Online Appendix for

Relationship stickiness, international trade, and economic uncertainty

Julien Martin* Isabelle Mejean† Mathieu Parenti‡

October 23, 2023

O Online Appendix

O.1 Details on the data

The dataset used in the analysis includes information on individual transactions, such as the seller identifier, buyer identifier, product category (at the 8-digit level of the European combined nomenclature), date (month and year), and the value of the shipment in euros. However, it does not provide information on whether the transactions are arm's length or intrafirm.[1] The dataset covers the period from 1993 to 2017, but the analysis focuses on various sub-periods. The main reason for analyzing sub-periods is due to changes in the product category nomenclature over time, which complicates the definition of product markets when estimating stickiness. To address this issue, a harmonization algorithm,

*UQAM and CEPR, Email address: martin.julien@uqam.ca.
†Sciences Po and CEPR, Email address: isabelle.mejean@sciencespo.fr.
‡Paris School of Economics, INRAE, Université Libre de Bruxelles, CESifo, and CEPR, Email address: mathieu.parenti@psemail.eu.

[1]We discuss how intrafirm trade may affect our results in section O.7.

described in Behrens et al. (2018), is applied to recover time-invariant product categories. This algorithm minimizes information loss when applied over shorter time horizons. In the baseline specification, the measurement of relationship stickiness uses data from 1996 to 2006. To assess robustness over time, an alternative period of 2011-2017 is also considered. For the analysis of the impact of uncertainty on the dynamics of firm-to-firm relationships, data from 1996 to 2010 is used.

The dataset encompasses all the relationships involving French exporters and European buyers within each product category.[2] The analysis focuses on trade with the eleven historical members of the European Union. The decision to exclude the new member states from the analysis is based on two reasons. Firstly, information on individual importers prior to their entry into the EU is not available, limiting the completeness of the dataset. Secondly, the entry of new member states into the European Union lead to significant trade adjustments, potentially resulting in substantial changes in firm-to-firm relationships. The large churning in relationships could introduce biases in the estimates of relationship stickiness, warranting their exclusion from the analysis.

Table O.1 presents comprehensive statistics related to the data used to calculate the baseline measure of relationship stickiness. The observations in the dataset are based on individual transactions, with each transaction representing a unique combination

[2]The dataset used in the analysis does not include exports from the smallest French exporters. These exporters are allowed to complete a simplified form that does not specify the product category. As a result, transactions involving these exporters are excluded from the analysis. In 2007, the simplified reporting regime applied to 21,616 exporters out of a total of 66,131, accounting for approximately 2% of transactions and 0.5% of the total value of French exports. Furthermore, there was a significant increase in the declaration threshold for firms to fill in the most detailed customs form in 2011. This change further motivates the selection of the preferred time period for analysis as 1996-2006, prior to the implementation of the threshold increase.

of a seller, a buyer, a product, and a specific month. Columns (1)-(3) of the table provide statistics on the number of transactions and the number of firms involved in these transactions. Columns (4)-(6) offer additional statistics on the dimensionality of the graph when treating multi-product importers and exporters as independent units. This approach considers firms that engage in transactions involving multiple products as distinct entities.

In the cross-section, the dataset exhibits a bipartite graph structure that connects individual seller×product pairs to individual importers. This graph structure predominantly aligns with a many-to-one matching pattern. Specifically, at a given point in time (defined by a particular month in a specific year), we observe most buyers purchasing a specific product from a single seller, while sellers simultaneously serve multiple importers, even within a single country. This observation is visually presented in Figure O.1, which showcases the distribution of the number of sellers interacting with a given importer during a specific month (top panel) and the distribution of the number of partners from the same country with whom a French exporter engages (bottom panel). More than 90% of importers have only one French supplier for a given product within a particular month. Even when focusing on importers with more than 50 transactions, this proportion remains high, exceeding 80%. Conversely, 26% of French exporters sell the same product to multiple partners within the same month, and this proportion increases to 55% when considering partners located in different countries.[3] Given this underly-

[3]This finding contrats with the results in Bernard et al. (2018) who employ qualitatively similar data and observe many-to-many relationships between exporters and importers. One potential explanation for this discrepancy, apart from the different country coverage, is that their analysis does not condition on a specific product, whereas our approach does. Indeed, our data reveals that buyers often engage with several French exporters within a given month, albeit for different products (refer to Figure O.3 and the comparison with Figure O.1, where the former aggregates partners across products within a

ing data structure, the model in section 3 assumes a many-to-one matching structure, where importers interact with a single supplier at any given moment. If an importer is observed engaging with two different exporters within a month, we consider those two transactions as occurring simultaneously.

The distribution of the number of transactions by buyers is highly skewed, as illustrated in Figure O.2. Only 8% of importers are observed over more than 20 transactions with French firms, but they account for more than 85% of trade. The dynamics of their relationships with French firms should thus provide insightful information. At the other end of the spectrum, 44% of buyers are engaged in only one transaction with a French seller over the ten-year period. These buyers make surprisingly tiny transactions: they account for only 1.5% of the value of trade. There are good reasons to believe that a substantial share of these transactions correspond to non-market transactions, such as samples sent by exporters to prospective clients. We thus decided to exclude these one-shot buyers when estimating stickiness.

Figure O.4 illustrates the distribution of the number of French partners with whom individual buyers interact throughout their entire presence in the dataset. Buyers that appear only once in the data are excluded from this analysis. The figure reveals that approximately 67% of buyers have a single partner in France, while less than 7% have three or more partners (as indicated by the circles line). It is important to note that interacting with a single partner in France is more likely for firms engaged in a small number of transactions. To provide further insights, the figure presents three additional distributions depicting the number of partners per buyer for importers involved in at least 5, 10, or 50 transactions. Even within the subset of importers observed in at

firm while the latter counts the number of partners for a specific product). Once we condition the analysis on a particular product, the occurrence of purchasing from multiple French exporters becomes exceedingly rare.

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

	# transac.	# sellers	# buyers	# sellers ×products	# buyers ×products	# buyer×seller ×products
	(1)	(2)	(3)	(4)	(5)	(6)
EU12	101,379,585	109,522	1,583,220	1,340,346	14,195,710	19,383,546
Belgium	19,872,676	74,924	185,596	637,007	2,488,213	3,596,690
Denmark	1,938,872	23,057	26,962	126,801	249,992	352,214
Germany	19,426,804	61,159	349,803	495,009	2,621,373	3,537,033
Greece	2,003,763	20,238	31,828	139,837	302,191	419,877
Ireland	1,293,531	16,414	15,925	88,334	182,032	270,832
Italy	12,662,419	51,963	280,641	381,644	2,144,174	2,792,808
Luxemburg	3,086,374	31,580	19,028	199,820	402,186	560,297
Netherlands	6,158,922	44,031	90,507	267,196	772,004	1,099,336
Portugal	4,833,183	33,528	67,248	238,463	762,041	1,024,489
Spain	12,581,119	53,471	237,767	419,964	1,928,424	2,490,565
UK	10,487,916	49,325	151,545	360,504	1,321,563	1,923,611

Notes: This table is based on French customs firm-to-firm data covering the period from 1996 to 2006. The first row presents aggregated statistics for all countries, while the subsequent rows provide country-specific statistics. Each column represents a different aspect of the data:

- Column (1): Total Count of Transactions. A transaction is defined as a trade flow occurring within a specific month and year, involving a particular seller-buyer pair and product.
- Column (2): Number of Exporters - This column indicates the count of unique exporters observed throughout the entire period.
- Column (3): Number of Importers.
- Column (4): Count of Distinct Seller-Product Pairs.
- Column (5): Number of Unique Buyer-Product Pairs.
- Column (6): Cumulative Count of Seller-Buyer-Product Triplets.

least 50 transactions, around one-third of buyers consistently interact with the same exporter, indicating a level of loyalty in their firm-to-firm relationships. This observation aligns with the notion that certain firm-to-firm relationships in international markets exhibit a high degree of stickiness. The empirical analysis aims to investigate whether this stickiness is systematically associated with specific products or sectors. It seeks to explore whether relationship stickiness is linked to the uniqueness or specificity of certain products or industries.

O.2 Measuring the duration of a trade relationship

The duration is measured as the number of months between the first and the last uninterrupted transactions involving a French exporter and a foreign importer. However, this

Figure O.1: ***Distribution of the number of partners, per buyer/seller and date (month×year)***

