

Learning and the Value of Trade Relationships*

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

March 23, 2016

Abstract

How valuable are long-term supplier relationships? To address this question, this paper explores relationships between U.S. importers and their suppliers abroad. We establish several facts: almost half of U.S. imports involve relationships three years or older, relationship survival and traded quantity increase as a relationship ages, and long-term relationships were more resilient in the 2008-09 financial crisis. We present a model of importer learning and calibrate it to detailed panel data on U.S. importer relationships. Model estimates imply large differences in the value of relationships across countries. Counterfactuals show that relationship-level responses to shocks are central to aggregate trade dynamics.

Keywords: International Trade, Firm Relationships, Learning, Institutions

JEL-Codes: F11, F14, L14, D22

*We thank Bernardo Blum, Fabrice Defever, and Olga Timoshenko for helpful discussions, as well as seminar participants at the Federal Reserve Board, the World Bank, the 2015 Rocky Mountain Empirical Trade conference, the 2015 Society for Economic Dynamics conference, the Washington Area International Trade Symposium, the European Bank for Reconstruction and Development, Nottingham University, Stockholm University, Uppsala University, UQAM, University of Montreal, Carnegie Mellon University, the Southern Economic Association meetings, the Stanford RDC Conference, and the Federal Reserve Board International Trade Conference for helpful comments. This work was undertaken while Monarch was 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.

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

[‡]Federal Reserve Board. Contact: t.schmidteisenlohr@gmail.com

1 Introduction

Every arm's-length shipment depends on a relationship between a firm and its supplier. But not all relationships are equal; many die quickly while others survive, grow and last for decades. Since long-term relationships likely differ in their contributions to the level and dynamics of economics activity, it is important to understand their prevalence, their size and how they develop. Panel data on relationships between importers and exporters present a unique opportunity to explore these issues in detail.¹

Exploiting the universe of U.S. import data, this paper establishes a set of facts about long-term relationships and quantifies their value to firms. Almost half of all U.S. arms-length imports are concentrated in importer-exporter relationships that are three years or older. Strikingly, the distribution of relationship ages differs widely across countries. While some countries (like Spain and China) deliver most of their U.S. exports through short-term relationships, others (like Japan and Germany) send much more through long-term relationships.² In addition, as relationships age, the amount that is traded increases, as does the likelihood that the relationship will survive another year. Long-term relationships are also more resilient in the face of aggregate shocks: Although new relationship formation dropped significantly during the Great Recession of 2008-09, the probability that a long-term relationship survived was unaffected.

These findings— especially those showing increased survival and trade flows for older relationships— inspire a model where importers learn about their foreign suppliers over time. An importer has an initial belief about the probability they will receive a usable product from a new supplier, based on two country-specific parameters: the share of reliable suppliers and the quality of contract enforcement. Through repeated interaction, the probability of receiving low

¹For recent work, see, for example, Eaton et al. (2014), Macchiavello and Morjaria (2015), and Antras and Foley (2015). The former studied trade between Colombia and the U.S. in a search model. The latter two papers analyzed learning within relationships in two particular cases: Kenyan rose exports and a U.S. food exporter, respectively.

²More generally, there are substantial differences between countries in the distributions of relationships across age cohorts, in terms of both the number of relationships and the value of trade, as we show in detail below.

quality inputs from one's supplier declines, leading to an increase in the quantity traded and a greater probability of survival. A key element of the model is that the speed of learning can differ with country characteristics: Learning occurs much faster in a successful transaction with a supplier in a country without good contract enforcement, or with a low share of reliable suppliers. From the model, we derive an intuitive expression for quantifying the value of a relationship, the *inherent value*: What would an importer starting a new relationship pay to have the accumulated partner-specific knowledge embodied in a long-term relationship?

The model is calibrated to U.S. relationship-level import data. We estimate separate parameters for each of the top 20 trading partners of the U.S., using moments from each country's distributions of relationships and trade across age cohorts. The parsimonious model generates age distributions that closely match those in the data. Furthermore, the obtained parameters capturing contract enforcement line up well with a comparable measure from the World Bank.

With these estimated parameters that govern the survival and growth of relationships, we can calculate the value of relationships for different countries. We find that long-term relationships are, on average, 7.2 times more valuable than new relationships, with a maximum of 13.8 (Spain) and a minimum of 3.0 (Philippines). By examining relationship dynamics, we find that one year of experience with a Chinese supplier is much more valuable than one year with a German supplier. Because relationships with Chinese suppliers often fail in the first year, the learning boost from a single successful transaction is large. However, the value of an older relationship is higher with a German supplier, as there is a higher long-run survival probability for relationships with German suppliers than Chinese.

