

Learning and the Value of Relationships in International Trade*

PRELIMINARY DRAFT. DO NOT CIRCULATE!

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

February 13, 2015

Abstract

We use U.S. import data to explore the characteristics of importer-exporter relationships. While many relationships split up after only one period, others are very persistent. Even when a relationship breakup occurs, importers overwhelmingly replace imports from that partner by buying from a supplier familiar to them. Furthermore, 43% of new product purchases come from firms that the importer interacted with in an earlier year. These results indicate the presence of large matching frictions in international trade. We develop a model to study and quantify the role of reputation and learning in explaining the observed patterns of importer-exporter relationships. Predictions of the model regarding the correlation between switching behavior and source country institutions, the number of export partners, and the number of products purchased are borne out in the data.

*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. 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

Some results are under disclosure review.

Parts of the text that refer to these results have been replaced by

...

These results should be available in the near future.

1 Introduction

What is the contribution of successful, sustained firm-to-firm relationships to international trade? It is plausible that relationships between exporters and importers can have major implications for international trade flows. Successful relationships may lead to expansion for firms on both sides, while failed relationships involve disrupted production possibilities in addition to wasted money and effort. In fact, breaking relationships might be so costly for an importing firm that it might choose to remain with an exporting partner even if the particulars of the transaction are not ideal. Newly available data capturing firms on both sides of an international trade transaction demonstrates that there is substantial variation in the duration of relationships in international trade. While a large number of importers buy from a specific exporter only once, there are many trade relationships that are long-lasting, with the same importer and exporter trading with each other over many years.

Recent work (e.g. Monarch (2014) and Eaton et al. (2014)) has suggested different mechanisms that may explain the longevity of importer-exporter relationships. Firms may keep trading with each other because there is ongoing learning about a trading partner, because there are high fixed costs of finding another partner, or simply because there are cost differences across potential suppliers. But the extent to which these respective mechanisms and relationship sustainability contributes the observed patterns of trade relationships and to trade flows more broadly remains an open question. To shed more light on these issues, this paper employs detailed data from the U.S. Census to study importer-exporter trade relationships in great detail. We analyze a model of learning in importing 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.

We begin by generating a set of new stylized facts to guide a model of exporter-importer relationships¹. First, we study the overall duration composition of trade rela-

¹All 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.

tionships. We look at importer-exporter relationships in terms of value and in terms of numbers. While ... percent of total (arm's length) trade takes place in relationship that have lasted for at least three years, more than ... of the overall trade relationships in a given year are new. Another relevant dimension is the type of product traded. For this, we compare differentiated with non-differentiated products, building on the classification developed by Rauch (1999). Perhaps surprisingly, non-differentiated products tend to be traded in ... relationships than differentiated products. One reservation may be that a large fraction of differentiated product trade takes place among related parties that is likely to be long-term. The length of relationships is also systematically related to source country institutions. The ... the rule of law in the supplier's country, the longer the average importer-exporter relationship. Finally, the duration of trade relationships can be explained by firm size. The ... a firm is importing overall, the longer on average its relationships with its suppliers. This could be the result of an advantage of ... firms in keeping relationships alive. Further analysis should reveal the extent to which these two factors contribute to the observed pattern.

We next generate several new results on relationships at the importer-exporter-product level. The data shows that 48.8% of these relationships continue from year-to-year. Additionally, even when switching does occur, 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. Underscoring the importance of familiarity even further is the finding that 43% of such new product purchases also come from export partners used to buy other HS10 products in the past. A key insight of our work is to use variation in the behavior of multi-product importers along these lines to separately identify the process of dynamic learning about a supplier (which occurs across multiple periods and products) from the static search cost explanation.

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, and which supplier to use for both new and existing products. We analyse these choices at different levels of aggregation: country level aggregates of switching behavior, firm level decisions, and firm-product level decisions. 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 initially, even with perfect information about the price of the traded product.

In a preliminary calibration exercise, we present first results where we estimate key parameters of the model through a method of simulated moments exercise. In the final section, we present evidence for the additional predictions implied by the model regarding the relationship between institutional quality and switching, as well as link between the total number of suppliers used and the probability of switching.

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. Recent work has made use of the U.S. Customs database used in this paper, which provides information about U.S. importers and their foreign exporting partners. 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. Other work takes advantage of two-sided trade data 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. Our work also fits into the literature on multi-product firms in international trade, including Bernard et al. (2010) and Bernard et al. (2011). In this project, we combine a theory

of trade network formation, multi-product importers and dynamic learning behavior by importers about the quality of buyers.