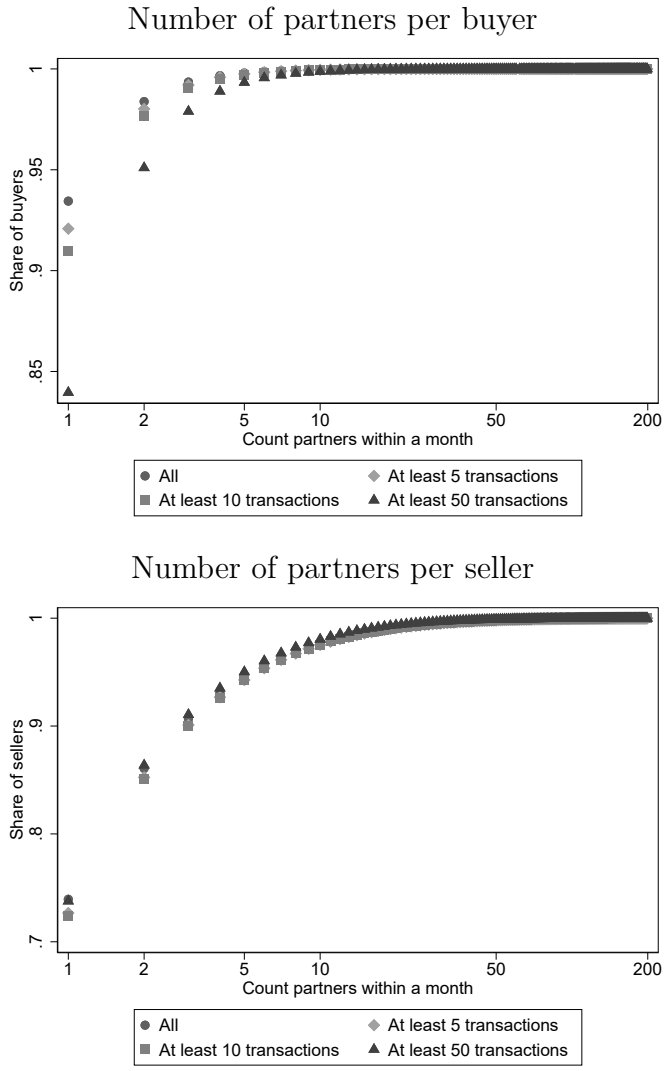

Notes: Cumulated distributions are presented in the bottom panel for the number of partners a French exporter engages with in a specific country. In the top panel, the cumulative distribution illustrates the number of partners a foreign buyer (×product) interacts with within a given month. The calculation of partner counts is based on a sub-sample of importers (or exporters) involved in a minimum of two transactions throughout the analysis period. Additionally, separate distributions are provided for importers/exporters with at least 5, 10, and 50 transactions.

Figure O.2: ***Distribution of the number of transactions, per buyer×product***

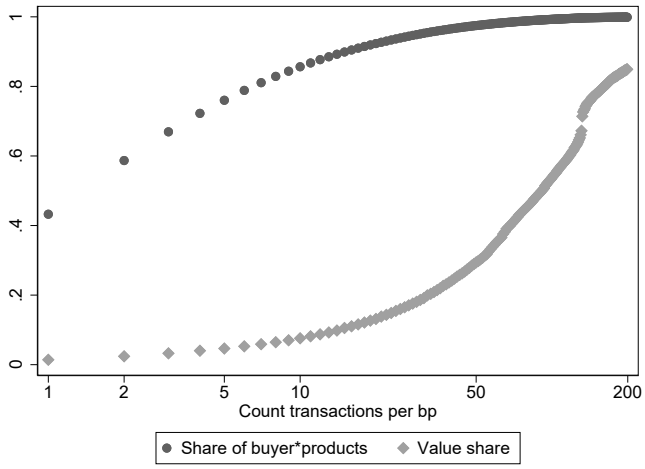

Notes: The graph illustrates the cumulative distribution of the number of transactions per foreign buyer (×product). Each transaction corresponds to the purchase of a specific good, from a particular seller within a given month. The light grey line represents the share of buyers within the population, while the dark line represents the corresponding share in the overall value of exports.

definition poses several challenges that we will now discuss. In cases where an importer interacts with multiple exporters within a single month, determining the continuity of the relationship becomes somewhat ambiguous as the sequence of these transactions cannot be precisely defined. To address this, we consider these transactions to occur simultaneously. If the importer had an existing interaction with one of the exporters before, we consider the relationship to continue. Furthermore, if, in the following period, the firm is observed interacting with one or several of these partners again, we include the transaction that took place simultaneously with other transactions. In unreported results, we have verified that our estimates remain virtually unchanged when we exclude importers from the estimation sample who eventually engage with multiple sellers within a month.

The bilateral nature of our data further complicates the concept of a continuous relationship. We lack observations on transactions between foreign importers and their non-French suppliers. However, this censoring does not directly impact our measurement

Figure O.3: ***Distribution of the number of partners, per buyer/seller and date (month×year), without conditioning on a particular product***

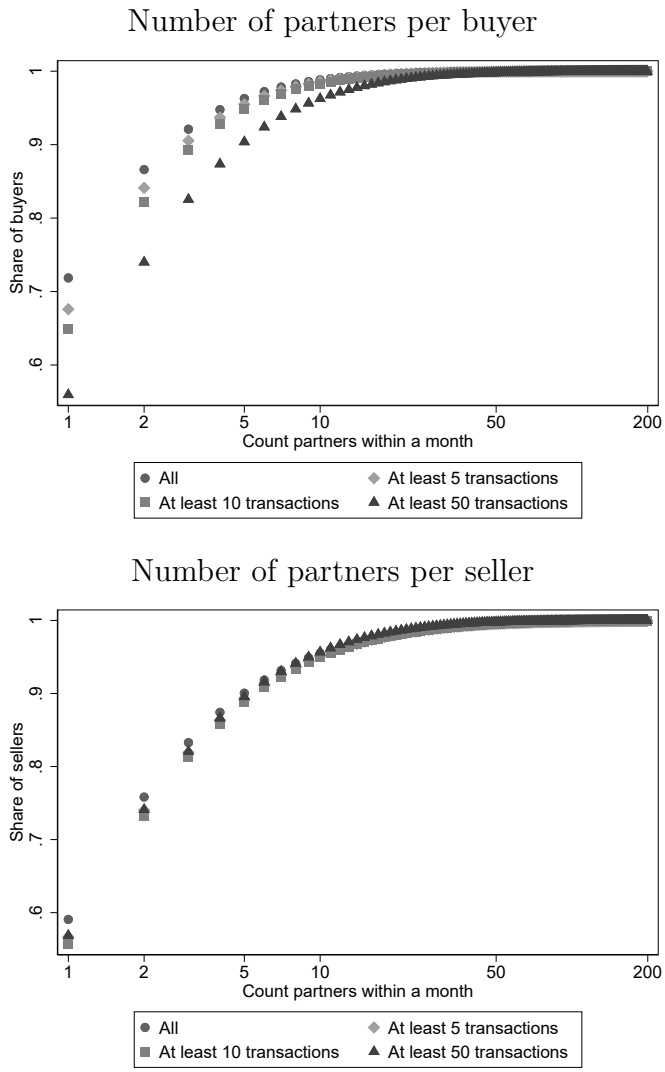

Notes: The graph displays cumulative distributions for two statistics: the number of partners a French exporter interacts with in a specific country (bottom panel) and the number of partners a foreign buyer interacts with within a particular month (top panel). Both statistics are calculated considering all products within a firm. The number of partners is determined based on a sub-sample of importers (resp. exporters) involved in at least two transactions throughout the analysis period. Additionally, separate distributions are shown for importers (resp. exporters) with at least 5, 10, and 50 transactions.

Figure O.4: ***Distribution of the number of French partners, per buyer×product***

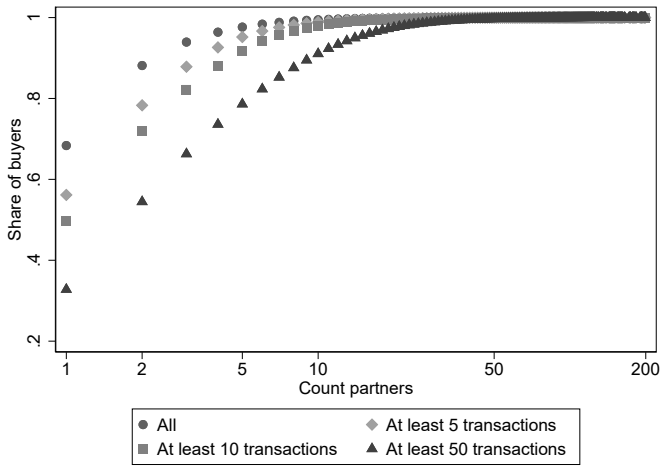

Notes: The graph illustrates the cumulative distribution of the number of partners for each foreign buyer (×product). In this context, a partner refers to a French exporting firm. The calculation of the number of partners is performed using a sub-sample of importers involved in at least two transactions (referred to as "All"), as well as subsets with at least 5, 10, and 50 transactions.

of duration, which relies on the time elapsed between the first and last transactions involving two firms, regardless of whether the relationship terminates due to the importer switching to another French supplier, a non-French partner, or ceasing to purchase the product altogether. Nevertheless, it is worth noting that some durations may be overestimated if the buyer switches to a non-French seller before reverting back to their previous partner. In such cases, we observe two consecutive transactions that we assign to a single continuous relationship. To address this concern, we propose an alternative estimation of stickiness. In this approach, duration is computed as the number of transactions within a continuous relationship, rather than the number of months between any two transactions.

The aforementioned issue is closely tied to the frequency of transactions. While defining continuous relationships would be straightforward with monthly transaction frequencies, utilizing actual transaction data introduces significant heterogeneity in transaction

frequencies. On average, there is a 33% probability of a transaction occurring in a given month, corresponding to an average of one transaction every three months. However, 25% of buyers purchase French products more frequently than once every two months, while firms in the first quartile of the distribution make purchases less than once every 10 months (refer to Table O.2 in the appendix). Our baseline approach for measuring durations treats a relationship involving two transactions spaced six months apart in the same manner as a relationship consisting of seven monthly transactions. In section O.5, we address the possibility that the heterogeneous frequencies across products might impact our results. To do so, we replicate the methodology for estimating stickiness using an alternative duration measure that considers the number of transactions within a specific relationship, rather than the elapsed months.

When working with durations, it is common to encounter censoring issues. In our analysis, we address right censoring by excluding transactions that start within two years before the conclusion of our estimation sample. This exclusion ensures that we have sufficient information to accurately measure the duration of relationships. In a robustness check, we additionally estimate relationship stickiness using durations measured in a sample where we exclude left-censored relationships. By doing so, we focus solely on relationships for which we have complete information and avoid any potential biases that may arise from incomplete data on the starting point of these relationships.