Another way to quantify the value of relationships is to show how different trade flows would be if there were no long-term relationships. Our experiment projects the response of trade flows to a reset of trading relationships, meaning all accumulated knowledge about suppliers is wiped out. Trade first declines sharply and then moves back to steady-state levels only slowly, taking several

years. The effect of the shock differs across countries: resetting all relationships is more costly for Japan than for China because, in Japan, long-term relationships have higher survival probabilities and therefore contribute more to overall trade. Across all countries in our model, the total trade loss in the transition back to a steady state averages 1.5 times the annual steady-state value, with the Philippines losing the least (0.17) and Mexico losing the most (2.45).

We perform two additional experiments to show how trade reacts to other types of shocks. Our second counterfactual demonstrates the effects of a one-time increase in the number of new relationships for each source country. In the case of China, this provides a temporary boost to trade that quickly falls off in later years. In contrast, more new U.S.-Japan relationships lead to a protracted increase in trade over the course of a few years. This is because relationships die less quickly in Japan, so for some years the increase in trade from learning dominates the reduction in trade from relationship attrition. Our third experiment quantifies the role of relationships in the 2008-09 trade collapse, in which new relationship formation declined dramatically. By feeding in observed relationship creation, we illustrate its dynamic effects on trade volumes. Fewer younger relationships translate into fewer long-term (and high-volume) relationships in later years— a tilting of the age distribution toward younger relationships. This implies that imports from countries like China, where relationships die more quickly, fall more at the time of the shock. For the same reason, however, trade recovers faster in those countries.

There is a burgeoning literature that, like our paper, uses “two-sided” international trade data to study relationships. Blum et al. (2013) use linked data on importers and exporters in Latin American to demonstrate that firms that export only occasionally often export the same goods to the same importers across multiple exporting spells. Eaton et al. (2014) study relationships between Colombian exporters and U.S. importers. They calibrate a search and matching model with learning to match exporter decisions, including sales, number of clients, and transition probabilities. Kamal and Sundaram (2013) use the same data to determine how likely Bangladeshi textile producers are

to follow other exporters in the same city in exporting to a particular partner. Monarch (2014) finds that most of the time, U.S. importers remain with their Chinese exporting partners. Two-sided trade data is also used to study the effects of firm heterogeneity on trade: Bernard et al. (2014) use Norwegian trade data to develop a model of relationship-specific fixed costs of exporting. Carballo et al. (2013) look at relationships in several Latin American countries and develop a model to analyze the role of competition. Heise (2015) studies the effects that firm-to-firm relationships have on price rigidity and exchange-rate pass-through using U.S. importer data and finds that prices grow within a relationship as trade increases. We complement this literature by focusing on cross-country heterogeneity in relationship patterns and using it to quantify the value of relationships.

A set of papers directly studied the role of long-term relationships for international trade with more disaggregated data.³ Egan and Mody (1992) provide survey evidence that importers from developed countries initiate emerging-economy trade relationships with very small purchases. Rauch and Watson (2003) rationalize this finding in a model with importer learning. In their model, relationship persistence is the result of a dynamic tradeoff between the per-period costs of a supplier and the reliability of other potential partners. Macchiavello and Morjaria (2015) study Kenyan rose exporters and find that the value of a relationship increases with age. They also show that in long-run relationships, buyers have already learned the type of the seller, and therefore no costly signaling is necessary in times of crisis. Our project augments this literature by using the universe of U.S. arm's-length transactions.

Our model directly builds on Araujo et al. (2012), who study learning by exporters and test reduced-form predictions with Belgian firm-level data.⁴ In the

³Earlier work on buyer-supplier relationships in international trade centered on the study of networks: Rauch (2001) surveys the potential for transnational cultural networks to reduce barriers to entry. Rauch and Watson (2004) present a model where economic agents use their networks to produce or export more efficiently or to become an intermediary. Krautheim (2012) models how information-sharing in networks may affect the fixed costs of exporting and thereby the measured effects of distance on trade. Chaney (2014) studies a dynamic model where firms search for additional trade opportunities through their network.

