Relationship stickiness and economic uncertainty

Preliminary and incomplete.

Julien Martin[*] Isabelle Mejean[†] Mathieu Parenti[‡]

September 21, 2020

Abstract

This paper examines how the degree of stickiness in business relationships influences the real impact of aggregate uncertainty. We first develop a new index of relationship stickiness (RS) for more than 5,000 HS6 products based on the duration of firm-to-firm trade. The RS measure is derived from a simple search model in which a higher degree of stickiness translates into a lower probability of switching and longer firm-to-firm trade relationships, conditional on match quality. Using firm-to-firm export data, we measure the duration of individual relationships and estimate RS for each HS6 product. We then show that RS shapes the dynamics of firm-to-firm relationships in response to uncertainty shocks. Uncertainty shocks induce a significant and larger decrease in the rate at which new firm-to-firm relationships are formed in high-RS product categories. This implies that aggregate uncertainty is especially costly for firms engaged in global value chains and that uncertainty shocks will impact these relationships mainly along the intensive margin.

[*]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 Introduction

For a firm, changing suppliers can involve substantial costs. As a consequence, firm-to-firm relationships often display a strong degree of stickiness.[1] One dimension of this stickiness comes from the product itself that the firms exchange. While switching costs might be limited for products anonymously purchased on spot markets, their size may be non-negligible in presence of relationship-specific investments, of customization costs, or in markets displaying informational or contractual frictions.

In this paper, we argue that relationship stickiness is important to understand the impact of uncertainty on economic activity. From Dixit (1989) and Dixit and Pindyck (1994), it is well known that uncertainty may affect investment behavior in creating an option value of waiting. In uncertain times, firms prefer to delay investments involving sunk costs until more information is revealed. As long as stickiness reflects important relationship-specific sunk costs, the impact of uncertainty should be magnified by relationship specificity. We test this assumption using a panel of firm-to-firm trade data involving French exporters and their European partners.

To this end, we first develop a new measure of products' relationship stickiness.[2] Our measure builds on the idea that the *duration* of firm-to-firm trade relationships conveys information on the relational specificity of products. The measure is backed by a simple theoretical framework of firm-to-firm input trade. In this model, firms receive offers randomly and decide to switch to a new input supplier whenever its offer is sufficiently below the price charged by the buyer's existing partner. In this environment, larger switching costs and/or more frictions contribute to lengthening existing firm-to-firm relationships. The duration of firm-to-firm relationships, conditional on the quality of their match, emerges as a relevant moment that can be exploited to recover a product-level measure of stickiness.

For estimation, we use French firm-to-firm export data, which panel dimension allows to follow importers over time, and compute the duration of each of their relationships with a French firm. The high degree of granularity in the data is exploited to control for heterogeneity in the quality of a match. Once match quality is controlled for, one can use the variability of average durations across products to recover a measure of relationship stickiness (RS) for more than 4,000 HS6 products.

[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.

[2] See the related literature section for a discussion of and a comparison with existing measures.

We present a body of evidence supporting the view that the recovered measure of relationship stickiness does capture some relational specificity at the product level. First, we show that the measure is stable when computed from French exports toward different destinations or over different time periods.[3] Second, we find that the stickiness measure varies substantially across sectors. Specialty chemicals or parts and accessories that entail large customization costs are found among the most sticky products, while motor vehicles or men suits are among the least sticky products. Third, we show that our measure correlates with existing proxies for relationship specificity found in the literature. For instance, it correlates with the ones developed by Rauch (1999) and Nunn (2007), but it displays substantially more variability, notably within the group of manufacturing products. More sticky products also turn out to be more complex products, products with a smaller elasticity of substitution, and products that are more upstream in the value chain. Fourth, we find that the measure relates to product-level outcomes in a way that is consistent with three theoretical results from the literature: sectors with more sticky products display a higher share of intrafirm trade as predicted by Antràs and Chor (2013); relational stickiness interacts with institutional quality to shape countries comparative advantages as in Nunn (2007); trade of more sticky products is more sensitive to distance, consistent with the view that the information and monitoring costs related to distance are exacerbated by relationship stickiness (Rauch, 1999; Head and Ries, 2008).

Armed with this new measure, we study how uncertainty shocks affect the dynamics of trade relationships in more or less sticky product markets. Uncertainty has been shown to be a potential threat to economic growth as it reduces firms' incentive to invest (Bloom, 2009). Our analysis focuses on one particular type of investments, namely investments associated with the firm's international expansion. International trade is indeed associated with substantial sunk costs that affect trade at the extensive margin. For this reason, this dimension of firms' activity is strongly affected by uncertainty shocks (Novy and Taylor, 2019). We provide evidence consistent with this view using the formation of new firm-to-firm relationships as an outcome variable.

We bring together micro data on firm-to-firm relationships and macro data on uncertainty to quantify the trade impact of high uncertainty episodes. Our analysis uses the "World Uncertainty Index" developed by Ahir et al. (2019) to recover external measures of uncertainty shocks. The database covers 143 countries, including 12 countries in our sample, from 1996 onwards and measures uncertainty at a quarterly frequency using text mining techniques

[3] We also show that the measure correlates with the duration of trade relationships between countries, suggesting that our duration-based measure captures product characteristics that are not specific to France.

applied to the quarterly Economist Intelligence Unit (EIU) country reports. Based on this database, it is possible to construct a panel of "uncertainty shocks" that we merge with product-level information on the number of new (and disrupted) relationships involving French firms and their European partners, at a quarterly frequency.

The number of new trade relationships involving French exporters and European importers is consistently found lower in periods of high policy uncertainty. Quantitative effects somewhat vary depending on the specification, with a contemporaneous effect varying between -1 and -9%, and a persistence over at least six months. More interestingly, we show that the impact of uncertainty is especially pronounced in product markets that feature a high degree of stickiness. The impact of uncertainty on the establishment of new firm-to-firm relationships varies between zero and minus two percent when moving along the distribution of RS, from the least to the most sticky markets. In high stickiness product markets, uncertainty episodes are associated with a pronounced reduction in the number of new trade relationships, a contraction of disrupted trade relationships and an increase of trade at the intensive margin. This is consistent with the view that such markets feature high costs of switching from a supplier to another. Because the option value of staying with an existing partner increases in period of high uncertainty, we observe less turnover in firm-to-firm relationships in risky time in markets characterized by such high switching costs. Stickiness in firm-to-firm relationships is a significant driver of the response of the economy to policy uncertainty. A corollary of our results is that uncertainty is especially costly for firms engaged in global value chains, whose products are characterized by a high degree of stickiness (Antràs and Chor, 2013).

The heterogeneous impact of uncertainty across product categories sheds lights on the importance of relationship-specific investments in international trade. The associated costs are likely to matter for a range of economic outcomes. The elasticity of trade flows is likely to depend on the stickiness of trade relationships. This implies that the responses of trade to exchange rate shocks or trade policy might differ along this dimension of heterogeneity. The international transmission of shocks shall also be impacted by the degree of stickiness, with implications for the level and the comovement of economic fluctuations. We hope that the measure developed in this paper will stimuate research along these questions.

Related literature. The paper mostly contributes to two strands of the literature, namely the trade literature on relation-specific investments and the macroeconomic literature on uncertainty and the business cycle. The importance of stickiness in an international context has been underlined repeatedly. 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). Last, the purpose and design of trade policy (shall) depend on the specificity of business relationships (Antràs and Staiger, 2012).[4]

In this literature, relationship specificity is usually proxied using either the measure developed by Rauch (1999), or the measure developed by Nunn (2007).[5] Our contribution to this literature is a novel measure of product relationship specificity at a disaggregated level recovered from the duration of firm-to-firm trade relationships.

The measures developed by Rauch (1999) and Nunn (2007) rely on a characterization of the markets on which products are traded. 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. Nunn's measure uses Rauch's classification to assess the specificity of inputs entering production processes, a good being called more "specific" when its production is more intensive in differentiated inputs. Whereas such classifications have proved useful, we propose a measure computed at a finer level of disaggregation, and which captures the impact of a wider set of product-market characteristics contributing to stickiness.

From this point-of-view, the closest paper to ours 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.[6] 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.[7]

[4] Outside of the trade literature, input specificity has also been shown to be a significant determinant of the propagation of shocks in value chains (Barrot and Sauvagnat, 2016).

[5] 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 (Antràs et al., 2012).

[6] In this lighter procedure, relationship stickiness is evaluated in relative terms along the distribution of products. As a consequence, we can not directly interpret our estimates in terms of a monetary switching cost, which Monarch (2014) can do.

