Learning and the Value of Relationships in International Trade*

Ryan Monarch[†] Tim Schmidt-Eisenlohr[‡]

March 30, 2015

Abstract

This paper explores the characteristics of importer-exporter relationships and their role in international trade using detailed U.S. import data. We generate new stylized facts that highlight the relevance of firm-to-firm relationships for international trade. While most relationships are short-term, the value of trade is spread evenly across new, medium-term, and long-term relationships. Relationship length is closely tied to product characteristics, source country institutions, and firm size. In addition, established relationships are central to firms expanding their product scope or adjusting their supply network. We estimate a model of learning about the reliability of trading partners in international trade that is consistent with these findings. Counterfactuals based on a calibrated version of the model suggest that relationships have a large quantitative role in explaining international trade flows and that disruptions of relationships have long-lasting adverse effects.

Keywords: international trade, firm relationships, learning, institutions

JEL-Codes: F11, F14, L14, D22

*We thank seminar participants at the Federal Reserve Board and the World Bank for helpful comments. This work was undertaken while under Special Sworn Status at the U.S. Census Bureau. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau, the Federal Reserve Board, or anyone affiliated with the Federal Reserve System. All results have been reviewed to ensure no confidential information is disclosed. Clint Carter and William Wisniewski were extremely helpful with data requests and disclosure processes. All errors are ours.

†Federal Reserve Board. Contact: ryan.p.monarch@frb.gov

‡University of Illinois at Urbana-Champaign. Contact: t.schmidteisenlohr@gmail.com

1 Introduction

A successful relationship between an importing firm and an exporting firm is at the core of every transaction in international trade. Earlier work in international trade has argued that relationships between importers and exporters are key to understanding international trade flows.[1] More recently, researchers have used importer-exporter transaction data that allows for a more detailed analysis of trade relationships over time.[2] Such work has taken important steps towards testing the underlying mechanisms driving the formation and durability of these relationships, including the role of learning the long-term compatibility of one's trading partner. However, the relevance of the distinction between short and long-term relationships, and the implications of successful, sustained relationships for the evolution of trade volumes are still not well understood. In our view, there are two key questions. First, how important is learning relative to other explanations in understanding the dynamics of trade relationships and traded quantities within relationships over time? Second, what is the quantitative importance of learning in international trade and how does this depend on country-specific institutions or other factors?

To answer these questions, we exploit detailed data provided by the U.S. Census, analyze a model of learning in importing, and calculate model predictions in order to evaluate the quantitative importance of learning. Our empirical results demonstrate that successful relationships between firms are relevant for international trade, both in the aggregate and for individual firm decisions. We use a model based on previous work by Araujo et al. (2012) to clarify the mechanisms at work, and calibrate it to assess the quantitative importance of learning about trading partners. Our counterfactual experiments confirm the importance of sustained relationships to trade: a significant share of U.S. imports can be explained by the quality of relationships U.S. firms enjoy with companies in other countries.

[1] Greif (1993), Rauch (2001) and Rauch and Watson (2004), for example, provided evidence for the idea that trade networks are central to solve enforcement problems across borders.

[2] See, for example, Eaton et al. (2014), Macchiavello and Morjaria (forthcoming) and Antràs and Foley (forthcoming). The former study trade between Colombia and the U.S. in a search model. The latter two papers analyze learning within relationships in two particular cases; Kenyan rose exports and a U.S. food exporter, respectively.

We first generate a set of new stylized facts about exporter-importer relationships[3]. Although 36% of total (arm's length) trade takes place in relationships that have lasted for at least three years, the overall age composition of trade relationships tends to skew towards new relationships. Drilling down further into the data, we find that the greatest fraction of non-differentiated product imports is traded in longer-term relationships (three or more years), while differentiated products are primarily traded in shorter-term relationships. In addition, the length of relationships is also systematically related to source country institutions. Generally, the stronger the rule of law in the supplier's country, the longer the average importer-exporter relationship. Finally, the duration of trade relationships is also influenced by firm size- larger firms tend to have longer supplier relationships.

Relationships also matter for individual importer sourcing decisions. Importers tend to buy from many suppliers at the same time. This pool of exporting partners affects the sourcing of new products by an importer- 43% of such new product purchases also come from export partners used to buy other HS10 products in the past. Having relationships with many export partners also affects future choices of whom to buy from. Although relationships are likely to continue from year-to-year (Monarch (2014)), when importers change their partner, 52% of all supplier switching decisions are to *familiar* partners, meaning to partners that were used in small amounts for the same HS10 product previously, or that were used for other HS10 products.

These stylized facts lead us to put forth a model that incorporates dynamic learning about suppliers, fixed costs of searching for a new supplier, and price differences across suppliers. The framework borrows from the model of exporter learning by Araujo et al. (2012), but adds a number of additional features to make it compatible with the empirical findings discussed above. Our setup leads to predictions about how importer decisions to stay or switch, as well as the role of source country institutions on the longevity of relationships. Switching arises endogenously in the model, due to learning occurring over multiple periods. This implies the lowest cost producer is not always the main supplier

[3] All statements about total trade and total number of relationships in the following refer to unrelated-party trade or arm's length trade, though in principle, related party trade is an additional dimension available for study.

initially, even with perfect information about the price of the traded product.

We calibrate the model for U.S. importer supplier decisions made over 10 years from 20 countries. For this, we match the age distribution of relationships in the data to that generated by the model. The trade shares that the model predicts, moments that are not explicitly targeted, are very close to the data. This indicates that the path of traded quantities over time within relationships is captured surprisingly well by our parsimonious model. We then go on to calculate counterfactuals using our model. We find that if all existing relationships in steady state were wiped out, there would be a dramatic decline in international trade flows. Recovery back to steady state would be a slow process, taking several years. Another experiment studies a surge in the creation of new trade relationships. While new relationships have a sizable impact upon their creation, because of learning it takes some years for their full effect on trade to unfold.

Previous work on the topic of dynamic buyer-supplier relationship formation in international trade centers on the study of networks: Rauch (2001) surveys the potential for transnational cultural networks to help smooth international trade and reduce barriers to entry, while Rauch and Watson (2004) present a general equilibrium model through which economic agents can use their supply of networks to either produce/export more efficiently or to become an intermediary. The network explanation for trade flows is also at the heart of the work of Chaney (2014), who finds that exporting firms rely on networks to search for additional trade opportunities.

Our paper is also related to studies about the role of new entry versus experience in international trade. Egan and Mody (1992) also study the role of linkages between buyers and sellers, and the role of long-term relationships. Macchiavello and Morjaria (forthcoming) study Kenyan rose exporters and find, as in Araujo et al. (2012), that the value of the relationship increases in age. They also show that in very long-run relationships, buyers already learned the type of the seller and therefore no costly signaling is necessary in times of crisis. Blum et al. (2013) study exporting spells of firms in Chilean data. In their data, large firms that always export but enter and exit specific markets and small exporters that sometimes only sell domestically. Impullitti et al.

(2013) develop a model of entry and exit, where firms pay sunk costs to start exporting and face persistent productivity shocks. Besedeš (2008) tests a number of predictions for relationships in international trade, and finds that reliable suppliers lead to longer relationships and larger export orders, while only a small fraction of relationships end as a result of switching behavior.

There is additionally a burgeoning literature that uses "two-sided" international trade data to study the nature of importer-exporter matching. Eaton et al. (2014) study the relationship between Colombian exporters and the number of U.S. importers they partner with over time and calibrate a search and matching model to match exporter decisions, including sales, number of clients, and transition probabilities. Kamal and Krizan (2012) use U.S. Census trade transaction data to document trends in the formation of importer-exporter relationships. Kamal and Sundaram (2013) use the same U.S. import data to determine how likely textile producers in Bangladeshi cities are to follow other exporters in their same city to export to a particular partner, while Monarch (2014) finds that U.S. importers are very likely to remain with their Chinese exporting partners. The two-sided trade data is also used to study the effects of heterogeneity on trade: Bernard et al. (2014) develop a model of relationship-specific fixed costs to exporting using Norwegian buyer-supplier trade data. Carballo et al. (2013) look at importer-exporter relationships in several Latin American countries and develop a model to analyze the role of competition.

The rest of the paper is organized as follows. In Section 2, we describe the main features of the importer-exporter database we use. Section 3 presents empirical findings about U.S. importer relationships with foreign partner firms. Section 4 describes the model we use that is inspired by the empirical work discussed above, and presents separate predictions that can be tested. Section 5 describes the model calibration and counterfactuals. Section 6 concludes.

2 Data

The data come from the Longitudinal Foreign Trade and Transaction Database (LFTTD), collected by U.S. Customs and Border Protection and maintained by the U.S. Census Bureau. Every transaction of a U.S. company importing or exporting a product requires the filing of *Form 7501* with U.S. Customs and Border Protection, and the LFTTD contains the information from each of these forms.[4] There are typically close to 50 million transactions per year. In this paper, we utilize the import data, which includes quantity and value exchanged for each transaction, HS 10 product classification, date of import and export, port information, country of origin, and a code identifying the foreign exporting partner. Known as the *manufacturing ID*, or *MID*, the foreign partner identifier contains limited information on the name, address, and city of the foreign supplier.[5] Monarch (2014) and Kamal et al. (2015) found substantial support for the use of the MID as a reliable, unique identifier, both over time and in cross-section. Bernard et al. (2010), Kamal and Krizan (2012), Pierce and Schott (2012), Kamal and Sundaram (2013), Dragusanu (2014), and Eaton et al. (2014) have all used this variable in the context of studying U.S. firm relationships in international trade.