⁴Our work is also related to the wider literature on learning and entry. Impullitti et al. (2013) develop a model of entry and exit, where firms pay sunk costs to start exporting and

tradition of the seminal work of Jovanovic (1982), a number of other papers consider firm learning in different contexts, such as its relationship with exports (Albornoz et al. (2012)), foreign demand (Ruhl and Willis (2014)), growth (Arkolakis et al. (2015)), and international prices (Bastos et al. (2015)).

Finally, our work adds to the wider literature on trade dynamics. Alvarez and Lucas (2007) develop a dynamic solution to the general-equilibrium trade model of Eaton and Kortum (2002), while Eaton et al. (2011) embed this framework in an international real business cycle model. Several papers study how firm behavior shapes aggregate trade dynamics. Ruhl (2008) looks at the extensive margin of trade and its implications for trade elasticities. Atkeson and Burstein (2010) look at the dynamic product and process innovation decision of exporters. Alessandria and Choi (2014) show that long-term gains of tariff cuts are larger when not only sunk costs of exporting but also continuation costs are taken into account. Alessandria et al. (2014) consider the sluggish response of trade volumes to changes in trade barriers or price shifts by modelling the buildup of a firm's exporting capacity over time.

In sum, our paper contributes to the literature in three ways. First, it provides facts on long-term relationships over time and across countries using the universe of U.S. import data. Second, it extends the model in Araujo et al. (2012) to study how relationships in different countries mature, and how they drive trade flows. Finally, we quantify the value of relationships across countries.

The rest of the paper is organized as follows. Section 2 describes the importer-exporter data. Section 3 presents empirical findings on trade relationships. Section 4 discusses the model and presents testable predictions. Section 5 describes the calibration exercise and counterfactuals. Section 6 concludes.

face persistent productivity shocks. Besedeš and Prusa (2006) study the effect of product differentiation on U.S. importer relationships with their suppliers. Besedeš (2008) 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. Timoshenko (2015a) looks at Colombian firm-level data to distinguish between the role of sunk costs and learning for the persistence of exporting. Timoshenko (2015b) finds that new exporters do more product switching and develops a model of learning that rationalizes this finding.

2 Data

The data come from the Longitudinal Foreign Trade and Transaction Database (LFTTD), which is collected by U.S. Customs and Border Protection and maintained by the U.S. Census Bureau. Every transaction in which a U.S. company imports or exports 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.⁵ 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, Harmonized System (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.⁶ Monarch (2014) and Kamal et al. (2015) find substantial support for the use of the MID as a reliable, unique identifier, both over time and in cross-section. Pierce and Schott (2012), Kamal and Sundaram (2013), Eaton et al. (2014), and Heise (2015) have all used this variable in the context of studying U.S. firm relationships in international trade.

For our analysis, we eliminate related-party transactions, as U.S. firms who are importing from foreign affiliates will likely have very different relationship dynamics than those involved in arm's-length transactions. U.S. importers whose domestic operations are classified as wholesale or retail are also dropped. We follow the methods of Bernard et al. (2009) 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. All of our results come from the U.S. import data on relationships from 1997 through 2011 (the most recent year of data availability).

⁵Approximately 80 to 85 percent of these customs forms are filled out electronically (Krizan (2012)).

⁶Specifically, 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.

Finally, some definitions: an **importer** is a U.S. importing firm, while an **exporter** is a non-U.S. firm identified by the MID as exporting to the U.S. A **relationship** is an observation of an importer-exporter combination. The **age** of a relationship is how many consecutive years that relationship has appeared in the U.S. import data, with the first observation of a relationship considered to be age 0.⁷ We call the distribution of relationship counts over different ages the **count distribution**, and the distribution of trade values across relationships of different ages the **value distribution**.

3 Relationships in International Trade

In this section, we present findings from our relationship-level data. We first show that long-term relationships are a meaningful component of international trade. Almost half of U.S. imports occur within older relationships, with distributions for individual source countries varying greatly. As a relationship ages, trade increases and the relationship itself becomes more likely to survive an additional year. Also, older relationships proved more resilient in the 2008-09 Great Recession. Finally, we show that long-term relationships are more likely to arise in countries with good institutions, higher per-capita GDP, or that are members of the Organization for Economic Cooperation and Development (OECD), as well as among larger exporters and older importers.

