balassa Index	$\mathbf{1}_{Balassa>1}$
Rule of law					
× *RS*	0.349***		0.408***	0.286**	0.022***
	(0.053)		(0.053)	(0.120)	(0.006)
× Nunn specif.		0.811***	0.978***	0.316*	0.027*
		(0.100)	(0.117)	(0.168)	(0.015)
× Upstreamness			0.034	0.013	0.002
			(0.021)	(0.024)	(0.002)
Fixed effects		$country(122)$ and $sector(4,326)$			
Observations	296,187	296,187	292,938	527,284	527,284
R-squared	0.605	0.606	0.610	0.012	0.139

Clustered (country) standard errors in parentheses with *, **, *** denoting significance at the 10, 5 and 1% levels.

Table 7: ***Gravity for trade in goods with sticky relationship***

	(1)	(2)	(3)	(4)	(5)	(6)
	Value	UV	Value	UV	Value	UV
Distance (log)	-0.571***	0.467***	-0.699***	0.101***	-0.990***	0.087***
	(0.020)	(0.006)	(0.021)	(0.006)	(0.023)	(0.007)
RS	1.008***	-0.166***				
	(0.093)	(0.041)				
- × dist.	-0.151***	0.042***	-0.143***	0.020***	-0.113***	0.066***
	(0.012)	(0.005)	(0.012)	(0.004)	(0.010)	(0.005)
Upstreamness	1.008***	-0.166***				
	(0.041)	(0.043)				
- × dist.	0.010**	-0.084***	0.021***	0.028***	-0.012*	0.047***
	(0.005)	(0.000)	(0.005)	(0.002)	(0.007)	(0.002)
Fixed effects	exporter-importer		exp + imp +sect		exp.×sect +imp.×sect	
Observations	5,703,782				5,473,330	
R-squared	0.164	0.178	0.285	0.654	0.578	0.770

Clustered (country) standard errors in parentheses with *, **, *** denoting significance at the 10, 5 and 1% levels.

correlation between country-specific measures is around .5. More importantly the correlation between the measure computed pooling all countries is highly correlated with the country level estimates (almost .8). Note that the pairwise correlation are not computed over the exact same sample as some products are not exported to every country.

Time consistency. A second way to assess the stability of the measure is to split the sample in two periods (1997-2003 and 2003-2007). In this exercise, we exclude countries which are not in the EU in the first half of the sample. We now have three RS measures (RSall, RS9703, RS0307) which are computed over the all years, the first half of the sample period, and the second half of the sample period. The measure sounds quite stable over time. The correlation between the measure computed over the first half of the sample and the one computed over the second half of the sample is 0.6. Furthermore, the correlation between these two measures and the measured computed over the whole sample reaches .8.

5 Trade, uncertainty and stickiness

Trade and uncertainty. Economic and policy uncertainty have been shown to have a significant impact on the real economy (Bloom, 2009). While the bulk of the literature on uncertainty focus on the domestic market, several papers have emphasized the importance of uncertainty to understand trade patterns such a firms' decision to export (see eg. Novy and Taylor, 2014; Handley and Limão, 2017, 2015; Pierce and Schott, 2016).

We conjecture that uncertainty may affect trade patterns by impeding the creation of new business relationships. If there are some costs to establish a new trade relationship and / or if trade partners are locked-in these relationships, then firms may delay the formation of new relationships in periods of high uncertainty. We further conjecture that these effects are more pronounced when the traded products entail stickier relationships. We test these hypothesis by combining our measure of stickiness with a measure of policy uncertainty.

Empirical strategy. We want to test whether the impact of uncertainty on the probability to form a new trade relationship is stronger for trade involving more sticky relationships. We run the following regression:

$$\#\ new\ relations_{pct} = \alpha Uncertainty_{ct} + \beta RS_p + \gamma RS_p \times Uncert_{ct} + FE + \epsilon_{pct} \quad (4)$$

The left hand side variable ($\#\ new\ relations_{pct}$) is simply the number of new interactions between a French seller and a foreign buyer in a destination and for a given product category.

We focus on years 2001-2010. We identify an interaction as new if the seller and the buyer have not interacted between 1995 and the date of the new relationship.