[7] The present 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

Beside constructing a new measure of business stickiness, our paper contributes to the literature on uncertainty and economic growth. Following the seminal contribution by Dixit and Pindyck (1994), a large theoretical and empirical literature has emerged, that studies the consequences of uncertainty in macroeconomics. At the microeconomic level, uncertainty is empirically shown to affect the relationship between patenting and firms' productivity (Bloom and Reenen, 2002), the responsiveness of investment to demand shocks (Bloom et al., 2007), or hiring decisions (Schaal, 2017). In the aggregate, the level of policy uncertainty affects aggregate output and employment (Bloom, 2009). Closer to us is the literature on uncertainty and trade. Novy and Taylor (2019) link uncertainty to the volatility of international trade. A series of papers discuss the reduction in policy uncertainty induced by Portugal's accession to the European Community (Handley and Limao, 2015), and China's entry into the WTO (Handley and Limao, 2017; Pierce and Schott, 2016), and how it explains the boom in exports after entry. Several papers have also explored the impact of Brexit-driven uncertainty on trade. Graziano et al. (2018) document significant extensive and intensive responses of product-level trade flows to changes in uncertainty tied to the Brexit. In comparison with this literature, we provide further evidence that uncertainty affects trade at the firm-to-firm extensive margin and that the effect is more pronounced in stickier product markets.

The rest of the paper is organized as follows. Section 2 describes the firm-to-firm data used throughout the paper, and provides stylized facts on the structure and duration of firm-to-firm relationships. Section 3 derives our measure of relationship stickiness from a parsimonious search model, explains how it is estimated and discusses how it compares with alternative measures used in the related literature. Section 4 is devoted to the empirical investigation of the trade impact of policy uncertainty. Finally, section 5 concludes.

2 Data

2.1 Dataset

Both our measure of relationship stickiness and the empirical analysis based on this measure take advantage of a panel of firm-to-firm trade data provided to us by the French customs and described in Bergounhon et al. (2018). The dataset covers each single export transaction between French firms and their individual partners in the European Union. Importantly, the data 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

of relationship stickiness, controlling for individual characteristics.

shipment (in euros). The dataset covers the period from 1993 to 2017 but the analysis exploits various sub-periods. The main reason why we do not work on the whole panel is that the nomenclature for product categories, which is exploited to characterize product markets by their stickiness, changes over time.[8] As a consequence, we use the harmonization algorithm described in Behrens et al. (2018) to recover time-invariant product categories. The induced information loss is minimized when the algorithm is applied over shorter horizons, which explains that we work on various sub-periods.[9] In the baseline specification, relationship stickiness is measured using data from 1996 to 2006, with left-censoring controlled for using pre-1996 data. We also check the robustness over time, using the 2011-2017 period as an alternative.[10] Armed with the corresponding measures of relationship stickiness, we then assess the role of uncertainty on the dynamics of firm-to-firm relationships using data over 2000-2010 and 2011-2017.

For each product category, we observe all firm-to-firm relationships involving a French exporter and a European buyer, over time.[11] 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 101 millions firm-to-firm transactions. Table 1 provides descriptive statistics on the dimensionality of the data, in the overall European Union as well as in each destination. We observe almost 110,000 different French exporters over the period, that interact with 1.7 million 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 19.5 millions firm-to-firm relationships, that thus interact over 5 transactions, on average.[12]

[8]Another constraint that we take into account while selecting the sub-periods of analysis is a discontinuity in the data between 2010 and 2011 attributable to a change in the declaration threshold above which firms report the type of products being exported.

[9]See the discussion in Bergounhon et al. (2018) of the trade-off incurred when working with such harmonization algorithms, between working on longer periods, and maintaining the granularity of the (harmonized) product classification.

[10]The 2011-2017 sample comes with some caveats as the threshold to declare the product category of exports has jumped from 150K euros per year to 460K euros in 2011. This raises concerns to properly identify the beginning or the end of trade relationships.

[11]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 21,616 exporters (out of 66,131) accounting for 2% of transactions and .5% of the value of French exports.

[12]The data do not allow to distinguish arm's length and intrafirm trade. Some of these firm-to-firm relationships may thus involve affiliates of the same multinational firm.

2.2 Structure of the firm-to-firm data

While the dimensionality of the data is important, not all transactions are used in the estimation sample. In particular, the estimation of relationship stickiness neglects all buyers that we observe over a single transaction throughout the period under study. These represent as much as 44% of the importers in the raw data (see Figure 1) and 1.5% 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.[13] 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 foreign buyers also display a lot of heterogeneity regarding the "intensity" of their relationships with French exporters. This heterogeneity is illustrated in Figure 1. 15% of foreign buyers are observed over two transactions, with 20% 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, about 15% of foreign buyers are observed over more than 10 different transactions, often involving several partners.[14] These firms are good candidates to observe the duration of their relationships with French firms.

Figure 2 shows the distribution of the number of French partners, individual buyers interact with, over the whole timespan of their presence in the dataset.[15] 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 use information on the number of partners per buyer, 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 third 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

[13] Indeed, we do observe that the time distribution of these one-time transactions is not homogenous across years. In particular, the share of one-transaction buyers is significantly larger in the first and last periods of the sample.

[14] Almost one out of ten buyers observed over more than 10 transactions always interact with the same French seller.

[15] Here and in the rest of the paper, statistics are based on the sub-sample that excludes buyers appearing just once in the data.

relationships in international markets are especially sticky. The question that the empirical analysis addresses is whether this is systematically related to the specificities of some products or sectors.[16]

We terminate this raw description of the data structure with a last stylized fact later used to motivate our econometric model. 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 purchasing a particular product from a single seller while instead sellers simultaneously serve several importers (even within a country). This is illustrated in Figure 3 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 55% when we pool partners located in different countries.[17]

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 relationships with French firms.[18] Table 2 provides statistics on this variable. In this table, we take the perspective of a particular importer of a good and compute, for each of these individuals, the mean duration of its relationships with French firms ("Mean duration"),

[16]Some of these faithful relationships may take place within multinational companies. As explained in footnote 12, the data do not allow to control for this possibility and we cannot exclude the corresponding relationships from the estimation of relationship stickiness as a consequence. We however check ex-post that estimated relationship stickiness indeed correlates with an (external) measure of intra-firm trade.

[17]This is in contrast with Bernard et al. (2018) 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.

[18]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. To control for right-censoring, we neglect from the analysis all firm-to-firm relationships starting during the last year of the estimation sample.

the frequency of its transactions with French firms ("Frequency of transactions") and the probability that, conditional on switching from one French exporter to another, it returns back with an exporter it had already interacted with in the past ("Proba Recall").

Statistics show considerable variability across importers, in terms of the mean duration of their relationships with French firms. The median importer in our data interacts with a given French firm over a period of ten months. A small number of importers however display substantially stickier relationships, as shown by the mean of the distribution that is substantially above this number, at 18 months. While these statistics are calculated for each importer, averaging across all relationships it has with French suppliers, 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.*

Another important feature of the data, which the model also takes into account, is that the frequency of transactions displays significant heterogeneity across importers (second line in Table 2). On average, the probability of a transaction occuring in a given month is equal to .332, which corresponds to a transaction every 3 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 10 months. 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. This means that we do not keep the whole history of an importer's partners in memory and instead suppose that a new relationship with a French seller that the importer had already interacted with before is equivalent to a new relationship with a new partner. Abstracting from 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 (See the last line in Table 2).

3 Measuring relationship stickiness

3.1 Empirical strategy

Section 2 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. In this section, we explain how to build and estimate a measure of relationship specificity at the product-level from the duration of firm-to-firm relationships in trading this product. The model is a simple search model between sellers and buyers of a given product. It is assumed that products systematically vary in terms of their degree of business stickiness because of heterogeneous search frictions and heterogeneous costs associated with switching from one supplier to the other. 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 products or any other product-specific characteristics. We remain purposely agnostic on the exact micro-foundations at the root of such stickiness and we consider instead these features as given and use the cross-section of products to quantify its relative size across product categories.