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 the five broad empirical findings about U.S. importer relationships with foreign partner firms that form the backbone of our project. 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 reduced form tests we run to examine these predictions. 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.² 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.³ 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. (2009), 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.

²Approximately 80-85% of these customs forms are filled out electronically (Krizan (2012)).

³Specifically, 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.

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 exporters who are importing from separate branches of the same firm will likely have very different relationship dynamics than arm’s-length exporters.

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 level⁴. We additionally distinguish between new, short-term, and long-term relationships, where a new relationship is one that is not found in any previous year, a short-term relationship is one lasting 1-2 (consecutive) years and a long-term relationship is one lasting 3 or more years⁵.

3 Empirical Findings

3.1 Relationship Length

Table 1 presents a breakdown of U.S. imports in 2007 based on the length of relationships. As can be seen, the largest fraction of trade takes place among ... relationships, but both ... relationships and ... relationships account for non-trivial fractions of trade. On the other hand, from a count basis, the ... of relationships within a year are new. We take the numbers from this table to be indicative of two new stylized facts. First, learning is not just about source or destination countries (Albornoz et al. (2012)) or about improving productivity (De Loecker (2013)), but also about trading partners: firms are trying out new suppliers, learning about them, and in many cases abandoning them just as rapidly. Second, most of trade is occurring from ... relationships, indicating that the

⁴Thus within this framework, it is possible for an importer-exporter to have multiple relationships with each other.

⁵The 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

... importers appear to be also the best at maintaining relationships, whether because of better screening before initiating relationships or learning faster which firms are better partners.

Panel A of Table 2 presents the same information on relationships along a different dimension. The first is to look at relationships based on how differentiated the product being purchased is. 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). The exact split does not matter for any of the ordinal rankings below. The highest share of trade in ... products is clustered in long-term relationships, while ... products are predominantly in shorter relationships. The same result holds for the fraction of relationships as well. To some extent, this is quite surprising. In particular, the fact that ... [in process for disclosure]

Another dimension in which the structure of relationships differs is across source countries exporting to the U.S. Information about relationship duration for U.S. importers and firms 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 ... relationships with U.S. firms. The ... of imports from China, as well as the ... of relationships between Chinese and U.S. firms, are among ... importer-exporter pairings for an industry. This contrasts with U.S. imports from Germany or the U.K., which are on average, ... A second stylized fact from Panel B is the effect of distance on relationship length. The closest foreign countries to the U.S. (Mexico and Canada) tend to have ... relationships then their counterparts of similar institutional quality that are farther away.

Finally, we examine relationships based on the total size of 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 C of Table 2 shows that the largest firms form ... relationships, while small importers are predominantly engaging in ... transactions.

There are several explanations why U.S. importers may choose to stay with the same partner over time. These include avoiding the cost of searching for a new partner, favor-

able 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 and importantly, learning by the importer about the quality and reliability of the supplier.

While none of these mechanisms besides the price are directly observable, we can use our rich data to differentiate between them based on their differential predictions for the dynamics and patterns of trade flows over time. At the same time, we design the features of the model to allow for the types of distinctions outlined above: different behavior across products, institutional quality of the source country, distance, and the asymmetric ability to search effectively for U.S. importers.

3.2 Multiple-Partner Sourcing

In this subsection we describe new findings on the prevalence of sourcing from multiple partners at the same time. We shift the focus now from HS2 industries into more specific 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. However, learning about the quality of an exporting partner can obviously happen across products as well, a dimension not captured in the relationship breakdowns above. Figure 1 demonstrates that such a channel is likely to matter: there is significant variation in the number of products imported. Specifically, 66.1% of U.S. importing firms import more than one product, and these firms account for 98.3% of U.S. imports. Furthermore, 10% of firms in the data import more than sixteen HS10 products.

Connecting these findings back to relationship information demonstrates that, as predicted, many U.S. importers rely on multiple sources for their imports. For U.S. importing firms, the average number of 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.

3.3 Dynamic trade relationships and Partner Familiarity

The results on the distribution of relationship length above demonstrate that some trade relationships are extremely stable while others are much more transitory. In this subsection, we go into more depth about how year-to-year supplier decisions are made. We show statistics that reveal the decision of which exporter to purchase from is not random, even if a relationship is ended. Consistent with the statistics on relationship length, we find high persistence in trade relationships, a trend that is consistent both with substantial fixed costs of searching for a trading partner and information asymmetries about the reliability of new suppliers.