Most Trade is in Long-Term Relationships Table 1 presents a breakdown of U.S. arm's-length imports in 2011 by the age of relationships. Two points stand out. First, the largest fraction of trade occurs in long-term relationships—those that have four or more consecutive transactions—while new relationships account for only about one-fifth of U.S. imports. At the same time, new relationships account for the vast majority of total relationships in 2011.

⁷The distinction between consecutive and non-consecutive years of a relationship makes little difference to any of the findings below. Results using non-consecutive years of a relationship are available upon request.

The complete distributions for 2011 are shown in Figure 1: both the count distribution and the value distribution decrease nearly monotonically with age.⁸ In line with Table 1, the value distribution is much more skewed toward higher ages.

These facts mask heterogeneity in the age structure of trade at the source-country level. Table 2 presents a summary of the count distributions and value distributions for the top 20 trading partners of the U.S. There is a wide range for the share of long-term trade across countries: 33 percent of U.S. arm's-length imports from Spain occur in long-term relationships, but 67 percent of imports from Taiwan do. Taiwan also has roughly double the share of long-term relationships as Spain. There is a rough positive correlation between the share of long-term relationships and the share of trade in long-term relationships, as shown in Figure 2. Later in this section, we address the question of why certain countries tilt toward long-term relationships while others tilt towards shorter relationships.

Trade Increases as Relationships Age The differences between the count distributions and the value distributions imply that, on average, older relationships trade more. We find that one explanation for this is that the amount traded within a surviving relationship increases over time.

To show this, we generate cohorts of newly formed relationships for each year from 1998 to 2006 and follow them over time.⁹ Table 3 presents the results of the following regression with relationship fixed effects:

$$Value_{mx,t} = \sum_{k=1}^K \beta_{A,k} \mathbb{1}[Age_{mx,t} = k] + \beta_m Y_m + \beta_x Y_x + f_{mx} + u_{mx,t} \quad (1)$$

where $Value$ is the log value of trade between a U.S. importer m and an exporter x at time t , Age is the age of that relationship at time t (and Age_0 - the first

⁸Any relationship 11 years or older is classified as 11 years old. The uptick in traded value for relationships for the final segment is explained by this simplification; extending the maximum age of a relationship forward more years smoothes the tail of the distribution.

⁹A new relationship in 1998 is one not found in 1997, our first year of data. The most recent year of data is 2011.

year a relationship is found- is omitted), and f_{mx} is a relationship fixed effect. Y_m and Y_x are sets of importer or exporter controls, which include the following variables: firm size (proxied by total imports or exports), the number of HS10 products traded, and the total number of relationships.¹⁰ The regression is run separately for relationships that last a total of $K = 5, 7, 10,$ and 13 years.

Trade increases as a relationship ages. As the relationship moves close to its end, trade falls off, but for relationships that last for (at least) the maximum number of measurable years, values essentially plateau at level higher much higher than the initial level of trade. The results are presented graphically in Figure 3 by graphing the $\beta_{A,k}$ terms for $K = 5, 7, 10,$ and 13-year relationships.

The Survival Probability Increases as Relationships Age We next explore the survival probability of a trading relationship. Figure 4 shows the conditional survival probability of a relationship by age: what share of relationships of age k in 2010 survive into 2011? This shows that the older a relationship is, the higher the likelihood that relationship survives an additional year.

We can also look for determinants of long-term relationships with formal survival analysis. To this end, we estimate a proportional hazard model that accounts for the effects of importer size, exporter size, and source country institutional quality.¹¹ As in the previous subsection, this requires following cohorts of new relationships over time, defining the hazard as the disappearance of the relationship.

Echoing the picture in Figure 4, Figure 5 shows that there is little decline in the survival function- the probability a relationship survives at least to age k - after age 4. This finding is not driven by *firm-level* entry and exit. When estimating the probability of relationship survival on the sample of firms existing for 3 years or 7 years (approximating a balanced panel of firms) we find qualitatively similar results for the survival model. Since we also control

¹⁰Heise (2015) performs a similar exercise using the LFTTD data.

¹¹We use the Rule of Law index from the World Bank's World Governance Indicators, which originated in the work of Kaufmann et al. (2010). This index reflects perceptions about contract enforcement, property rights, courts, and crime within a country.

for total importer and exporter firm size, firm-level trends are not driving these findings. Nor is it the case that relationships with larger initial trade volumes are driving these results; including first-year relationship size in the hazard model specification does not alter the finding that the probability of survival is lowest for the youngest relationships.