We also follow Bernard et al. (2009) methods for cleaning the LFTTD. Specifically, we drop all transactions with imputed quantities or values (which are typically very low-value transactions) or converted quantities or values. For the statistics below, we also eliminate related-party transactions, as U.S. plants who are importing from foreign affiliates will likely have very different relationship dynamics than those involved in arm's-length transactions.

Finally, some definitions: an **importer** is a U.S. importing firm, while an **exporter** is a non-U.S. firm identified by the MID exporting to the U.S. A **relationship** is an observation of an importer-exporter-industry combination, where industry is measured at the HS2 level.[6] We additionally distinguish between new, medium-term, and long-

[4] Approximately 80-85% of these customs forms are filled out electronically (Krizan (2012)).

[5] Specifically, the MID contains the first three letters of the producer's city, six characters taken from the producer's name, up to four numeric characters taken from its address, and the ISO2 code for the country of origin.

[6] Thus within this framework, it is possible for an importer and exporter to have multiple relationships

term relationships, where a new relationship is one that is not found in any previous year, a medium-term relationship is one lasting 1-2 (consecutive) years and a long-term relationship is one lasting 3 or more years.[7]

3 Empirical Findings

In this section, we present suggestive evidence for the existence of a new channel of learning in international trade. In particular, firms participating in international activity not only actively update their information about source or destination countries (Albornoz et al. (2012)) or undertake productivity-enhancing improvement to underlying practices (De Loecker (2013)), but also are constantly updating information about trading partners and determining whether they are in a healthy relationship. We demonstrate the existence of such learning in two ways: first, with a "bird's-eye" view of relationships in international trade. We find that the largest number of relationships are new, while, in value, new, medium-term and long-term relationships contribute about equally to international trade. Second, we use a "ground-level" approach to studying importer behavior. Here, the evidence points to the importance of experimenting with new suppliers, updating beliefs, and making future supply decisions based on information gleaned through these interactions. The set of stylized facts we establish guide our model of relationship-level learning that we use to estimate the value of relationships in international trade.

3.1 Relationship Length

Table 1 presents a breakdown of U.S. imports in 2007 based on the length of relationships. There are two main findings illustrated by the table. First, although the largest fraction of trade takes place among medium-term relationships, trade is fairly evenly distributed between new, medium and long-relationships. Newly formed importer-exporter relation-

with each other.

[7]The distinction between consecutive and non-consecutive years of a relationship makes very little difference to any of the findings below, as relationships disappearing and reappearing is not common in the data

ships account for one-quarter of U.S. imports. Second, new relationships account for the majority of total relationships in 2007, with only about 10% having existed for three or more years. Taken together, the numbers from this table are indicative of a correlation between relationship longevity and the value of trade: although it is rare for long-term relationships to be established, if a relationship persists, it is likely to consist of elevated trade flows.

Table 2 presents the same relationship classification along a number of different cuts of the data. The first looks at how product differentiation affects relationship length. For this we calculate the fraction of products within HS2 codes that are classified by Rauch (1999) as being differentiated. We then divide HS2 products into non-differentiated (below 50 percent of HS10 codes differentiated) and differentiated (the remaining industries).[8] The highest share of trade in non-differentiated products is clustered in long-term relationships, while trade in differentiated products tend to skew more towards shorter relationships. Longer-term relationships are also a higher share of the total number of relationships in non-differentiated products. To some extent, this is quite surprising. In particular, the fact that many differentiated goods require substantial relationship-specific investments may have led one to expect more stickiness in relationships for these products. One hypothesis is that since non-differentiated products have spot prices, there is no need to conduct intensive search for lower-price exporters, and firms simply continue buying from their current partner. In the market for a differentiated product, search is meaningful and more commonly leads to changing partners. Alternatively, while relationship-specific investments may be important for differentiated goods, failure to provide products of sufficient quality to the buyer may also be more prevalent in these industries, increasing separation rates.

Another dimension in which the structure of relationships differs is across different size categories U.S. importing firms. We split U.S. importers into three size bins with equal numbers of firms in each, based on the total value of imports for that firm. Panel B of Table 2 shows that the largest firms form longer-lasting relationships, while small importers are predominantly engaging in single-year transactions. Additionally, most of

[8] The exact split does not matter for any of the ordinal rankings of categories described below.

the imports by large firms are in longer relationships, the opposite of smaller firms. Large firms are well-known to be more active in international trade and to benefit the most from trade openness. Our data shows that they also tend to have the longest relationships. This represents a new dimension in which large firms are more successful in international trade: the ability to match effectively with an exporter over time is correlated with a firm's overall size as an importer.

Finally, we illustrate that the structure of relationships differs across source countries exporting to the United States. Information about relationship duration for U.S. imports from selected major trading partners is found in Panel B of Table 2. One takeaway from this decomposition is that countries that score generally lower on "Rule of Law" indices provided by the World Bank and others tend to have shorter-lasting relationships with U.S. firms. For example, one-third of imports from China, and two-thirds of relationships between Chinese and U.S. firms, are among newly formed importer-exporter pairings. This contrasts with U.S. imports from Germany or Japan, which are on average, longer lasting. The scatterplot in Figure 1 also demonstrates this finding: the higher the rule of law, the smaller the fraction of new relationships, and the larger the fraction of long-term relationships. Panel B also shows that among countries with lower measures of the rule of law, distance seems to have an impact on relationship length. For example, Mexico and Venezuela tend to have longer-lasting relationships then their counterparts of similar institutional quality that are farther away. The trend is not apparent among more developed economies.

Continuing in this vein, we test whether source countries with better institutions tend to have a higher share of maintained relationships over time. Our estimating equation takes the following form:

$$ShareStay_c = \alpha + \beta_1 \lambda_c + \gamma X_c + \nu_c \tag{1}$$

Are U.S. import relationships are more likely to persist in countries with better legal institutions (λ_c), due to better enforcement of contracts? The variable *ShareStay* is the fraction of importer-exporter relationships that are maintained between the U.S. and from country c averaged over the years 2002-2008, while X_c is a vector of country-specific controls, including log per capita GDP in country c and private credit coverage. To measure the quality of institutions λ, we use a collection of institutional quality variables taken from the World Bank World Development Indicators. First is the Strength of Legal Rights index, which measures the degree to which collateral and bankruptcy laws protect the rights of borrowers and lenders in a country. Since this variable is not directly a measure of contracting rights,we also include both the number of procedures required to enforce a contract, and the ordinal ranking of countries by such a procedure. Neither of these two institutional quality variables has significant changes over our sample period from 2002-2008, so we simply use 2004 values for each, as well as for per-capita GDP, while averaging the ratio of staying partners for each country over time.[9] As can be seen in Table 3, higher legal rights and fewer procedures in exporting are both indicative of a larger share of firms remaining with their partner over time.

3.2 Importer Sourcing Behavior

Having described the role of relationship longevity in international trade, we next illustrate a number of findings concerning how U.S. importers source. While the results above describe aggregate splits of U.S. imports, this section will present other suggestive evidence for importer learning from an importer-level perspective. For this reason, we shift the focus now from HS2 industries into more disaggregated HS10 products, and now define a **product relationship** as a transaction (or set of transactions) between an importer and an exporter for a particular HS10 product.

For a particular HS10 product, many importers use more than one exporter. The average number of suppliers for a firm-product purchase is 3 (with a median of 1), and 36.5% of firm-product combinations involve the use of more than one supplier. Thus learning about the quality of an importer-exporter match is likely to occur for many different partners simultaneously.

[9]The results are the same if we use yearly measures of *ShareStay* for each country c with individual year observations of $PCGDP$ and λ.

However, learning about an exporting partner can obviously happen across products as well, a dimension not captured in the relationship breakdowns above.[10] Indeed, U.S. firms rely on many different sources for their imports: at this level, the average number of foreign partners is 23.3, and the median is 4. 72.4% of firms have more than one partner. There is also additional variation in the number of countries a firm is buying from: the median number of countries is 3, and the average is 11.4.

We next show that long-lasting relationships increase not only the value of trade, but also an importing firm's product scope. Overall, 72.2% of importing firms have at least one "new product" each year, where a new product is one that is purchased by an importer in a year that was not purchased the previous year. Having a long-lasting relationship is key to explaining the purchase of new products, as 43.9% of such purchases come from existing importer-exporter relationships.