Suppose that an importer is willing to purchase a certain product. Every period, it receives with probability λ an offer $\tilde{p}$ from a new input supplier and 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 an offer that is not only better but also covers its switching cost. Namely, the firm decides to switch whenever $\tilde{p} < \frac{p}{\gamma}$ where $\gamma > 1$ is a parameter governing the size of switching costs. This occurs with a probability $\lambda H_P(p/\gamma)$.[19]

Under these conditions, the length of a buyer-seller relationship, conditional on its price follows a geometric law with mean

[19] Here, it is implicitly assumed that p is 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. Alternatively, one may argue that the new offer induces the importer and its existing partner to renegotiate "on-the-match" (see Postel-Vinay and Robin (2002) for an application of this assumption in the context of frictional labor markets). While such assumption would deliver that firm-to-firm prices tend to decrease with the age of the buyer-seller relationship (consistent with Fontaine et al. (2020)), it would not affect the expected duration of a firm-to-firm relationship, the object of interest in this paper. The reason is that in our model as in this framework, the importer always ends up interacting with the firm with the lowest cost of serving her, which does not depend on the supplier's price offer but only her ability to beat potential competitors.

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

This generalizes in continuous time where offers follow a Poisson process and the probability of receiving an offer during an infinitely small period of time dt is λdt. The duration $\mathcal{T}$ of a relationship at price p then follows an exponential law $\mathcal{E}$ with parameter $\lambda H_P(p/\gamma)$ denoted by

$$\mathcal{T}|p \sim \mathcal{E}[\lambda H_P(p/\gamma)]$$

The expected duration of a relationship is thus the inverse of the probability of switching. It is a function of the firm's existing deal p, the product-specific 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 quality-adjusted 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.

Confronting this model to the data forces us to make two additional parametric assumptions. First, we assume that the distribution of quality-adjusted prices is inverse-Pareto with shape parameter k. Second, the importer's demand curve is assumed to be iso-elastic with $\sigma > 1$ being the price elasticity of demand.[20] Under these assumptions, one can write the distribution of durations conditional on the size r of the transaction, instead of the (unobserved) price offered by the supplier:

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

where r_{min} is the lower bound of the distribution of transactions and $\eta \equiv \frac{\gamma^k}{\lambda}$. From now on, we will interpret η as a product-specific indicator of relationship stickiness, capturing various forces that tend to lengthen firm-to-firm relationships, conditional on a match. In the context of our model, longer durations conditional on a match can be the outcome of less frequent offers (a low λ), large switching costs (a high γ) or little dispersion in the distribution of price offers (a high k). While the parametric assumptions required to recover equation (2) are obviously important, we argue that the model's insight, namely that such product-specific

[20] Taken together, these two assumptions imply that the distribution of observed transactions between buyers and sellers is approximately Pareto for large transactions. This agrees for instance with the canonical model of firm heterogeneity under monopolistic competition (Melitz and Redding (2014))

characteristics tend to increase durations, conditional on a match, is likely to hold under alternative parametric assumptions.

Our data is a vector of realized durations for all relationships involving a European buyer and a French exporter. To recover the parameters of equation (2), we use the statistical properties of the product-specific empirical distribution of these random variables. Under the model's assumptions, the expected duration of a relationship conditional on transactions r falling in the quantile of order q of its product-specific distribution writes:[21]

$$\begin{aligned}\mathbb{E}\left[\mathcal{T} \mid R \in R_q\right] &= \mathbb{E}\left[\eta \left(\frac{R}{r_{min}}\right)^{-\frac{k}{\sigma-1}} |R \in R_q\right] \\ &= \eta \ \ln\left[\frac{\mathbb{P}(R \geq r_{q-1})}{\mathbb{P}(R \geq r_q)}\right] \end{aligned} \tag{3}$$

where R_q denotes the q^{th} quantile of the distribution:

$$R_q := [r_{q-1}, r_q] \equiv \left\{ r \mid \bar{H}_R^{-1}\left(\frac{q-1}{Q}\right) \leq r \leq \bar{H}_R^{-1}\left(\frac{q}{Q}\right)\right\}$$

and $\bar{H}_R(r) \equiv \mathbb{P}(R > r)$. Equation (3) being log-linear in η and η being specific to product categories but constant across countries and transaction size, it is enough to use a fixed effect model to recover an estimate of the product-specific index of relationship stickiness, up to a constant. See the details of the empirical implementation in Appendix A.2.

3.2 Stylized facts on relationship specific indicators

Using the strategy described in Section 3.1, we recover the relative level of stickiness for 5,186 HS6 products. Figure 4 shows the distribution of estimates. We see that there are significant differences in the level of relationship stickiness across HS6 products, with a mean at 2.86, a median at 2.97 and an interquartile range of 0.63. Since we estimate the log of the η parameter, an interquartile of .63 means that the expected duration of trade flows, conditional on the quality of the match, is 1.88 times larger at the 75th percentile of the product-distribution than at the 25th. This is sizable. Note that the precision of estimates however varies across products, as shown by the grey area in Figure 4.[22] We take this margin of error into account

[21] The second line uses the property of the Pareto distribution. If X is Pareto distributed with shape parameter κ and locus x_m, then $\frac{q}{Q} = 1 - \left(\frac{X_q}{X_m}\right)^{\kappa}$, where Q is the number of cut points, and X_q the value for the q^{th} cut-point. See details in Appendix A.2.

[22] Given the size of standard errors, listing products found at each tail of the distribution might not be very informative. A glance at the data reveals that among the most relationship-specific products are a number of industrial chemical, pharmaceutical and mineral products. These are goods that are highly differentiated, are

in the empirical analysis using a parametric bootstrap procedure. For each product, we make 400 draws in a Gaussian distribution whose mean is the estimated stickiness of the product and the standard deviation its associated standard error. We then run 400 regressions using the relationship stickiness generated in these draws. The coefficients and their standard errors are obtained by taking the mean and standard deviation of these estimates.

Table 4 shows how our measure correlates with other product-specific attributes used in the literature. The first column reports the pairwise correlation coefficients while column (2) reports the coefficients of a regression of our RS measure on all other characteristics. The degree of product stickiness is positively correlated with alternative measures of product specificity used in the literature, most notably Rauch (1999) and Nunn (2007). Differentiated products tend to be more relationship-specific, as shown by the positive correlation with the dummy for differentiated products recovered from Rauch (1999) and the negative correlation with elasticities of substitution estimated in Imbs and Mejean (2015). More complex goods also involve more stickiness, as shown by the positive correlation of our indicator with both Nunn (2007) and Hausmann and Hidalgo (2014). Finally, we find a positive correlation between the level of upstreamness of a product and its degree of stickiness. This correlation suggests products far from the final demand entail more buyer-specific investment, more elaborated contracts or more customization than products dedicated to final consumption. This is consistent with Antràs and Chor (2013) view that global value chains entail substantial locked-in effect.

Although these correlations all have the expected sign, the linear combination of existing indicators only explains 7% of the heterogeneity recovered from our estimation (column (2)). The reason is that the RS indicator is extremely heterogeneous, including within particular industries.[23] The high disparity of stickiness is further illustrated in Figure 5, which compares average RS measures across categories of the broad economic classification (BEC).[24] Final consumption goods like cars or consumer goods display on average a low level of stickiness

often customized to the particular needs of the firm's client and tend to be purchased frequently, thus generating long-lasting relationships between the firm and its client. These facts may seem surprising as chemicals are often considered as homogenous products. It is worth stressing that the chemical industry is split between commodity and specialty chemicals, the latter being chemicals that are tailored for each client. At the other side of the distribution, one finds a number of final good products that are usually produced in large quantities and sold in anonymous markets (e.g. Men's suits), some non-differentiated primary goods (Ferro-alloys or Raw Sild), and a number of capital goods such as machines used in the textile industry that are purchased infrequently and are not subject to relationship stickiness as a consequence.

[23] A typical example is the degree of stickiness in the mineral product industry that encompasses some of the most and the least relationship-specific products.

[24] BEC categories corresponding to less than 5 hs6 products are excluded from the figure.

while parts and accessories or food processed for the industry display higher average RS indices. Here as well however, we see significant level of dispersion within these categories.

One may wonder the extent to which the high dispersion in the recovered distribution of RS indices is not the consequence of the estimation strategy identifying relationship-stickiness using a fixed effect, which is likely to absorb part of the noise in the data. While this possibility cannot be entirely ruled out, we argue that what we recover is not pure noise as the estimated indicators appear relatively stable both across countries and over time. **TO BE UPDATED** We have indeed estimated a country-specific distribution of RS indicators using the same empirical strategy but country-specific datasets.[25] The pairwise correlation between country-specific measures is around .5. More importantly the correlation between the baseline distribution recovered from the pooled sample and the country-specific estimates is significant, between 69 and 74%. Stability of our estimates is also assessed over time, by estimating relationship stickiness using the 2011-2017 period. Here as well, the correlation is significant, at .6. Note that our baseline distribution of RS indices is the one estimated over 1996-2006 since the underlying customs data are of better quality in the earlier years of the sample.[26]

3.3 External validity tests

We conclude the description of the proposed measure of relationship-specific indices by reproducing some results in the literature showing a significant impact of relationship-specific investments on various trade outcomes. Whereas the literature is based on alternative measures of relationship-specificity, we show that the same qualitative results are confirmed using our indicator instead. We think of these exercises as useful sanity checks that the measure we will later exploit to assess the impact of uncertainty in sticky product markets indeed captures what it is meant to. We summarize insights recovered from these exercises in the main text and report the detailed results in Appendix A.3.