In the model, the duration of a relationship is represented as a function of the probability of a switch, which refers to an importer splitting from its current partner to begin interactions with a new one. However, in the actual data, these two concepts do not align perfectly due to the heterogeneity in transaction frequencies. This disparity is highlighted in Table O.2, which compares statistics on (i) the mean duration of a buyer's relationships with French suppliers, (ii) the inverse of the probability of the buyer switching to a new supplier, and (iii) the inverse of the conditional probability of switching

given a trade transaction. If buyers consistently purchased French products at regular intervals, such as every month, the three statistics would convey the same information. However, as indicated in the fourth line of Table O.2, transaction frequencies are neither close to 1 nor homogeneous across buyers. On average, the probability of a transaction occurring in a given month is .369, corresponding to a transaction every three months. However, 25% of buyers purchase French products more frequently than once every two months, while firms in the first quartile of the distribution make purchases less than once every 10 months. Due to these heterogeneous frequencies, the three available duration measures are not equivalent. Generally, it can be demonstrated that the mean duration falls between the two switching probabilities. In the data, the three statistics exhibit a correlation of more than 50%, suggesting that the heterogeneity in transaction frequencies does not completely distort the distribution of trade durations across buyers and products.

Table O.2: ***Descriptive statistics on alternative measures of the duration of firm-to-firm relationships***

	Mean	Median	P25	P75
Mean duration	23	28	5	32
$1/\mathbb{P}$(switch)	28	19	8	37
$1/\mathbb{P}$(switch|Trade)	8	3	2	7
Frequency of transactions	0.369	0.250	0.130	0.500
Proba Recall	0.053	0.000	0.000	0.000

Notes: This table presents statistics on alternative measures of durations. The first line represents our benchmark measure, which is the average number of months between the first and last transactions within a continuous relationship for a specific pair of firms ("Mean duration"). "$1/\mathbb{P}$(switch)" is the reciprocal of the switching probability, calculated as the number of switching episodes divided by the total number of months a buyer appears in the data. "$1/\mathbb{P}$(switch|Trade)" is the reciprocal of the conditional switching probability given a transaction, computed as the number of switching episodes divided by the total number of transactions. The "Frequency of transactions" is obtained by dividing the total number of transactions by the overall duration of the importer's presence in the data. This measure represents the monthly probability of a transaction. Lastly, "Proba Recall" quantifies the probability that, during a switching episode, the buyer starts interacting with an exporter with whom they had a previous relationship. These statistics are calculated for each individual importer and then averaged across buyers, using the dataset covering the period from 1996 to 2006.

In our definition of a relationship, we consider two continuous periods with the same

firms interacting, interrupted by another relationship involving the same importer but a different French exporter, as two separate relationships. This means that we do not retain the complete history of the importer's partners. By abstracting from the entire history of the buyer's interactions with French sellers, we simplify the analysis significantly. Furthermore, in the data, the probability of a "recall," which refers to a buyer switching back to a supplier they have interacted with before, is very small. This is evident from the last line in Table O.2. Given the limited occurrence of such recalls, we can safely ignore them for the purposes of our analysis.

We now turn to two stylized facts on the duration of firm-to-firm relationships. Figure O.5 displays the distribution of durations in the sample used for analysis. The mode of the distribution is at one month which corresponds to an importer that interacts with a French firm over a single month, before eventually switching to another French or non-French supplier. 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. This may explain that the distribution is shifted to the right when we weight relationships by their value (blue bars in Figure O.5).

Table O.3 provides evidence of a positive correlation between the duration of trade relationships and the mean size of transactions, which we use as proxy for the quality of the match. This correlation is observed across buyers within a particular product category and within a buyer, across different suppliers encountered throughout their interactions with French firms. The empirical framework includes controls for the correlation between relationship duration and transaction size. By controlling for this correlation, the framework aims to isolate the product-specific attributes that contribute to different degrees of stickiness.

Figure O.5: ***Distribution of the durations in firm-to-firm relationships***

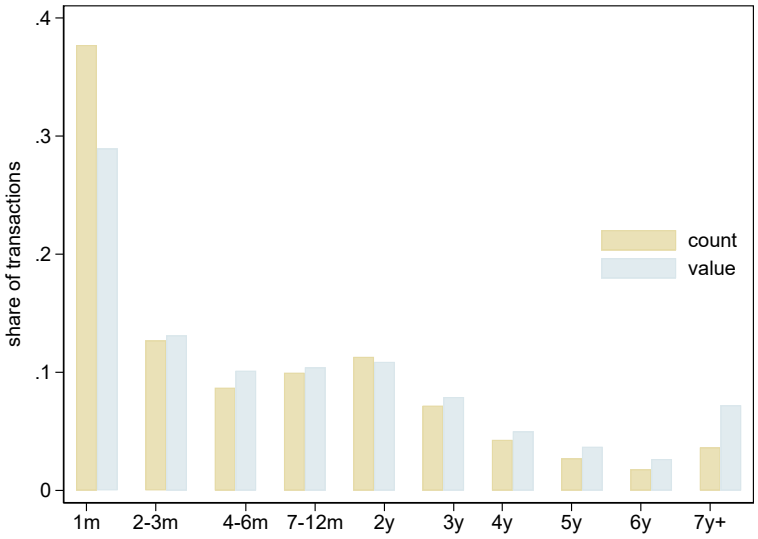

Notes: The distribution of durations is presented in terms of two measures: as a share of the total number of relationships ("count") and as a share of the aggregate export value ("value"). These statistics are derived from the analysis of the 19.5 million firm-to-firm relationships identified during the period of 1996-2006.

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

	(1)	(2)	(3)
	Log of duration		
Log of mean exports	.026***	.069***	.239***
	(.000)	(.000)	(.000)
Observations	14,954,675	14,954,569	10,400,818
R^2	0.001	0.151	0.372
Fixed effects		Product	Product× buyer

Notes: This table presents the correlation between the duration of a relationship and a measure of the average size of transactions occurring within that relationship. The statistics are computed using the dataset that covers the period from 1996 to 2006. The standard errors are indicated in parentheses. Significance levels: ∗ 10%, ∗∗ 5%, ∗ ∗ ∗ 1%.

O.3 Theoretical appendix

In this section, we provide a detailed description of the model solution used in our baseline estimation, as well as its extension to incorporate uncertainty. The simulation parameters used in our analysis are listed in Table O.4 at the end of this section.

O.3.1 Solution under a multiplicative switching cost

To solve the model, a functional form for the switching cost, denoted as $C(\gamma; p)$, needs to be specified. In the benchmark empirical model, the firm's reservation price, denoted as $p^*(\gamma; p)$, is assumed to be proportional to the price of its current supplier:

$$p^*(\gamma; p) = \frac{p}{\gamma}$$

Figure O.6 illustrates the functional form for the corresponding switching cost $C(\gamma; p)$.[4] Under the calibration, the switching cost is decreasing in p, which comes from $V()$ being sufficiently convex.

O.3.2 Solution under a fixed switching cost

To examine the impact of the parametric assumption regarding the form of the switching cost on the qualitative results, we compare the predictions of the baseline model to an alternative framework where the switching cost remains invariant to the buyer's reservation price. In this alternative framework, the switching cost is defined as $C(\gamma, p) = \gamma - 1$, and the optimal switching strategy is no longer multiplicative in p. Consequently, the relationship between durations and prices is no longer log-linear.

[4] The calibration in this figure is performed to match the median duration of firm-to-firm relationships at the median of the simulated price distribution. Detailed information on the calibrated parameters can be found in Table O.4.

Figure O.6: ***Switching cost as a function of prices***

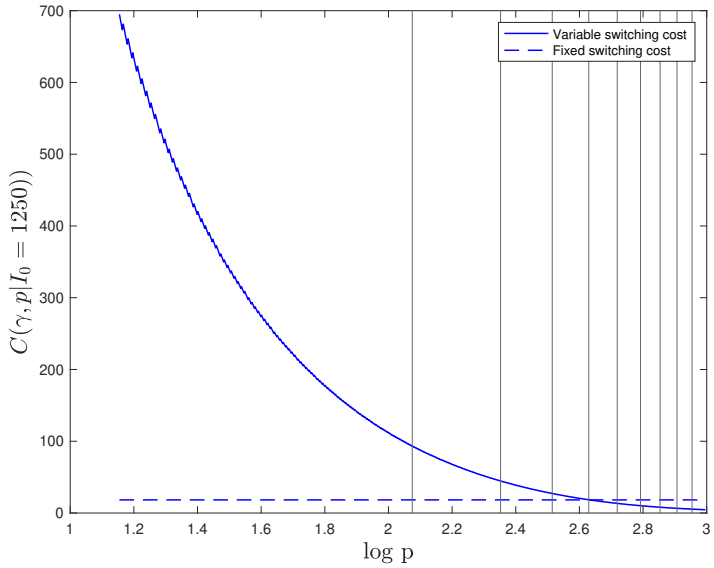

Notes: The figure provides a comparison of the distributions of switching costs in the multiplicative and fixed cost models. The vertical lines represent the first to ninth deciles of the price distribution. The calibration of the model is conducted to ensure that the median duration of relationships in our data is matched.

Figure O.7 illustrates the solutions of both models under different degrees of stickiness. The model with a fixed switching cost exhibits more convexity compared to the baseline log-linear case. However, the qualitative impact of increasing stickiness (i.e., increasing γ) remains the same as in the baseline case, resulting in a shift of the distribution of expected durations. It is important to note that the measure of stickiness employed in the empirical analysis relies on the ranking of expected durations across the distribution of price deciles. Despite the differences in the parametric assumptions, the information obtained from the empirical strategy remains valuable, as it captures the relative stickiness of different products and sectors.