The product-level measure of stickiness is the one described in the rest of the paper. The last variable of interest is a country-time measure of uncertainty. We use the measure of policy uncertainty developed by Baker et al. (2016). This measure has several advantages: it is a measure of uncertainty computed with the same methodology across countries, it is available over the same time interval as our dataset, it is available at a monthly frequency. Figure 5 reports the measure of policy uncertainty for the four countries we consider during the analysis, namely Germany, Italy, the UK, and Spain. From these series, we define uncertainty episodes as periods where the uncertainty index is one standard deviation above its average level.[11] We see that the uncertainty index and the uncertainty episodes are correlated but are far from coinciding across countries.

In equation 4, we expect α to be negative meaning that there are less business creations during uncertainty episodes. The main coefficient of interest is γ. We expect it to be negative meaning that the negative impact of uncertainty on business creation is stronger for stickier products.

Last, we control for different sets of fixed effects. In all specifications, we control for product-month fixed effects. These allow us to account for product specific patterns in business creation due to seasonality. We further control for country fixed effects to account for time unvarying characteristics that may be correlated with the average level of uncertainty in a country and the dynamism of new trade relationships. In the more demanding specification, we add country-time fixed effects. This absorbs our measure of policy uncertainty but there is still room to identify the interaction between uncertainty and relationship stickiness.

Results. Table 8 presents the results of the joint impact of uncertainty and product stickiness on the formation of new trade relationships. As expected, we find that French exporters initiated a lower number of trade relationships during uncertainty episodes. In column (2), we show that the negative impact of uncertainty is exacerbated for product categories entailing stickier relationships. This result survives if one includes country-time fixed effects. It is also robust to an alternative definition of uncertainty episodes - uncertainty episodes being defined as periods during which the uncertainty index deviates by more than 1.64 standard deviation.

These results point to a precise mechanism through which uncertainty impacts the real economy. We show that uncertainty affects the creation of new business relationships. We further show that this effect is stronger for product categories entailing stickier trade relation-

[11] We test the robusness of our results to alternative definition of uncertainty episode.

