Relationship Stickiness, International Trade, and Economic Uncertainty

Julien Martin
UQAM & CEPR

Isabelle Mejean
CREST- Ecole Polytechnique & CEPR

Mathieu Parenti
SBS-EM, ECARES, Université libre de Bruxelles & CEPR

January 2021

ECARES working paper 2021-03

ECARES
ULB - CP 114/04
50, F.D. Roosevelt Ave., B-1050 Brussels BELGIUM
www.ecares.org

Relationship stickiness, international trade, and economic uncertainty*

Julien Martin† Isabelle Mejean‡ Mathieu Parenti§

December 18, 2020

Abstract

This paper examines how the degree of stickiness in business relationships influences the real impact of aggregate uncertainty. We first develop a novel 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 stylized search model in which a higher degree of stickiness implies a lower probability of switching and longer firm-to-firm trade relationships, conditional on match quality. Relationship stickiness 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. The relationship between uncertainty and separation rates also varies along the distribution of RS indices, the probability of a trade relationship ending

*We are grateful to François de Soyres, Swati Dhingra, Julian di Giovanni, Alessandro Ferrari, Stefania Garetto, Keith Head, Sebastian Heise, Andrei Levchenko, Florian Mayneris, Ryan Monarch, Stephen Terry, and Frederic Warzynski, as well as participants at various seminars and conferences including the CEPR-ERWIT conference, the CEPR-CEBRA-Bank of Lithuania 2020 conference, the World Bank, the Bank of Canada, the Aarhus trade workshop, the Sauder-Vancouver School of Economics seminar, the Geneva trade and development workshop, the CIRANO macro workshop, the Warwick seminar, the SCSE conference, the St Petersburg conference on trade and urban economics, the University of Cergy-Pontoise, the University of Zurich, and Banque de France. I. Mejean gratefully acknowledges support from a public grant overseen by the French National Research Agency (ANR) as part of the "Investissements d'Avenir" program (Idex Grant Agreement No. ANR-11-IDEX-0003- 02Labex ECODEC No. ANR-11-LABEX-0047 and Equipex reference: ANR-10-EQPX-17 "Centre d'accès sécurisé aux données" CASD) as well as the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 714597). J. Martin is holder of the UQAM research chair on the local impact of global firms. M. Parenti gratefully acknowledges support from the Fund for Scientific Research (FNRS Grant PDR/GTO T.0180.19).

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

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

§ECARES (ULB) and CEPR, Email address: mparenti@ulb.ac.be.

being significantly reduced in sticky-product markets in uncertain times. These results provide evidence that trade of sticky products is characterized by wait-and-see behaviors during uncertainty episodes.

1 Introduction

How do firm-to-firm relationships shape the response of international trade to macroeconomic shocks? One defining characteristic of trade relationships is their degree of stickiness. Standard trade models assume transactions take place in spot markets where relationships can be initiated or replaced at virtually no cost. Yet increasing evidence shows a number of product markets display a large degree of stickiness that results in persistent firm-to-firm relationships.[1] Responding to macroeconomic shocks is likely to be more difficult for firms engaged in such sticky relationships. One type of shock in which stickiness might be especially relevant is uncertainty shocks. Following Dixit and Pindyck (1994), uncertainty indeed affects investment behaviors by creating an option value of waiting. Such transmission of economic shocks to the real economy is likely to be especially pervasive in the presence of relationship stickiness.

In this paper, we study this question and make two main contributions. First, we provide a novel measure of relationship stickiness recovered from unique firm-to-firm trade data. Second, we study the transmission of uncertainty shocks to international trade in sticky-product markets. Our measure of relationship stickiness (RS) exploits a rich panel of all firm-to-firm relationships involving French exporters and their European partners over the 2002-2006 period. We use the mean duration of trade relationships to back out a product-level measure of relationship stickiness. Armed with this measure and external time-series of aggregate uncertainty in various destination countries, we then study how trade adjusts to uncertainty shocks and how the mechanism of the adjustment varies depending on the degree of product stickiness. Uncertainty leads to a drop in the creation of new relationships and, in sticky-product markets, a reduced separation rate within firm-to-firm relationships. Such adjustments are consistent with wait-and-see behaviors in periods of high uncertainty in markets characterized by high relationship-specific investments. When comparing uncertainty shocks with GDP shocks, we find such behaviors to be specific to uncertainty episodes.

Our measure of relationship stickiness 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

[1] The literature uses several terms 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 characteristics as relationship stickiness.

switching costs and more frictions contribute to lengthening existing firm-to-firm relationships, conditional on the quality of a match. The duration of relationships is thus a relevant moment that can be exploited to recover a product-level measure of stickiness.[2] For estimation, we use French firm-to-firm export data, whose panel dimension allows us to follow importers over time, and compute the duration of each of the importers' relationships with a French firm. The unique level of disaggregation of the data allows us to control for individual characteristics that affect the quality of a match and contribute to the dispersion in the duration of relationships, within a product category. After controlling for the quality of a match, we can use the variability of average durations *across products* to recover a measure of relationship stickiness (RS) for more than 5,000 HS6 products.

We present a body of evidence supporting the view that the recovered measure of relationship stickiness does capture relational specificity at the product level. First, we show the measure is stable when computed from French exports to different destinations or over different time periods.[3] Second, we find the stickiness measure varies substantially across sectors and across products within a sector. Specialty chemicals or parts and accessories that entail large customization costs are found among the most sticky products, whereas motor vehicles or men's 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 measures developed by Rauch (1999) and Nunn (2007), but it displays substantially more variability, notably within the group of manufacturing products. More sticky categories 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 Levchenko (2007) and Nunn (2007); and 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 novel measure, we study how relationship stickiness affects the adjustment

[2]The measure is thus backed out from the ex-post duration of firm-to-firm relationships. As a consequence, it is agnostic about the exact microeconomic features that lead to such stickiness, which can be explained by a combination of switching costs, relationship-specific sunk investments, information asymmetries, or informational frictions.

[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.

of trade flows to uncertainty shocks. Uncertainty has been shown to be a potential threat to economic growth, because 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. Indeed, international trade is associated with substantial sunk costs that affect trade at the extensive margin. For this reason, uncertainty shocks have a large effect on this dimension of firms' activity (Novy and Taylor, 2019). We provide evidence consistent with this view using highly disaggregated trade data.

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 applied to the quarterly Economist Intelligence Unit (EIU) country reports. Based on this database, we can 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.

Periods of high policy uncertainty systematically display a reduced number of new trade relationships involving French exporters and European importers. Quantitative effects somewhat vary depending on the specification, with a contemporaneous effect estimated 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 -2% and -11% when moving from the first to the third quartile of the distribution of RS, from the least to the most sticky markets. We also provide evidence of separation rates being affected by uncertainty. Here, the qualitative impact varies along the RS distribution. More uncertainty is associated with a higher probability of trade relationships ending in less sticky markets, whereas the separation rate is instead muted where relationships are stickier, on average.[4]

In high-stickiness product markets, uncertainty episodes are thus associated with wait-and-see behaviors that materialize in the data through muted adjustments at the extensive margin. This finding is consistent with the view that such markets feature high costs of

[4] At first view, such a result may appear mechanical, because product markets that exhibit longer relationships are likely to adjust less at the extensive margin. However, the identification strategy adopted to estimate this effect uses the variability across relationships *within* a product. The systematic cross-product heterogeneity in separation rates that contributes to the dispersion in RS indicators is thus controlled for and does not contribute to the identification of a trade response to uncertainty shocks.

switching from one supplier to another. Because the option value of staying with an existing partner increases in periods of high uncertainty, we observe less turnover in firm-to-firm relationships in risky times 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.[5]

We conclude the analysis by studying the implications of these results for trade growth. Consistent with the previous literature, we estimate a -12 pp response of product-level trade growth to uncertainty episodes. The vast majority of this effect is driven by a lower net entry of firm-to-firm relationships, which is especially pronounced in sticky markets. Interestingly, we can compare our results with those associated with a shock to the level of growth in the destination. Episodes of low growth are also associated with significantly less product-level trade. But around 50% of the effect is driven by the intensive margin, especially in sticky-product markets. This finding echoes the argument in Antras (2020) that severe but temporary shocks such as the 2008-09 trade collapse or the COVID-crisis are unlikely to change firms' sourcing strategies due to stickiness in trade relationships (in particular in global value chains). Our finding that adverse shocks lead to a slow-down in the creation of new relationships and, for shocks to the level of demand, a sharp drop along the intensive margin, but that they are not associated with a strong increase in the disruption of firm-to-firm trade relationships in sticky products, is consistent with this view.