O.3.3 Log-normal price distribution

We also conducted simulations using an alternative log-normal distribution for prices. This alternative distribution was calibrated based on the findings of Head et al. (2014),

Figure O.7: ***Expected durations in the baseline model and in a model with a fixed switching cost***

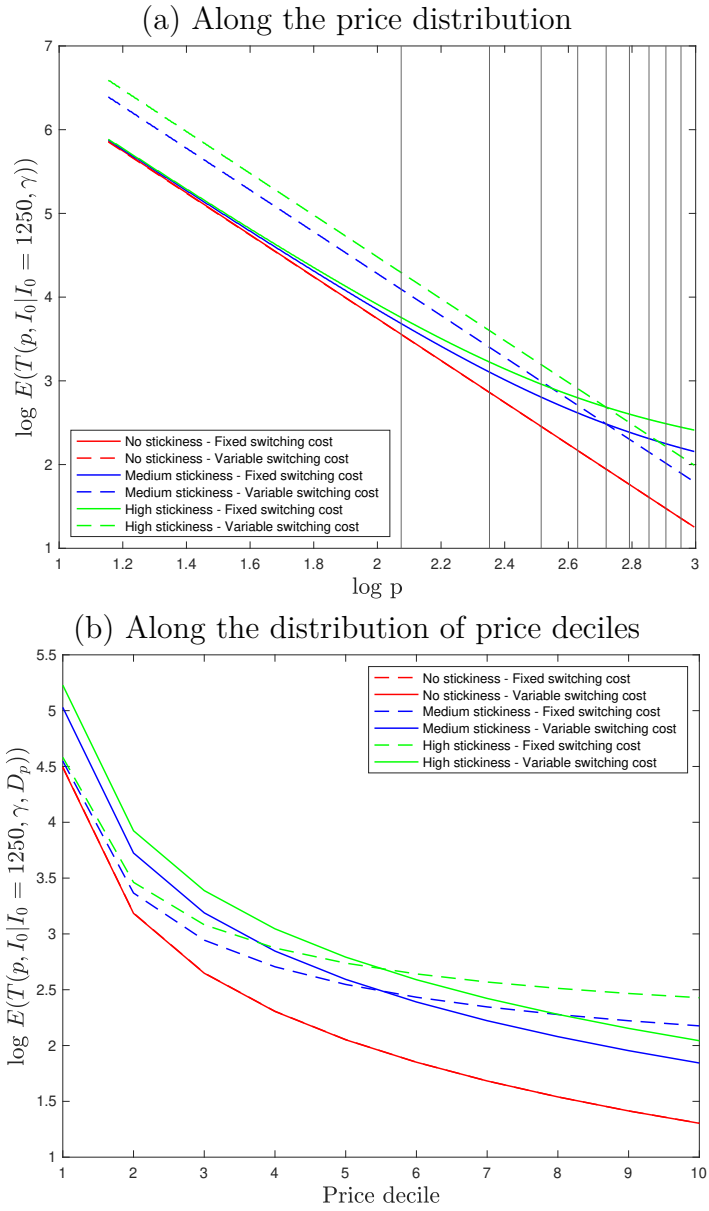

Notes: In the figure, we can observe a comparison of the model solutions under two different assumptions for the switching cost: the baseline variable switching cost (solid lines) and the fixed switching cost (dotted lines). The blue lines represent the baseline calibration that matches the median duration of relationships in the data. The top panel of the figure displays prices on the x-axis, while the bottom panel shows price deciles. The red lines represent a scenario with no switching cost ($\gamma = 1$), where the relationship between durations and prices is flat. This implies that the duration of relationships remains constant regardless of the price of the current supplier. The green lines depict a "sticky" scenario where the duration at the median is set to 18 months, which corresponds to the 75th percentile of the product-level distribution of median durations.

Figure O.8: ***Robustness to the distribution assumption: Pareto versus log-normal***

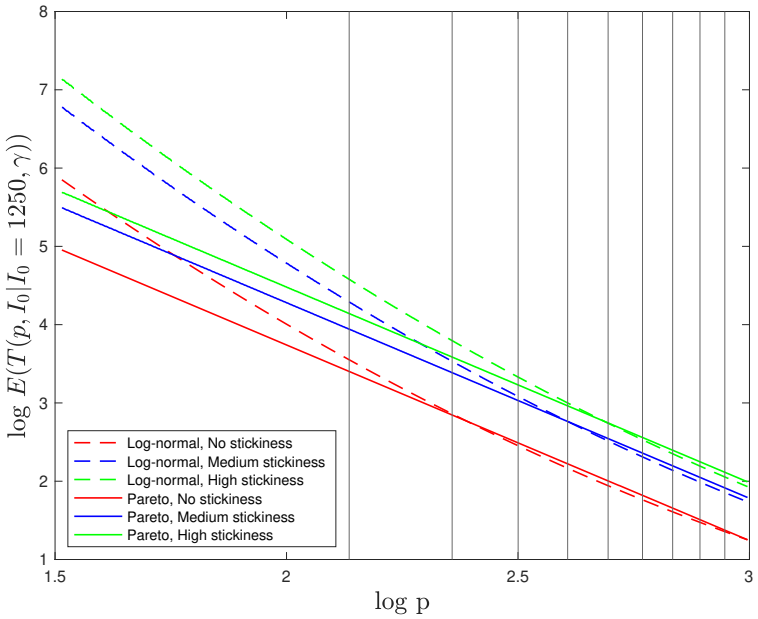

Notes: The figure compares the solutions of the model under two different assumptions for the distribution of prices: the Pareto distribution (solid lines) and the log-normal distribution (dashed lines). The various colors in the figure represent different calibrations for three levels of stickiness. The vertical lines in the figure represent the first to ninth deciles of the price distribution.

who estimated shape parameters for both Pareto and log-normal distributions using French export data. They determined that a log-normal distribution with a shape parameter of 0.8 provided a good fit to the data. In our simulations, we adjusted the location parameter of the log-normal distribution to match the expected duration at the median of the price distribution observed in our data, which was 12. This adjustment ensured consistency with the observed durations while allowing us to examine the model's performance under an alternative pricing distribution.

Figure O.8 compares the results obtained under the Pareto and log-normal distributions. The log-linear relationship between durations and prices is lost when moving away from the Pareto case, as expected. However, even under the log-normal distribution, the convexity is relatively moderate, indicating that fitting a linear relationship

to the actual distribution generated by the log-normal assumption does not introduce a significant error. The intercept of the linear relationship remains informative about the level of stickiness determined by the model's parameters.

O.3.4 A risk-averse manager under mean-preserving uncertainty

In Section 3.2, we extended the model to incorporate uncertainty shocks. More specifically, we introduced a non-mean-preserving uncertainty shock leading to downside risk. We now test for the robustness of our predictions using an alternative assumption, in which the shock is mean-preserving. In solving the model under the assumption of a mean-preserving uncertainty shock, we incorporate risk aversion. We assume that the utility of profits, denoted as $u(\pi(p, I))$, is concave in the level of profits, represented as $\pi(p, I)$. Specifically, we use the functional form $u(\pi(p, I)) = \pi(p, I)^{1-\rho}$, where ρ represents the degree of relative risk aversion. In our case, we set ρ to a value of 0.9. By introducing risk aversion into the utility function, we capture the firm's aversion to uncertainty and its willingness to trade off expected profits with reduced risk. This allows us to examine the impact of risk aversion on optimal switching decisions and the dynamics of the model. The remaining parameters of the model are kept unchanged, except for the switching costs, which are adjusted to match the observed durations at the median of the price distribution.

In Figure O.9, the bottom panel presents the results of simulations conducted under different levels of uncertainty, while the top panel represents the same comparative statics in the baseline model with downside risk. These results allow us to compare the effects of uncertainty in both models. Similar to the baseline model, we find that higher levels of uncertainty are associated with longer durations on average, particularly in sticky product markets. This confirms that the impact of uncertainty discussed in Section 5 is not limited to shocks affecting only the first moment of the profit function. The presence

Table O.4: ***Parameter values for the simulations***

Value function	
Outside option V_0	10
Discount rate β	0.953
Meeting probability λ	0.286
Separation rate δ	0.05
Median duration - No stickiness	7
$\rightarrow \gamma$ (Fixed switching cost)	1.001
$\rightarrow \gamma$ (Variable switching cost)	1.001
Median duration - Medium stickiness	12
$\rightarrow \gamma$ (Fixed switching cost)	19.28
$\rightarrow \gamma$ (Variable switching cost)	1.242
Median duration - High stickiness	15
$\rightarrow \gamma$ (Fixed switching cost)	31
$\rightarrow \gamma$ (Variable switching cost)	1.345
AR(1) process	
Auto-correlation α	0.9
Drift ψ	0
μ	50
Volatility σ - No uncertainty	0
Volatility σ - Low volatility	10
Volatility σ - High volatility	150
Lower bar $\underline{I}$	100
Upper bar $\bar{I}$	1250
Initial Income State I_0	1250

of risk aversion in the model allows us to capture the real effects of uncertainty on firms' optimal switching decisions and relationship durations.

O.4 Further results on relationship stickiness

This section present additional results on relationship stickiness.

- Figure O.10 displays the estimates of relationship stickiness obtained from two sample periods: 1996-2006 and 2011-2017. The correlation coefficient between the two measures is 0.55.