Antràs and Chor (2013) argue that the locked-in effect induced by relationship-specific investments may have consequences for firms' propensity to vertically integrate. In their property-rights model, downstream firms have an incentive to integrate suppliers because of contractual frictions in the procurement of a customized component later integrated in

[25] Here, we focus on those countries that are important destinations for France exports, namely Belgium, Germany, Italy, Spain, and the UK, as the empirical strategy requires to observe a sufficient number of firm-to-firm relationships for each product category.

[26] As mentioned in Section 2, the dataset is censored because the Customs do not request firms to declare the product exported, below a certain threshold. Since the threshold has been tripled in 2011, censoring is significantly larger in the most recent period.

the production. A corollary of such framework is that vertical integration should be more pervasive in product markets that display more intense locked-in effects. We provide evidence that this is the case in Table A.2 using the prevalence of intra-firm trade as a measure of vertical integration. Using US trade data, we show that the share of intra-firm in overall trade is significantly larger for products displaying high RS indicators. Alone, RS explains 10 percent of the cross-product dispersion in the data.

Nunn (2007) and Levchenko (2007) argue that high relationship-specific investment goods require sounds institutions, in the form of quality of contract enforcement, property rights, shareholder protection, etc. As a consequence, institutions can shape the geography of trade as other sources of comparative advantages. In Table A.3, we replicate the empirical exercise in Nunn (2007) using more disaggregated data and our measure of relationship-stickiness. We further control for the relation-specificity measure developed by Nunn (2007) to identify an effect beyond and above Nunn's. Both his and our measures point in the same direction, namely a specialization of countries with better institutions into the production of goods that are more relationship-specific.

Finally, the impact of relationship-specificity is investigated in a gravity context. Namely, Table A.4 interacts distance with our measure of relationship stickiness in an otherwise standard gravity equation for bilateral trade. We find that the negative impact of distance on trade flows is stronger for product categories that exhibit a higher degree of relationship stickiness. While a structural interpretation is not possible in this reduced-form context, several theoretical mechanisms can help rationalize the evidence. First, the increased distance elasticity may explain by information frictions being large in more sticky markets, which on the one hand increases the cost of switching to a new supplier and on the other hand reinforces the geographic concentration of trade (Rauch, 1999). An alternative interpretation is that stickier relationships are associated with higher monitoring costs which increase with distance (Head and Ries, 2008).

4 Trade, uncertainty and stickiness

With the measure of relationship-stickiness at hand, we now turn to the paper's core question, namely how product stickiness shapes the impact of economic uncertainty.

4.1 Empirical strategy

The assumption tested is that uncertainty affects trade patterns by impeding the creation of new business relationships. If there are sunk costs associated with establishing 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 hypotheses by combining our index of stickiness with measures of policy uncertainty.

The Poisson specification estimated in the analysis takes the following generic form:

$$E(X_{pct}|Uncertainty_{ct}, RS_p, FE) = exp\left(\alpha Uncertainty_{ct} + \beta RS_p + \gamma RS_p \times Uncert_{ct} + FE\right) \tag{4}$$

where the left-hand side variable which is computed for each country c, product p and period t can be the overall value of French exports, the number of new firm-to-firm relationships or the number of relationships ending between period t and $t+1$. $Uncertainty_{ct}$ is a measure of policy uncertainty for country c and period t, which is described below. The tested assumption is that uncertainty impedes trade and the α coefficient is thus expected to be negative. By digging into various margins of international trade, we can further investigate the mechanisms. In particular, we expect the impact of uncertainty to mostly work through the extensive margin, i.e. the net creation of firm-to-firm relationships, which is associated with considerable sunk costs and should thus react to uncertainty. Finally, we also expect the impact of uncertainty to be particularly pronounced in high relationship-stickiness sectors, thus the interaction term between uncertainty and the RS indicator ($RS_p \times Uncert_{ct}$). As explained below, the regressions include various vectors of fixed effects, that allow to control for unobserved heterogeneity between products, countries and periods.

Uncertainty episodes are measured at the country and quarterly levels using the "World Uncertainty Index" (WUI) developed in Ahir et al. (2019). Using text-analysis of the Economist Intelligence Unit country reports, they construct uncertainty index series for 143 countries, at a quarterly frequency from 1996 onwards. The approach to construct the WUI is to count the number of times uncertainty is mentioned in the EIU country reports and scale these numbers by the total number of words in each report. Based on the WUI series for 12 countries in our sample, we define uncertainty episodes as periods where the uncertainty index is one standard deviation above its average level.[27] The corresponding series are matched with the firm-to-firm trade data for 2000-2010 and 2011-2016, aggregated to the quarterly frequency to fit with the WUI data. We define new interactions as the first transaction involving a particular pair of firms, going back to data from 1995 to avoid left-censoring. Disrupted trade relationships are defined as transactions involving a seller and a buyer that never interact again over the

[27] We have also tried with uncertainty episodes defined as periods in which the index is 1.64 standard deviation above its average. Results are virtually unchanged. In Table 5, we also present results that directly use the level of the index to measure uncertainty.

rest of the period of data availability.

4.2 Results

Benchmark results are presented in Table 5 where we use the number of new firm-to-firm trade relationships as left-hand side variable. Results consistently show that episodes of high uncertainty are associated with significantly less creations of new firm-to-firm relationships, especially in product markets displaying stickier relationships. In column (1), the cross-product heterogeneity in the elasticity of trade to uncertainty is identified with product×quarter fixed effects to control for seasonality and country×period fixed effects so that the identification is across products. The negative coefficient on the interaction suggests that, in comparison with normal times, periods of high uncertainty are characterized by significantly less new firm-to-firm relationships in high RS product markets. The result is confirmed in Column (2), when the identification is within a product×period, across countries that do or do not experience high uncertainty episodes. In comparison with others, destinations that feature high uncertainty episodes are characterized by a significantly lower rate of creations of new firm-to-firm relationships, especially in high relationship-specific product markets. In quantitative terms, specification (1) implies that the number of new relationships for a product at the first quartile of the distribution of the RS indicator drops by about 2% in periods of high uncertainty. For a more sticky product, at the third quartile of the distribution of RS, the number of new relationships drops by 11% in periods of uncertainty.[28]

Columns (3) and (4) use the value of the World Uncertainty Index rather than dummies for high uncertainty episodes. The finding that the negative effect of uncertainty on new relationships is stronger for more sticky products is robust to using this alternative measure of uncertainty. Finally, column (5) tests the robustness of the result to the period of analysis, using 2011-2016 instead of the 2000-2010 period used in the baseline. As mentioned in Section 2, the quality of the data is somewhat reduced in the most recent period as left-censoring induced by the declaration threshold for product categories in the customs data has severely increased in 2011. As a consequence, the measure of new firm-to-firm relationships is noiser. Still, results based on these data are consistent with the baseline, although the impact of uncertainty is no longer significant for the average RS product.

In Table 6, we dig deeper into the trade effect of uncertainty, by looking at its correlation over space and over time. In columns (1)-(2), we test for spillover effects of uncertainty shocks,

[28] The first quartile is 2.61 and the third quartile is 3.23. For the first quartile, we thus compute: $E(X|Uncertainty = 1)/E(X|Uncertainty = 0) - 1 = exp(.37 - 2.61 \times 0.15) - 1$ ==-0.021. The same formula for the third quartile gives -0.108.

on other destination countries. In theory, an uncertainty shock in country c could have two opposite effects on trade towards alternative destinations. Trade could be *diverted* to these destinations, so that the decline in the creation of new trade relationships in the destination hit by the shock would be compensated by an increase in the rate at which French exporters establish new relationships with importers from other countries. But the spillover could also be negative, with uncertainty in one country reducing firms' incentive to invest in new trade relationships in other destinations as well. This is expected to be the case if firms perceive uncertainty to potentially spread to the rest of the EU. To address this question empirically, the specification in Column (2) of Table 5 is augmented with a variable summing uncertainty shocks occurring in all alternative EU destinations but the country itself, interacted with the RS indicator. Results suggest that the overall spillover effect is negative, i.e. shocks in another EU destination induce a significant decline in the creation of new trade relationships in the country under study. As expected, the effect is smaller than the impact of the shock itself and is significantly larger in product markets displaying higher relationship stickiness. Finally, the same qualitative results are found in Column (2), using identification across products within a destination$\times$period.