Related literature. The paper contributes primarily to two strands of the literature, namely, the trade literature on relation-specific investments, and the macroeconomics 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), and associated welfare gains (Chor and Ma, 2020). 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 optimal design of trade policy depends on the specificity of business relationships (Antràs and Staiger, 2012).[6]

In this literature, relationship specificity is usually proxied using either the measure developed by Rauch (1999), or the measure developed by Nunn (2007).[7] Our contribution to

[5] 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).

[6] 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). And Hémous and Olsen (2018) study how stickiness in business relationships can affect innovation incentives.

[7] Other related measures have been developed, including the Herfindahl index of intermediate input use

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 that 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 less computationally demanding procedure to estimate these costs, which allows us to recover them for a wide range of products.[8] Furthermore, we work with highly disaggregated seller-buyer relationships observed over various destinations. Doing so allows us to purge our measure from country-specific costs and obtain a measure of relationship stickiness at the fine product level.[9]

Our paper also 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

(Levchenko, 2007), the share of wholesalers importing a product (Bernard et al., 2010), suppliers' R&D expenses and the number of patents that they issued (Barrot and Sauvagnat, 2016), or the distance to final demand (Antràs et al., 2012). Chor and Ma (2020) develop a measure of contractibility in the spirit of Nunn (2007).

[8] In this procedure, relationship stickiness is evaluated in relative terms along the distribution of products. As a consequence, we cannot directly interpret our estimates in terms of a monetary switching cost, which Monarch (2014) can do. We have checked that, for the 44 products that we can match with his estimates, the correlation between his and our measures is positive.

[9] 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 (e.g., Besedes and Prusa, 2006). Whereas most of the papers in the literature look at the duration of trade flows at the product-level, Schmidt-Eisenlohr and Monarch (2015) and Heise (2016) instead exploit firm-to-firm data, as we do. Schmidt-Eisenlohr and Monarch (2015) show the survival probability of seller-buyer relationships increases with their size and age, using matched US importer-exporter data. And Heise (2016) studies the systematic relationship between exchange-rate pass-through and the duration of firm-to-firm relationships. Instead, we use the duration of seller-buyer relationships in international markets to back out a measure of relationship stickiness, controlling for individual characteristics.

patenting and firms' productivity (Bloom and Reenen, 2002), the responsiveness of investment to demand shocks (Bloom et al., 2007), or hiring decisions (**?**). 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, 2017a; 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), Ahmad et al. (2020), and Exton and Rigo (2020) 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. Consistent with Carballo et al. (2018), we find uncertainty affects trade flows through the extensive margin.

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 export transaction between French firms and their individual partners in the EU. 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 shipment (in euros). The dataset covers the period from 1993 to 2017, but the analysis exploits various sub-periods. The main reason 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.[10] 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.[11] 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.[12] 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.[13] 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 B.5 in Appendix provides descriptive statistics on the dimensionality of the data, in the overall EU 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 in over five transactions, on average.[14]

2.2 Structure of the firm-to-firm data

Although 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 one-shot

[10] 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.

[11] 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.

[12] 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 change in threshold raises concerns properly identifying the beginning or the end of trade relationships.

[13] Whereas the dataset is exhaustive, exports from the smallest French exporters cannot be exploited, because these firms are allowed to complete 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.

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

buyers represent as much as 44% of the importers in the raw data (see Figure B.1 in Appendix) and 1.5% of the value of trade. The reason we abstract from these one-time transactions is that our empirical strategy is based on the duration of a relationship that is not defined when the importer is observed just once. One may argue that neglecting these transactions biases our measure of stickiness up because 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.[15] Our strategy thus consists of ignoring these transactions to estimate the degree of stickiness at the product level and using them ex-post to check whether they indeed are concentrated in those sectors that we find display less stickiness. As expected, we find one-shot buyers are more prevalent in markets featuring a lower degree of stickiness.

The remaining foreign buyers also display considerable heterogeneity regarding the "intensity" of their relationships with French exporters. This heterogeneity is illustrated in Figure B.1. Fifteen percent 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.[16] These firms are good candidates for observing the duration of their relationships with French firms.

Figure 1 shows the distribution of the number of French partners with whom individual buyers interact over their entire time in the dataset.[17] Overall, 67% of buyers have a single partner in France, whereas 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 observe a third of "loyal" buyers that repeatedly interact with the same exporter. Such behavior is consistent with the idea that some firm-to-firm relationships in international markets are especially sticky. The question that the empirical analysis addresses is whether this stickiness is systematically related to the specificities of some products or

[15] 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.

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

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

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

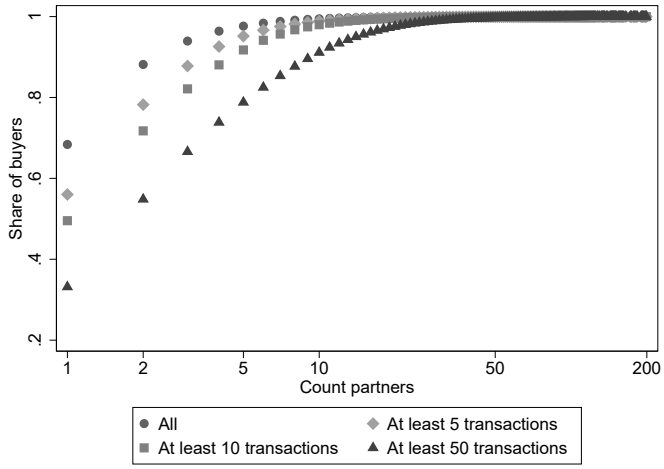

Notes: Cumulated distribution of the number of partners, per foreign buyer (×product). A partner is a French exporting firm. The number of partners is calculated over the sub-sample of importers that are involved in at least two transactions ("All") and at least 5, 10 and 50 transactions.

sectors.[18]

We terminate this raw description of the data structure with a last stylized fact later used to motivate our econometric model. Namely, we 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, whereas sellers simultaneously serve several importers (even within a country). This finding is illustrated in Figure 2, which shows the distribution in the number of sellers interacting with a given importer during a particular month (top panel) and the distribution in the number of partners from the same country a French exporter is interacting with (bottom panel). More than 90% of importers have only one French supplier for a given product within a given month. Even when we concentrate on importers that we observe over many (i.e., at least 50) transactions, this proportion is high, above 80%. Instead, 26% of French exporters sell the same product within the same month to several partners located in the same country, the proportion increasing to 55% when we pool partners located

[18]Some of these faithful relationships may take place within multinational companies. As explained in footnote 14, the data do not allow us to control for this possibility. We cannot exclude that relationship stickiness emerges as a consequence of intra-firm trade (or vice-versa). We show estimated relationship stickiness is indeed correlated with an (external) measure of intra-firm trade.

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

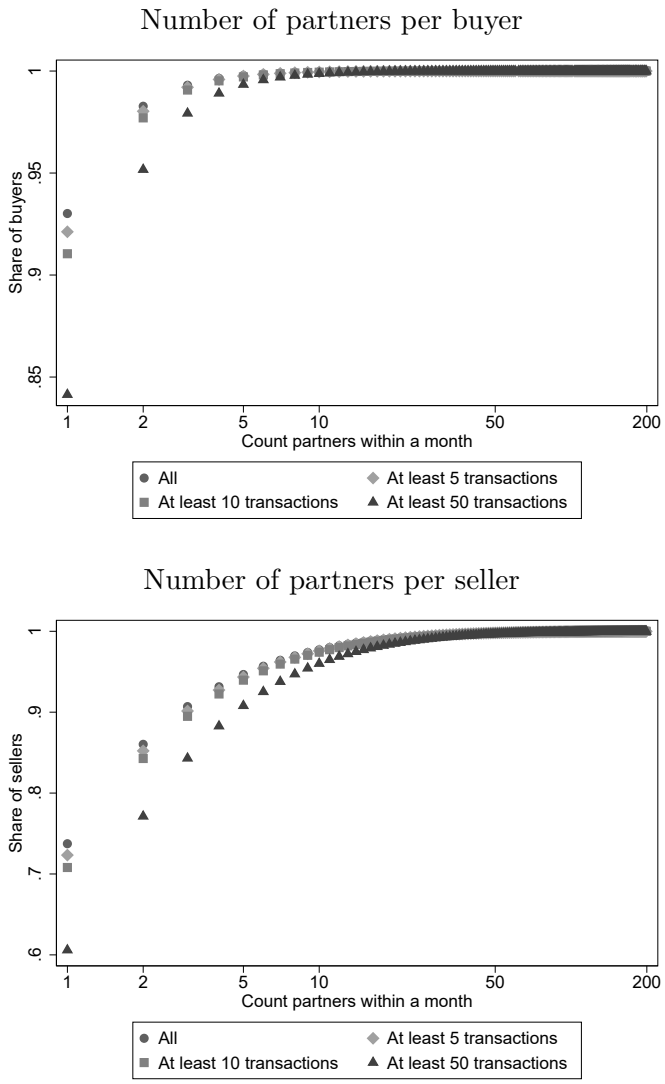

Notes: Cumulated distributions of the number of partners a French exporter interacts with in a given country (bottom panel) and the number of partners a foreign buyer (×product) interacts with within a particular month (top panel). The number of partners is calculated over the sub-sample of importers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