Additionally, even when an importing firm changes its primary supplier, it typically replaces the value of imports from a partner that it has purchased from before. The importer can know such a minority supplier either because she previously bought a small share of an HS10 product from that source or through her purchase of a different HS10 product from that firm in a previous year. Both types of familiarity turn out to be important. We find that 25.9% of switching is to partners that were used in the minority for the same HS 10 product, while an additional 26.6% of all partner switches are to a supplier that a U.S. importer has bought a different HS10 product from. Thus over half of all partner switching is to what can broadly be called "familiar" partners.[11] This constitutes robust evidence that familiarity with a supplier is central to the buying decisions of an importer.

Taken together, the "ground-level" findings can be summarized as follows: importers often have multiple partners, and rely on information about those partners when adjusting either the products they import, or the mix of suppliers they use. Together with the

[10] Specifically, 66.1% of U.S. importing firms import more than one product, and these firms account for 98.3% of U.S. imports. 10% of firms import more than sixteen HS10 products.

[11] If we eliminate those cases where each type of familiar switch is impossible, i.e. excluding firms that only used a single partner for an HS 10 product from the first definition, and excluding one-product importing firms from the second definition, the share of switching to familiar partners rises to 69.9%.

results in Section 3.1, we believe the data presents striking evidence of the existence of learning about foreign firms.

But what is the contribution of such learning to trade flows? In order to answer this question, we must account for numerous other explanations for why U.S. importers may choose to stay with the same partner over time. In addition to learning about the quality and reliability of a supplier, these include avoiding the cost of searching for a new partner, favorable pricing terms, the gains from experience of a long-standing supplier with respect to the customization of the product to the specific needs of the importer. While none of these mechanisms besides the price are directly observable, the goal of the theoretical design is to include features of the model that allow us to differentiate between these explanations using our rich relationship data, based on their differential predictions for the dynamics and patterns of trade flows over time. At the same time, the model is consistent with the distinctions outlined above: different behavior across products, institutional quality of the source country, and the asymmetric ability to search effectively for large U.S. importers.

4 Model

In the following we outline a model about learning in exporter-importer relationships. We derive several testable predictions and calibrate the key parameters of the model using moments from the U.S. import data. The model builds heavily on work by Araujo et al. (2012). While their analysis focused on the problem of exporters, we flip their framework to study the related decisions of importers.

4.1 Basic Setup

One importer is matched to one exporter. The importer has all bargaining power and offers the exporter a quantity-price pair. The exporter can accept or reject the offer. Assume that all transactions are done cash-in-advance, meaning the importer has to pay the exporter before goods are sent.[12] In this case, there is a risk that the exporter defaults on the contract and does not deliver the goods after receiving the payment. However, an exporter can only do so when an opportunity for cheating arises. This is more likely the worse the legal institutions in the source country.[13] Let λ measure the quality of legal institutions, so that an opportunity to cheat arises with probability $1 - \lambda$. Assume that a fraction $\widehat{\theta}$ of suppliers are patient whereas the remainder of them are myopic. As in Araujo et al. (2012), we assume that the difference in the discount rates is so large that patient suppliers always want to keep a trade relationship alive, whereas myopic firms try to deviate from the contract whenever they get an opportunity to do so.

Buyer Behavior Since there are two types of suppliers in the economy, learning plays a central role. Initially, buyers believe (correctly) that the probability any seller of a product will fulfill the contract is equal to the population mean $\widehat{\theta}$. Every period that a relationship survives, they update their beliefs according to Bayes Rule. Remember that a myopic supplier defaults whenever there is an opportunity (probability $(1 - \lambda)$). If a buyer has successfully purchased from a seller for k periods, the posterior probability that the seller is patient can be derived as:

$$\theta_k = \frac{\widehat{\theta}}{\widehat{\theta} + \left(1 - \widehat{\theta}\right) \lambda^k} \tag{2}$$

Importantly, the probability only depends on the length of time that a buyer has been buying from the same seller. It is easy to see that for large k, θ_k converges to 1, that is the buyer is almost certain that the seller is of the patient type.

Profits from a relationship of age k For a buyer-seller relationship of age k, the buyer receives the goods from the supplier in two scenarios: either the seller is patient (an event with expected probabiliy θ_k), or that the seller is myopic but does not face any

[12] Note that surveys suggest that most trade is done on open account terms, which represents the opposite from cash-in-advance. This assumption will therefore be relaxed in the later analysis.

[13] Alternatively, one could assume that an exporter always has the opportunity to cheat and that the exporter can go to court to enforce contracts. Legal institutions would then determine the probability that enforcement is successful.

opportunity to renege (an event with expected probability $(1-\theta_k)\lambda^k$). We denote this probability by $\tilde{\theta}_k = (\theta_k + (1-\theta_k)\lambda)$. Thus expected profits when buying from a supplier that an importer has traded with for k periods are:

$$\mathrm{E}[\pi_k] = (\theta_k + (1-\theta_k)\lambda) R(q) - cq - f = \tilde{\theta}_k R(q) - cq - f \tag{3}$$

The buyer can sell the goods for revenue $R(q)$ if they are successfully delivered by the supplier, which occurs with probability $\tilde{\theta}_k$. As the buyer has all bargaining power, she pays the seller the marginal cost of production c for each unit purchased q, even if the product is not delivered.[14] Finally, the importer has to pay a per period cost of sustaining the trade relationship f. The firm can always decide to cancel the trade relationship and receive profits of zero.

Distributions of relationships and trade across cohorts In the following, we derive the fraction of relationships in different age cohorts which we will use in our calibration. Furthermore, we also show how trade shares across cohorts are calculated in our model. Assume that there is a constant and exogenous probability that relationships get destroyed between periods that is given by $\delta \in (0,1)$. Then, the probability a relationship of age k survives to another period is $surv_k = (1-\delta)\tilde{\theta}_k$. A relationship is hence alive after k periods with probability $alive_k = alive_{k-1} * surv_{k-1}$. Assuming the probability of a new relationship being alive is 1, we have:

$$alive_k = (1-\delta)^k \left(\lambda^k(1-\hat{\theta}) + \hat{\theta}\right) \tag{4}$$

We define the steady state to be such that the number of relationships is constant, that is in each period the number of new relationships equals the number of dying relationships. This means that the share of all existing relationships that are age k can be calculated

[14]The importer needs to pay at least the production costs cq in order for patient sellers to participate. The participation constraint of patient sellers is independent of $\hat{\theta}$ and λ and is given by: $T - cq \geq 0$, where T is the payment from the buyer to the seller.

as:

$$
\begin{aligned}
ageshare_k &= alive_k / \sum_{s=0}^{\infty} alive_s & (5) \\
&= (1-\delta)^k \left(\lambda^k (1-\hat{\theta}) + \hat{\theta}\right) \left[\frac{1-\hat{\theta}}{1-(1-\delta)\lambda} + \frac{\hat{\theta}}{\delta}\right]^{-1},
\end{aligned}
$$

as $\sum_{s=0}^{\infty}(1-\delta)^s \left(\lambda^s(1-\hat{\theta}) + \hat{\theta}\right) = \frac{1-\hat{\theta}}{1-(1-\delta)\lambda} + \frac{\hat{\theta}}{\delta}$.

The above assumed an infinite number of possible relationship lengths, but we can also compute these shares given a maximum possible age for a relationship. Such a calculation is useful for matching with our data on relationships described above. Let $ageshare_{k|K}$ denote the fraction of relationships in total relationships that survived for k periods if no relationship is allowed to be active for more than K periods. Then we have:

$$
\begin{aligned}
ageshare_{k|K} &= \frac{alive_k}{\sum_{s=0}^{\infty} alive_s - \sum_{s=K}^{\infty} alive_s} & (6) \\
&= \frac{(1-\delta)^k \left(\lambda^k(1-\hat{\theta}) + \hat{\theta}\right)}{\frac{1-\hat{\theta}}{1-(1-\delta)\lambda} + \frac{\hat{\theta}}{\delta} - (1-\delta)^K \left[\lambda^K \frac{1-\hat{\theta}}{1-(1-\delta)\lambda} + \frac{\hat{\theta}}{\delta}\right]}.
\end{aligned}
$$

This object is now directly comparable to the tables on relationship age shares in Section 3, and depends only on underlying parameters of the model.

Similarly, we can calculate the share of different cohorts in total trade in steady state as:

$$
tradeshare_k = \left(\mathrm{E}[R_k] alive_k\right) / \sum_{s=0}^{\infty} \left(\mathrm{E}[R_s] alive_s\right). \tag{7}
$$

Here, trade shares by cohort depend on specific assumptions on demand, since expected revenues $\mathrm{E}[R_k]$ are typically a non-linear function of the payment probability $\tilde{\theta}_k$. For example, we show below that under CES demand expected revenues $\mathrm{E}[R_k] = \left(\tilde{\theta}_k\right)^{\sigma} B$, where B is a collection of parameters.