In columns (3)-(4) of Table 6, we study the dynamics of the trade impact of uncertainty. To this aim, the benchmark specification is augmented with lags of the uncertainty variable, interacted with the RS indicator. The impact of uncertainty is shown persistent over time, for at least four quarters.

Up to now, the analysis has focused on the creation of new firm-to-firm relationships as an outcome variable. The reason is that the trade literature has extensively discussed the prevalence of fixed costs as a barrier to the international development of firms. To the extent that some of these costs are sunk, the extensive margin of trade should be especially sensitive to uncertainty. In Table 7, we complete the analysis with regressions studying how other margins of international trade adjust to uncertainty. In columns (1) and (2), we reproduce the analysis in Table 5 using the number of disrupted trade relationships as an outcome variable. A disrupted relationship is defined as a transaction involving a seller and a buyer that will never interact again in the future. To the extent that uncertainty episodes create an option value of waiting, exporters should be reluctant to end a costly trade relationship when uncertainty is high. Results suggest that it is indeed the case for products displaying high relationship stickiness. Namely, episodes of high uncertainty are associated with significantly more disruption in trade, but the coefficient on the interaction is significantly negative, meaning that high RS products instead display significantly less disrupted relationships in uncertain time. More specifically, the estimates imply that uncertainty leads to a 2% increase in the number of disruptions for a product at the first quartile of the distribution of RS, and a 2% decrease in the number of disruptions for a product at the third quartile of the distribution of RS. Uncertainty thus has an asymmetric impact on products depending on their level of relationship stickiness. The fact that sticky relationships are not more likely to be disrupted in period of high uncertainty is consistent with **Antras** argument that in GVC, most of the trade decline is along the intensive margin - precisely because of the high cost of destroying specific relationships.

Finally, columns (4) and (5) in Table 7 present results using the log of product-level bilateral trade as left hand-side variable. Here, expected results are more ambiguous. The intensive margin of trade is not expected to adjust to more uncertainty while the overall impact of uncertainty at the extensive margin is ambiguous since exporters limit their investments in new trade relationships whereas the number of disrupted reationships can increase or decrease depending on the product. Results show that the overall impact of uncertainty is negative in the mean product market but becomes positive for high enough RS products.

We end up the analysis by examining what our results imply regarding the dynamics of trade in period of high uncertainty. We follow Bricongne et al. (2012) and decompose the mid-point growth in product-level trade of French exports into an intensive margin component, the growth of trade within continuing relationships, and two extensive margin components, the start of new relationships and the ending of existing relationships. Unlike Bricongne et al. (2012) but consistent with our previous investigations, our unit of analysis is at the seller-buyer level. The entry term is thus the start of a new relationship between a French seller and a foreign buyer for a given product. The seller may thus have been active in the market during the previous period.

The results of the decomposition are reported in Table 8. The first column shows uncertainty has a negative impact on the growth of product-level trade that is magnified in stickier markets. Quantitatively, the mid-point growth of trade flows is 4.5% lower in period of high uncertainty for the median product, with the growth contraction varying between 4.1% and 4.7% when moving from the first to the last decile of the distribution of relationship stickiness. These aggregate patterns however hide interesting heterogeneity regarding the margin of the adjustment, as shown in the next 3 columns. First, negative adjustments at the intensive margin (a slowdown in the volume of trade) are mainly driven by high stickiness products. Second, whereas more stickiness magnifies the slowdown of the creation of new relationships in uncertain time, the impact of uncertainty shocks on relationships ending changes sign when moving along the RS distribution. In less sticky markets, the decline in the formation of new relationships is reinforced by an acceleration of relationships ending. Instead, the number of terminated relationships is lower in uncertain periods for high-enough RS products. This pattern is consistent with the interpretation of extensive margin adjustments in terms of the option value of waiting. In sticky markets, where the option value of waiting is particularly large because of larger sunk costs, uncertain periods are characterized by a freezing of decisions whereby firms neither start nor end business relationships. When such decisions are frozen, the only margin through which trade can adjust is the intensive one. Note that all these results are robust to introducing product-country fixed effects as shown in tables 5-8.

The results of this decomposition offer an explanation to the importance of intensive margin adjustments during the crisis. For instance, Bricongne et al. (2012) show most of the trade drop during the 2008 crisis has occurred along the intensive margin, contrary to what most models of trade predict. Pol Antras (add ref) hypothesizes that this pattern can be explained by the importance of lock-in effects in GVCs. We provide evidence that relationship stickiness does limit the adjustments along the extensive margin in period of uncertainty, which may thus explain what occurs during the crisis that are often associated with high uncertainty levels.

5 Conclusion

In this paper, we discuss the extent to which stickiness in firm-to-firm relationships can amplify the real impact of uncertainty with a particular emphasize on international trade. This is a topical question for at least two reasons. On the one hand, uncertainty is prevalent in the current international context. Firms engaged in international markets have to cope with the uncertainty induced by negotiations over the Brexit, the trade war involving the United States and most of its partners and now the consequences of the Covid-19 pandemy. On the other hand, sticky trade relationships are prevalent within Global Value Chains. The fragmentation of production processes generates locked-in effects as a consequence of relationship-specific investments (Antràs and Chor, 2013). Firms engaged in GVCs often need to customize product or adjust their logistics chain to the particular needs of the firm located downstream. This generates a substantial degree of persistence in firm-to-firm relationships.

In this paper, we study the interaction between these two phenomenons from an empirical standpoint. We exploit highly detailed firm-to-firm data involving French firms and their partners in the European Union, covering a period of more than twenty years. Using these data, we first construct an indicator of relationship stickiness at the product-level. The empirical strategy consists in comparing the average duration of firm-to-firm relationships conditional on a match quality and derive from this an ex-post measure of stickiness. The strategy is backed out from a structural model of firm-to-firm matching in product markets. In the model, the duration of a firm-to-firm relationship is a function of firms' attributes, notably the price adjusted for quality that the supplier can offer to her client, as well as product-specific attributes, the cost of switching from a supplier to another, the extent of search frictions or the dispersion of offers across potential suppliers. These product-market attributes contribute to lengthening firm-to-firm relationships conditional on the quality of a match. This is what is captured by the relationship stickiness indicator.

We discuss how the relationship stickiness indicator recovered using this strategy compares with alternative measures of relationship-specific investments. Relationship stickiness is found larger in more differentiated product markets, for stages of production localized more upstream in value chains and for more complex goods, which production is likely to involve more relationship-specific investments. Consistent with the previous literature, we show that qualifying products by their position in the RS distribution helps explain the share of intra-firm flows in the value of bilateral trade and the patterns of specialization of countries displaying different levels of institution quality.

Armed with the RS indicator, we estimate the propagation of uncertainty shocks to the real economy. We estimate a significant impact of uncertainty on the extensive margin of trade. Episodes of high uncertainty are characterized by significantly less new firm-to-firm relationships, the impact being significantly stronger in product markets displaying stickier firm-to-firm relationships. This is consistent with sticky relationships generating high sunk costs, which firms are reluctant to pay in periods of high uncertainty. The propagation of policy uncertainty to the real economy is thus intimately linked to the type of relationships in which sellers and buyers are engaged. The modern organization of production into fragmented processes increases the sensitivity of the economy to uncertainty shocks.

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.

Ahir, Hites, Nicholas Bloom, and Davide Furceri, "World Uncertainty Index," 2019.

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

_ 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.

Barrot, Jean-Noel 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*, 2018, *100* (3), 424–439.

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.

_ and John Van Reenen, "Patents, Real Options and Firm Performance," *Economic Journal*, March 2002, *112* (478), 97–116.

Bloom, Nick, Stephen Bond, and John Van Reenen, "Uncertainty and Investment Dynamics," *Review of Economic Studies*, 2007, *74* (2), 391–415.

Bricongne, Jean-Charles, Lionel Fontagné, Guillaume Gaulier, Daria Taglioni, and Vincent Vicard, "Firms and the global crisis: French exports in the turmoil," *Journal of International Economics*, 2012, *87* (1), 134–146.