Long-Term Relationships Are More Resilient During the financial crisis of 2008-09, international trade plummeted. Besides the large decline in aggregate trade, the crisis also triggered changes to the number and the age structure of relationships.

The first line of Table 4 shows that the number of relationships declined: Compared to 2007, there were 3 percent fewer relationships in 2008 and an additional 7 percent fewer in 2009. Although both relationship formation and survival declined, the major decrease was in new relationship formation. In 2008, there were 7 percent fewer relationships formed than in 2007; in 2009, there were 15 percent fewer. By contrast, the number of continuing relationships was approximately constant from 2007 to 2009 (rising by 3 percent in 2008 and falling by the same amount in 2009).

Table 5 confirms that long-term relationships were much more likely to outlast the crisis. It presents the results of a regression of an indicator for 2008-09 and an interaction with institutional quality on country-age bin survival probabilities.¹² We see that the oldest relationships were no less likely to survive during the crisis years than before or after. Even though 3-5 year relationships were less likely to survive, these relationships still tended to persist if they were in countries with fairly good institutions. Thus, long-term relationships not only trade more and have higher survival probabilities, but are also better at withstanding external headwinds that destroy young relationships.

Determinants of Long-Term Relationships Having shown the importance of long-term relationships in international trade, we now consider the

¹²The regression includes age fixed effects. We again use all country-age-year observations in a year, as long as that country has more than 100 relationships.

determinants of relationship ages across countries. We do this in the context of reduced-form regressions that control for firm and product characteristics.

Our specification is as follows:

$$RelLength_{mx} = \beta_m Y_m + \beta_x Y_x + \beta_c Y_c + f_i + u_{mx}, \quad (2)$$

where *RelLength* is the age of a relationship in 2011 between U.S. importer *m* and exporter *x*. Y_m and Y_x are sets of importer or exporter controls, which include the following variables for importers and exporters: firm age (where firm “birth” is proxied by the first year of appearing in U.S. import data), firm size (proxied by total imports), number of relationships, and number of products traded. Regressions also include importer-industry fixed effects f_i .¹³ Y_c is a set of country controls that includes rule of law (again from Kaufmann et al. (2010)), per capita GDP, distance from the U.S., a trade agreement dummy, and an OECD membership dummy.

Columns (1) to (5) of Table 6 present the results. The regressions show that U.S. importers have longer relationships with exporters from countries with better institutions (higher rule of law). To address concerns about industry composition, we rerun the regressions restricting the sample to only machinery and electric technological products (HS 84-86) or to textile products (HS 50-63) and find very similar results (available upon request). Interestingly, despite its general importance for trade levels and patterns, distance is uncorrelated with relationship length. Similarly, Free Trade Agreements (FTAs) do not seem to affect the age of relationships. In contrast, source countries with higher GDP per capita tend to have longer relationships— an effect that disappears when including the OECD dummy.

An alternative specification is one in which source-country controls are replaced by source-country fixed effects, f_c , which is found in Column (6) of Table 6. Figure 6 plots these fixed effects against the rule of law measure. The two series line up well, providing further evidence for a positive relationship between

¹³ Since importers can import from more than one industry, we assign each importer the industry that it imports the most from.

rule of law and the length of relationships.

Guided by these facts, we develop a model of relationships and learning, calibrate its parameters, and analyze the value of relationships.

4 Model

This section outlines a model of learning in trade relationships. We derive predictions consistent with our within-relationship findings on trade flows and survival, as well as the country-level determinants of long-term relationships. The model directly builds on Araujo et al. (2012), adjusted to study learning by importers rather than exporters.

The section proceeds in five steps. First, it presents the basic setup and explains how learning works. Second, it derives results on relationship survival and the average age of relationships. Third, it introduces a constant elasticity of substitution (CES) demand structure and analyzes firm decisions. Fourth, it obtains expressions for the count and value distributions that are central to our calibration. Finally, it presents our measure for the value of a relationship.