Figure O.9: ***Impact of more uncertainty on expected durations under various degrees of stickiness***

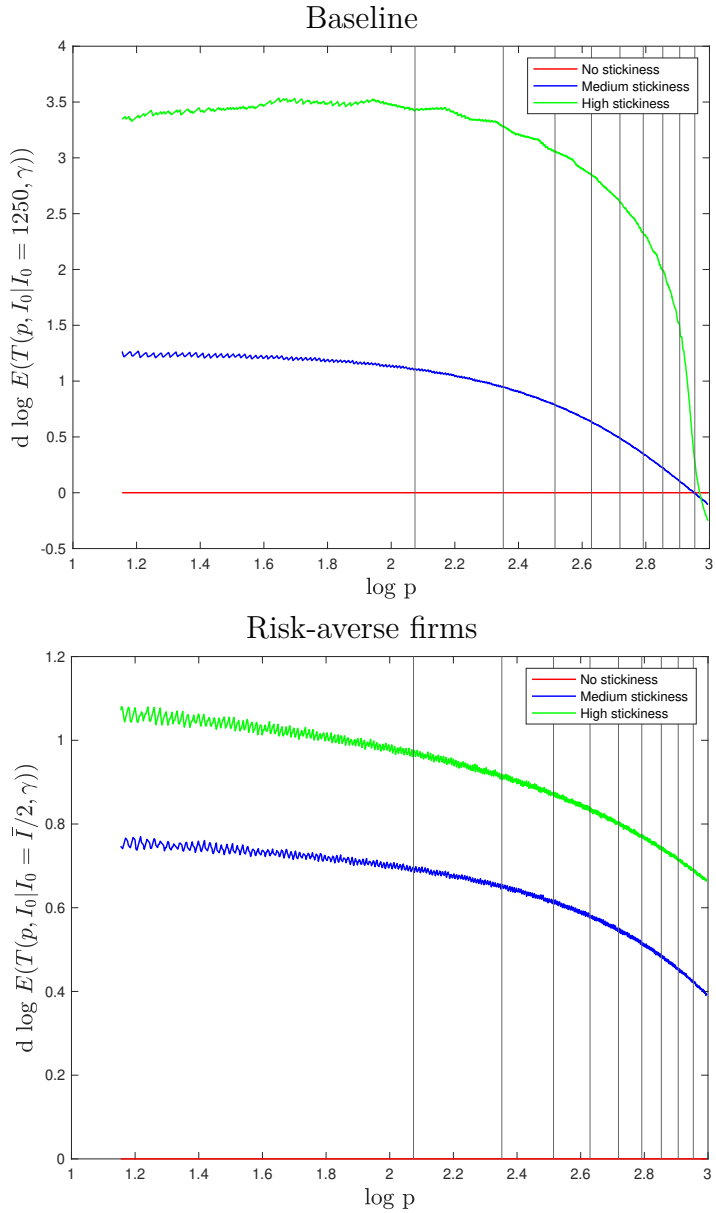

Notes: The figure depicts the percentage change in the expected duration of relationships, contrasting the states of no uncertainty and high uncertainty, across product markets with varying degrees of stickiness. The red line corresponds to a market with no stickiness, the blue line represents a moderate level of stickiness, and the green line represents a high degree of stickiness. In the top panel, which reflects the baseline calibration, the uncertainty shock affects an economy characterized by a high demand level, leading to the emergence of downside risk. Conversely, the bottom panel explores a scenario of a mean-preserving shock, assuming that firm managers exhibit risk aversion.

Figure O.10: ***Comparison of estimated stickiness indicators, across periods***

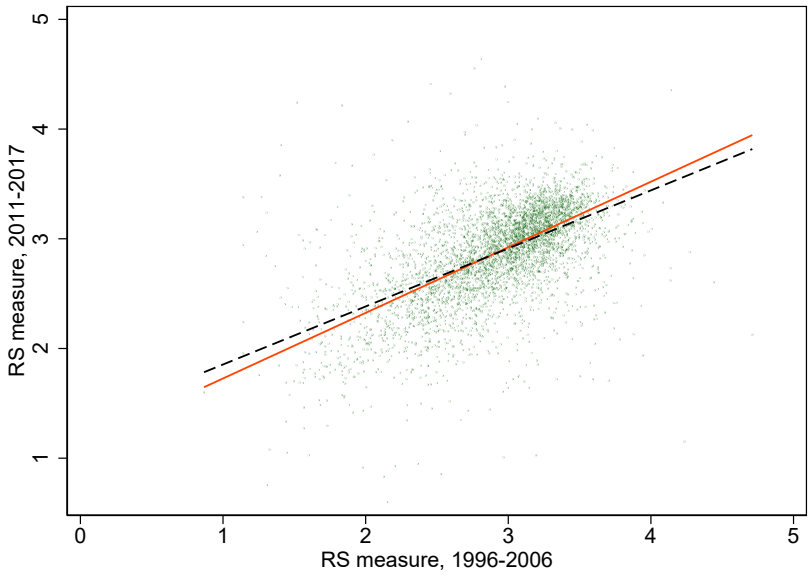

Notes: The figure represents a scatter plot of the baseline RS indicator recovered from the 1996-2006 period (x-axis) against the measure estimated from 2011-2017 (y-axis). The red and dotted lines correspond to the fitted lines, recovered from an unweighted linear fit (dotted line) or a linear regression in which products are weighted by the inverse of the estimated standard error recovered in the baseline case (red line).

- Figure O.11 depicts the dispersion in relationship stickiness across different Broad Economic Categories (BEC). The graph highlights that passenger cars exhibit the lowest level of stickiness, indicating a higher likelihood of switching suppliers compared to other BEC categories. On the other hand, parts and materials demonstrate the highest stickiness, suggesting a greater tendency for firms to maintain long-term relationships in this category. It is worth noting that there is significant variation in stickiness within each BEC category. This indicates that even within a specific industry or product category, there are firms that exhibit different levels of stickiness in their relationships with suppliers.

Figure O.11: ***Relationship stickiness across Broad Economic Categories***

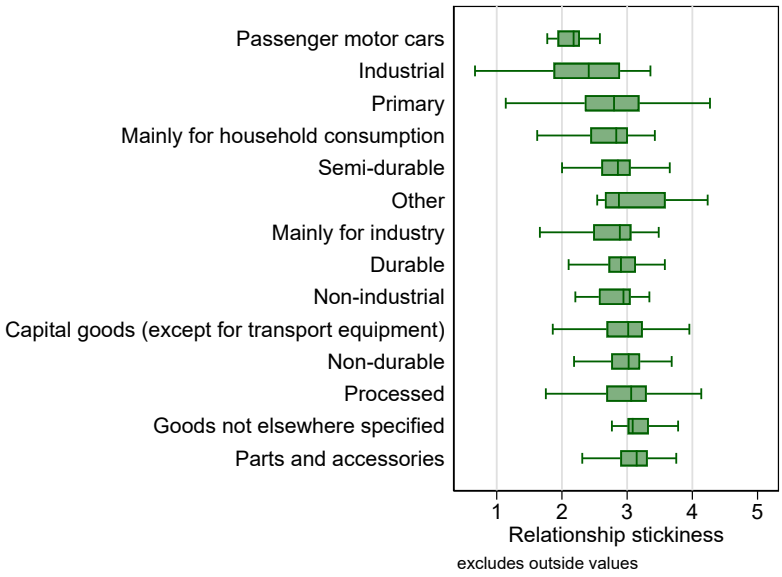

Notes: The figure represents a boxplot of relationship stickiness across Broad Economic Categories (BEC).

O.5 Alternative measures of relationship stickiness

Table O.5 presents a set of robustness exercises that aim to assess the stability of the estimated Relationship Stickiness (RS) indicators. These exercises involve alternative definitions of sales quantiles in equation (3). The table summarizes the results obtained from using quintiles instead of deciles, country-specific quantiles, or a subset of the top quantiles of the distribution of transaction sizes. When using country-specific quintiles, the support and shape of the distribution of productivities are not constrained to be homogeneous across destinations within a product. Focusing on the top percentiles restricts the analysis to the right tail of the distribution of firms' sizes, which aligns better with the model's Pareto assumption as discussed in Head et al. (2014). The table reports that, in all three cases, the distribution of RS indices obtained from these alternative definitions is highly correlated with the baseline estimate, with correlations exceeding 90%. Additionally, the inter-quartile ranges of the RS indices remain comparable to

those obtained in the baseline estimation. These findings suggest that the estimated RS indicators are robust and not heavily influenced by the specific choice of sales quantiles used in the analysis.

In an attempt to align the semi-elasticity of the duration with respect to the size of the quintile to unity, as suggested by the theoretical model, we conducted a robustness exercise by imposing a coefficient of one.[5] The results remained virtually unchanged. This suggests that this component of the estimated equation does not capture a substantial portion of the heterogeneity present in the data. We also conducted an additional robustness exercise by controlling for country fixed effects. Again, the results remained largely unchanged, with the estimated RS indicators highly correlated with the baseline estimates. Furthermore, we examined the importance of controlling for the size of the transaction by running a specification without size control. In this case, the estimates were also highly correlated with the baseline estimates, indicating that the size variable contributes little explanatory power to the relationship stickiness measures.

Another potential concern regarding the measurement of durations relates to the infrequent occurrence of transactions. It is possible that products that are purchased less frequently may artificially exhibit longer durations, thereby inflating the measured stickiness. To address this concern, we conducted a robustness check that takes an extreme approach. Instead of measuring durations based on the timespan between the first and last transactions, we measured "durations" using the number of transactions observed within a continuous relationship. Remarkably, the correlation between the stickiness indicator obtained from this alternative measure and the baseline estimate is 83%, indicating a substantial level of agreement. Furthermore, the inter-quartile range of the

[5]The structural model implies that the semi-elasticity of the duration with respect to the mass of firms in the corresponding quantile of the distribution should be equal to 1 (as shown in equation (3)). However, in our baseline estimation, we found a semi-elasticity of 0.15, significantly lower than unity.

stickiness distribution remains comparable between the two measures (see Table O.5). This suggests that the potential mis-measurement of durations resulting from the heterogeneous frequency of purchases is not a significant concern in our analysis.[6] To further reinforce this point, we conducted additional checks. First, we computed the duration of a relationship regardless of potential interruptions, which yielded a correlation of 0.81 with our baseline measure. Second, we calculated the duration of a relationship at the 4-digit level of the HS nomenclature and found a correlation of 0.71 with the baseline measure. These additional analyses provide further support for the robustness of our findings.