4.2 CES Demand

We now solve the model under CES demand for the final good, q, where the demand function is $q = Ap^{-\sigma}$ with elasticity of substitution σ.[15] Expected importer profits in a relationship that has lasted k periods are given by equation (3). Profit maximization implies the optimal price:

$$p_k = \frac{\sigma}{\sigma - 1} \frac{1}{\tilde{\theta}_k} c \tag{8}$$

Expected profits are:

$$\mathrm{E}[\pi_k] = \frac{1}{\sigma} \tilde{\theta}_k A \left[\frac{\sigma}{\sigma - 1} \frac{1}{\tilde{\theta}_k} c \right]^{1-\sigma} - f \tag{9}$$

Revenues and learning For later reference, it is useful to derive expected revenues as a function of $\widetilde{\theta}$:

$$\mathrm{E}[R_k] = (\tilde{\theta}_k)^{\sigma} B, \tag{10}$$

with $B = A \left[\frac{\sigma}{\sigma-1} c\right]^{1-\sigma}$. Note that the marginal cost term c can be exporter-specific, meaning B can be as well. This means that an importer's expected revenues from a relationship depend on how long the relationship has lasted as well as the input cost.

4.3 Simulations

In this subsection, we demonstrate some of the features of the model with a numerical example. We simulate trajectories for profits, prices, and other key variables in the CES-version of the model.

The buyer is comparing profits from using two different sellers. Note that without information $\tilde{\theta}_k$ in the model, profits are maximized by simply using the buyer with least

[15] Such a demand follows from maximizing a CES utility function of the form $U = \left(\int_{\Omega} q(\omega)^{\frac{\sigma-1}{\sigma}} d\omega \right)^{\frac{\sigma}{\sigma-1}}$ subject to a budget constraint and an ideal price index $P = \left(\int_{\omega \in \Omega} p(\omega)^{1-\sigma} d\omega \right)^{\frac{1}{1-\sigma}}$. A is then given by $A = P^{\sigma} Q$, where $Q = \int_{\Omega} q(\omega) d\omega$.

cost. However, by allowing for dynamic adjustment of partners, we can endogenize the decision to switch partners, whereby an importer might first prefer to use a buyer of higher cost that it has better information about, switching only once it learns enough about the other buyer to be sure they will not default. Consider the case with two sellers: Seller 1 has a lower cost, but the buyer has been buying from Seller 2 for 3 periods longer than Seller 1. Setting $\lambda = 0.6$, the share of good sellers $\widehat{\theta} = 0.6$, costs $\{c^1, c^2\} = \{1, 1.2\}$, and $k_2 = k_1 + 3$, we can obtain the graphs found in Figure 3.

In Panel A, the solid line represents the supplier with high cost that possesses a better reputation, by virtue of the fact that $\tilde{\theta}_{k_2}$ is higher. This is because the high-cost supplier has been used for longer. Panels B and C show that eventually, as information improves about the low cost seller, the buyer can charge lower prices for the final good, generating higher revenues. Indeed, Panel D demonstrates that by period 7, the reputation of the low cost seller improves enough such that there are higher profits from utilizing that seller, thereby inducing dynamic switching behavior.

The model is thus consistent with the empirical findings about importer sourcing behavior described in Section 3: importers can have multiple suppliers at the same time, and can switch between suppliers as they learn more about newer exporters. Indeed, growing familiarity with new suppliers leads to endogenous switching behavior in the model.

We also see from the model that higher values of λ lead to more surviving relationships over time. Figure 4 demonstrates that beliefs about the quality of a supplier evolve slower in countries with high levels of λ (Panel A). At the same time, the direct effect of λ leads to less default in countries with good enforcement at every point in time. Figure 4 shows that overall the latter effect dominates, that is, the probability of a relationship lasting to any length k is higher in high rule-of-law countries.

5 Matching and Counterfactuals

5.1 Matching with the data

To evaluate the model and calculate welfare effects, we need to calibrate the key parameters of the model δ, $\hat{\theta}$ and λ. For this, we implement a simple algorithm that minimizes the distance between key moments in the data and our model. The calibration is quite intuitive. Consider Figure 2. It shows the fraction of importer-exporter pairs that have been together for a given number of periods (1 to 20), given the probability of survival above and fixed parameter values. This is a declining function as some relationships may end because the supplier cheats and because with $\delta > 0$ some relationships die for exogenous reasons. Note that, in order to calculate this steady state distribution, we assume that the number of new entrants is constant over time. We obtain it by normalizing the probabilities that firms are alive at different points in time by the sum of all these probabilities (see equation (5) in the theory section).

To calibrate the model, we employ an algorithm that searches for the parameter vector (δ, $\hat{\theta}$, λ) that minimizes the sum of squared differences between the age shares predicted by the model and those found in the data. That is, it solves the following problem:

$$\arg\min_{\delta,\hat{\theta},\lambda} Error = \sum_{t=1}^{N} (ageshare - \widehat{ageshare})^2,$$

where $ageshare$ is the value predicted by the model and $\widehat{ageshare}$ is taken from the data.

5.2 Calibration

We now present results from a calibration exercise and study counterfactuals. For this, we run the matching procedure discussed in section 5.1 for a set of 20 U.S. trading partners.[16] The matching procedure delivers the following parameter estimates: λ has a mean of 0.299

[16] These are China, Hong Kong, Taiwan, Italy, Germany, United Kingdom, Canada, India, France, Japan, South Korea, Mexico, Thailand, Indonesia, Chile, Spain, Brazil, Netherlands, Philippines and Vietnam.

and a standard deviation 0.12; $\hat{\theta}$ has a mean of 0.444 and a standard deviation 0.128; δ has a mean value of 0.313 and a standard deviation 0.061, where the means and the standard deviations are calculated across the individual parameters estimated for each of the 20 countries.[17]

In the following, we illustrate how the parameters can be put back into the model to check its fit and to run counterfactuals. For this, we use the mean values for the three parameters reported above. We begin by calculating the expected payment probabilities $\tilde{\theta}_k = \theta_k + \lambda(1 - \theta_k)$ within a relationship that has lasted for k periods. This is shown in panel A of figure 5. As time goes by, the belief about the supplier being a good type θ_k, and therefore the expected payment probability $\tilde{\theta}_k$ converges to 1. Also note how the speed of learning is declining over time, with most learning taking place in the first few periods. To clarify the working of the model, it is interesting to turn different components on and off, as in panels B-D of figure 5. Panel B where δ is set to zero is identical to panel A, as the death shock does not affect learning. In panel C, enforcement λ is set to zero meaning all impatient suppliers instantaneously reveal their type. After period 1, there is no further learning and the payment probability jumps to 1. In panel D all suppliers are patient, so the payment probability is equal to 1 throughout.

We can now calculate the age shares and trade shares given the parameters - remember that the former were used for the calibration, so a close match between data and the model should be expected. In the current simulations, we assume that relationships end after 20 periods. We can extend this to more periods but attrition is so high that this does not matter quantitatively. The predicted age shares are depicted in panel A of figure 6. As expected, the fraction of relationships within a cohort in total relationships declines with its age. This is the case for two reasons. First, there is realized supplier impatience, which results in relationship death. Second, the exogenous death probability δ leads to shrinking cohort sizes over time. Again, we can study the different components in turn. In panel B, where δ is set to zero, the distribution of relationships over cohorts becomes much more uniform. In panel C, λ is again equal to zero and learning takes place only in

[17]We only report the means and the standard deviations as additional data points have not been disclosed yet. We will provide these country by country in the next draft.

the first period. The distribution is governed by the death shock δ after the first period. Panel D shuts down learning completely, so the distribution is entirely determined by the death shock. The distribution therefore is very similar to the one in panel C.

Next, we calculate the trade shares of different cohorts. Remember from the theory section that these depend on a combination of the fraction of total relationships a cohort represents and the quantity a firm within a cohort trades.The former is declining with age whereas the latter increases in the length of a relationship. To see the latter, recall that there is a direct mapping between $\hat{\theta}$ and q in the model and notice that $\hat{\theta}$ in figure 5 is always weakly increasing over time. The trade shares are shown in figure 7. Panel A shows the full model again. Notice that because of the increasing quantities traded by older cohorts, the trade share distribution has much less weight on the new relationships and more weight on the older relationships than the distribution of the number of relationships depicted in figure 6, panel A. Panel B shows what happens when δ equals zero. Then, the quantity effect dominates and older cohorts have a larger share in trade than younger ones. In panel C there is again no enforcement. Then, there is a little trade in the first period as the probability of payment is very low. However, from period 2, buyers perfectly know the type of their supplier and start buying large quantities. The trade shares are still declining in age after that as the exogenous death shock is still active. In panel D there is no learning, so all cohorts sell the same quantity per relationship. Trade shares and age shares are the same in this case.

The calibration exploits the distribution of relationships across cohorts to identify the three parameters of interest δ, $\hat{\theta}$ and λ but does not use any information on trade shares. It is natural to check for both the fit of the model with the matched moments as well for its ability to predict trade shares. Table 4 compares data moments with moments generated from the calibrated model. Not surprisingly, there is a very close fit in the share of relationships between the model and the data. To evaluate the fit of the trade shares we have to assume an elasticity of substitution σ as this governs how differences in $\hat{\theta}$ map into quantities. For the current comparison we use $\sigma = 3$. The match between trade shares in the model and the data is surprisingly good even though none of the moments were targeted.