We begin by looking at the probability that an importing firm keeps buying a product from the same main supplier year to year. Define a U.S. firm as “staying” with a supplier if it obtains the largest share of its purchases of a product from the same firm for two consecutive years.⁶ It is critical to note that this definition of staying is *firm-HS10 product* specific- i.e. one firm with many export products could have multiple stay/switch decisions within one period. Additionally, a firm could buy a product only in one year and that firm-product observation would not enter into this data of relationship dynamics.⁷

According to this definition, 48.8% of U.S. importers stay with the same partner from one year to the next. But even among those who change their main supplier, it is very common to lean on experience from previous interactions with partners that were used as minority suppliers. That is, even when switching, importers tend to buy from firms they are familiar with.

There are two ways in which an importing firm can be “familiar” with a supplier of a given HS10 product in our data, other than it being her current main source. First, the importer can know about another supplier through her purchase of a different HS10 product from that firm in a previous year. Second, the importer can know about a supplier because she previously bought a minority share of the HS10 product from that source.

⁶The average share of trade from this “major partner” used by a U.S. firm-product combination is 85%, with a median of 100%.

⁷We study new product purchasing decisions later in this section.

Both types of familiarity turn out to be important. We find that 26.6% of all partner switches (again defined as a U.S. firm-HS10 product combination buying from a new partner) are to a supplier that a U.S. importer has bought a different HS10 product from. An additional 25.9% of switching is to partners that were used in the minority for the same HS 10 product. Thus over half of all partner switching is to what can broadly be called “familiar” partners. Furthermore, if we eliminate those cases where each type of familiar switch is impossible, i.e. excluding one-product importing firms from the first definition, and excluding firms that only used a single partner for an HS 10 product from the second definition, the share of switching to familiar partners rises to 69.9%. This constitutes robust evidence that familiarity with a supplier is central to the buying decisions of an importer.

Familiarity could also matter for the purchase of “new” HS10 products. Define a new purchase of an HS10 product as a U.S. importer buying an HS10 product that it had not purchased in the previous year. 72.2% of importing firms buy at least one new product each year. Again, we find familiarity to be a key explanatory factor in these purchases: 43.9% of new products come from partners used one year previously for a different HS10 product.

To summarize, there are a number of key results from the data that the dynamic model of import sourcing we work with should match:

1. The length of trading relationships differs across importing firms, products, and source countries.
2. Importing firms have multiple export partners, even for the same HS10 product.
3. Importers have strong links to their chosen export partners over time.
4. Exporter choice is heavily influenced by prior experience with that partner, both within and across products.
5. New product purchases are governed in part by exporter usage in other products.

4 Model

In the following we outline a model about learning in exporter-importer relationships. We derive several testable implications 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 use their framework to study the related decisions of importers. We follow their basic setup closely before extending it to allow for multi-product firms as well as for differences in production costs.

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, that is the importer has to pay the exporter before goods are sent.⁸ 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 the more likely, the worse the legal institutions in the source country.⁹ Let λ measure the quality of legal institutions, so that an opportunity to cheat arises with probability $1 - \lambda$. Assume that a fraction $\hat{\theta}$ 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 As there are two types of suppliers in the economy, learning plays a central role. Initially, buyers believe (correctly) that the probability that any seller of a product is patient and will fulfill the contract is equal to the population mean $\hat{\theta}$. Every period that a relationship survives, they update their beliefs according to Bayes Rule.

⁸Note 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.

⁹Alternatively, 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.

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:

$$\theta_k = \frac{\hat{\theta}}{\hat{\theta} + (1 - \hat{\theta}) \lambda^k} \quad (1)$$

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 good type.

The Static Case In the following, we introduce a fixed cost of keeping a trade relationship from one period to the next. Denoted by f , this cost is paid at the beginning of the period, before the optimal import bundle is chosen. Then, expected importer profits when buying from a supplier that she traded with for k periods are:

$$\pi_k = \max \{(\theta_k + (1 - \theta_k) \lambda) R(q) - cq - f, 0\} \quad (2)$$

The buyer can sell the goods for revenue R if they are successfully delivered by the supplier. This happens with probability $\theta_k + (1 - \theta_k) \lambda$ (that is, if the seller is non-myopic, or if the seller is myopic, but no opportunity to default occurs). As the buyer has all bargaining power, she pays the seller the marginal cost of production c for each unit purchased q . Finally, the importer has to pay the per period cost of sustaining the trade relationship f . The firm can always decide to cancel the trade relationship and receive profits of zero.