Dixit, Avinash, "Entry and exit decisions under uncertainty," *Journal of political Economy*, 1989, *97* (3), 620–638.

Dixit, Avinash K. and Robert S. Pindyck, *Investment under Uncertainty* number 5474. In 'Economics Books.', Princeton University Press, 1994.

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.

Fontaine, François, Julien Martin, and Isabelle Mejean, "Price discrimination within and across EMU markets: Evidence from French exporters," *Journal of International Economics*, 2020, p. 103300.

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.

Graziano, Alejandro, Kyle Handley, and Nuno Limao, "Brexit Uncertainty and Trade Disintegration," NBER Working Papers, National Bureau of Economic Research, Inc December 2018.

Handley, Kyle and Nuno Limao, "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 CEPII's distances measures: The GeoDist database," Working Papers 2011-25, CEPII research center December 2011.

Melitz, Marc J and Stephen J Redding, "Heterogeneous firms and trade," in "Handbook of international economics," Vol. 4, Elsevier, 2014, pp. 1–54.

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 Taylor, "Trade and Uncertainty," *Review of Economics and Statistics*, Forthcoming 2019.

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, "A concordance between ten-digit U.S. Harmonized System codes and SIC/NAICS product classes and industries," Finance and Economics Discussion Series 2012-15, Board of Governors of the Federal Reserve System (US) 2012.

— and — , "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.

Schaal, Edouard, "Uncertainty and Unemployment," *Econometrica*, November 2017, *85* (6), 1675–1721.

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

A Appendix

A.1 Measuring the duration of a trade relationship

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 transactions. This is illustrated in Table A.1 which compares statistics on i) the mean duration of a buyer's relationships with French suppliers, ii) the inverse of the probability of

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

	# transac. (1)	# sellers (2)	# buyers (3)	# sellers ×products (4)	# buyers ×products (5)	# buyer×seller ×products (6)
EU15	101,085,679	109,456	1,743,157	1,331,702	14,348,859	19,504,028
Belgium	20,093,986	75,611	220,839	644,380	2,567,705	3,680,980
Germany	19,591,647	61,949	380,942	500,587	2,690,140	3,609,025
Italy	12,766,637	52,825	302,048	386,961	2,185,160	2,835,711
Spain	12,696,214	54,079	259,753	424,676	1,973,209	2,537,203
UK	10,592,077	49,920	173,118	364,629	1,368,087	1,971,993
Netherlands	6,313,236	45,401	110,954	274,736	815,679	1,145,419
Portugal	4,940,157	34,244	77,370	242,825	785,200	1,048,799
Luxemburg	3,161,404	32,178	25,376	204,952	420,501	579,303
Austria	2,392,499	23,368	44,254	133,799	349,275	448,760
Greece	2,040,793	20,829	36,768	142,327	314,962	433,051
Sweden	2,029,067	20,934	36,153	119,912	270,737	358,207
Denmark	1,993,252	23,877	34,368	130,478	264,146	366,991
Ireland	1,391,572	18,062	23,445	95,108	205,661	297,275
Finland	1,083,138	14,499	17,769	78,293	138,397	191,311

Notes: This table is based on French customs firm-to-firm data for 1996-2006. The first line corresponds to all countries while the rest of the table gives statistics for individual countries. 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.

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

	Mean	Median	P25	P75
Mean duration	18	10	3	25
Frequency of transactions	0.332	0.222	0.095	0.500
Proba Recall	0.013	0.000	0.000	0.000

Notes: This table gives statistics on the distribution of durations, frequencies and recall probabilities, across importers connected to French firms. The "Mean duration" of a buyer's relationships with French firms is computed as the mean number of months between its first transaction with a given supplier and the next transaction with a different partner. The "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 are calculated on the dataset covering the 1996-2006 period.

Figure 1: ***Distribution of the number of transactions, per buyer×product***

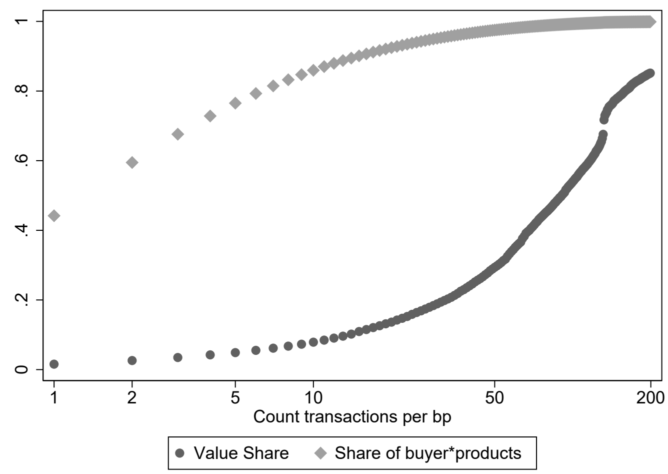

Notes: Cumulated distribution of the number of transactions per foreign buyer (×product). A transaction is the purchase of a particular good, to a given seller, in a given month. The light grey line corresponds to the share in the population of buyers and the dark line measures what this represents in the overall value of exports.

Figure 2: ***Distribution of the number of French partners, per buyer×product***

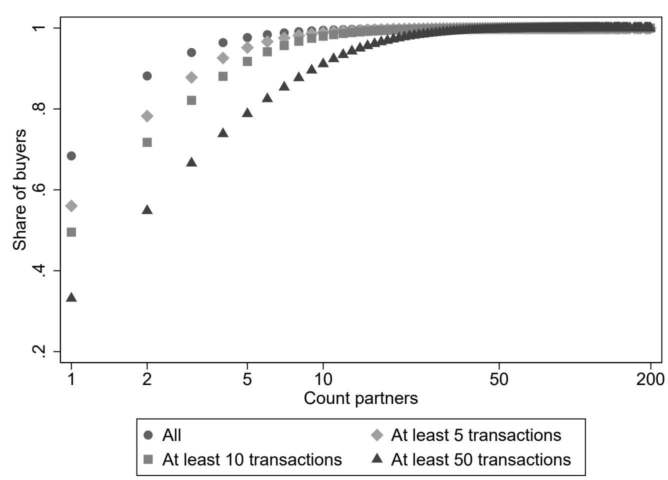

Notes: Cumulated distribution of the number of partners, per foreign buyer (×product). A partner is a French exporting firm. The number of partners is calculated over the sub-sample of importers that are involved in at least two transactions ("All") and at least 5, 10 and 50 transactions.

Figure 3: ***Distribution of the number of partners, per buyer/seller and date (month×year)***

Number of partners per buyer

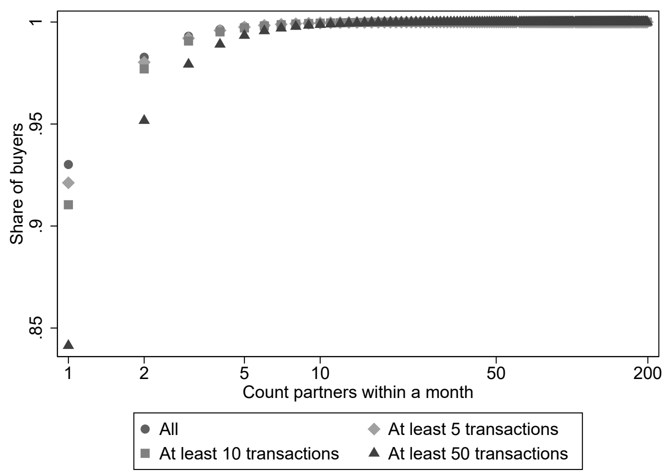

Number of partners per seller

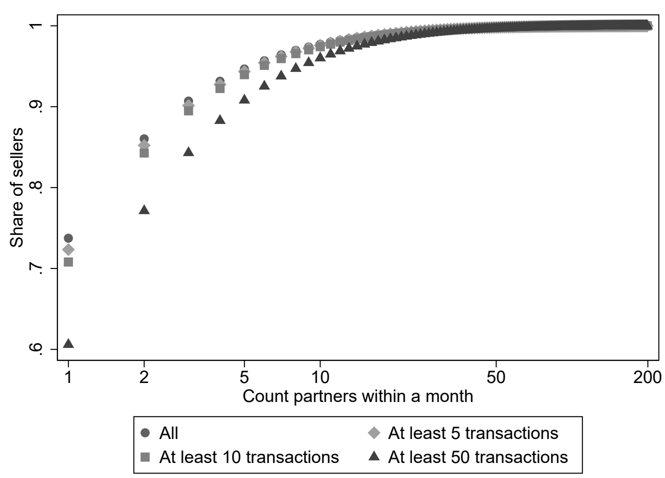

Notes: Cumulated distributions of the number of partners a French exporter interacts with in a given country (bottom panel) and the number of partners a foreign buyer (×product) interacts with within a particular month (top panel). The number of partners is calculated over the sub-sample of importers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