5.3 Counterfactuals

The model studied above creates a role for relationship longevity in international trade flows- in particular, trade between two firms increases as a relationship strengthens. Given the underlying parameters of the model that we estimated, we are now able to calculate several counterfactuals to quantify this value of relationships for trade.[18]

Reset of All Trade Relationships We begin with an experiment that illustrates the positive benefits from long-lived trade relationships. Assuming an initial steady state with 1000 relationships, we shock the model by wiping out all established relationships. Final good producers still match with exporters to buy their goods, but the relationship distribution is condensed to be 1000 new relationships. There is then a transition back to the steady state distribution of relationships over time. As the impulse response function in Figure 8 shows, the average age of relationships gradually rises, as do the traded quantities. On impact, traded quantities first drop by $\frac{2}{3}$, and slowly move back to the steady state value of 780.6, returning to the original total traded amount after about 20 periods. In our calibration, 1000 relationships in steady state generate 3.43 times more trade than the same number of new relationships.

One period of increased entry The next experiment studies the effect of a one-off increase in the number of new exporting firms. Results are shown in Figure 9. As above, we begin in the same steady state as before with 1000 relationships. Now we shock the model so that 200 additional relationships are created in a single period- that is, over and above the 192.3 required to sustain the steady state number of relationships. With the 200 additional relationships, traded quantities increase by 5.8 percent on impact. This initial increase is represented by the dotted line in the graph. Even though some of the 200 relationships die over time, traded quantities first rise as a result of learning. They peak at 844.9 after 2 periods, or an increase of 8.2 percent over steady state. After that, trade slowly moves back to its steady state level.

[18] For the counterfactuals, we assume that the maximum duration of relationships is 80 periods.

Permanent increase in entry The third experiment, shown in Figure 10, studies a permanent increase in the number of new relationships per period by 50. It starts again in a steady state with 1000 relationships and a per period entry of 192.3. After 20 periods, the number of new relationships created per period increases permanently by 50 to 242.3. This leads to a rise in trade from 780.6 to 792.0, shown by the dotted line. As can be seen, there is an additional amount traded beyond the direct amount from this increase in the extensive margin, again due to endogenous trade increases governed by learning. Over time, the economy moves to a new steady state with a total trade quantity of 983.6 and 1260 relationships, an increase of 26.0 percent in both measures.

In sum, our counterfactuals illustrate how the model can be employed to run informative experiments. They show that learning and relationships can have first-order effects on the levels and the dynamics of trade and highlight the importance of distinguishing between permanent and transitory shocks to the economy and to trade relationships.

6 Conclusions

This paper employs rich data on U.S. imports to analyze importer-exporter relationships. We present a set of new stylized facts and develop a model, building on Araujo et al. (2012), to clarify the different mechanisms at play and quantify the role of learning in explaining the patterns in the U.S. data. While still preliminary, our results suggest that long-run relationships are key to international trade. We identify several factors that affect the length of relationships, including product characteristics, the size of total imports and the quality of legal institutions in the source country. We also show that importers rely on preexisting relationships when purchasing new products or adjusting which suppliers they buy from. These findings are in line with a model of importer-exporter relationships where learning about the reliability of the trading partner represents a central aspect of a firm's decision. We estimate this model in order to quantify the contribution of successful, sustained relationships for international trade. Our counterfactual experiments demonstrate that long-term relationships have a significant role in explaining trade flows, and their disruption can have long-lasting negative effects.

References

Albornoz, Facundo, Héctor F Calvo Pardo, Gregory Corcos, and Emanuel Ornelas, "Sequential Exporting," *Journal of International Economics*, 2012, *88* (1), 17–31.

Antràs, Pol and C. Fritz Foley, "Poultry in Motion: A Study of International Trade Finance Practices," *Journal of Political Economy*, forthcoming.

Araujo, Luis F, Giordano Mion, and Emanuel Ornelas, "Institutions and Export Dynamics," Working Paper 220, National Bank of Belgium 2012.

Bernard, Andrew B, Andreas Moxnes, and Karen Helene Ulltveit-Moe, "Two-sided Heterogeneity and Trade," Working Paper 20136, National Bureau of Economic Research 2014.

_ , J Bradford Jensen, and Peter K Schott, "Importers, Exporters and Multinationals: A Portrait of Firms in the US that Trade Goods," in "Producer dynamics: New evidence from micro data," University of Chicago Press, 2009, pp. 513–552.

_ , _ , Stephen J Redding, and Peter K Schott, "Wholesalers and Retailers in International Trade," *American Economic Review: Papers and Proceedings*, 2010, *100*, 408–413.

Besedeš, Tibor, "A Search Cost Perspective on Formation and Duration of Trade," *Review of International Economics*, 2008, *16* (5), 835–849.

Blum, Bernardo S, Sebastian Claro, and Ignatius J Horstmann, "Occasional and Perennial Exporters," *Journal of International Economics*, 2013, *90* (1), 65–74.

Carballo, Jerónimo, Gianmarco IP Ottaviano, and Christian Volpe Martincus, "The Buyer Margins of Firms' Exports," Discussion Paper 9584, CEPR 2013.

Chaney, Thomas, "The Network Structure of International Trade," *American Economic Review*, 2014, *104* (11), 3600–34.

Dragusanu, Raluca, "Firm-to-Firm Matching Along the Supply Chain," 2014. Harvard University, mimeo.

Eaton, Jonathan, Marcela Eslava, Cornell J Krizan, Maurice Kugler, and James Tybout, "A Search and Learning Model of Export Dynamics," 2014. Mimeo.

Egan, Mary Lou and Ashoka Mody, "Buyer-Seller Links in Export Development," *World Development*, 1992, *20* (3), 321–334.

Greif, Avner, "Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders' Coalition," *The American Economic Review*, 1993, pp. 525–548.

Impullitti, Giammario, Alfonso A Irarrazabal, and Luca David Opromolla, "A Theory of Entry into and Exit from Export Markets," *Journal of International Economics*, 2013, *90* (1), 75–90.

Kamal, Fariha and Asha Sundaram, "Buyer-Seller Relationships in International Trade: Do Your Neighbors Matter?," 2013. Mimeo.

_ and C.J. Krizan, "Decomposing Aggregate Trade Flows: New Evidence from US Traders," Working Paper 12-17, U.S. Census Center for Economic Studies 2012.

_ , _ , and Ryan Monarch, "Identifying Foreign Suppliers in U.S. Merchandise Import Transactions," Working Paper 15-10, U.S. Census Center for Economic Studies 2015.

Krizan, C.J., "The Nuts and Bolts of International Trade Transactions: A Data User's Primer," 2012. U.S. Census Center for Economic Studies, mimeo.

Loecker, Jan De, "Detecting Learning by Exporting," *American Economic Journal: Microeconomics*, 2013, *5* (3), 1–21.

Macchiavello, Rocco and Ameet Morjaria, "The Value of Relationships: Evidence from a Supply Shock to Kenyan Rose Exports," *American Economic Review*, forthcoming.

Monarch, Ryan, "It's Not You, It's Me: Breakups in U.S.-China Trade Relationships," Working Paper 14-08, U.S. Census Center for Economic Studies 2014.

Pierce, Justin R and Peter K Schott, "The Surprisingly Swift Decline of U.S. Manufacturing Employment," Working Paper 18655, National Bureau of Economic Research 2012.

Rauch, James E, "Networks versus Markets in International Trade," *Journal of International Economics*, 1999, *48* (1), 7–35.

_ , "Business and Social Networks in International Trade," *Journal of Economic Literature*, 2001, pp. 1177–1203.

_ **and Joel Watson**, "Network Intermediaries in International Trade," *Journal of Economics & Management Strategy*, 2004, *13* (1), 69–93.