4.1 Basic Setup and Learning

An importer wishing to initiate a trade relationship encounters an exporter. The importer has all bargaining power and offers the exporter a quantity-price pair. The exporter can accept or reject the offer. The importer needs to pay fraction α of the agreed total price in advance.¹⁴

If exporter produces the goods, they must decide how much effort to put into their production. When producing, the exporter either exerts full effort

¹⁴Survey- and country-level customs data show that open account— that is, payment after delivery— is the predominant payment form in international trade. This may well be consistent with our assumption: Even a small advance payment of 5 or 10 percent is sufficient to generate the results. Firms that pay 10 percent or less in advance would likely indicate that they are buying on open account. Alternatively, one could assume that verification of the quality of supplied inputs takes time. Then, even under full open account, the importer may pay the exporter before realizing the bad quality of a delivery.

or no effort. With effort, production costs are c_1 per unit of output, whereas without effort production costs are zero. Goods are only useful for the importer when the exporter exerted effort in their production. The importer needs to hire c_2 worker units to assemble one unit of inputs into a final good. Workers for assembly need to be hired in advance— that is, when the goods are ordered from the exporter and before input quality has been revealed.

The ability of an exporter to put zero effort into production depends on local contract enforcement. The better the enforcement, the harder it becomes to shirk. We model this by assuming that there is a country-specific quality of enforcement, λ . With probability λ , the exporter is forced to put in high effort. With probability $1 - \lambda$, the exporter has an opportunity to cheat and put zero effort. If cheating occurs, the importer ends the relationship. Relationships can also dissolve for exogenous reasons with probability $\delta \in (0, 1)$.

Assume that a fraction $\hat{\theta}$ of suppliers are reliable (meaning they always put forth high effort), whereas the remainder are myopic. As Araujo et al. (2012) do, we assume that differences in the discount rates are so large that patient suppliers always make effort, whereas myopic firms try to deviate from the contract and put no effort whenever they get an opportunity to do so.

Bayesian updating 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 fulfills the contract is equal to the population mean, $\hat{\theta}$. Every period that a relationship survives, buyers update their beliefs according to Bayes Rule. If a buyer has successfully purchased from a seller for k periods, the posterior probability that the seller is reliable can be derived as

$$\theta_k = \frac{\hat{\theta}}{\hat{\theta} + (1 - \hat{\theta}) \lambda^k}. \quad (3)$$

Importantly, the probability only changes with 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 reliable.

The delivery probability Consider again a relationship of age k . The buyer will receive the goods from the supplier next period under two scenarios: either the seller is reliable (an event with expected probability θ_k), or the seller is myopic but does not face any opportunity to produce with low effort (an event with expected probability $(1 - \theta_k) \lambda$). Thus after k successful transactions, the delivery probability is:

$$\tilde{\theta}_k = (\theta_k + (1 - \theta_k) \lambda) = \left(\frac{\hat{\theta}}{\hat{\theta} + (1 - \hat{\theta}) \lambda^k} \right) (1 - \lambda) + \lambda. \quad (4)$$

The delivery probability is a key object in the model. It is increasing in k , the age of the relationship, as well as in $\hat{\theta}$, the fraction of reliable suppliers in the source country.¹⁵

There are two competing effects of the contract enforcement parameter, λ , on the delivery probability. Better enforcement leads both to a higher initial delivery probability $\tilde{\theta}_0$, but also to slower learning, meaning lower subsequent delivery probabilities.¹⁶ The differences can be seen in Figure 7, which illustrates the delivery probability for both high and low λ . For young relationships, the direct effect of λ dominates, and the delivery probability is higher for higher λ . For older relationships, the negative effect on $\tilde{\theta}_k$ from slower learning is more important, generating a probability that is lower for higher λ .

4.2 Relationship Survival and the Average Relationship Age

A relationship of age k survives for another period if it is not hit by the exogenous dissolution shock and if there is a successful delivery. It therefore survives with (conditional) probability: $surv_k = (1 - \delta) \tilde{\theta}_k$. Using this expression and Equation (4), it is straightforward to see that, matching our empirical finding, as a relationship ages ($k \uparrow$), the conditional probability of survival increases.

¹⁵ $\frac{\partial}{\partial \hat{\theta}} \left(\frac{\hat{\theta}}{\hat{\theta} + (1 - \hat{\theta}) \lambda^k} \right) = \frac{\lambda^k}{(-\hat{\theta} \lambda + \hat{\theta} + \lambda)^2} > 0$.

¹⁶ $\tilde{\theta}_0 = \hat{\theta} + (1 - \hat{\theta}) \lambda$.