The net present value of future profits Assume that with probability δ a relationship is separated for exogenous reasons, such as supplier exit. Further, assume that this shock takes place between trading periods, so that the Bayesian updating is not affected by this variable. We can then derive the net present value of a given trade relationship

as:

$$\Pi = \pi_0 + \sum_{i=1}^{\infty} \left\{ \left(\prod_{j=1}^i (1 - \delta)[\lambda(1 - \theta_j) + \theta_j] \right) \pi_i \right\} \quad (3)$$

$(1 - \delta)[\lambda(1 - \theta_j) + \theta_j]$ is the probability that a relationship active in period $j - 1$ survives to period j . $\prod_{j=1}^i (1 - \delta)[\lambda(1 - \theta_j) + \theta_j]$ is therefore the probability that a relationship that is formed in period 1 is still active in period i .¹⁰

4.2 Matching with the data

To evaluate the model and calculate welfare effects at a later point, 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 calibrations 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.

To calibrate the model, we employ an algorithm that searches for the parameter vector $(\delta, \hat{\theta}, \lambda)$ 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.

¹⁰We do not include an additional discount factor for the importer as this would complicate the presentation without adding any additional insights. It would of course be straightforward to include one.

4.3 Model extensions

In this section we discuss extensions to the baseline model. While several of the predictions still need to be taken to the data, we provide some preliminary evidence that is consistent with these extensions in the empirical section.

The case of two suppliers Suppose there are two suppliers of the same product, each of whom has been used by the importer. To distinguish between them we add a superscript $s \in \{1, 2\}$ to the relevant variables. Suppliers can differ in their production cost c^s . Furthermore, the buyer may have different posterior beliefs about their reliability $\{\theta_{k_1}^1, \theta_{k_2}^2\}$, where the length of purchasing time k is seller-specific. For now, assume that both suppliers come from the same base population and that they face the same enforcement probabilities. That is $\hat{\theta}^1 = \hat{\theta}^2$ and $\lambda^1 = \lambda^2$. This implies that differential beliefs about the supplier types can only arise if the importer has been buying from the two firms for a different number of periods, i.e. $k_1 \neq k_2$.

Suppose that the importer has a longer relationship with firm 1, meaning $k_1 > k_2$. It follows directly that $\theta^1 > \theta^2$, i.e. the buyer has a better opinion about seller 1's reliability. Suppose also, that seller 2 has a better technology that allows her to produce the product at a lower production cost $c^2 < c^1$. In this simple case we can make the following prediction:

Prediction 1 *For sufficiently large δ and fixed q , an importer may buy from a higher cost exporter if she has a longer relationship with that firm than with an alternative lower cost supplier.*

This prediction is quite straightforward. First, denote the likelihood of delivery by $\tilde{\theta}^s = \theta^s + (1 - \theta^s)\lambda$. Now, consider the limiting case of $\delta \rightarrow 1$. Then, the net present value of future profits collapse to the static one-period profits π_0^s . In that case, based on equation (2) an importer buys from the higher cost supplier 1 if:

$$\Delta c q = (c^1 - c^2) q < R(q) \left[\tilde{\theta}^1 - \tilde{\theta}^2 \right] \quad (4)$$

If production cost differences Δc are not too large, the familiarity effect ($\tilde{\theta}^1 > \tilde{\theta}^2$) dominates and the importer buys from the better known firm.

Searching for Suppliers We now turn to the dynamic aspects of searching for suppliers. As we saw in Prediction 1, a firm may decide to buy from a higher cost supplier in the short run. However, the importer may decide to try out a new supplier and keep ordering from it for a while to see whether she is reliable.

A challenge in analyzing this question is how much an importer has to order to make the supplier reveal its type. Given the dynamic nature of the relationship, we should expect such a constraint to be related to the maximum growth rate in the ordered quantities over time. More precisely, the discounted present value of future gains from trade has to be dominated by the one period deviation payoff for the myopic suppliers. It is straightforward to see that for sufficiently high discount rates of the myopic firms, this constraint can be arbitrarily weakened. Instead of analyzing this aspect explicitly, in the following, we assume that it is sufficient to order a very small amount of ϵ from the supplier to test its reliability. However, finding a new trading partner is costly. Whenever a firm wants to test a new supplier, it has to pay a fixed cost f_N . This could be a pure search cost. An alternative interpretation would be that this cost captures the fact that firms actually have to order more than ϵ in order to trigger defaults from myopic firms.