in different countries.[19,20] Given this data structure, the model in section 3 assumes many-to-one matching with importers interacting with a single supplier at a point in time. When the same importer is seen interacting with two different exporters within a month, we consider the two transactions to be sequential with a duration of half a month for the first transaction. In a robustness check, we have also tried removing from the estimation sample the 10% of importers that we eventually see interacting with multiple sellers within a month. Results available upon request are virtually unchanged.

2.3 Duration of firm-to-firm relationships

Using the time series of each buyer's interactions with French firms, we are able to construct the main statistic at the root of the estimation, namely, the duration of a buyer's relationships with French firms.[21] Figure 3 and Table B.6 provide statistics on this variable. In Table B.6, 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"), 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 to an exporter it had already interacted with in the past ("Proba Recall"). Figure 3 instead takes the perspective of a single firm-to-firm ($\times$ product) relationship and computes durations, the distribution of which is illustrated in the figure.

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 10 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. Although these statistics are calculated for

[19] 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 (see Figure B.2 and the comparison with Figure 2, the later counting the number of partners within a product while the former cumulates partners across products within a firm). Once we condition on a given product, purchasing from multiple French exporters becomes very rare.

[20] The many-to-one structure of firm-to-firm trade relationships at the product level is also documented by Monarch (2014) using data on U.S. imports from China. Monarch (2014) further shows that U.S. importers tend to import from a single country, as do Belgian importers as documented by Muûls (2015).

[21] 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.

Figure 3: ***Distribution of the durations in firm-to-firm relationships***

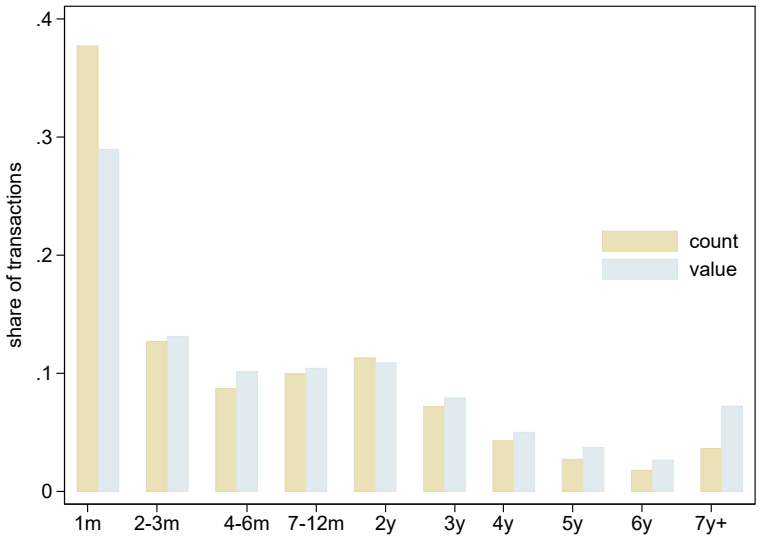

Notes: Distribution of durations, as a share of total number of relationships ("count") and as a share of aggregate export value ("value"). Statistics are recovered from the 19.5 million firm-to-firm relationships identified over the 1996-2006 period.

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; that is, we directly exploit the durations represented in Figure 3. Here, the mode is at one month which corresponds to an importer that interacts with two different French firms over two consecutive months for the same product. These very short relationships represent less than 40% of the population, whereas roughly 30% of firm-to-firm relationships last more than a year. Part of this heterogeneity is the consequence of heterogeneous match qualities, an importer being more likely to switch if her current match is not satisfactory. Results in Table B.7 show 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 correlation is true both across buyers within a product and within a buyer, across the different suppliers it meets throughout their 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 feature of the data is that the frequency of transactions displays significant heterogeneity across importers (second line in Table B.6). On average, the probability of a transaction occurring in a given month is .33, which corresponds to a transaction every three

months. However, 25% of buyers purchase French products more than once every two months, whereas in the first quartile of the distribution, firms purchase products less than once every 10 months. In the model and in statistics in Table B.6 and Figure 3, 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.[22] Therefore, we do not keep the whole history of an importer's partners in memory and instead suppose 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 of switching to a new supplier greatly simplifies the analysis. Moreover, the probability of a "recall", that is, of a buyer switching back to a supplier it knows from before, is very small in the data (see the last line in Table B.6).

3 Measuring relationship stickiness

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. We then benchmark our measure against various indicators in the literature to gauge its relevance.

3.1 Empirical strategy

The model is a simple search model between sellers and buyers of a given product. We assume 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 by any other product-specific characteristics. We remain purposely agnostic on the exact micro-foundations at the root of such stickiness. We consider instead these features as given and use the cross-section of products to quantify their relative size across product categories.

[22] We also tried to measure duration as the numbers of non-interrupted transactions between two firms. In a such case, if transactions occur in January and March, the number of transactions is two, whereas the number of months is three. The index of stickiness computed with this alternative definition is correlated at .7 with our baseline measure.

Suppose 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 switch and benefit from this offer. Suppose $\tilde{p}$ 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 on 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. That is, the firm decides to switch whenever $\tilde{p} < \frac{p}{\gamma}$, where $\gamma > 1$ is the wedge between the current price and the buyer's reservation price. This wedge may stem from switching costs being sunk upon starting a new relationship, but it could also reflect search costs incurred to be able to draw a price offer with probability λ. Our representation of these costs is deliberately very stylized, which, as we show below, allows us to draw a simple link between their magnitude and the expected switching probability, and thus the expected duration of a relationship. Indeed, switching suppliers then occurs with probability $\lambda H_P(p/\gamma)$.[23]

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

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

This formula 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 that has met a more competitive supplier is more

[23] Here, we implicitly assumed p is determined prior to the arrival of a new offer, that is, 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). Although such an assumption would suggest 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.

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.

To bring the model to the data, we make two additional parametric assumptions. First, we assume 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.[24] 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 $\eta \equiv \frac{\gamma^k}{\lambda}$. Hereafter, we 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). Although the parametric assumptions required to recover equation (2) are obviously important, we argue the model's insight, namely, that such product-specific characteristics tend to increase durations, conditional on a match, is likely to hold under alternative parametric assumptions.

Our data are 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 follows:[25]

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

[25] 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 is the value for the q^{th} cut-point. See details in Appendix A.2.

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

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 $\bar{H}_R(r) \equiv \mathbb{P}(R > r)$. Because equation (3) is log-linear in η and η is 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 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. Because we estimate the log of the η parameter, an interquartile of .63 means 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, which is sizable. Note the precision of estimates, however, varies across products, as shown by the grey area in Figure 4.[26] Empirically, products for which the number of fim-to-firm relationships is large tend to display tighter confidence intervals, the correlation between these two variables being high at 75%. This feature is expected as the empirical strategy implicitly relies on the law of large numbers to smooth the impact of heterogeneity in durations

[26] 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 goods are highly differentiated, are 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, because chemicals are often considered as homogenous products. Note 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.