In addition to the various measures of stickiness, we also computed a simple measure of average durations. The correlation between this average duration measure and our baseline measure of stickiness is high, at 0.87. However, we made an interesting observation regarding the sources of dispersion in average durations. We found that approximately one fifth of the across-product dispersion in average durations can be explained by the dispersion of sales within a specific product category. This indicates that there is variation in average durations even within the same product category, likely reflecting differences in match qualities between firms. In contrast, the dispersion in sales explains only a small fraction (1/20th) of the empirical variance in our measure

[6]It is important to note that this robustness check primarily addresses concerns about systematic differences in the frequency of purchases across products. However, another potential concern arises from within-product heterogeneity across firms. While a significant portion of this variability is smoothed out when averaging durations within quantiles, there may still be concerns if the measurement error associated with infrequent transactions is correlated with the transaction size. This correlation could arise if there are fixed costs per shipment, as argued by studies such as Hornok and Koren (2015) and Blum et al. (2019). Despite these potential concerns, the stability of our results across various definitions of durations and quantiles suggests that they are not significant in practice.

of stickiness. This comparison highlights the effectiveness of our empirical strategy in controlling for within-product dispersion, which we argue captures the heterogeneity in match qualities. Considering these results, we believe that the model-driven measure of stickiness used in our paper offers several advantages over a model-free measure such as the average duration of trade relationships. Our approach explicitly accounts for the within-product dispersion and provides a more nuanced understanding of the underlying mechanisms driving stickiness in international trade relationships.

To assess the stability of our relationship stickiness (RS) estimates, we conducted several tests using different estimation samples. Our hypothesis is that our estimation strategy captures the ex-post impact of product-specific attributes, implying that the RS estimates should be consistent regardless of the time period or country sample used. First, we estimated country-specific distributions of RS indicators using the same empirical strategy but focusing on important destination countries for French exports, including Belgium, Germany, Italy, Spain, and the UK. The correlation between the baseline distribution obtained from the pooled sample and the country-specific estimates is high, averaging around 77%. This suggests that the estimated RS indicators exhibit substantial consistency across countries. We also examined the stability of our estimates over time by estimating relationship stickiness using the 2011-2017 period. Despite potential data limitations in the more recent years due to increased censoring, the correlation between the RS estimates from this period and the baseline distribution is significant at 0.64. It is worth noting that our baseline distribution of RS indices is based on the 1996-2006 period, which benefits from better-quality customs data. Furthermore, we extended our estimation strategy to a completely different dataset, specifically the panel of firm-to-firm trade flows from Colombian exports. Although we could not estimate RS for the exact same set of products due to data limitations, we estimated equation (3) for 377 HS4 products with more than 100 transactions, pooling all destination countries

Table O.5: ***Alternative estimates of stickiness***

	(1) corr. w/ baseline	(2) IQR	(3) # of products
Baseline	100%	0.59	5,077
Robustness to the definition of size quintiles:			
Quintiles	98.8%	0.59	5,054
HS6-iso2 specific quantiles	99.3%	0.60	5,054
Top quantiles	90.2%	0.63	4,985
Imposing alpha==1	99.0%	0.58	5,054
Without size control	99.9%	0.58	5,054
Robustness to the specification:			
Average duration	86.5%	11.9	5,077
Adding country fixed effects	99.9%	0.60	5,054
# of transactions rather than months	83.4%	0.58	5,054
Dealing with left censoring	96.0%	0.59	5,045
Duration (including interruption)	80.5%	0.53	5,054
Relationship at the 4-digit level	71.6%	0.60	5,076
Stability over space and over time:			
Focus on Belgium	79.2%	0.71	4,819
Focus on Germany	80.5%	0.69	4,862
Focus on Italy	75.6%	0.71	4,691
Focus on Spain	77.3%	0.72	4,748
Focus on UK	74.2%	0.73	4,572
Sample 2011-2017	64.2%	0.55	4,564
Using Colombian data (HS4 level)	45.4%	0.55	377

Notes: The table presents a comparison of the baseline set of relationship stickiness (RS) estimates with several robustness exercises discussed in the text. The table includes three columns: (1) correlation coefficient with the baseline, (2) inter-quartile range, and (3) number of estimated coefficients.

together. The correlation between the recovered estimates and the mean RS per HS4 product in our baseline estimates is significant and positive, at 0.45. This suggests a consistent pattern of relationship stickiness across different datasets, even when considering a different exporting country. Overall, our stability analysis reveals high levels of correlation among RS estimates obtained from various datasets, including those from a different country. These positive correlations provide empirical support for our interpretation of RS as capturing the consequences of structural factors that lead to significantly different mean durations in relationships across various product categories.

O.6 External validity checks

Relationship stickiness and intrafirm trade. In their study, Antràs and Chor (2013) propose a property-rights model where relationship-specific investments create a "locked-in" effect, that pushes firms to integrate their suppliers. Specifically, downstream firms have an incentive to integrate suppliers due to contractual frictions in procuring customized components that are later integrated into production. According to this framework, the prevalence of vertical integration is expected to be higher in product markets with stronger locked-in effects.

Table O.6 investigates whether the extent of intrafirm trade in US product-level trade data varies systematically across different relationship stickiness indicators. The analysis focuses on the correlation between the relationship-specific indicator and the share of intrafirm trade in US exports (columns (1)-(2)) and US imports (columns (3)-(4)). To account for other product-level characteristics that are known to be correlated with relationship stickiness (as discussed in Section 4.2), columns (2) and (4) include additional controls. The share of intrafirm trade is derived from data provided by the Bureau of Economic Analysis for the year 2002. Intrafirm trade is identified based on 6-digit NAICS categories, which are merged with the HS6 nomenclature (version 2002) using the correspondence established by Pierce and Schott (2012).

The results indicate a positive and statistically significant correlation between a product's level of relationship stickiness and its share of intrafirm trade. Moreover, relationship stickiness accounts for approximately 11% of the variation observed in the share of intrafirm trade across different product categories. These findings suggest that relationship stickiness plays a meaningful role in explaining the prevalence of intrafirm trade.

Relationship stickiness and comparative advantages: In their studies, Nunn (2007) and Levchenko (2007) argue that goods requiring high relationship-specific investments are more likely to thrive in countries with strong institutions, including effective contract enforcement, secure property rights, and shareholder protection. These institutions, in turn, shape the geography of trade, similar to other sources of comparative advantage. To test the validity of our relationship stickiness measure, we replicate an exercise similar to Nunn (2007) using more disaggregated data obtained from the UN-COMTRADE database at the 6-digit level of the Harmonized System (HS) nomenclature, which we merge with our own measure of relationship stickiness.

Table O.7 presents the results of these regressions. In each regression, we control for the measure of relationship specificity developed by Nunn (2007). In the first three columns, we follow Nunn's approach and explain the value of a country's exports at the product level using an interaction term between the quality of the country's institutions (measured by Kaufmann et al. (2010)) and the degree of relationship stickiness of the product. In columns (4) and (5), we depart from Nunn's specifications and consider measures of specialization that account for product-country pairs with zero trade flows, such as the Balassa index and a dummy variable identifying Balassa indices above 1.[7]

The findings confirm the results of Nunn (2007) that countries with strong contract enforcement tend to specialize in the production of more relationship-specific goods. In columns (3) and (5), we show that both Nunn's measure and our measure of product stickiness have explanatory power in this regression. However, when we use the Balassa index as a measure of comparative advantage, the interaction term with Nunn's measure becomes statistically insignificant, while our relationship stickiness indicator remains

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

positively associated with greater trade from countries with robust enforcement laws.

Overall, these results provide support for the validity of our relationship stickiness measure and its ability to capture the role of institutions in shaping trade patterns, consistent with the findings of Nunn (2007) and Levchenko (2007).

Relationship stickiness and the distance effect: To assess how relationship stickiness interacts with standard determinants of international trade and shapes the geography of trade, we employ the gravity equation framework. We interact the distance variable with our relationship stickiness measure, and the results are presented in Table O.8. Bilateral trade data at the HS6 level are obtained from the BACI database for the year 2005 (Gaulier and Zignago, 2010), while distance is measured as the weighted distance between countries' main cities based on Mayer and Zignago (2011). Additionally, we control for product upstreamness in value chains and its interaction with distance.

The consistent findings across specifications indicate that the effect of distance on trade is amplified in product markets characterized by higher relationship stickiness. This result holds regardless of the fixed effects structure, including in the most demanding specification in column (4). The elasticity of trade with respect to distance also appears to increase for more upstream goods, although the effect is sensitive to the choice of fixed effects. Interpreting the magnified impact of distance for high-relationship stickiness products within this reduced-form framework is challenging. One possible interpretation is that information frictions are more significant in these markets, leading to higher switching costs and concentration of trade in closer geographic locations (Rauch, 1999). Another explanation could be that stickier relationships are associated with higher monitoring costs, which tend to increase with distance (Head and Ries, 2008).