The updated single-period profit equation from purchasing from a supplier s is thus:

$$\pi^s = \max_{q \geq \epsilon} \tilde{\theta}^s R(q) - c^s q - f_N \mathbb{1}[s = \text{new}] \quad (5)$$

There are three key aspects to determining which supplier an importer prefers:

- The dynamic reputation of seller s , $\tilde{\theta}^s$, compared to other sellers s' .
- The cost paid to seller s , c^s , compared to other sellers s' .
- The fixed cost of buying from a seller s if they have not been used before, f_N .

Which suppliers will a firm drop over time and which will it keep? Note that an importer does not automatically drop a new supplier with a higher marginal cost than the current one. If the exogenous shock δ is sufficiently high and the baseline share of patient suppliers sufficiently low, it may well be worth it to have a reserve supplier even if that firm is less efficient. We have shown in the Section 3 that, consistent with this intuition, many firms use multiple exporters for their products. It also speaks to the level of persistence in relationships over time which we find in the data:

Prediction 2 *Relationship persistence between buyers and suppliers can be high, even with price dispersion among export choices.*

This prediction still needs to be fully quantified to assess how much persistence can be generated by this mechanism for reasonable values of δ and marginal cost dispersions. Note that, to some extent, this is a corollary to Prediction 1. Trade with low-cost unknown parties may not be profitable. If learning is slow, the exogenous death rate δ is high and there is a substantial cost of trying out a new supplier, it may not be worthwhile to change from a familiar source to a new one.

Consider two suppliers, one with whom the importer had at positive experiences in the past and one who is new and therefore completely unfamiliar. Then, ceteris paribus, given the Bayesian updating, the importer should buy from the familiar exporter rather than the new one. This is captured in the next prediction:

Prediction 3 *Switching is likely to occur to partners that are already known through prior purchases.*

Additionally, the model predicts the growth patterns of new relationships:

Prediction 4 *New relationships are likely to start small, then grow faster, compared to preexisting relationships.*

This prediction is the same as in Araujo et al. (2012), in their case for exporters learning about importers. There are two potential explanations. First, following the logic in Araujo et al. (2012), the importer may choose optimal quantities based on her belief about a specific exporter. This leads to an increasing path of purchases over time. Alternatively, the importer may keep buying from a familiar source while buying small amounts from

a new firm until enough information has been gathered. A theoretical challenge in the second approach is how to pin down how much an importer needs to order from a supplier to generate learning. Above, we simply assume some ϵ quantity is necessary in order to start or maintain the learning process. That is, we abstract away from optimal import share calculations on behalf of the buyer.

Separately from our findings discussed above, the simple model discussed here lends itself to two additional tests. It is informative to consider a case where the quality of legal institutions λ can vary by country. As can be seen from equation (1), the speed of learning decreases in the strength of contract enforcement. This is quite intuitive. Myopic suppliers can only deviate when contract enforcement fails, which happens more often in countries with bad legal institutions. We should therefore expect firms importing from exporting countries with better institutions to have more persistent relationships and to switch less.

Prediction 5 *The share of firms switching should be lower in countries with better institutions.*

This prediction is also a byproduct of the model designed in Araujo et al. (2012), though they do not have dual-sided firm data in order to calculate the share of switching.

Secondly, the model implies that the more partners an importer uses overall, the more likely a switch becomes. We have shown this to be the case in the simple case with two suppliers, whereby switching is more likely if the importer has ongoing relationships with each supplier. This result should be easily generalizable to the case of more suppliers.

4.4 Multi-Product Importers

We next consider multi-product importers. Let us go back to the case of buying from a single supplier. When would an importer buy multiple products from the same firm? To study this problem, we follow the multi-product firm literature by assuming that every producer has a core product. Adding additional products moves the firm away from its core competency and therefore increases production costs. Assume therefore that additional products have higher marginal costs by factor $\gamma > 1$. Further, assume

that learning about a supplier's type happens at the firm and not at the product level. Therefore, buying multiple products does not increase the speed of learning. Assume also that now there is an additional fixed cost f_p that has to be paid per product bought. Under these assumptions, it is straightforward to calculate profits from buying product n of a supplier as: Profits from product n are thus:

$$\pi_s(n) = (\theta_k + \lambda(1 - \theta_k)) R(q) - c\gamma^n q - f_p \quad (6)$$

Product $n > 1$ is bought whenever the current profit term is greater than 0. This can be solved for.

$$\theta_k > \frac{1}{1 - \lambda} \left[\frac{c\gamma^n q + f_p}{R(q)} - \lambda \right] \quad (7)$$

We can also solve equation (6) for the number of products sold. This delivers:

$$n = \frac{\ln \left[\tilde{\theta}_k R(q) - f_p/cq \right]}{\ln \gamma} \quad (8)$$

Prediction 6 *The more familiar an importer is with a supplier, the more products she buys from that firm ($\partial n/\partial \theta > 0$).*

4.5 Simulations

In this subsection, we demonstrate some of the features of the model with a numerical example. Specifically, we solve the model under CES demand for the final good, and simulate trajectories for profits, prices, and other key variables in the model. This allows us to demonstrate the prediction results more clearly.