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

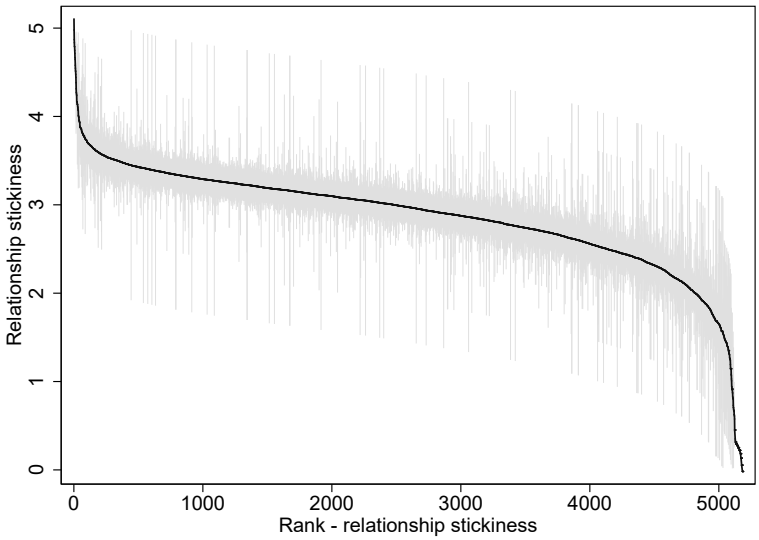

Notes: The figure shows the distribution of estimated relationship stickiness indicators (solid line) and their 10% confidence interval (grey area). The distribution covers 5,186 HS6 products.

within deciles of the product-specific subsamples. The larger the number of relationships, the better the approximation. Importantly, the number of underlying observations does not affect the point estimates, which are orthogonal to the number of observations. We take this margin of error into account in the empirical analysis using a parametric bootstrap procedure.

Table 1 shows how our measure correlates with other product-specific attributes used in the literature. The first column reports the pairwise correlation coefficients, and 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). We find little correlation with Nunn (2007) indicator. The reason is the difference in the level of aggregation. To give a concrete example, in his index, the car industry has a high level of input specificity. Consistent with this measure, we find most parts and accessories in this industry have a high level of stickiness. But final products like cars do not, which lessens the correlation. Finally, we find a positive correlation between the level of upstreamness of a product and its degree of stickiness. This

Table 1: ***Correlation with other measures***

Measure	Corr(η,.)	OLS η
	(1)	(2)
$\mathbf{1}_{differentiated}$ (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: This table reports the pairwise correlation coefficients (column (1)) and the multivariate correlations (column (2)) bewteen estimated RS indices and various characteristics of these products. Robust standard errors in (). Significance levels: * 10%, ** 5%, *** 10%.

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, which 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 12% of the heterogeneity recovered from our estimation (column (2)). The reason is that the RS indicator is extremely heterogeneous, including within particular industries.[27] The high disparity of stickiness is further illustrated in Figure B.3, which compares average RS measures across categories of the broad economic classification (BEC).[28] Final consumption goods like cars or consumer goods display, on average, a low level of stickiness, whereas parts and accessories or food processed for the industry display higher average RS indices. Here as well, however, we see a significant level of dispersion within these categories.

In Table B.8, we present the results of a regression of our measure of stickiness on various statistics directly calculated from the firm-to-firm data. We first show products that with a higher share of buyers that appear only once in our data and with a higher share of buyers that interact with several suppliers within a month exhibit lower levels of stickiness. These correlations are consistent with the view that if buyers can afford to enter markets for short

[27] 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.

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

periods of time or if they trade with many suppliers, they can more easily switch from one supplier to another.[29] We then explore two predictions of our model, namely, that stickiness should be higher in markets with more frictions and where the distribution of sales in more concentrated. As a proxy of sales distribution, we measure the concentration of exports across French firms for each product category. As expected, product categories with a higher level of sales concentration exhibit a higher level of stickiness. The explanatory power of sales concentration is low (R^2 of 1%), however, suggesting the distribution of sales is not the main force explaining heterogeneity in stickiness across products. As a rough measure of frictions, we compute the dispersion of unit values recovered from the whole distribution of transaction-level unit values within product categories. The correlation between stickiness and price dispersion is positive and significant (R^2 of 3%), consistent with the view that products with higher levels of frictions are more sticky. Last, we look at the correlation between the activity of wholesalers in a product category and the level of stickiness. We find a negative and significant correlation between stickiness and the share of wholesalers, which is consistent with the view that the least sticky products involve less customization and specificity and are thus more likely to be exported by intermediaries.

A final check of the distribution of estimates consists of studying their stability across estimation samples. Our argument is indeed that the estimation strategy allows us to capture the ex-post impact of product-specific attributes. If this is the case, we should expect our estimates to be roughly consistent regardless of the time period or the country sample used to estimate relationship stickiness. We therefore estimated a country-specific distribution of RS indicators using the same empirical strategy but estimating RS for each destination.[30] 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 high, between 69% and 74%. We also assess the stability of our estimates over time, by estimating relationship stickiness using the 2011-2017 period. Here as well, the correlation is significant, at .6. Note our baseline distribution of RS indices is the one estimated over 1996-2006, because the underlying customs data are of better quality in the earlier years of the sample.[31]

[29] We also compute stickiness on a sample in which we exclude buyers interacting with several suppliers within a month. Stickiness measured on this sample has a correlation of .87 with our baseline measure.

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

[31] As mentioned in section 2, the dataset is censored because the Customs does not request that firms declare the product exported, below a certain threshold. Because the threshold tripled in 2011, censoring is significantly larger in the most recent period.

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 the same qualitative results are confirmed using our indicator instead. We think of these exercises as useful sanity checks that the measure we 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 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 the production. A corollary of such a framework is that vertical integration should be more pervasive in product markets that display more intense locked-in effects. We confirm this result in Table A.2 using the prevalence of intra-firm trade as a measure of vertical integration. Using US trade data, we show intra-firm trade as a share of overall trade is significantly larger for products displaying high RS indicators. Alone, RS explains 10% of the cross-product dispersion in the data.

Nunn (2007) and Levchenko (2007) argue high-relationship-specific investment goods require sounds institutions, in the form of quality of contract enforcement, property rights, shareholder protection, and so on. As a consequence, institutions can shape the geography of trade as other sources of comparative advantages do. 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 measure and ours point in the same direction, namely, a specialization of countries with better institutions into the production of goods that are more relationship-specific.

Finally, we investigate the impact of relationship stickiness 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 the negative impact of distance on trade flows is stronger for product categories that exhibit a higher degree of relationship stickiness. Although 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 be explained by information frictions being large in stickier 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).

These sanity checks all point into the same direction. Our indicator of relationship stickiness seems to capture meaningful variability across (disaggregated) product markets, which correlates with external indicators in a way that is consistent with our interpretation.

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 adjustment of trade to economic-uncertainty shocks.

4.1 Motivation, data and empirical strategy

In the tradition of Dixit and Pindyck (1994), the transmission of uncertainty shocks to the real economy is interpreted as a consequence of firms' investment decisions, with firms being reluctant to invest in uncertain times.

In the international economics literature, the real option of waiting is often driven by the presence of sunk costs to establish a new relationship. In Handley and Limao (2017b), trade-policy driven uncertainty lowers the probability of entry in a destination in presence of a sunk entry cost. In Fillat and Garetto (2015), a higher sunk cost of entry widens the band of inaction: it reduces the probability that a domestic firm enters a foreign market while reducing the exit probability for already exporting firms. Uncertainty may also affect entry and exit decisions in international markets if international trade is subject to search frictions and the search process is costly. The mechanism would thus be similar to Schaal (2017) who show that in presence of search costs and search frictions in the labor market, uncertainty reduces firms' decision to hire new workers and to fire current ones.

In this literature, the real option of waiting increases with sunk costs and search frictions, conditional on a level of uncertainty. Importantly, these parameters are precisely those that drive the dispersion in our measure of relationship stickiness. More specifically, relationship stickiness is higher for products that have high sunk costs and/or for which there are a high level of frictions. We thus conjecture the option value of waiting – to start or to end a relationship – should be increasing along the distribution of RS indices.

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 is a measure of trade computed for each country c, product p, and period t. Importantly, we systematically control for product or product×period fixed effects so that the identification exploits the variability across destinations and eventually over time, within a product. This dimension of heterogeneity has been largely neglected when estimating relationship stickiness and is thus useful to separately identify the response of trade to uncertainty shocks, conditional on the level of stickiness.

Our measure of stickiness is estimated, and we use the standard errors associated with the RS measure to account for this first-stage error. For each product, we make 400 draws of RS in a Gaussian distribution calibrated to the mean and standard deviation of the corresponding estimated standard errors. We then run 400 regressions using the relationship stickiness generated in these draws. The coefficients and their standard errors reported in estimation tables are obtained by taking the mean and standard deviation of these estimates.