These sanity checks consistently support our main findings. Our relationship stickiness measure captures meaningful variability across disaggregated product markets, and its correlation with external indicators aligns with our interpretation. It provides further

Table O.6: ***Share of intrafirm trade and relationship stickiness***

	(1)	(2)	(3)	(4)
	Share of intrafirm			
	exports		*imports*	
RS (η)	0.214***	0.167**	0.158***	0.136***
	(0.029)	(0.045)	(0.022)	(0.033)
Nunn' measure		0.400***		0.197***
		(0.066)		(0.048)
Upstreamness		0.059***		0.013
		(0.016)		(0.012)
Elasticity (σ)		0.001		-0.005
		(0.006)		(0.004)
Observations	435	375	435	375
R-squared	0.105	0.154	0.106	0.111

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

evidence that relationship stickiness plays a role in shaping trade patterns.

O.7 Robustness on relationship stickiness and uncertainty

Table O.9 and Table O.10 present an extensive robustness analysis of the results discussed in Section 5.2. Both tables follow a similar structure. Table O.9 focuses on testing the robustness of the results concerning the effect of uncertainty on entry behavior, while Table O.10 examines the robustness of exit patterns.

Non-durable goods. In columns (1)-(2), we examine the robustness of the baseline regression by excluding durable products from the sample. This exclusion is motivated by the findings of Novy and Taylor (2019), who highlight that durable goods exhibit the highest sensitivity to uncertainty shocks. In their model, firms can postpone orders of durable products during periods of high uncertainty by relying on their existing inventories. The effect of such behavior on our results is not entirely clear, as it primarily affects adjustments at the intensive margin. However, it is possible that firms may still delay the search for new trading partners even if they have inventories to rely on.

Table O.7: ***Institutional comparative advantage***

	(1)	(2)	(3)	(4)	(5)
	log(exports)			Balassa Index	$\mathbf{1}_{Balassa>1}$
Rule of law					
$\times$ *RS*	0.334***		0.392***	0.198***	0.020***
	(0.053)		(0.048)	(0.070)	(0.006)
$\times$ Nunn specif.		0.823***	1.105***	0.382	0.045**
		(0.099)	(0.145)	(0.298)	(0.021)
$\times$ Upstreamness			0.068	0.037	0.007
			(0.044)	(0.070)	(0.005)
Fixed effects			*country*(122) *and sector*(4,326)		
Observations	291,157	291,157	287,929	518,744	518,744
R-squared	0.605	0.606	0.610	0.012	0.141

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

Table O.8: ***Gravity for trade in goods with sticky relationship***

	(1)	(2)	(3)	(4)
Distance (log)	-0.552***	-0.189***	-0.338***	-0.756***
	(0.015)	(0.029)	(0.031)	(0.034)
- $\times$ RS		-0.130***	-0.122***	-0.078***
		(0.011)	(0.011)	(0.010)
- $\times$ Upstreamness		0.010**	0.020***	-0.014**
		(0.005)	(0.005)	(0.007)
RS	-0.259***	0.790***		
	(0.012)	(0.089)		
Upstreamness	0.060***	-0.022		
	(0.006)	(0.042)		
Fixed effects				
Exporter	✓	✓	✓	
Importer	✓	✓	✓	
Product			✓	
Exporter$\times$Product				✓
Importer$\times$Product				✓
Observations		5,568,400		5,342,338
R-squared	0.165	0.166	0.284	0.578

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

Figure O.12: ***Uncertainty episodes and GDP growth***

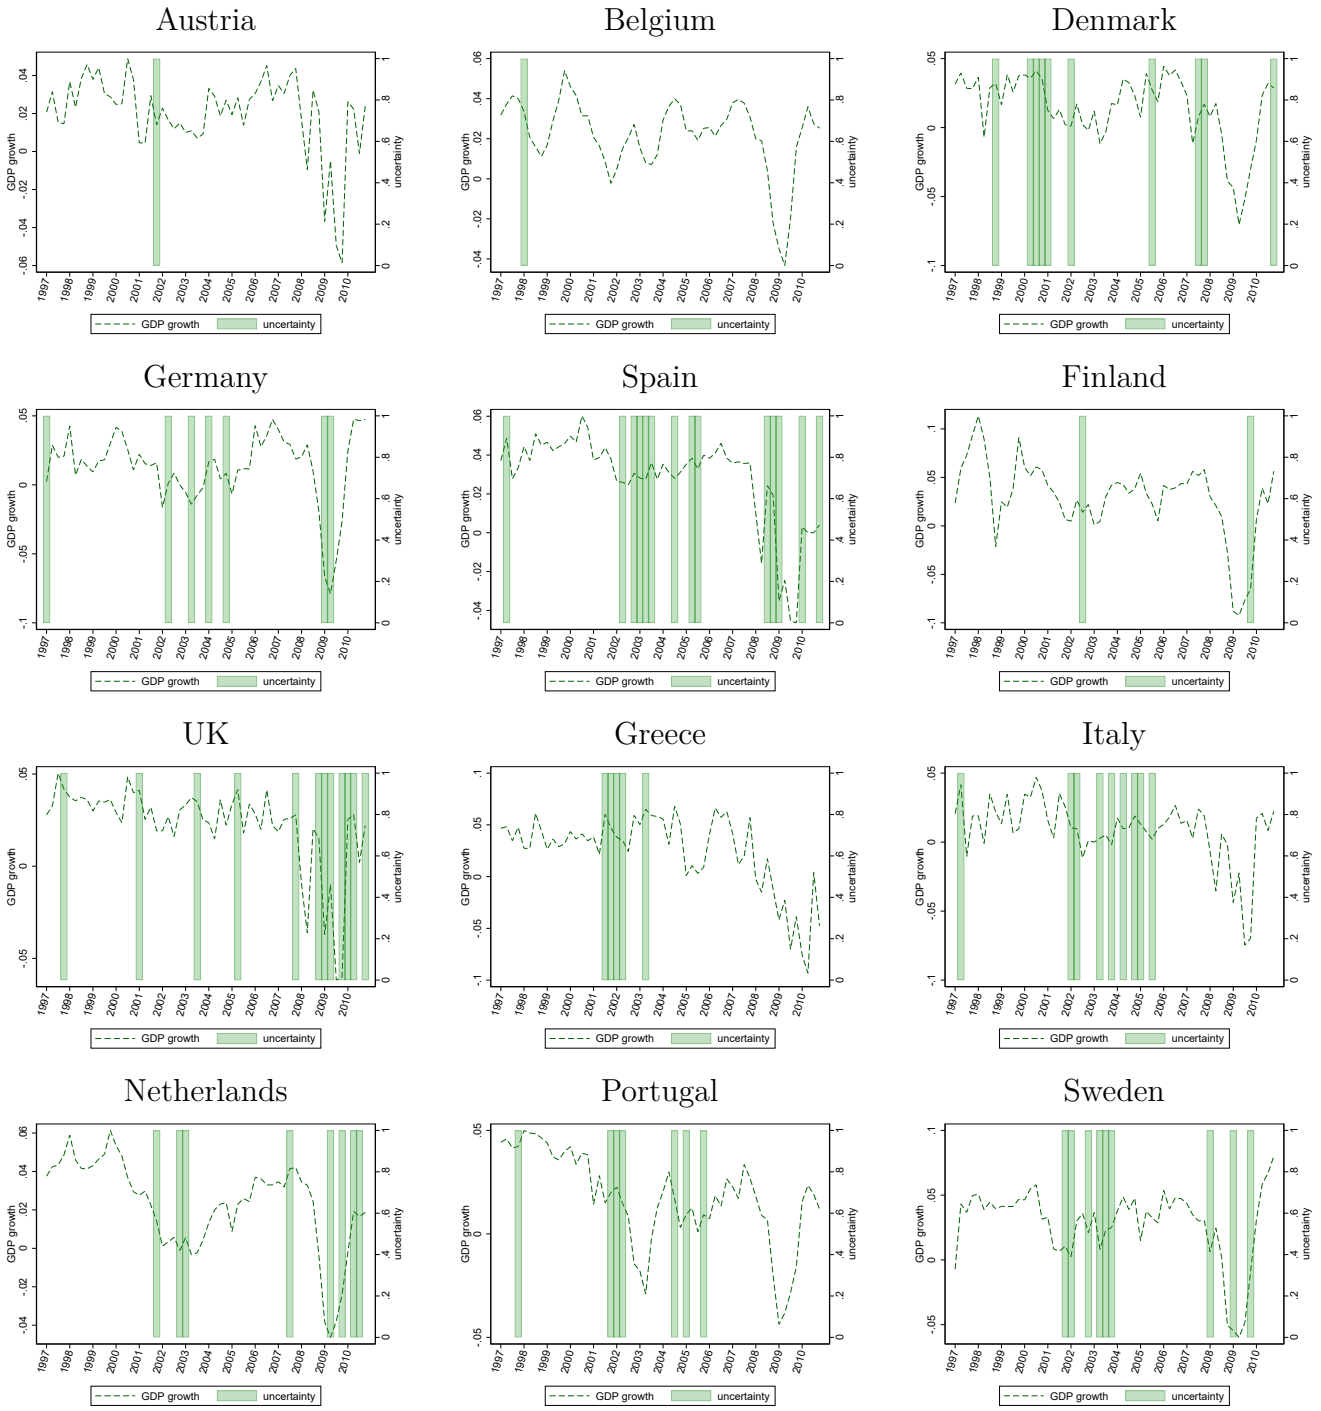

Despite these potential complexities, we find that our results remain robust when excluding durable products from the analysis. However, it is worth noting that focusing on non-durable products shifts the elasticities of separations in response to uncertainty episodes towards more positive values. This suggests that uncertainty has a stronger impact on separations in the absence of durable goods.