Figure 4: ***Distribution of RS estimates***

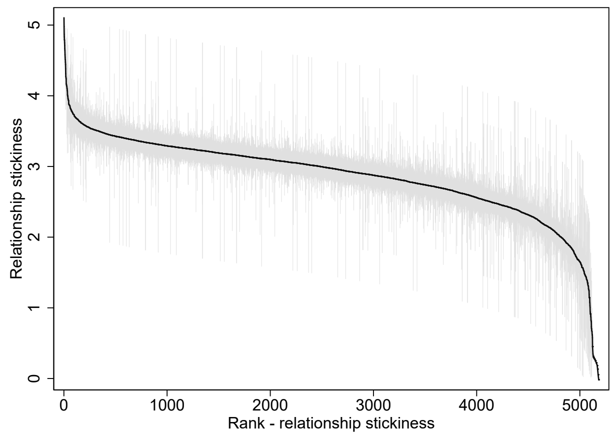

Notes: The figure shows the distribution of estimated relationship stickiness indicators (solid line) and their 10% confidence interval (grey area). The distribution covers 5,186 HS6 products.

Table 3: ***Duration and the size of trade flows***

	(1)	(2)	(3)
		Log of duration	
Log of mean exports	.041***	.070***	.237***
	(.000)	(.000)	(.001)
Observations	6,904,758	6,904,585	3,331,224
R^2	0.003	0.151	0.242
Within R^2	0.003	0.007	0.057
Fixed effects		Product	Product× buyer

Notes: This table correlates the duration of a relationship with a measure of the size of this transaction. Statistics are calculated on the dataset covering the 1996-2006 period.

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

Measure	Corr(η,.)	OLS η
$\mathbf{1}_{differentiated}$ (Rauch)	.08***	.06**
Share of not homogen. products (Nunn)	.04**	$-.02$
Upstreamness (Antras et al.)	.14***	.21***
Elasticity of subs. (Imbs & Mejean)	-.6***	-.16***
Product complexity (Haussman & Hidalgo)	.16***	.09***
Observations		3,863
R^2	-	.12

Notes:

Figure 5: ***Relationship stickiness across Broad Economic Categories***

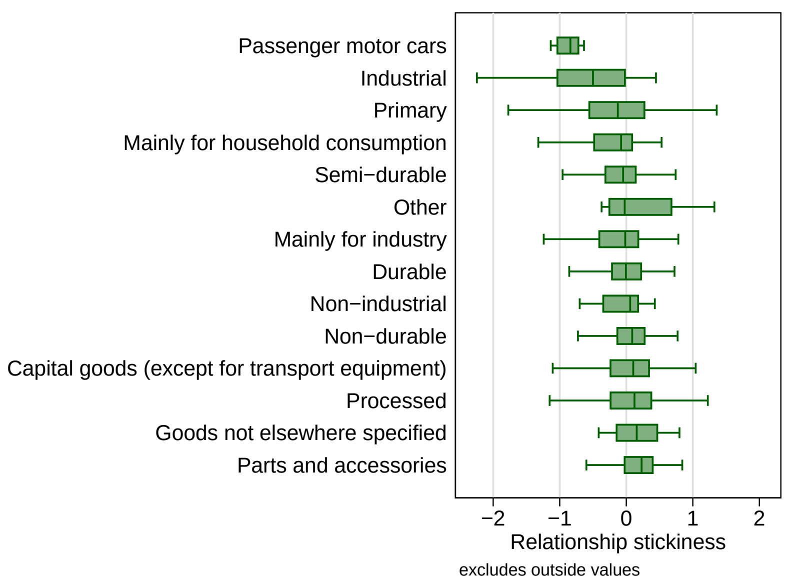

Table 5: ***Uncertainty and the creation of new trade relationships: Baseline results***

	(1)	(2)	(3)	(4)	(5)
Dep. var:			*# new trade relationships*		
Uncertainty shock dummy	0.37***				
	(0.008)				
- × RS index	-0.15***	-0.12***			-0.13***
	(0.003)	(0.002)			(0.008)
Uncertainty index			1.27***		
			(0.023)		
- × RS index			-0.51***	-0.40***	
			(0.008)	(0.007)	
Observations	3,302,770	3,302,770	3,302,770	3,302,770	1,500,366
Period	2000-2010	2000-2010	2000-2010	2000-2010	2011-2016
Fixed Effects					
Product × quarter		✓		✓	✓
Product × period	✓		✓		
Country	✓		✓		
Country × period		✓	✓	✓	

Table 6: ***Uncertainty and the creation of new trade relationships: Spillovers and persistence***

	(1)	(2)	(3)	(4)
Dep. var:	*# new trade relationships*			
Uncertainty	0.27***		0.25***	
	(0.007)		(0.006)	
- × RS index	-0.12***	-0.12***	-0.10***	-0.08***
	(0.002)	(0.002)	(0.002)	(0.002)
Uncertainty other countries	-0.01***			
	(0.002)			
- × RS index	-0.002**	-0.003***		
	(0.001)	(0.001)		
Uncertainty × Lag 1			0.19***	
			(0.005)	
- × Lag 2			0.25***	
			(0.005)	
- × Lag 3			0.19***	
			(0.005)	
- × Lag 4			0.22***	
			(0.006)	
Uncertainty × RS Index × Lag 1			-0.08***	-0.07***
			(0.002)	(0.002)
- × Lag 2			-0.09***	-0.07***
			(0.002)	(0.002)
- × Lag 3			-0.08***	-0.06***
			(0.002)	(0.002)
- × Lag 4			-0.08***	-0.04***
			(0.002)	(0.002)
Observations	3,637,726	3,637,726	3,636,211	3,637,726
Fixed Effects				
Product × period			✓	
Product × quarter	✓	✓		✓
Country	✓		✓	
Country × period		✓		✓

Table 7: ***Uncertainty and trade: Other margins of adjustment***

	(1)	(2)	(3)	(4)
Dep. var:	*# disrupted trade relationships*		*Export Value*	
Uncertainty shock dummy	0.23***		-0.03***	
	(0.007)		(0.003)	
- × RS index	-0.08***	-0.03***	0.01***	0.01***
	(0.003)	(0.007)		(0.001)
Observations	2,546,156	2,546,156	5,687,280	5,687,280
Period	1996-2006	1996-2006	1996-2010	1996-2010
Fixed Effects				
Product × quarter		✓		✓
Product × period	✓		✓	
Country	✓		✓	
Country × period		✓		✓

Table 8: ***Uncertainty and trade: margin decomposition***

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dep. var:	Growth	=Start	+ End	+ Intensive	Growth	=Start	+ End	+ Intensive
Uncertainty shock	-0.03***	-0.01***	-0.02***	0.00	-0.03***	-0.02***	-0.01***	0.00
	(0.003)	(0.002)	(0.002)	(0.001)	(0.003)	(0.002)	(0.002)	(0.001)
- × RS index	-0.005***	-0.004***	0.001***	-0.003***	-0.002***	-0.003***	0.002***	-0.002***
	(0.001)	(0.001)	(0.001)	(0.000)	(0.001)	(0.001)	(0.001)	(0.000)
RS index	-0.02***	-0.05***	0.03***	0.00***				
	(0.000)	(0.001)	(0.000)	(0.000)				
Observations	3,538,965	3,538,965	3,538,965	3,538,965	3,538,965	3,538,965	3,538,965	3,538,965
Fixed Effects								
Product × country					✓	✓	✓	✓

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 second 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 .332, which corresponds to a transaction every 3 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 10 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 switching probabilities. In the data, the three statistics are correlated at more than 50% meaning that heterogeneity in the frequency of transactions does not completely distort the distribution of trade durations, across buyers and products.