In equation (4), $Uncertainty_{ct}$ represents 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, that is, the net creation of firm-to-firm relationships, which is associated with considerable costs and should thus react to uncertainty. Finally, we also expect the impact of uncertainty to be particularly pronounced in high relationship-stickiness sectors, which we capture with the interaction term between uncertainty and the RS indicator ($RS_p \times Uncert_{ct}$).

A legitimate concern is that the analysis relies on data similar to those used to estimate relationship stickiness. Note, however, that whereas stickiness is estimated by exploiting differences in duration across product categories, our analysis of the impact of uncertainty includes product fixed effects, and thus exploits a different dimension of the data. Furthermore, if our measure of stickiness were contaminated by uncertainty - stickiness reflecting, for instance, a higher sensitivity to uncertainty - the interaction between stickiness and uncertainty wouldn't be significant.[32]

We measure uncertainty episodes at the country and quarterly levels using the "World Un-

[32] One can also make the point that, as long as uncertainty increases the duration of relationships more for the more sticky products, the product ranking of stickiness will not be affected.

certainty Index" (WUI) developed in Ahir et al. (2019). Using text analysis of the Economist Intelligence Unit country reports, they construct an 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 when the uncertainty index is one standard deviation above its average level.[33] 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. Figure B.4 shows a heat map of the corresponding uncertainty series. Although the panel solely covers EU countries, the data display a significant degree of heterogeneity across countries and over time. In some specifications, we also control for GDP growth shocks in the destination to separately identify the impact of uncertainty and other aspects of business cycles. Market-price GDP growth data are taken from Eurostat's national accounts indicators. The GDP shocks are computed as periods in which GDP growth is lower than the average minus one standard deviation of GDP growth.

4.2 Uncertainty and the extensive margin

We first estimate the impact of uncertainty shocks on the formation of new firm-to-firm relationships. Here, new relationships are defined as the first transaction involving a particular pair of firms, going back to data from 1993 to avoid left-censoring. Results are presented in Table 2. They consistently show that episodes of high uncertainty are associated with significantly fewer 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 fewer new firm-to-firm relationships in high-RS product markets. The result is confirmed in column (2), where the identification is within a product×period and 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)

[33] We have also tried with uncertainty episodes defined as periods in which the index is 1.64 standard deviations above its average. Results are virtually unchanged. In Table 2, we also present results that directly use the level of the index to measure uncertainty.

Table 2: ***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		✓		✓	✓

Notes: Poisson estimations with high-dimensional fixed effects. Uncertainty shocks is a dummy equal to 1 in periods when uncertainty in the destination country is above-average uncertainty plus one s.d. of uncertainty. *RS* is our measure of relationship stickiness which is not centered (Mean: 2.9, P05: 1.8, P95: 3.5). Bootstrapped standard errors are in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

implies the number of new relationships for a product in the first quartile of the distribution of the RS indicator drops by about 2% in periods of high uncertainty as shown in Figure 5. For a more sticky product, in the third quartile of the distribution of RS, the number of new relationships drops by 11% in periods of uncertainty.[34]

Columns (3) and (4) use the value of the WUI 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. 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

[34]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.

Figure 5: ***Impact of an uncertainty shock on new relationships, along the distribution of RS***

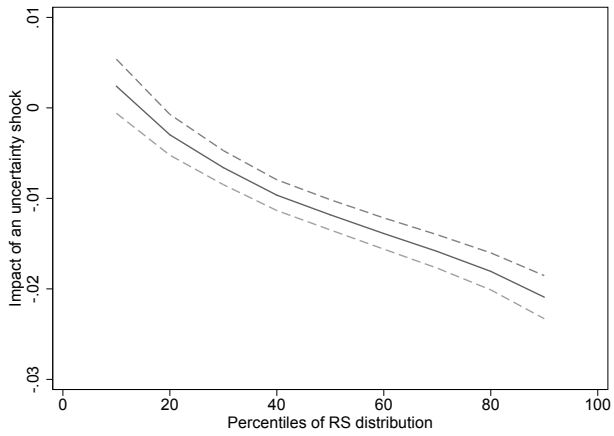

Notes: This figure is recovered from results in Table 2, column (1). It shows the percentage-point impact of an uncertainty shock, on the number of new firm-to-firm relationships.

period, because left-censoring induced by the declaration threshold for product categories in the customs data severely increases 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.

In Table B.9, we dig deeper into the trade effect of uncertainty by looking at its correlation over space and time. In columns (1)-(2), we test for spillover effects of uncertainty shocks, on other destination countries. In theory, an uncertainty shock in country c could have two opposite effects on trade with 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. We expect this scenario 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 2 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 the overall spillover effect is negative, that is, 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. The same qualitative results are found in column (2), using identification across products within

Table 3: ***Uncertainty and the disruption of trade relationships***

	(1)	(2)
Dep. var:	*# disrupted trade relationships*	
Uncertainty-shock dummy	0.23***	
	(0.007)	
- × RS index	-0.08***	-0.03***
	(0.003)	(0.007)
Observations	2,546,156	2,546,156
Period	1996-2006	1996-2006
Fixed Effects		
Product × quarter		✓
Product × period	✓	
Country	✓	
Country × period		✓

Notes: Poisson estimations with high-dimensional fixed effects. Uncertainty shocks is a dummy equal to 1 in periods when uncertainty in the destination country is above-average uncertainty plus one s.d. of uncertainty. *RS* is our measure of relationship stickiness, which is not centered (Mean:2.9, P05: 1.8, P95: 3.5). Bootstrapped standard errors are in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

a destination×period. In columns (3)-(4) of Table B.9, we finally study the dynamics of the trade impact of uncertainty. To this aim, we augment the benchmark specification with lags of the uncertainty variable, interacted with the RS indicator. The impact of uncertainty is shown to be 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 affect the value of waiting, we also expect to see an impact on separation rates. Firms should indeed be reluctant to end a costly trade relationship when uncertainty is high. We investigate this possibility in Table 3. Here, the left-hand-side variable is the number of disrupted relationships observed over a particular period. The variable is constructed using relationships that we observe for the last time over a particular month. The number of disrupted relationships over a quarter and destination is the number of such relationships cumulated over the last quarter. The assumption is that, if these relationships had not ended, firms would have been active at least once in the next quarter. This assumption

Figure 6: ***Impact of an uncertainty shock on disrupted relationships, along the distribution of RS***

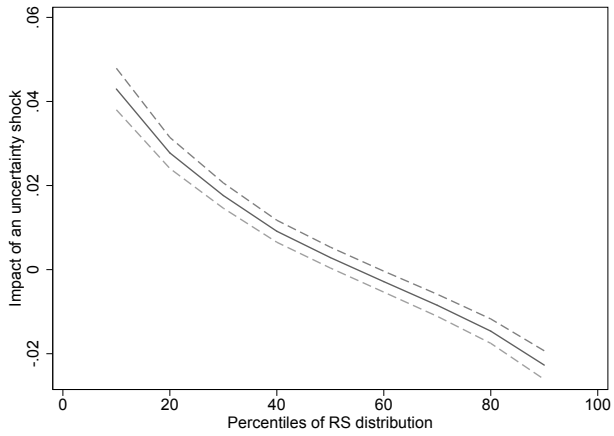

Notes: This figure is recovered from results in Table 3, column (1). It shows the percentage-point impact of an uncertainty shock, on the number of disrupted firm-to-firm relationships.

is reasonable because transactions take place once every three months, on average (Table B.6).

Results in Table 3 suggest the impact of uncertainty shocks on separation rates is indeed significant, although the sign of the relationship varies along the distribution of relationship stickiness. As illustrated in Figure 6, uncertainty periods are associated with significantly more firm-to-firm separations, in product markets displaying little stickiness. In the first quartile of the distribution, the impact is thus positive and significant, at 2%. In the third quartile, the effect is instead negative, at -2%.

The fact that stickiness increases the real option of waiting is consistent with standard models of investment choice under uncertainty. However, that uncertainty increases separation rates for the least sticky products suggests instead that our measure of uncertainty also has a negative impact on the level of expected profits. One can rationalize this positive effect through two alternative interpretations. First, our measure of uncertainty may be negatively correlated with expected demand, which would affect expected profits. Second, even if uncertainty shocks were pure second-moment shocks, they could still have a level impact in presence of managers' risk aversion or non-constant elasticities of profits to demand shocks. Whereas a first-moment negative shock on profits would only reinforce the negative impact of uncertainty on entry, it could act as a countervailing force on exit: wait-and-see behaviors would still dominate for high-enough levels of stickiness, while separation rates could increase for low RS products.