A Figures

Figure 1: Relationship Longevity and Source Country Rule-of-Law

Panel A: Fraction of New Relationships

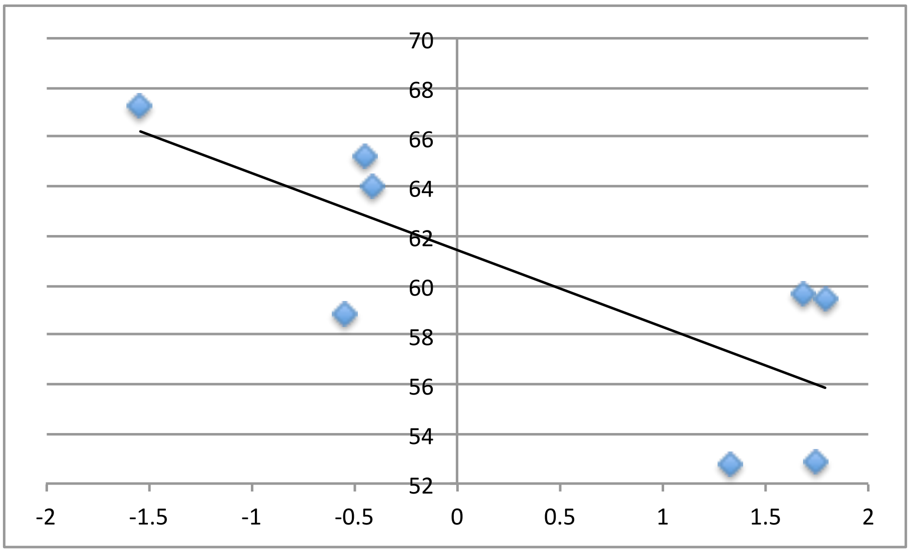

Panel B: Fraction of Long-Term Relationships

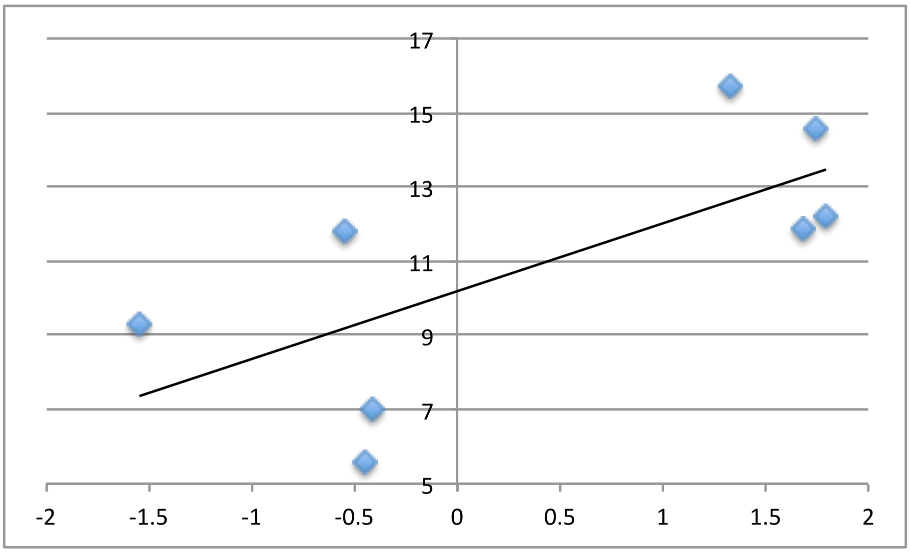

Note: In the above graphs, the x-axis is the fraction of importer-exporter relationships that are either new (Panel A) or long-term (Panel B), graphed against the rule of law measurement for that source country. The countries (from left to right) are Venezuela, Mexico, China, Vietnam, Mexico, Japan, Canada, United Kingdom, Germany, and Canada. See also Table 2 Panel C.

Figure 2: Age shares of importer-exporter relationships

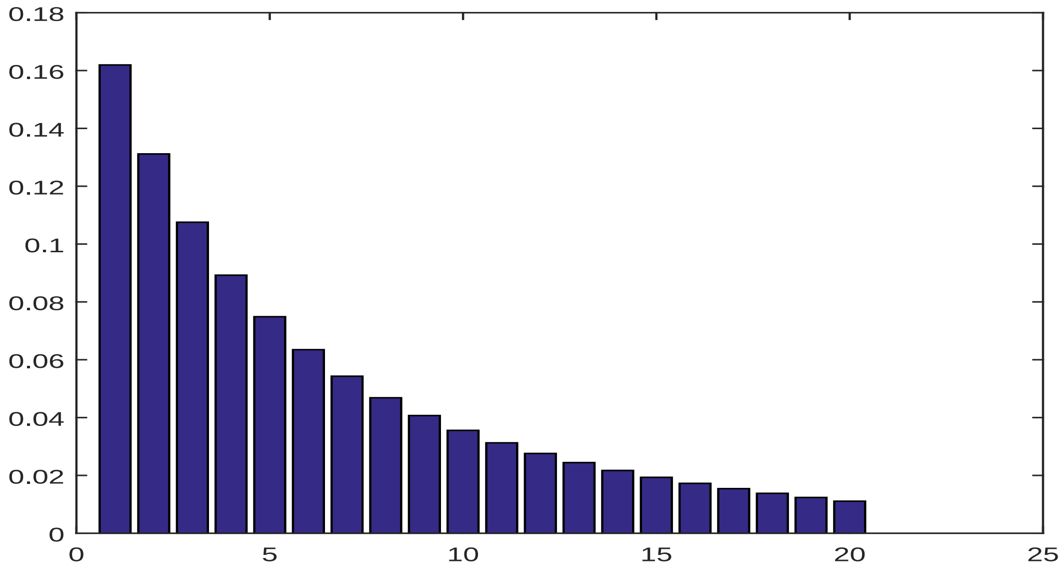

Notes: This graph illustrates the distribution of buyer-seller relationship ages implied by Equation (6), with parameter values $\delta = 0.1$, $\lambda = 0.5$, $\widehat{\theta} = 0.8$, and $K = 20$.

Figure 3: Relationship Longevity, Cost and Supplier Choice

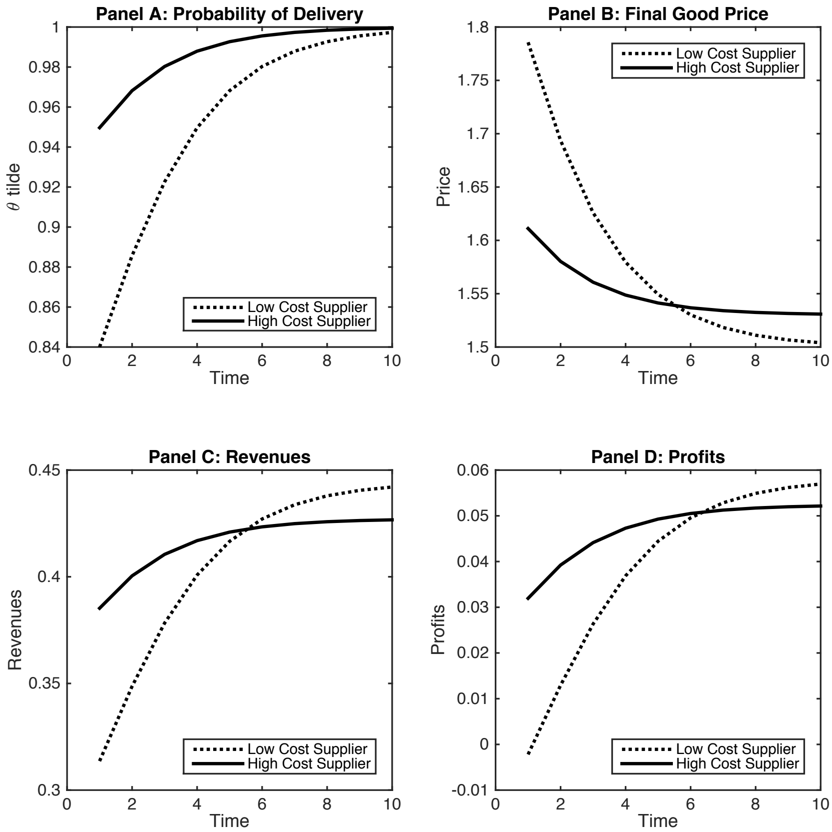

Note: These simulations are for final good producers facing CES demand and having expected profits according to Equation (3). The choice is which supplier to use, a high-cost seller who is in a long-term relationship with the buyer (solid line) or a low-cost seller who is unfamiliar with the seller (dotted line).

Figure 4: Model Simulations- Different λ

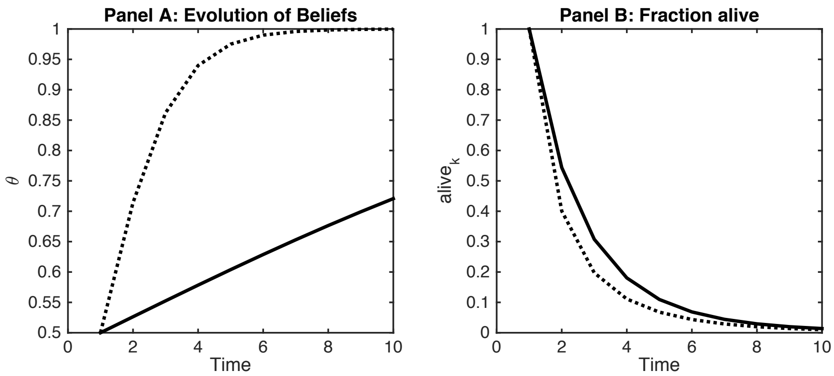

Note: These graphs illustrate the effect of different levels of institutional quality (λ) on objects in the model. The simulations are the importer's belief that the exporter is patient (θ_k) as the relationship ages, and the respective age distribution of relationships ($alive$) for $\lambda = 0.4$ (the dotted line) and $\lambda = 0.9$ (the solid line).