The unconditional probability that a relationship is still alive after k periods is the probability it has survived all previous periods: $alive_k = alive_{k-1} surv_{k-1}$. Since the probability of any potential relationship being alive at least 0 periods is 1 (that is, $alive_0 = 1$), this delivers

$$alive_k = (1 - \delta)^k \left(\lambda^k (1 - \hat{\theta}) + \hat{\theta} \right). \quad (5)$$

From this, we can derive the average relationship age as:

$$age = \frac{\sum_{k=0}^{\infty} alive_k k}{\sum_{k=0}^{\infty} alive_k} = \frac{\sum_{k=0}^{\infty} (1 - \delta)^k \left(\lambda^k (1 - \hat{\theta}) + \hat{\theta} \right) k}{\sum_{k=0}^{\infty} (1 - \delta)^k \left(\lambda^k (1 - \hat{\theta}) + \hat{\theta} \right)}$$

Echoing our empirical result that countries with better institutional quality tend to have older relationships, we can derive the following proposition:

Proposition 1 *A country's average relationship age increases in the share of reliable suppliers in a country $\hat{\theta}$, and in the enforcement probability λ , for $\lambda > 1 - \frac{\delta}{(1-\delta)\sqrt{\hat{\theta}}}$.*

Proof. See Appendix. ■

Trade relationships, on average, last longer if there is a larger fraction of patient firms to begin with. In addition, for sufficiently large λ , the average age also rises with the strength of enforcement. The relationship between these parameters and the model-based average age of a relationship is plotted in Figure 8.

4.3 Firm Optimization

Expected profits Taking into account the delivery probability, expected profits when buying from a supplier that an importer has traded with for k periods are

$$E[\pi_k] = \tilde{\theta}_k [R(q) - (1 - \alpha)T_k] - \alpha T_k - c_2 q,$$

where T_k is the agreed payment from the importer to the exporter. The buyer can assemble the final goods and sell them for revenue $R(q)$ if the intermediate inputs are of high quality. In that case, the buyer will also pay the remaining outstanding bill of $(1 - \alpha)T_k$. In any case, the importer will incur the prepayment to the importer, αT_k , and the prepayment to its own workers, $c_2 q$.

We assume that there are no financing costs and that patient suppliers do not discount the future, so advance payments are not valued more than post-delivery payments.¹⁷ The participation constraint of a reliable exporter (who always puts in effort) is given by $\alpha T_k + (1 - \alpha)T_k \geq c_1 q$.¹⁸ As the buyer has all bargaining power, this constraint binds in equilibrium, so that $T_k = T = c_1 q$. Plugging this back into expected profits delivers

$$E[\pi_k] = \tilde{\theta}_k [R(q) - (1 - \alpha)c_1 q] - \alpha c_1 q - c_2 q. \quad (6)$$

CES Demand In order to derive implications for trade revenues within a relationship, we need to introduce a demand structure. Assume that demand for the final good has the standard CES form: $q_{t,k} = A_t p_k^{-\sigma} = P_t^\sigma Q_t p_k^{-\sigma}$, with elasticity of substitution $\sigma > 1$.¹⁹ Expected importer profits in a relationship that has lasted k periods are given by Equation (6), where revenue at any time t is $p_k q_{t,k}$. Profit maximization implies the optimal price:

$$p_k = \frac{\sigma}{\sigma - 1} \left[(1 - \alpha)c_1 + \frac{\alpha c_1 + c_2}{\tilde{\theta}_k} \right]. \quad (7)$$

¹⁷It would be easy to relax either of these assumptions, but it would not add any value to the analysis. See Schmidt-Eisenlohr (2013) for a model of payment choices with positive interest rates.

¹⁸We assume that the buyer always offers a contract that is acceptable to patient and impatient sellers, so the patient-buyer participation constraint is the relevant one. In principle, there could be a separating case where the buyers offer a low total payment, T_K , and only attracts the impatient suppliers. However, as shown in Schmidt-Eisenlohr (2013), the pooling dominates under very weak conditions.

¹⁹Such a demand follows from maximizing a CES utility function of the form $Q_t = \left(\int_{\Omega} q(\omega)^{\frac{\sigma-1}{\sigma}} d\omega \right)^{\frac{\sigma}{\sigma-1}}$ subject to a budget constraint and an ideal price index $P_t = \left(\int_{\omega \in \Omega} p(\omega)^{1-\sigma} d\omega \right)^{\frac{1}{1-\sigma}}$.