Figure 5: ***Policy uncertainty, 2000-2015 (Baker, Bloom, Davis 2016)***

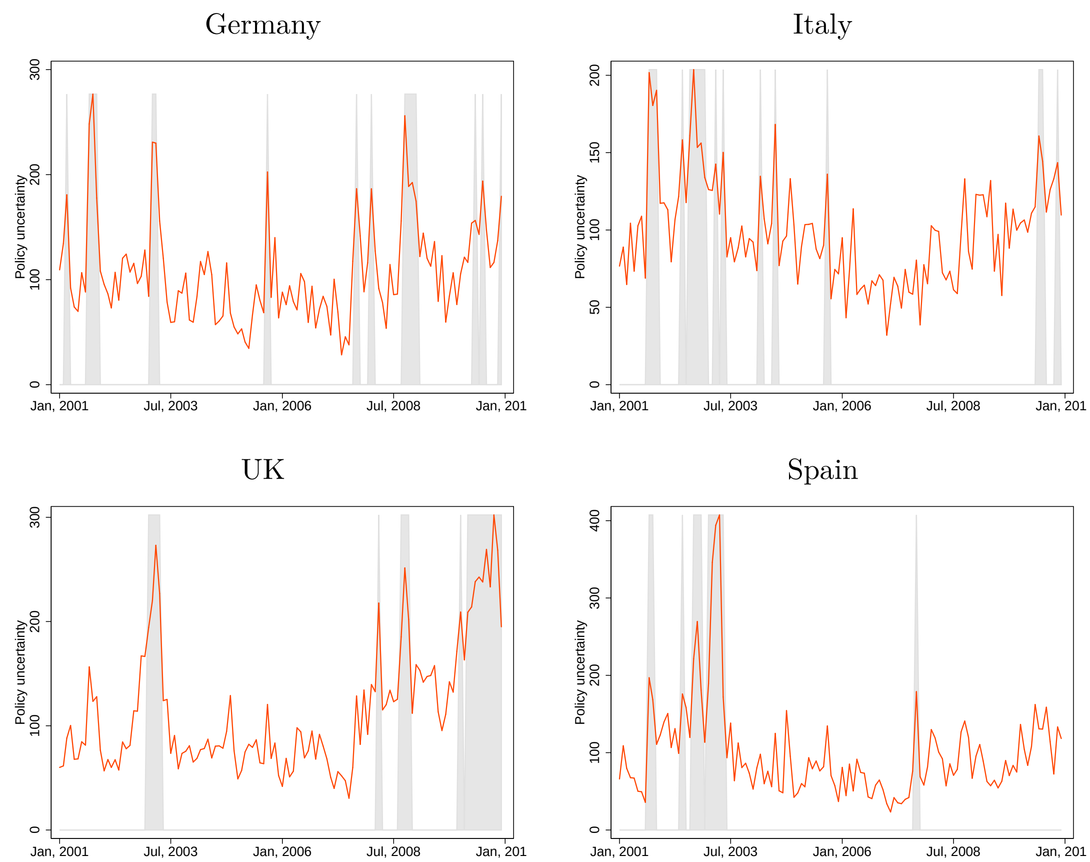

Table 8: ***Uncertainty, stickiness, and new relationships***

	(1)	(2)	(3)	(4)	(5)
Dep. var:			*# new trade relationships*		
Uncertainty	-0.05***	-0.06***		-0.09***	
	(0.005)	(0.005)		(0.007)	
$- \times RS$		-0.06***	-0.01***	-0.07***	-0.01***
		(0.010)	(0.002)	(0.014)	(0.003)
Fixed effects:					
Country	✓	✓		✓	
Product-month	✓	✓	✓	✓	✓
Country × Time			✓		✓
Method	Poisson	Poisson	LPM	Poisson	LPM
Threshold	1 s.d.	1 s.d.	1 s.d.	1.64 s.d.	1.64 s.d.
Obs.			1,986,261		

ships. The propagation of policy uncertainty to the real economy is thus intimately linked the type of relationships in which sellers and buyers are engaged.

6 Conclusion

7 Appendix

7.1 Equation (3)

That the distribution of productivities is Pareto (with parameters φ_{min} and k) means that

$$\mathbb{P}(\tilde{\varphi} \leq \varphi) = 1 - \left(\frac{\varphi}{\varphi_{min}}\right)^{-k}.$$

Now, under the assumption that a buyer faces an isoelastic demand with elasticity σ and simply resells the imported good with no value-added, we have $R(\varphi) = A\varphi^{\sigma-1}$ where A is a demand shifter. The c.d.f. of sales is then

$$\mathbb{P}(\tilde{R} \equiv R(\tilde{\varphi}) \leq R) \equiv \mathbb{P}(\tilde{\varphi} \leq \varphi) = 1 - \left(\frac{R^{\frac{1}{\sigma-1}}}{R_{min}^{\frac{1}{\sigma-1}}}\right)^{-k}$$

where $R_{min} = A\varphi_{min}^{\sigma-1}$. Sales are therefore distributed according to a Pareto with scale R_{min} and shape $\frac{k}{\sigma-1}$.

7.2 Empirical implementation

Instead of working with duration at the transaction level, we compute the average duration of a transaction within transaction-size deciles. The subscript p indexes the product category of the transaction.

$$\int_d^{d+1} \mathcal{T}(R)_{zpd} f(R) dR = \frac{(\gamma_p)^{k_p}}{\lambda} \frac{k_p}{\sigma - 1} \int_d^{d+1} \frac{1}{R} dR = \frac{(\gamma_p)^{k_p}}{\lambda} \frac{k_p}{\sigma - 1} log(\frac{R_{d+1}}{R_d}) \tag{5}$$

Because we have a Pareto distribution of sales: $\frac{d}{10} = 1 - \left(\frac{R_d}{R_{min}}\right)^{\frac{-k}{\sigma-1}}$. The integral over size deciles gives:

$$\int_d^{d+1} \mathcal{T}(R)_{zpd} f(R) dR = \frac{(\gamma_p)^{k_p}}{\lambda} log\left(\frac{10-d}{9-d}\right) = \eta_p log\left(\frac{10-d}{9-d}\right) \tag{6}$$

Note that $log\left(\frac{10-d}{9-d}\right)$ is not defined for $d = 9$. We have thus chosen to compute the average durations for the following size-intervals:

- $[P1, P10]$ (duration for transaction between the first percentile and the first decile)
- $[P10, P20], ..., [P80, P90]$ (duration between deciles)
- $[P90, P99]$ (duraction between the $8th$ decile (P90) and the last percentile (P99)

References

Acemoglu, Daron, Simon Johnson, and Todd Mitton, "Determinants of Vertical Integration: Financial Development and Contracting Costs," *Journal of Finance*, 06 2009, *64* (3), 1251–1290.

Antras, Pol and Robert W. Staiger, "Offshoring and the Role of Trade Agreements," *American Economic Review*, December 2012, *102* (7), 3140–3183.