In sum, results presented in this section provide strong evidence for a response of trade to

uncertainty shocks, at the extensive margin. Such adjustments are consistent with episodes of high uncertainty generating significant wait-and-see behaviors from firms that need to invest money in their international development, especially in sticky markets. In the next section, we conclude the analysis by studying the consequences of such behaviors for overall trade adjustments.

4.3 Trade adjustments to uncertainty and GDP shocks

We now examine what our results imply regarding the dynamics of trade in periods of high uncertainty. We follow Bricongne et al. (2012) and Carballo et al. (2018) by decomposing 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.[35] Trade growth is computed using year-on-year growth for each quarter to remove seasonality. We then regress product-level growth and its components on uncertainty and its interaction with relationship stickiness. Here, we further control for GDP shocks and its interaction so that we can compare the dynamics of trade with uncertainty and GDP shocks. For comparability purposes, the GDP shock variable is a dummy that takes the value of 1 when the growth in the destination is one standard deviation below its average over the period of estimation. Results are robust if we use the growth level instead.

We report the point estimates in Table 4, and plot the implied response of the different margins to uncertaity and GDP shocks in Figure 7. The top panel summarizes the response of trade growth to an uncertainty shock, and the bottom panel concerns a growth shock. Several interesting results emerge from the comparison of these graphs. First, both shocks are associated with significantly less trade growth. In comparison with normal times, episodes of high uncertainty are associated with 12 percentage points less growth, on average. The impact of a drop in the destination country's growth is larger, at -25 percentage points. For uncertainty shocks, the impact is roughly constant over the distribution of RS indices. Instead, the impact of a GDP shock is 6 percentage points larger at the 10^{th} than at the 1^{st} decile of the distribution. Second, the margins of adjustments to both shocks are significantly different. The response of trade to uncertainty shocks is mostly driven by the extensive margin, namely, the wait-and-see behaviors documented earlier, which is in line with Carballo et al. (2018).

[35] Unlike Bricongne et al. (2012) and Carballo et al. (2018) 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.

Figure 7: ***Impact of shocks on trade growth, along the distribution of RS***

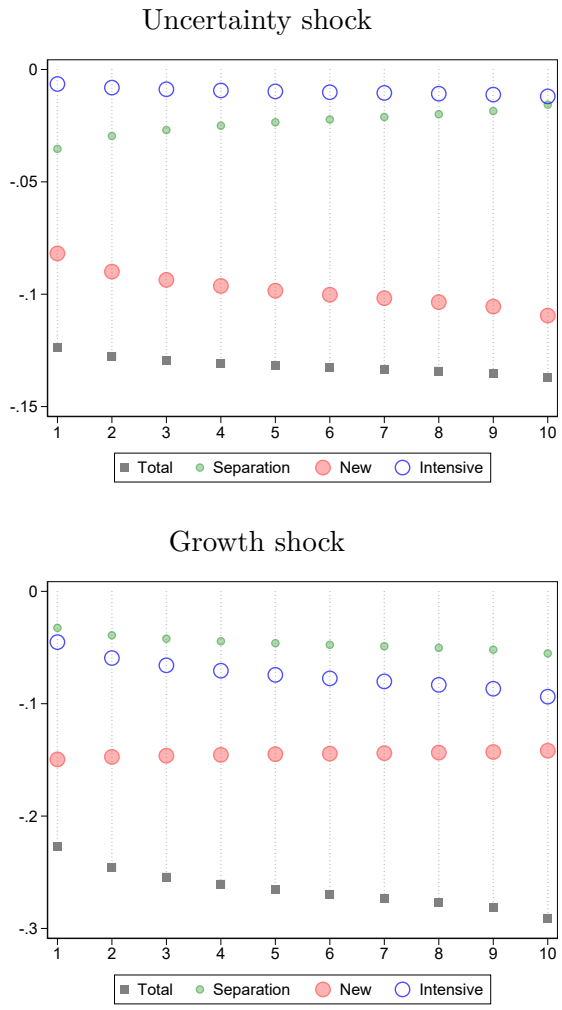

Notes: These figures summarized the response of product-level trade to an uncertainty shock (top panel) and a shock to the destination market's growth (bottom panel). Results are recovered from the estimation of the following equation:

$$Y_{pct} = \alpha Uncert_{ct} + \gamma RS_p \times Uncert_{ct} + \beta GDP_{ct} + \delta RS_p \times GDP_{ct} + FE + \varepsilon_{pct}$$

where the LHS variable is the mid-point growth rate or one of its components, $Uncert_{ct}$ and GDP_{ct} respectively denote an uncertainty and GDP shocks, and FE denotes fixed effects at the product×country level.

Table 4: ***Uncertainty and trade growth: margin decomposition***

	(1)	(2)	(3)	(4)
Dep. var:	Growth	=Start	+ End	+ Intensive
RS index	-0.07***	-0.21***	0.13***	0.01***
	(0.002)	(0.002)	(0.001)	(0.001)
Uncertainty shock	-0.11***	-0.05***	-0.06***	0.001
	(0.004)	(0.002)	(0.002)	(0.001)
- × RS index	-0.01***	-0.02***	0.01***	-0.003***
	(0.001)	(0.001)	(0.001)	(0.0004)
GDP shock	-0.14***	-0.16***	-0.002	0.02
	(0.007)	(0.004)	(0.003)	(0.003)
- × RS index	-0.04***	0.01***	-0.01***	-0.03***
	(0.002)	(0.001)	(0.001)	(0.001)
Observations	3,538,965	3,538,965	3,538,965	3,538,965

Notes: OLS estimation. Bootstrapped standard errors are in parentheses. Growth is the 12-month growth of product-level French exports to a destination. Start, end, and intensive are the different growth margins, namely the number of new seller-buyer relationships, the number of disrupted relationships, and the evolution of seller-buyer sales along the intensive margin. Uncertainty shocks is a dummy equal to one in periods when uncertainty in the destination country is above average uncertainty plus one s.d. of uncertainty. GDP shock is a dummy equal to one in periods when GDP growth is below average GDP growth minus one s.d. of GDP growth. *RS* is our measure of relationship stickiness, which is not centered (Mean: 2.9, P05: 1.8, P95: 3.5). Significance levels: $*$ 10%, $**$ 5%, $***$ 1%.

Instead, the intensive margin is a significant driver of the response of trade to GDP shocks. This finding is consistent with evidence in Bricongne et al. (2012) who show the trade collapse in 2008 was largely driven by the intensive margin. In the panel under study, the intensive margin contributes to one fifth to one third of the overall adjustment, depending on the degree of stickiness in the product market. Finally, the intensity of trade adjustments does vary along the distribution of stickiness. For uncertainty shocks, the heterogeneity mostly kicks in at the extensive margin. Consistent with results in section 4.2, results show strong evidence for muted extensive adjustments in sticky-product markets, with fewer entries somewhat compensated by fewer separations. When the shock hits GDP growth, the heterogeneity instead affects adjustments at the intensive margin, whereas the response of trade at the extensive margin is roughly constant. Intensive adjustments amount to 9.5 percentage points at the 10^{th} decile of the RS distribution, against only 4.5 p.p. at the 1^{th} decile.

We check the robustness of our results by using two alternative proxies for first and second

moment shocks. More specifically, we use the average stock returns and the average volatility of returns computed by Baker et al. (2020).[36] The results, presented in Table B.10, show an increase in the volatility of returns is associated to a decrease in entry – magnified for the most sticky products, and a increase in the disruption of trade relationships – dampened for the most sticky products. These findings are consistent with the wait-and-see behavior of sticky products when uncertainty increases.[37]

Antras (2020) argues severe but temporary shocks such as the 2008-09 trade collapse or the COVID-crisis do not fundamentally change firms' sourcing strategies and are likely followed by a fast recovery. Our findings are consistent with this view. The adverse effect of negative shocks is mainly driven by a reduction in the creation of new relationships and, for first-moment shocks, a drop in firm-to-firm trade along the intensive margin. By contrast, negative shocks have little impact, if any, on the disruption of sticky firm-to-firm trade relationships. Firms are thus likely to expand quickly in the aftermath of a negative shock by building on their existing relationships.

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 emphasis on international trade. This question is a topical one 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 Brexit, the trade war involving the US and most of its partners, and now the consequences of the COVID-19 pandemic. 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 global value chains often need to customize product, or adjust their logistics chain to the particular needs of the firm located downstream. This need generates a substantial degree of persistence in firm-to-firm relationships.