Excluding intrafirm trade. In columns (3)-(4), we address concerns about the potential impact of intrafirm trade on our estimates. Due to data limitations, we are unable to distinguish between arm's length transactions and intrafirm transactions in our dataset. However, it is possible that the response of both types of trade flows to uncertainty episodes could drive the observed heterogeneity that we attribute to relationship stickiness in Section 5.2.

To control for this potential confounding effect, we construct a sub-sample using external data from the INSEE-LiFi survey. This survey allows us to identify French exporters that are part of multinational firms. We then exclude from our analysis all their exports to countries where they have multinational linkages, either through having an affiliate or because their headquarters are located in those countries. By doing so, we remove all intrafirm trade flows from the estimation sample. However, it is important to note that this procedure may also remove transactions that are not intrafirm, as exporters with affiliates in a country may still export to non-affiliated partners.

Interestingly, the results in columns (3)-(4) are qualitatively and quantitatively unchanged in terms of the response of entries to uncertainty episodes. This suggests that the observed relationship between uncertainty and entry is not solely driven by intrafirm trade flows. However, it is worth noting that the elasticities at the exit margin are slightly shifted towards more positive values when intrafirm transactions are neglected. This is likely due to the over-representation of less sticky products in the estimation sample, as stickiness is correlated with the prevalence of intrafirm trade (as shown in

Table O.6).

Relationship stickiness: level vs. rank. In columns (5)-(6), we introduce a different approach to address concerns about the potential contamination of the level of relationship stickiness (RS) by uncertainty. Instead of using the level of RS, we focus on the ranking of products within the distribution of stickiness indicators displayed in Figure 2. The rationale behind this approach is that while the level of RS may be influenced by uncertainty, the ranking of products should be immune to this source of endogeneity, as long as uncertainty is common across products.

The results obtained from this alternative measure of stickiness are robust and actually indicate a stronger heterogeneity in adjustments to uncertainty compared to the baseline specifications.

Relationship stickiness: 2011-2017. In columns (7)-(8), we use the level of stickiness estimated over the period 2011-2017. This measure of relationship stickiness is based on a different sample compared to the one used to assess the role of uncertainty on the creation and destruction of trade relationships. The results obtained from this alternative measure of stickiness are consistent with the baseline specifications. In fact, these results indicate an even stronger heterogeneity in adjustments to uncertainty compared to the baseline specifications.

In summary, the results presented in this section confirm and reinforce the findings discussed in Section 5.2. These results provide further evidence regarding the sensitivity of trade to uncertainty episodes, the role of adjustments at the extensive margin, and the magnified impact of uncertainty in more sticky product markets.

Table O.9: ***Uncertainty and the creation of new trade relationships: Robustness***

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dep. var:	*# new trade relationships*							
Uncertainty episode	0.412***		0.153***		0.002		0.056**	
	(0.023)		(0.021)		(0.006)		(0.023)	
- × RS index	-0.163***	-0.144***	-0.066***	-0.048***				
	(0.008)	(0.008)	(0.007)	(0.008)				
- × RS Percentile					-0.001***	-0.001***		
					(0.000)	(0.000)		
- × RS11-17							-0.034***	-0.034***
							(0.008)	(0.009)
Specification	wo durables		wo MNEs		RS Pctiles		2011-2017	
Obs.	2,083,539	2,332,462	2,519,005	2,814,576	2,326,720	2,587,862	2,544,177	2,826,460
Fixed Effects								
Product × quarter		✓		✓		✓		✓
Product × period	✓		✓		✓		✓	
Country	✓		✓		✓		✓	
Country × period		✓		✓		✓		✓

Notes: The table presents Poisson estimations with high-dimensional fixed effects to examine the impact of uncertainty on new trade relationships. The main explanatory variable is an indicator for uncertainty episodes, which is equal to 1 during periods when uncertainty in the destination country exceeds the average uncertainty plus one standard deviation. The measure of relationship stickiness, denoted as RS, is included as a key independent variable. The first two columns (1)-(2) exclude observations related to bilateral exports of firms that have an affiliate or their headquarters in the destination country. Columns (3)-(4) focus on a sub-sample that excludes durable goods. In columns (5)-(6), the measure of relationship stickiness is replaced by the percentile ranking of products within the distribution of stickiness indicators (as shown in Figure 2). Columns (7)-(8) employ relationship stickiness estimated on a different sample period, specifically from 2011 to 2017. To facilitate comparison, the table also reports the estimated impact of an uncertainty shock for products at the 25th percentile (P25) and 75th percentile (P75) of the relationship stickiness distribution. All estimations control for GDP growth and its interaction with the relationship stickiness measure as additional covariates. Robust standard errors are in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

Table O.10: ***Uncertainty and the disruption of new trade relationships: Robustness***

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dep. var:	*# disrupted trade relationships*							
Uncertainty episode	0.300***		0.183***		0.160***		0.082***	
	(0.023)		(0.023)		(0.025)		(0.006)	
- × RS index	-0.094***		-0.069***		-0.051***		-0.041***	
	(0.008)		(0.008)		(0.008)		(0.009)	
- × RS 11-17					-0.041***	-0.030***		
					(0.009)	(0.009)		
- × RS Percentile							-0.001***	-0.001***
							(0.000)	(0.000)
Specification	wo durables		wo MNEs		2011-2017		RS Pctiles	
Obs.	1,494,042	1,553,904	1,812,912	1,882,808	1,685,927	1,747,973	1,847,772	1,915,448
Fixed Effects								
Product × quarter		✓		✓		✓		✓
Product × period	✓		✓		✓		✓	
Country	✓		✓		✓		✓	
Country × period		✓		✓		✓		✓

Notes: The table presents Poisson estimations with high-dimensional fixed effects to examine the impact of uncertainty on disrupted trade relationships. The main explanatory variable is an indicator for uncertainty episodes, which is equal to 1 during periods when uncertainty in the destination country exceeds the average uncertainty plus one standard deviation. The measure of relationship stickiness, denoted as *RS*, is included as a key independent variable. The first two columns (1)-(2) exclude observations related to bilateral exports of firms that have an affiliate or their headquarters in the destination country. Columns (3)-(4) focus on a sub-sample that excludes durable goods. In columns (5)-(6), the measure of relationship stickiness is replaced by the percentile ranking of products within the distribution of stickiness indicators (as shown in Figure 2). Columns (7)-(8) employ relationship stickiness estimated on a different sample period, specifically from 2011 to 2017. To facilitate comparison, the table also reports the estimated impact of an uncertainty shock for products at the 25th percentile (P25) and 75th percentile (P75) of the relationship stickiness distribution. All estimations control for GDP growth and its interaction with the relationship stickiness measure as additional covariates. Robust standard errors are in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

References

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

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.

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

Bernard, Andrew, Andreas Moxnes, and Karen Helene Ulltveit-Moe, "Two-sided Heterogeneity and Trade," *Review of Economics and Statistics*, 2018, *100* (3), 424–439.

Blum, Bernardo S., Sebastian Claro, Kunal Dasgupta, and Ignatius J. Horstmann, "Inventory Management, Product Quality, and Cross-Country Income Differences," *American Economic Journal: Macroeconomics*, January 2019, *11* (1), 338–388.

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.

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.

_ , Thierry Mayer, and Mathias Thoenig, "Welfare and Trade without Pareto," *American Economic Review*, 2014, *104* (5), 310–316.

Hornok, Cecília and Miklós Koren, "Administrative barriers to trade," *Journal of International Economics*, 2015, *96* (S1), 110–122.

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.

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.

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

Table O.11: ***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***
	(0.001)	(0.002)	(0.001)	(0.001)
Uncertainty shock	-0.11***	-0.04***	-0.07***	0.002
	(0.003)	(0.002)	(0.001)	(0.001)
- × RS index	-0.01***	-0.02***	0.01***	-0.004***
	(0.001)	(0.001)	(0.001)	(0.0004)
GDP shock	-0.15***	-0.16***	-0.01	0.02
	(0.006)	(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	3,568,288	3,568,288	3,568,288	3,568,288
Period		1996-2010		

Notes: OLS results are reported with bootstrapped standard errors shown in parentheses. The dependent variable, growth, represents the year-on-year growth rate of product-level French exports to the destination under consideration. The overall growth is decomposed into three different growth margins: Start refers to the number of new seller-buyer relationships, End represents the number of disrupted relationships, Intensive reflects the evolution of seller-buyer sales within existing trade relationships. The uncertainty shocks variable is a dummy equal to 1 during periods when uncertainty in the destination country exceeds the average uncertainty plus one standard deviation. The GDP shock variable is a dummy equal to 1 during periods when GDP growth is below the average GDP growth minus one standard deviation. *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%.

Table O.12: ***Uncertainty and trade growth: margin decomposition with an alternative measure of uncertainty***

Dep. var:	(1) Growth	(2) =Start	(3) + End	(4) + Intensive
RS index	-0.02	-0.32***	0.29***	0.01
	(0.017)	(0.012)	(0.008)	(0.006)
Vol. of returns	-0.18***	-0.09***	-0.09***	-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)	(0.003)	(0.002)
- × RS index	0.01***	-0.01***	0.01***	0.01***
	(0.002)	(0.001)	(0.001)	(0.001)
Observations	3,568,288	3,568,288	3,568,288	3,568,288

Notes: OLS results are reported with bootstrapped standard errors shown in parentheses. The dependent variable, growth, represents the year-on-year growth rate of product-level French exports to the destination under consideration. The overall growth is decomposed into three different growth margins: Start refers to the number of new seller-buyer relationships, End represents the number of disrupted relationships, Intensive reflects the evolution of seller-buyer sales within existing trade relationships. 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%.