As above, expected importer profits from using any seller s at time t are:

$$\pi_t^s = (\theta_k^s + (1 - \theta_k^s) \lambda) R(q) - c^s q - f$$

where k is the number of periods that the buyer has been buying from supplier s by time t .

We again use $\tilde{\theta}^s = (\theta_k^s + (1 - \theta_k^s) \lambda)$ to save notation. Analysis of the problem while allowing for CES demand ($q = A (p^s)^{-\sigma}$) for the final good of the producer is straightforward. Assume there are two potential sellers to choose from, with costs to the buyer c_1 and c_2 . The optimal price depends on which supplier s is used:

$$p_t^s = \frac{\sigma}{\sigma - 1} \frac{1}{\tilde{\theta}^s} c^s$$

This means that revenue from supplier s at time t is $R_t^s = A (p_t^s)^{1-\sigma}$, and profits are:

$$\pi_t^s = \frac{1}{\sigma} \tilde{\theta} A \left[\frac{\sigma}{\sigma - 1} \frac{1}{\tilde{\theta}^s} c^s \right]^{1-\sigma} - f \quad (9)$$

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

Setting $\lambda = 0.6$, the share of good sellers $\hat{\theta}=0.6$, 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 $\tilde{\theta}$ 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, increasing revenue at a faster pace. 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.

Furthermore, we can justify the purchases of more products from a seller using our framework of multi-product importers. Again as before, we have the profits from indi-

vidual product n from seller s as

$$\pi^s(n) = (\theta_k + \lambda(1 - \theta_k)) R(q) - c\gamma^n q - f_p$$

with products only being bought that satisfy the condition in Equation (7). We set the marginal cost of an extra product at $\gamma = 1.02$. Figure 4 demonstrates the evolution of the number of products purchased over time, again with more expensive seller initially selling more products to the buyer, up until the reputation of the cheaper seller improves enough.

5 Empirics

In this section, we describe the tests we undertake for testing the model above, using the LFTTD data described in Section 2.

First, we test whether source countries with better institutions tend to have a higher share of maintained relationships over time. Guided by the model, our estimating equation is:

$$ShareStay_c = \alpha + \beta_1 \lambda_c + \beta_2 PCGDP_c + \nu_c \tag{10}$$

Based on the model, we would expect a greater fraction of U.S. import relationships to persist in countries with better legal institutions λ , due to better enforcement of contracts and the low rate of learning that takes place in sources with better institutional quality. The variable *ShareStay* is the fraction of importer-exporter relationships that are maintained between the U.S. and from country c averaged over all the years of our sampling frame, while *PCGDP* is log per capita GDP in country c . We also include private credit coverage as an additional regressor. 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.¹¹ As can be seen in the top part of Table 3, the results are consistent with the predictions of the model: higher legal rights and fewer procedures in exporting are both indicative of a larger share of firms remaining with their partner over time.

In line with the results about familiarity, we run the same specification replacing the dependent variable with the share of firms that obtain new products from familiar exporters at country level. This variable is the fraction of all new products coming from a certain country that arrive from familiar partners. From the bottom panel of Table 3, we see again that better institutions lead to a greater share of new products being sourced from familiar partners.

It is also possible to test the extent to which switching is linked to the number of partners used. There are two avenues through which familiarity might be an important determinant of the decision. First, an exporter could be familiar as a minority partner for the same product. To explore this channel, we test whether having more partners within a firm-HS10 product code combination is correlated with whether a minority-to-majority switch occurs. We also include total size of imports in each firm-product combination as a regressor. The results in Table 4 confirm that having more partners for a product indeed means importers are more likely to switch. The second avenue for familiarity is buying a product from an export partner previously used for some other product. Here, we test whether a switch came from a familiar firm against the total number of partners used by a firm (rather than an firm-HS10 product combination). Table 5 confirms that these types of switches are indeed more likely to occur among firms that have more overall firm

¹¹The results are the same if we use yearly measures of *ShareStay* for each country c with individual year observations of *PCGDP* and λ .

partners. As above, firms with more overall partners are more likely to switch. However, including either by including total firm imports directly or splitting firms into size deciles, we see that given the same number of partners, larger-volume importers are actually more likely to remain with their partner over time.