We study the interaction between these two phenomena from an empirical standpoint. We exploit highly detailed firm-to-firm data involving French firms and their partners in the EU, covering a period of more than 20 years. Using these data, we first construct an indicator of relationship stickiness at the product level. The empirical strategy consists of comparing the

[36] We use the macro measures computed by Baker et al. (2020), because they have a better country coverage.

[37] An increase in the level of returns (a positive shock to the first moment) leads to higher export growth, in particular for the most sticky products, which is largely driven by the intensive margin.

average duration of firm-to-firm relationships conditional on a match quality and derive from this comparison an ex-post measure of stickiness. 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 fewer new firm-to-firm relationships, the impact being stronger in product markets displaying stickier firm-to-firm relationships. This finding 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 relationship in which sellers and buyers are engaged. The modern organization of production into fragmented processes increases the sensitivity of the economy to uncertainty shocks.

Finally, note we have focused here on trade adjustment to uncertainty shocks, but stickiness may be important for a range of macroeconomic outcomes. The responses of trade to exchange-rate shocks or trade policy might differ across products with different levels of stickiness (Heise, 2016). The degree of stickiness also affects the international transmission of shocks, with implications for the level and the comovement of economic fluctuations. We hope the measure developed in this paper will stimulate research along these questions.

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.

Ahmad, Saad, Nuno Limão, Sarah Oliver, and Serge Shikher, "Brexit Uncertainty and its (Dis)Service Effects," NBER Working Papers 28053, National Bureau of Economic Research, Inc November 2020.

Antras, Pol, "De-Globalisation? Global Value Chains in the Post-COVID-19 Age," Technical Report, Harvard University November 2020.

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.

Baker, Scott R., Nicholas Bloom, and Stephen J. Terry, "Using Disasters to Estimate the Impact of Uncertainty," NBER Working Papers 27167, National Bureau of Economic Research, Inc May 2020.

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.

Carballo, Jeronimo, Kyle Handley, and Nuno Limão, "Economic and Policy Uncertainty: Export Dynamics and the Value of Agreements," NBER Working Papers 24368, National Bureau of Economic Research, Inc March 2018.

Chor, Davin and Lin Ma, "Contracting Frictions in Global Sourcing: Implications for Welfare," Technical Report, Dartmouth November 2020.

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

Exton, Oliver and Davide Rigo, "The Role of Customer Base in Exporter Dynamics," 2020. Unpublished paper.

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.

Fillat, José L. and Stefania Garetto, "Risk, Returns, and Multinational Production," *The Quarterly Journal of Economics*, 2015, *130* (4), 2027–2073.

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.

_ and _ , "Policy Uncertainty, Trade, and Welfare: Theory and Evidence for China and the United States," *American Economic Review*, 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.

Heise, Sebastian, "Firm-to-Firm Relationships and Price Rigidity - Theory and Evidence," CESifo Working Paper Series 6226, CESifo 2016.

Hémous, David and Morten Olsen, "Long-term Relationships: Static Gains and Dynamic Inefficiencies," *Journal of the European Economic Association*, 2018, *16* (2), 383–435.

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

Kaufmann, Daniel, Aart Kraay, and Massimo Mastruzzi, "The worldwide governance indicators : methodology and analytical issues," Policy Research Working Paper Series 5430, The World Bank September 2010.

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.

Muûls, Mirabelle, "Exporters, importers and credit constraints," *Journal of International Economics*, 2015, *95* (2), 333–343.

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, that is, 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 difference 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 this buyer's switching to a new supplier, and (iii) the inverse of the probability of switching, conditional on trade. If buyers purchased French products at regular intervals, for example every month, the three statistics would deliver the same information. As shown in the second line of Table B.6, the frequency of transactions is neither close to 1, 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 three months. Twenty-five percent of buyers, however, purchase French products more than once every two months, whereas 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 between the two switching probabilities. In the data, the three statistics are correlated at more than 50%, meaning 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:

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

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
$1/\mathbb{P}$(switch)	9	20	9	41
$1/\mathbb{P}$(switch|Trade)	2	3	2	6

Notes: This table provides statistics on alternative measures of durations. The first line is our benchmark measure, defined as the mean number of months between the first transaction involving a particular buyer and one of its supplier and the next transaction in which the same buyer interacts with a different partner ("Mean duration"). "$1/\mathbb{P}$(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. "$1/\mathbb{P}$(switch|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 1996-2006.

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, recovering a relative measure of relationship stickiness using the following log-linear specification is thus possible:

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

where FE_p is a product fixed effect, and ϵ_{qpc} is the error term. Because 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,[38]; (ii) we assign all the trade relationships to a size-decile (specific to a product category); and (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 percentiles of the distribution plus a quantile defined by transactions between the 1st and the 10th percentile, and the quantile of transactions between the 90th and the 99th percentiles.

A.3 External validity tests

Relationship stickiness and intrafirm 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

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

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 intrafirm trade is computed from intra-firm trade data released by the Bureau of Economic Analysis for year 2002. Intrafirm 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% 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 result 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 Kaufmann et al. (2010), and the degree of the 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.[39] 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 both Nunn's and our measures of product stickiness have explanatory power in this regression. When we instead us the Balassa index 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 final 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 the distance effect is magnified in product markets that display more relationship stickiness. This result is true whatever the structure of fixed effects, including in the most demanding specification in column (4). The elasticity of trade to distance also seems 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 a 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 with higher monitoring costs, which increase with distance (Head and Ries, 2008).

[39] 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.

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

Table B.5: ***Summary statistics on the structure of the dataset***

	# transac.	# sellers	# buyers	# sellers ×products	# buyers ×products	# buyer×seller ×products
	(1)	(2)	(3)	(4)	(5)	(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, and 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.

B Additional results

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

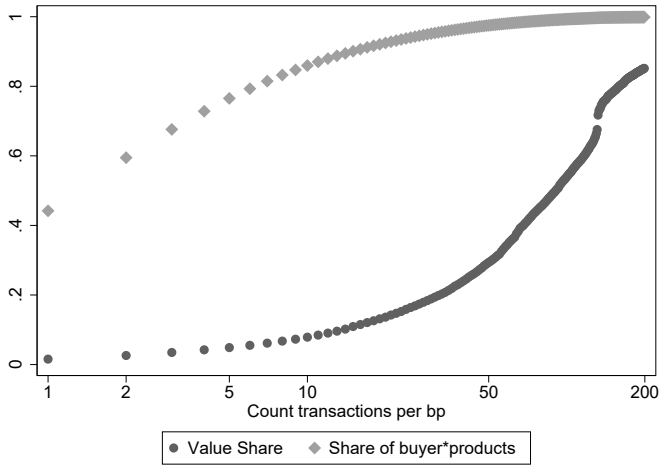

Notes: Cumulated distribution of the number of transactions per foreign buyer (×product). A transaction is the purchase of a particular good, to a given seller, in a given month. The light grey line corresponds to the share in the population of buyers and the dark line measures what this represents in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

Table B.6: ***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 B.2: ***Distribution of the number of partners, per buyer/seller and date (month×year), without conditionning on a particular product***

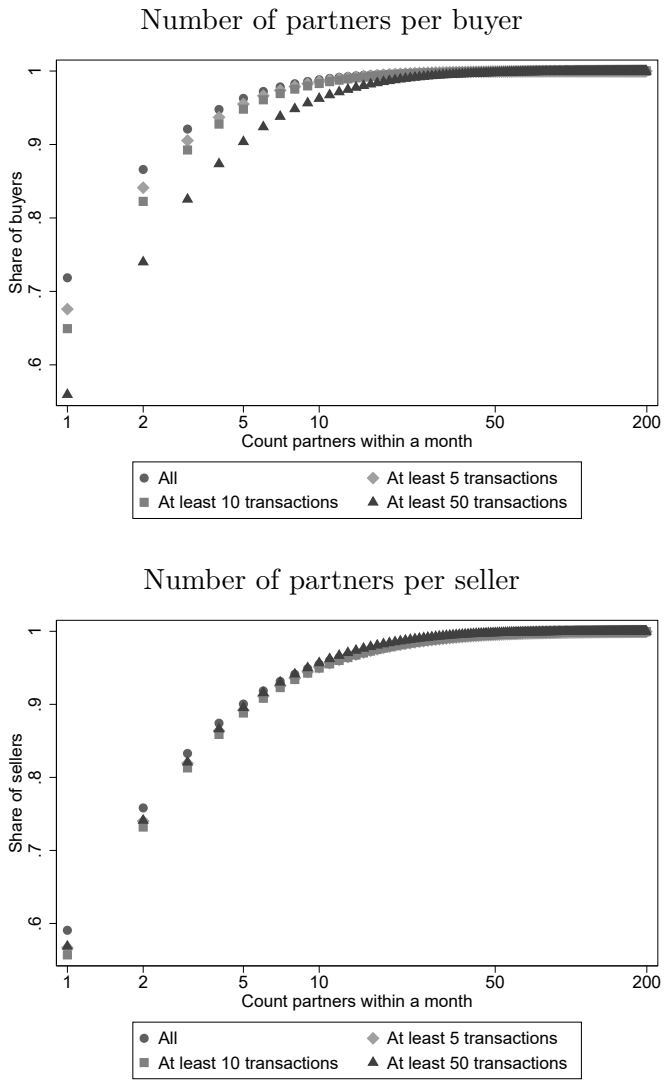

Notes: Cumulated distributions of the number of partners a French exporter interacts with in a given country (bottom panel), and the number of partners a foreign buyer interacts with within a particular month (top panel). Both statistics are calculated across products within a firm. The number of partners 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.