Figure 5: Predicted Payment Probabilities

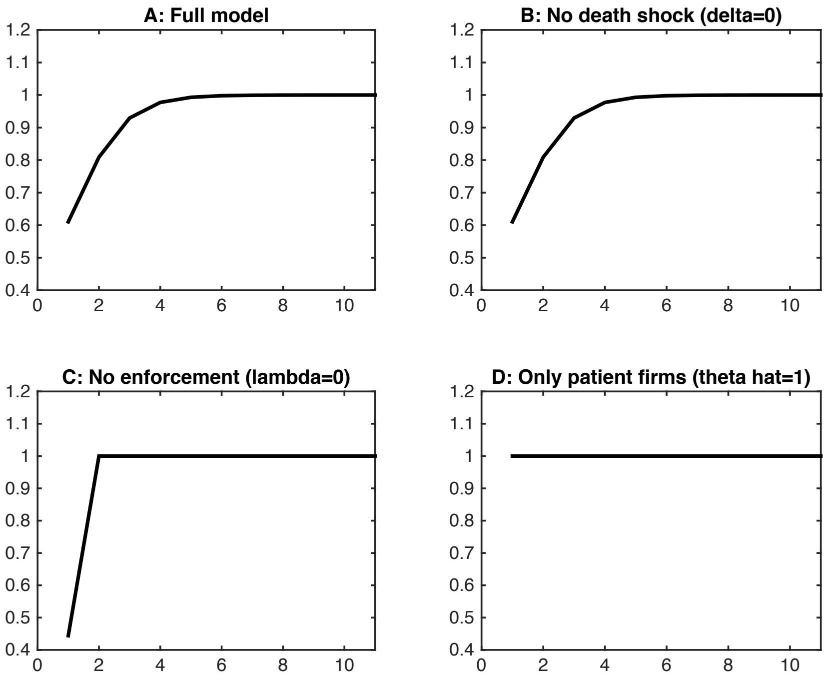

Note: These graphs illustrate the effects of varying model parameters on the predicted payment probability $\tilde{\theta}_k$ against the length of the relationship. The baseline model (Panel A) uses the parameters estimated from the data ($\lambda = 0.299, \widehat{\theta} = 0.444, \delta = 0.313$). Without a death shock δ (Panel B), the probability that a supplier over a relationship delivers the goods is unchanged, as learning is not affected by δ. With no enforcement of contracts (Panel C), any impatient seller instantly reneges on the contract, and learning ceases after the first period- if a relationship survives one period, the seller must be patient. If there are no myopic firms (Panel D), then there is never any learning, and the buyer believes that, death shock aside, their goods always arrive.

Figure 6: Predicted Age Shares

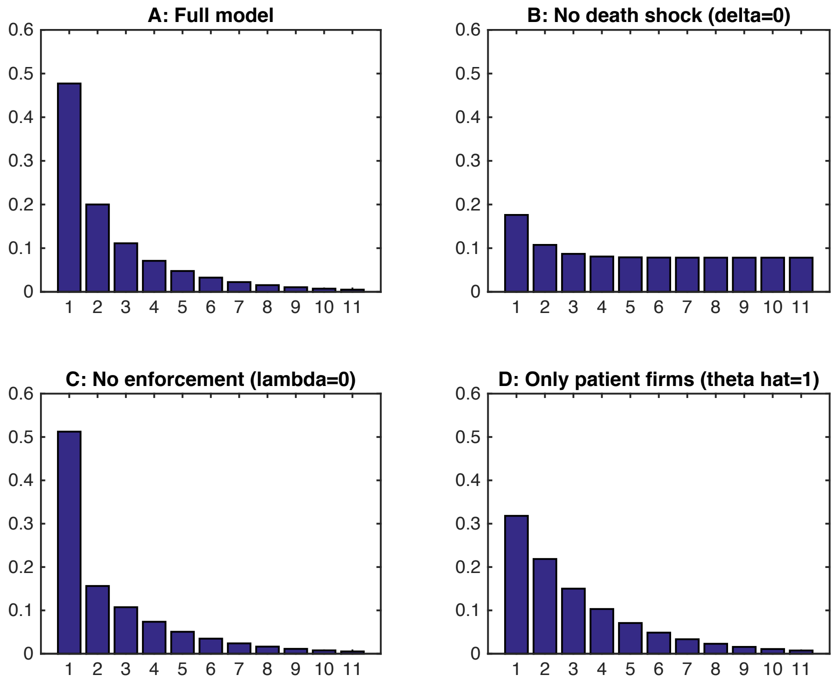

Note: These graphs illustrate the effects of varying model parameters on the predicted age distribution of relationships. In the baseline model with the estimated parameters (Panel A), longer relationships occupy a smaller fraction of total relationships at any one time. Without a death shock, the distribution of ages flattens over time, as longer relationships less likely to die off. With no enforcement of contracts (Panel C), the age distribution is entirely driven by the death shock δ after the first period. This mirrors the case of no myopic firms (Panel D), in which the entire distribution is driven only by δ from the beginning of a relationship.

Figure 7: Predicted Trade Shares, $\sigma = 3$

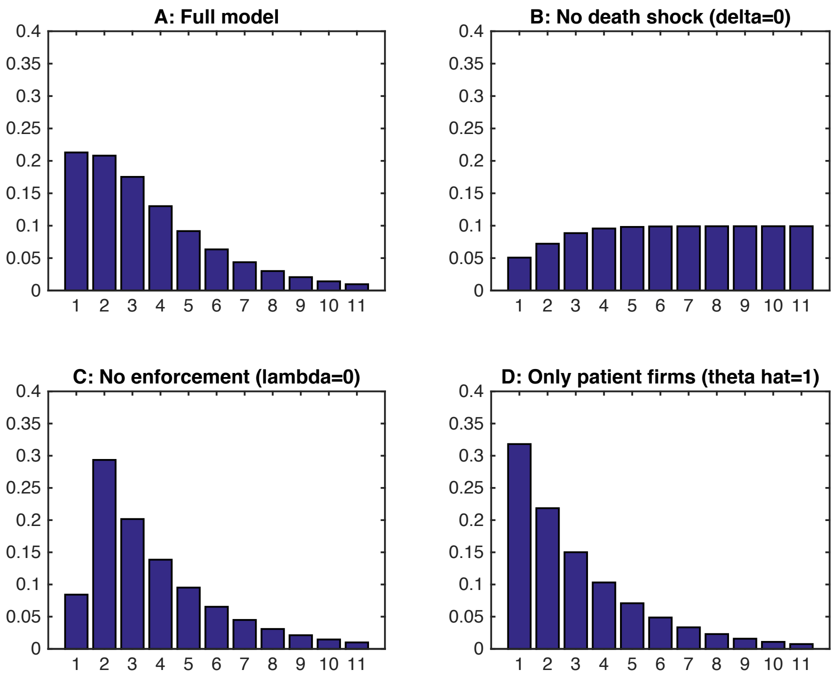

Note: These graphs illustrate the effects of varying model parameters on the predicted age distribution of trade under CES preferences for the importer's final good. This distribution (Equation (7)) depends both on the distribution of relationships by age and the quantity traded by any firm, which, in the case of CES, is implied by Equation (8). In the baseline model with the estimated parameters (Panel A), new relationships trade much less than older relationships, given the increasing quantities implied by the model. Without a death shock, the quantity effect outweighs learning, and older cohorts occupy an even higher fraction of total trade. With no enforcement of contracts (Panel C), little is traded until suppliers reveal their type, after which the distribution is entirely driven by the death shock δ after the first period. This again mirrors the case of no myopic firms (Panel D), in which the entire distribution is driven only by δ from the beginning of a relationship.

Figure 8: Experiment 1: Ending All Relationships

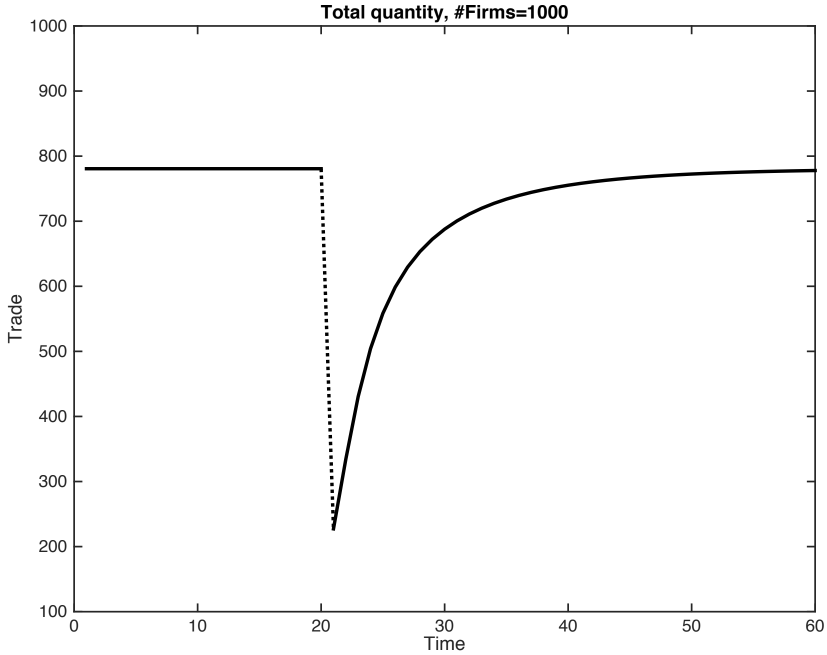

Note: This graph simulates the effect on total trade from a "reset" of all trade relationships back to new. The model is estimated for 1000 importer-exporter relationships, under the baseline parameters of the model ($\lambda = 0.299, \widehat{\theta} = 0.444, \delta = 0.313$).