6 Conclusions

This paper employs rich data on U.S. imports to analyze importer-exporter relationships. It presents a set of new stylized facts and develops a model, building on Araujo et al. (2012), to clarify the different mechanisms at work 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 identified several factors that affect the length of relationships, most importantly product characteristics, the size of total imports and the quality of legal institutions in the source country. All findings are in line with a model of importer-exporter relationships where learning about the reliability of the trading partner represents a central aspect. A more comprehensive quantification of the learning mechanism against alternative channels will constitute the main next step of our analysis.


References

- Albornoz, Facundo, Héctor F Calvo Pardo, Gregory Corcos, and Emanuel Ornelas**, “Sequential exporting,” *Journal of International Economics*, 2012, 88 (1), 17–31.
- 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,” Technical Report, 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**, “Multiple-product firms and product switching,” *American Economic Review*, 2010, 100 (1), 70–97.
- , – , and – , “Multiproduct Firms and Trade Liberalization,” *The Quarterly Journal of Economics*, 2011, 126 (3), 1271–1318.
- 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,” *Mimeo*, 2014.
- 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, Census Center for Economic Studies 2012.
- , – , and **Ryan Monarch**, “Identifying foreign suppliers in U.S. merchandise import transactions,” 2015. U.S. Census Center for Economic Studies, mimeo.

- 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.
- Monarch, Ryan**, “It’s Not You, It’s Me: Breakups in U.S.-China Trade Relationships,” Working Paper 14-08, Census Center for Economic Studies 2014.
- Pierce, Justin R and Peter K Schott**, “The surprisingly swift decline of US 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: Kernel Density, Number of Imported Products



Note: The graph shows the density of U.S. firms for each number of total HS 10 products imported. Product codes are taken from customs declarations by U.S. importers, and adjusted using the methodology described in Pierce and Schott (2009). The density is cut off for readability above 78 products, accounting for 99% of the total sample. The kernel density is computed using 1000 grid points within this range. 10% of U.S. importing firms import more than 16 HS 10 products.