_ , Davin Chor, Thibault Fally, and Russell Hillberry, "Measuring the Upstreamness of Production and Trade Flows," *American Economic Review*, 2012, *102* (3), 412–416.

Antràs, Pol and Davin Chor, "Organizing the Global Value Chain," *Econometrica*, November 2013, *81* (6), 2127–2204.

Baker, Scott R., Nicholas Bloom, and Steven J. Davis, "Measuring Economic Policy Uncertainty," *The Quarterly Journal of Economics*, 2016, *131* (4), 1593–1636.

Barrot, Jean-Noël and Julien Sauvagnat, "Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks," *The Quarterly Journal of Economics*, 2016, *131* (3), 1543–1592.

Behrens, Kristian, Brahim Boualam, Julien Martin, and Florian Mayneris, "Gentrification and Pioneer Businesses," 2018. Unpublished paper.

Bergounhon, Flora, Clémence Lenoir, and Isabelle Mejean, "A guideline to French firm-level trade data," 2018.

Bernard, Andrew, Andreas Moxnes, and Karen Helene Ulltveit-Moe, "Two-sided Heterogeneity and Trade," *Review of Economics and Statistics*, 2017, *Forthcoming.*

Bernard, Andrew B., J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott, "Intrafirm Trade and Product Contractibility," *American Economic Review*, May 2010, *100* (2), 444–448.

Besedes, Tibor and Thomas J. Prusa, "Product differentiation and duration of US import trade," *Journal of International Economics*, 2006, *70* (2), 339 – 358.

Bloom, Nicholas, "The Impact of Uncertainty Shocks," *Econometrica*, May 2009, *77* (3), 623–685.

Feenstra, Robert C. and Gordon H. Hanson, "Ownership and Control in Outsourcing to China: Estimating the Property-Rights Theory of the Firm*," *The Quarterly Journal of Economics*, 2005, *120* (2), 729–761.

Gaulier, Guillaume and Soledad Zignago, "BACI: International Trade Database at the Product-Level. The 1994-2007 Version," Working Papers 2010-23, CEPII research center October 2010.

Handley, Kyle and Nuno Limão, "Trade and Investment under Policy Uncertainty: Theory and Firm Evidence," *American Economic Journal: Economic Policy*, November 2015, *7* (4), 189–222.

— and — , "Policy Uncertainty, Trade, and Welfare: Theory and Evidence for China and the United States," *American Economic Review*, September 2017, *107* (9), 2731–2783.

Hausmann, Ricardo and Cesar Hidalgo, *The Atlas of Economic Complexity: Mapping Paths to Prosperity*, Vol. 1 of *MIT Press Books*, The MIT Press, January 2014.

Head, Keith and John Ries, "FDI as an outcome of the market for corporate control: Theory and evidence," *Journal of International Economics*, January 2008, *74* (1), 2–20.

Imbs, Jean and Isabelle Mejean, "Elasticity Optimism," *American Economic Journal: Macroeconomics*, July 2015, *7* (3), 43–83.

Levchenko, Andrei A., "Institutional Quality and International Trade," *Review of Economic Studies*, 2007, *74* (3), 791–819.

Mayer, Thierry and Soledad Zignago, "Notes on CEPIIŠs distances measures: The GeoDist database," Working Papers 2011-25, CEPII research center December 2011.

Monarch, Ryan, ""It's Not You, It's Me": Breakup In U.S.-China Trade Relationships," Working Papers 14-08, Center for Economic Studies, U.S. Census Bureau February 2014.

Novy, Dennis and Alan M. Taylor, "Trade and Uncertainty," NBER Working Papers 19941, National Bureau of Economic Research, Inc February 2014.

Nunn, Nathan, "Relationship-Specificity, Incomplete Contracts, and the Pattern of Trade," *The Quarterly Journal of Economics*, 2007, *122* (2), 569–600.

Pierce, Justin R. and Peter K. Schott, "The Surprisingly Swift Decline of US Manufacturing Employment," *American Economic Review*, July 2016, *106* (7), 1632–1662.

Postel-Vinay, Fabien and Jean Marc Robin, "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity," *Econometrica*, 2002, *70* (6), 2295–350.

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

Schmidt-Eisenlohr, Tim and Ryan Monarch, "Learning and the Value of Relationships in International Trade," 2015 Meeting Papers 668, Society for Economic Dynamics 2015.