Figure 9: Experiment 2: One Period of Higher Entry

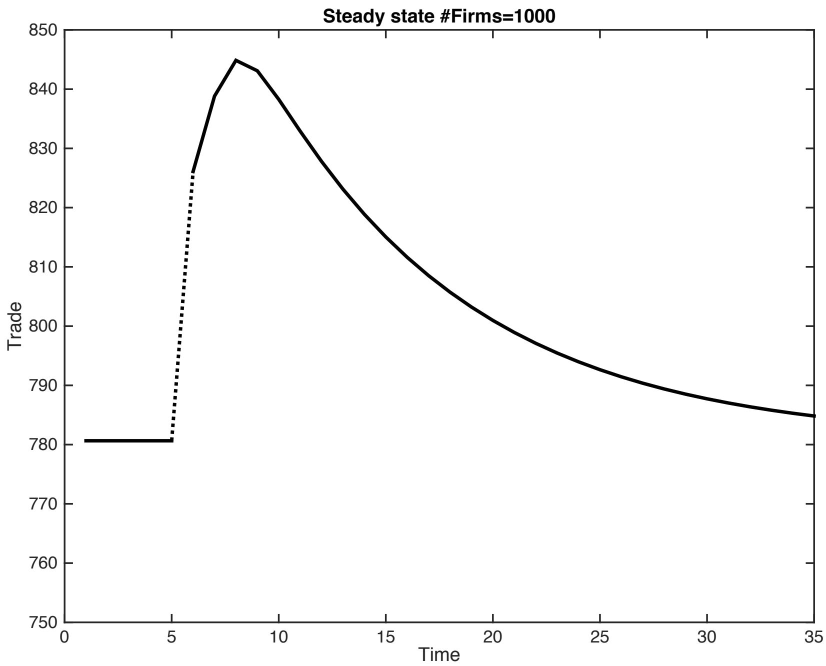

Note: This graph simulates the effect on total trade from a one-time increase in the number of relationships of 20%. Although some of the relationships die out over time, trade continues to increase within the first few periods based on learning about supplier quality. Trade eventually returns to its steady state level. The model is estimated for 1000 (steady-state) importer-exporter relationships, under the baseline parameters of the model ($\lambda = 0.299, \widehat{\theta} = 0.444, \delta = 0.313$).

Figure 10: Experiment 3: Permanent Increase in Entry

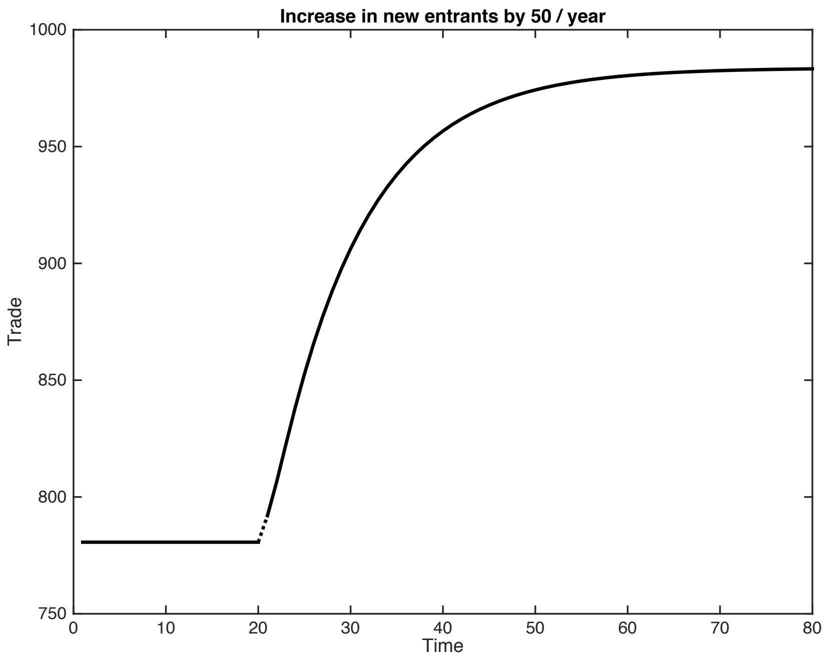

Note: This graph simulates the effect on total trade from a permanent increase in the steady-state number of relationships, achieved by a continual increase in the number of new relationships each year. There is an additional amount traded beyond the direct effect from additional relationships, which comes from the learning-based increases in quantity. Trade eventually returns to its steady state level. The model is estimated under the baseline parameters of the model ($\lambda = 0.299, \widehat{\theta} = 0.444, \delta = 0.313$).

B Tables

Table 1: Relationship Structure of U.S. Imports, 2007

	New	1-2 Years	3 or More Years
Share of Relationships	58.4	30.5	11.1
Share of Trade	25.3	38.7	36.0

For this table, a relationship is defined as a U.S. importing firm buying a product within an HS2 industry from a non-U.S. exporting firm. A new relationship is one that is not found in any previous year of data, while relationships of multiple years are defined similarly.

Table 2: Relationship Structure of U.S. Imports, Various Categories, 2007

Panel A: Product Differentiation

	% of Relationships			% of Trade		
	New	1-2 Years	3+ Years	New	1-2 Years	3+ Years
Differentiated	59.0	30.2	10.8	26.4	39.7	33.9
Non-Differentiated	54.4	32.0	13.5	23.1	36.7	40.1

For this panel, differentiated refers to an HS2 code where more than 50 percent of products are classified as differentiated. As above, a relationship is defined as a U.S. importing firm buying a product within an HS2 industry from a non-U.S. exporting firm.

Panel B: U.S. Firm Size

	% of Relationships			% of Trade		
	New	1-2 Years	3+ Years	New	1-2 Years	3+ Years
Small	74.4	20.7	4.8	68.6	24.4	7.1
Medium	62.4	28.5	9.2	48.7	35.0	16.3
Large	57.3	31.1	11.6	25.0	38.8	36.2

For this panel, firm size is defined as the total volume of imports across all products for that U.S. firm. Firms are split evenly between three categories of Small, Medium, and Large. As above, a relationship is defined as a U.S. importing firm buying a product within an HS2 industry from a non-U.S. exporting firm.

Panel C: Source Country

	% of Relationships			% of Trade		
	New	1-2 Years	3+ Years	New	1-2 Years	3+ Years
Venezuela	67.3	23.4	9.3	11.0	67.7	21.3
China	65.2	29.3	5.6	33.7	46.6	19.7
Mexico	58.9	29.4	11.8	14.7	31.3	53.9
Vietnam	64.0	29.0	7.0	27.5	39.2	33.2
Germany	52.9	32.5	14.6	26.8	30.0	43.2
U.K.	59.7	28.4	11.9	32.1	36.5	31.4
Canada	59.5	28.3	12.2	23.0	34.0	43.1
Japan	52.8	31.5	15.7	16.9	30.4	52.7

For this panel, as above, a relationship is defined as a U.S. importing firm buying a product within an HS2 industry from a non-U.S. exporting firm.

Table 3: Relationship between Institutions and Staying/ Switching Decision

Dependent Variable: Share of Importers Staying with Exporter Year-to-Year, 2002-2008

	(1)	(2)	(3)
Log Strength of Legal Rights	0.0489**		
	(0.02253)		
Procedures to Enforce a Contract		-0.00673***	
		(0.00195)	
Rank (Procedures)			-0.00106***
			(0.00027)
Log Per Capita GDP	0.02403***	0.01893***	0.01677**
	(0.00740)	(0.00743)	(0.00747)
Constant	0.16496**	0.52975**	0.37517***
	(0.05531)	(0.11278)	(0.07171)
N	151	152	152
R^2	0.14	0.18	0.20

Notes: The independent variables come from the World Bank's World Development Indicators. Strength of Legal Rights is an index from 0 to 10, and measures the degree to which collateral and bankruptcy laws protect the rights of borrowers and lenders in a country. Number of procedures to enforce a contract are the number of independent actions, mandated by law or courts, that demand interaction between the parties of a contract or between them and the judge or court officer. Per Capita GDP and Private Credit Coverage variables are also from the World Bank. Three asterisks implies significance at 1%, two asterisks implies significance at 5%.

Table 4: Relationship Structure of U.S. Imports vs. Model with $\sigma = 3$

	Share of Relationships			Share of Trade		
	New	1-2 Years	3+ Years	New	1-2 Years	3+ Years
Data	58.4	30.5	11.1	25.3	38.7	36.0
Model	50.9	31.9	17.2	28.7	34.1	37.1

Notes: The top row of the table is taken from Table 1. The bottom row of the table is the accompanying age shares estimated under the model, as described in Section 5.1. The baseline parameter values are $\lambda = 0.299$, $\widehat{\theta} = 0.444$, and $\delta = 0.313$. To calculate trade shares, we use CES preferences and $\sigma = 3$. Although the left panel is the targeted moment from our estimation, the right panel is not targeted.