Figure 2: Age shares of importer-exporter relationships






Figure 3: Model Simulations- Single Product Importers


Panel A: Reputation Information




Panel C: Revenue



Panel B: Final Good Prices

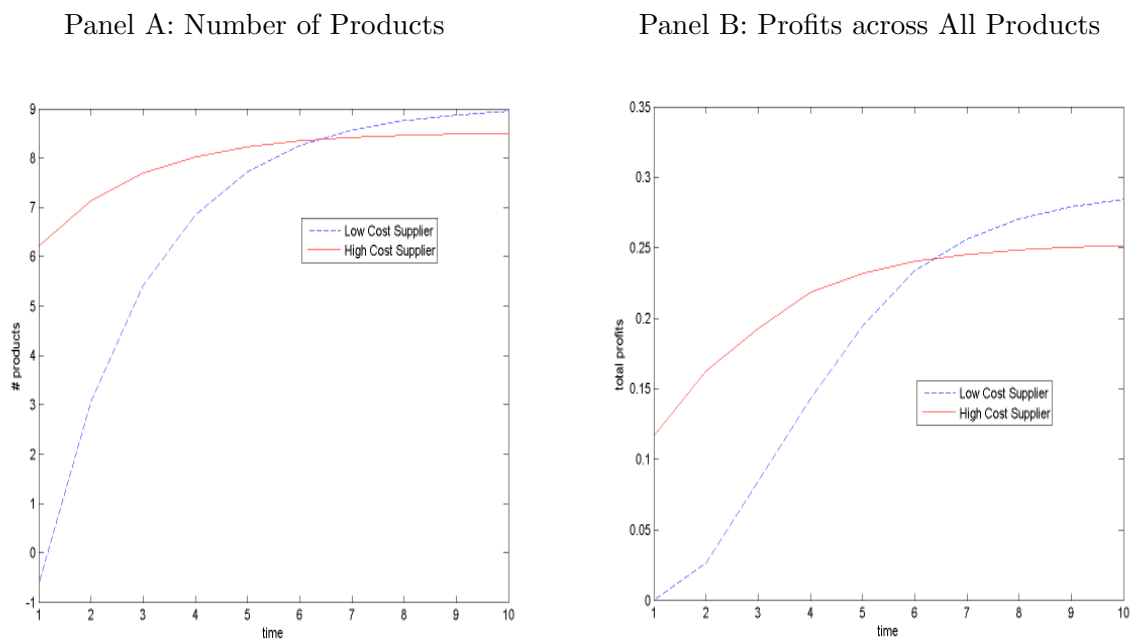


Panel D: Profits



Note: These simulations are for single product importers facing CES demand and having profits according to Equation (9) in Section 4.5. The choice is which supplier to use, a high-cost seller with a better starting reputation (solid line) or a low-cost seller with a less reputation (dotted line).

Figure 4: Model Simulations- Multiple Product Importers



Note: These simulations are for multiple product importers facing CES demand and having profits according to Equation (9) in Section 4.5. The choice is which supplier to use, a high-cost seller with a better starting reputation (solid line) or a low-cost seller with a lesser reputation (dotted line), as well as how many products to buy from each.

B Tables

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

	New	1-2 Years	3 or More Years
Share of Relationships
Share of Trade

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
Non-Differentiated

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: Source Country

	% of Relationships			% of Trade		
	New	1-2 Years	3+ Years	New	1-2 Years	3+ Years
Venezuela
China
Mexico
Vietnam
Germany
U.K.
Canada
Malaysia

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.

Panel C: U.S. Firm Size

	% of Relationships			% of Trade		
	New	1-2 Years	3+ Years	New	1-2 Years	3+ Years
Small
Medium
Large

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.

Table 3: Relationship between Institutions and Staying/ Switching Decisions*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*** (0.00740)	0.01893*** (0.00743)	0.01677** (0.00747)
Constant	0.16496** (0.05531)	0.52975** (0.11278)	0.37517*** (0.07171)
N	151	152	152
R ²	0.14	0.18	0.20

Dependent Variable: Share of Switches To Exporters Used for other Products, 2002-2008

	(4)	(5)	(6)
Log Strength of Legal Rights	0.06471** (0.02438)		
Procedures to Enforce a Contract		-0.00539** (0.00217)	
Rank (Procedures)			-0.00078** (0.00031)
Log Per Capita GDP	0.0264*** (0.00808)	0.02458*** (0.00836)	0.0237*** (0.00846)
Constant	-0.01302 (0.06102)	0.32657** (0.12559)	0.19129** (0.08084)
N	149	150	150
R ²	0.16	0.16	0.16

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: Staying/Switching Decisions, using Firm-Product Characteristics

	Stay (1)	Stay (2)	Switch to Minority (3)	Switch to Minority (4)
Total Partners	-0.00388*** (0.00014)		0.00163*** (0.00001)	
Log Total Partners		-0.19332*** (0.00132)		0.15586*** (0.00169)
Log Importer Size	0.02119*** (0.00044)	0.06504*** (0.00055)	0.08044*** (0.00081)	0.03624*** (0.00054)
Constant	0.26771*** (0.00467)	-0.7735*** (0.00568)	-0.60280*** (0.00857)	-0.26456*** (0.00621)
Year Fixed Effects	Yes	Yes	Yes	Yes
N	3,108,201	3,108,201	1,592,852	1,592,852
R ²	0.06	0.11	0.17	0.23

Notes: Three asterisks implies significance at 1%, two asterisks implies significance at 5%.

Table 5: Staying/Switching Decisions, using Firm Characteristics

	Stay (1)	Stay (2)	Other Product Switch (3)
Log Firm Total Partners	-0.07676*** (0.00074)	-0.05576*** (0.00045)	0.01910*** (0.00057)
Log Importer Firm Size	0.02870*** (0.00062)		0.00017 (0.00055)
Firm Size Deciles			
2		0.02097*** (0.00513)	
3		0.03723*** (0.00558)	
4		0.04520*** (0.00518)	
5		0.05996*** (0.00550)	
6		0.07523*** (0.00556)	
7		0.08581*** (0.00601)	
8		0.10629*** (0.00628)	
9		0.12123*** (0.00665)	
10		0.15504*** (0.00691)	
Constant	0.37481*** (0.00709)	0.17077*** (0.00615)	0.59676*** (0.00665)
Year Fixed Effects	Yes	Yes	Yes
N	3,108,201	3,108,201	1,592,852
R ²	0.06	0.11	0.17

Notes: Three asterisks implies significance at 1%, two asterisks implies significance at 5%.