A.2 Detailed on the estimation of relationship stickiness

As explained in Section 3.1, relationship stickiness is estimated by exploiting the following prediction of a search model with switching costs:

$$\ln \mathbb{E}\left[\mathcal{T} \mid R \in R_q\right] = \ln \eta + \ln \ln \left[\frac{\mathbb{P}(R \geq r_{q-1})}{\mathbb{P}(R \geq r_q)}\right] \tag{A.1}$$

where $\mathbb{E}\left[\mathcal{T} \mid R \in R_q\right]$ is the expected duration of a transaction, conditional on the transaction falling in the qth quantile of the distribution, $\eta \equiv \frac{\gamma^k}{\lambda}$ is the product-specific index of business stickiness and $\ln\left[\frac{\mathbb{P}(R \geq r_{q-1})}{\mathbb{P}(R \geq r_q)}\right]$ solely depends on the definition of quantiles.

The empirical counterpart of the left-hand side of this equation is the log of the mean duration of firm-to-firm relationships, in various size quantiles of the product- and country-specific distribution:

$$Dur_{qpc} \equiv \frac{1}{N_{qpc}} \sum_{sb \in R_{qpc}} Dur_{sb(c)p}$$

where $Dur_{sb(c)p}$ is the duration of the relationship involving buyer $b(c)$ located in country c, French exporter s and product p, and N_{qpc} is the number of such relationships in the quantile under study (R_{qpc}). Dur_{qpc} is the empirical counterpart of the expectation term in (3).

Based on this (noisy) measure of conditional expected durations, it is thus possible to recover a relative measure of relationship stickiness using the following log-linear specification:

$$\log Dur_{qpc} = FE_p + \alpha \ln \ln \left[\frac{\mathbb{P}(R \geq r_{q-1})}{\mathbb{P}(R \geq r_q)}\right] + \epsilon_{qpc} \tag{A.2}$$

where FE_p is a product fixed effect, and ϵ_{qpc} is the error term. Since the regression is estimated on the pooled sample of firm-to-firm relationships involving importers from various origin countries, we control for unobserved heterogeneity across countries using an additional set of country-specific fixed effects FE_c. The product fixed effect recovered from equation ((A.2)) can be interpreted as a measure of the relative stickiness of relationships in product market p.

To compute the mean duration conditional on a size quantile (Dur_{qpc}) we proceed as follows: i) we compute the size of a relationship as the average value of transactions involving a given seller-buyer pair, in constant euros,[29] ii) we assign all the trade relationships to a size-decile (specific to a product category), iii) we take the average duration within each bin. Each distribution is cut into 10 quantiles, the eight deciles in between the 10th and the 90th

[29]Nominal values are deflated by the French PPI constructed by INSEE.

percentiles of the distribution plus a quantile defined by transactions in between the 1st and the 10th percentile and the quantile of rtansactions lying between the 90th and the 99th percentiles.

A.3 External validity tests

Relationship stickiness and intra-firm trade. In Table A.2, we examine whether the prevalence of intra-firm trade in US product-level trade data is systematically different along the distribution of relationship stickiness indicators. Namely, we correlate the relationship-specific indicator with the share of intra-firm trade in US exports (Columns (1)-(2)) and US imports (Columns (3)-(4)). In columns (2) and (4) we further control for additional product-level characteristics that we know are correlated with RS (see Section 3.2). The share of intra-firm trade is computed from intra-firm trade data released by the Bureau of Economic Analysis for year 2002. Intra-firm trade is reported by 6-digit NAICS categories that are merged with the HS6 nomenclature (version 2002) using the correspondence developed by Pierce and Schott (2012). We find a positive and significant correlation between the level of relationship stickiness of a product and its share of intrafirm trade. Relationship stickiness explains around 10 percent of the dispersion in the share of intrafirm trade across product categories.

Relationship stickiness and comparative advantages: 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 recovered from the UN-COMTRADE database at the 6-digit level of the Harmonized Nomenclature which is merged with our own measure of relationship stickiness. The results are reported in Table A.3. 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 by an interaction term between the quality of the country's institutions, as measured by **???-> ask Mat**, and the degree of relationship-stickiness of the product. 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, namely the Balassa index and a dummy identifying Balassa indices above 1.[30] We confirm Nunn (2007) findings that countries with good contract enforcement specialize in the production of more relationship specific goods. In columns (3) and (5), we show that both Nunn's and our measures of product stickiness have explanatory power in this regression. When the Balassa index is instead used as a measure of comparative advantage, the interaction with Nunn's measure becomes unsignificant while our indicator remains positively associated with more trade from countries with good enforcement laws.

Relationship stickiness and the distance effect: In a last sanity check, we investigate how relationship-stickiness interacts with standard determinants of international trade to shape the geography of trade. Namely, we use the gravity equation and interact the distance variable with our measure of relationship stickiness. Results are summarized in Table A.4. Bilateral trade data at the hs6 level are taken from the BACI database for 2005 (Gaulier and Zignago, 2010). Distance is the weighted distance between countries' main cities from Mayer and Zignago (2011). Finally, we also control for the product upstreamness in value chains and its interaction with distance. Results consistently show that the distance effect is magnified in product markets that display more relationship stickiness. This is true whatever

[30]The Balassa index is computed using BACI multilateral data and is defined as the value of product-level exports originating from one particular source country over the value of worldwide exports in the same product category.

the structure of fixed effects, including in the most demanding specification in Column (4). The elasticity of trade to distance also seem to increase for more upstream goods, although the effect is sensitive to the structure of fixed effects. Interpreting the magnified impact of distance for high RS products is not possible in such reduced-form framework. A possible interpretation is that information frictions are more stringent 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). **could be cool here to add in the table the distance elasticity at the first and third quartile of the distribution of RS**

Table A.1: ***Descriptive statistics on alternative measures of the duration of firm-to-firm relationships***

	Mean	Median	P25	P75
Mean duration	18	10	3	25
$1/\mathbb{P}(\text{switch})$	9	20	9	41
$1/\mathbb{P}(\text{switch}\|\text{Trade})$	2	3	2	6

Notes: This table gives statistics on alternative measures of durations. The first line is our benchmark measure, defined by the mean number of months between the first transaction involving a particular buyer and one of its supplier and the next transaction the same buyer interact with a different partner ("Mean duration"). "$1/\mathbb{P}(\text{switch})$" is the inverse of the switching probability recovered as the number of switching episodes divided by the total number of months a particular buyer is present in the data. "$1/\mathbb{P}(\text{switch}|\text{Trade})$" is the inverse of the switching probability conditional on a transaction, computed as the number of switching episodes over the total number of transactions. Statistics are calculated for each importer before averaging across buyers, using the dataset covering the 1996-2006 period.

Table A.2: ***Share of intrafirm trade and relationship stickiness***

	(1)	(2)	(3)	(4)
	Share of intra-firm			
	exports		*imports*	
RS (η)	0.152***	0.092***	0.097***	0.062**
	(0.025)	(0.034)	(0.021)	(0.026)
Nunn' measure		0.409***		0.202***
		(0.066)		(0.049)
Upstreamness		0.067***		0.021*
		(0.016)		(0.012)
Elasticity (σ)		-0.002		-0.008**
		(0.006)		(0.003)
Observations	439	378	439	378
R-squared	0.074	0.133	0.055	0.075

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

Table A.3: ***Institutional comparative advantage***

	(1)	(2)	(3)	(4)	(5)
		log(exports)		Balassa Index	$\mathbf{1}_{Balassa>1}$
Rule of law					
$\times$ *RS*	0.196***		0.224***	0.110**	0.010***
	(0.034)		(0.030)	(0.047)	(0.003)
$\times$ Nunn specif.		0.812***	1.070***	0.367	0.041**
		(0.100)	(0.144)	(0.302)	(0.020)
$\times$ Upstreamness			0.077*	0.041	0.008
			(0.045)	(0.072)	(0.005)
Fixed effects		$country(122)$ *and* $sector(4,326)$			
Observations	296,185	296,185	292,957	527,406	527,406
R-squared	0.604	0.606	0.609	0.012	0.139

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

Table A.4: ***Gravity for trade in goods with sticky relationship***

	(1)	(2)	(3)	(4)
Distance (log)	-0.553***	-0.370***	-0.521***	-0.893***
	(0.015)	(0.019)	(0.020)	(0.025)
- × RS		-0.064***	-0.056***	-0.028***
		(0.006)	(0.006)	(0.006)
- × Upstreamness		0.002	0.012**	-0.021***
		(0.005)	(0.005)	(0.007)
RS	-0.198***	0.322***		
	(0.007)	(0.051)		
Upstreamness	0.044***	0.026		
	(0.005)	(0.040)		
Fixed effects				
Exporter	✓	✓	✓	
Importer	✓	✓	✓	
Product			✓	
Exporter×Product				✓
Importer×Product				✓
Observations		5,704,026		5,473,532
R-squared	0.164	0.164	0.285	0.578

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