Table B.7: ***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. Standard errors are in parentheses. Significance levels: ∗ 10%, ∗∗ 5%, ∗ ∗ ∗ 1%.

Figure B.3: ***Relationship stickiness across Broad Economic Categories***

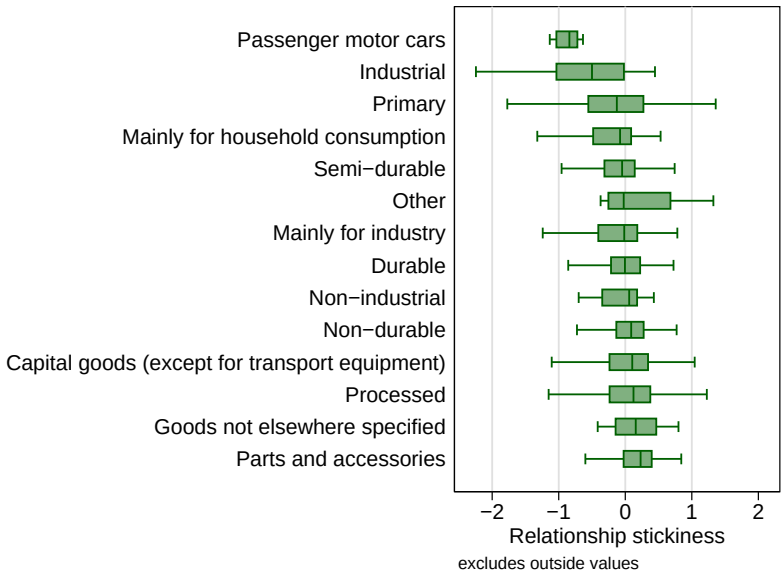

Notes: Boxplot of relationship stickiness across Broad Economic Categories (BEC). Link between the two measures based on the correspondence between 2003 BEC and HS 2002 nomenclature.

Table B.8: ***Correlation with moments of French data***

Reg. RS on:	β	$R2$
Share buyers active a single month	-0.95***	0.068
	(0.038)	
Share buyers matched with multiple suppliers	-3.78***	0.118
	(0.111)	
Product-level HHI	0.17***	0.010
	(0.018)	
Product-level dispersion in unit values	0.10***	0.03
	(0.006)	
Value share of wholesalers	-0.26***	0.032
	(0.016)	

Notes: Coefficients obtained from the regression of the measure of relationship stickiness on different statistics computed on French data. *Share buyers active a single month* is the share of buyers within a product category that import only once over the period. *Share of buyers with many suppliers* is the share of buyers within a product category that import a given product from several French suppliers within a given month. *Product-level HHI* is the Herfindahl index computed for each product category using firm-level exports. *Product-level dispersion in unit values* is computed as the interquartile range of unit values with product categories. *Value share of wholesalers* is the share of French exports within a product category that is made up by firms identified as retailers or wholesalers. Standard errors are in parentheses. Significance levels: $*$ 10%, $**$ 5%, $***$ 1%.

Figure B.4: ***Heat map of the uncertainty series***

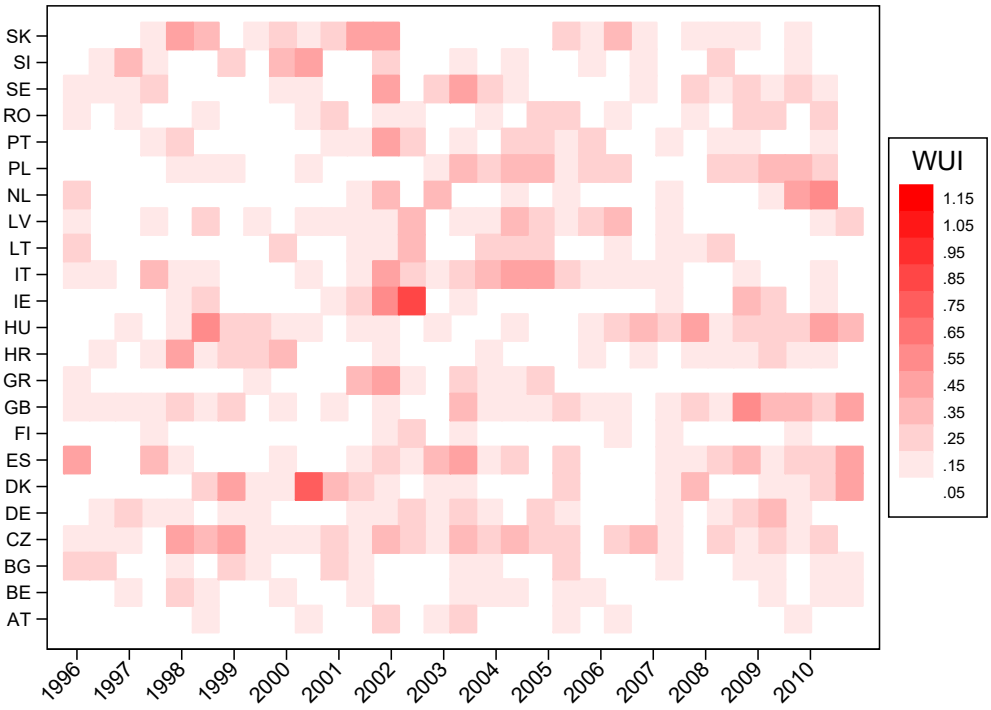

Notes: The figure shows a heat map of the World Uncertainty Index, in the panel under study.

Table B.9: ***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		✓		✓

Notes: Poisson estimations with high-dimensional fixed effects. Uncertainty shocks is a dummy equal to one in periods when uncertainty in the destination country is below average uncertainty minus one s.d. of uncertainty. *RS* is our measure of relationship stickiness. Trade diversion examined in columns (1) and (2). Persistence examined is columns (3) and (4). Standard errors are in parentheses. Significance levels: ∗ 10%, ∗∗ 5%, ∗ ∗ ∗ 1%.

Table B.10: ***Uncertainty and trade growth: margin decomposition***

	(1)	(2)	(3)	(4)
Dep. var:	Growth	=Start	+ End	+ Intensive
RS index	-0.02	-0.31***	0.27***	0.01
	(0.017)	(0.012)	(0.008)	(0.006)
Vol. of returns	-0.18***	-0.09***	-0.08***	-0.01***
	(0.006)	(0.004)	(0.003)	(0.002)
- × RS index	0.01***	-0.01***	0.02***	-0.00
	(0.002)	(0.001)	(0.001)	(0.001)
Level of returns	-0.01**	0.01***	-0.02***	-0.01***
	(0.005)	(0.003)	(0.002)	(0.002)
- × RS index	0.01***	-0.01***	0.01***	0.01***
	(0.002)	(0.001)	(0.001)	(0.001)
Observations	3,538,965	3,538,965	3,538,965	3,538,965

Notes: OLS estimation. Growth is the 12-month growth of product-level French exports to a destination. Start, end, and intensive are the different growth margins, namely the number of new seller-buyer relationships, the number of disrupted relationships, and the evolution of seller-buyer sales along the intensive margin. The level and volatility of stock returns are from (Baker et al., 2020). RS index is our measure of relationship stickiness, which is not centered (Mean: 2.9, P05: 1.8, P95: 3.5). Significance levels: $*$ 10%, $**$ 5%, $***$ 1%.