This expression is quite intuitive. The importer only pays $(1 - \alpha)c_1$ when the delivered inputs are of high quality, so these costs affect the optimal price-setting in the standard way. αc_1 and c_2 are the costs of prepayment to the exporter and the local workers, respectively, that are incurred independently of the success of the transaction. Optimal price setting allocates these costs evenly across all successful transactions. So the lower $\tilde{\theta}_k$ is, the more these costs increase the price demanded from final consumers. Importers do standard markup pricing over the effective costs per unit of output, $(1 - \alpha)c_1 + \frac{\alpha c_1 + c_2}{\tilde{\theta}_k}$. This implies that a successful transaction generates revenues of

$$R_{t,k}(\tilde{\theta}_k) = p_k q_{t,k} = A_t \left(\frac{\sigma}{\sigma - 1} \right)^{1-\sigma} \left[(1 - \alpha)c_1 + \frac{\alpha c_1 + c_2}{\tilde{\theta}_k} \right]^{1-\sigma}, \quad (8)$$

with expected revenues given by $\tilde{\theta}_k R_{t,k}(\tilde{\theta}_k)$. We can now derive the following proposition that matches our empirical finding:

Proposition 2 *The amount being traded, $R_{t,k}(\tilde{\theta}_k)$, within a relationship increases with the relationship's age, k .*

Proof.

$$\frac{\partial}{\partial k} R_{t,k}(\tilde{\theta}_k) = A_t \left(\frac{\sigma}{\sigma - 1} \right)^{1-\sigma} (1 - \sigma) \left[(1 - \alpha)c_1 + \frac{\alpha c_1 + c_2}{\tilde{\theta}_k} \right]^{-\sigma} (\alpha c_1 + c_2) (-1) \tilde{\theta}_k^{-2} \frac{\partial}{\partial k} \tilde{\theta}_k > 0,$$

as $\tilde{\theta}_k = \left(\frac{\hat{\theta}}{\hat{\theta} + (1 - \hat{\theta})\lambda^k} \right) (1 - \lambda) + \lambda$ is increasing in k for $\lambda < 1$. ■

Price index and CES Utility For our counterfactuals, it is necessary to derive the ideal price index, P_t , and the aggregate CES consumption, Q_t , so we can study their response to shocks. Let $N_{t,k}$ be the mass of firms at time t with relationships of age k (note that this notation implies $N_{t,k} = surv_{k-1} N_{t-1,k-1}$). Furthermore, the price index at time t is an aggregation of prices paid for all goods successfully supplied at time t , taking into account that delivery is only

successful with probability $\tilde{\theta}_k$. Thus, we can derive the ideal price index as

$$P_t = \left(\int_{\omega \in \Omega} p(\omega)^{1-\sigma} d\omega \right)^{\frac{1}{1-\sigma}} = \left(\sum_{s=0}^{\infty} \tilde{\theta}_s N_{t,s} p_s^{1-\sigma} \right)^{\frac{1}{1-\sigma}}. \quad (9)$$

We can also derive the aggregate consumption from CES imports as

$$Q_t = \left(\int_{\Omega} q(\omega)^{\frac{\sigma-1}{\sigma}} d\omega \right)^{\frac{\sigma}{\sigma-1}} = \left(\sum_{s=0}^{\infty} \tilde{\theta}_s N_{t,s} q_s^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}. \quad (10)$$

4.4 Count and Value Distributions

We can now derive model analogues to the count and value distributions that we studied in Section 3. In this section, we drop the t subscript, as we focus on the steady-state distributions.²⁰ In a steady state, the number of relationships is constant. Let $ageshare_{k|K}$ denote the fraction of relationships that survived for k periods if no relationship can be active for more than K periods. Then,

$$\begin{aligned} ageshare_{k|K} &= alive_k / \sum_{s=0}^K alive_s = \frac{alive_k}{\sum_{s=0}^{\infty} alive_s - \sum_{s=K}^{\infty} alive_s} \\ &= \frac{(1-\delta)^k (\lambda^k (1-\hat{\theta}) + \hat{\theta})}{\frac{1-\hat{\theta}}{1-(1-\delta)\lambda} + \frac{\hat{\theta}}{\delta} - (1-\delta)^K \left[\lambda^K \frac{1-\hat{\theta}}{1-(1-\delta)\lambda} + \frac{\hat{\theta}}{\delta} \right]}. \end{aligned} \quad (11)$$

We can directly compare this object to the relationship age shares reported in Section 3. Next, we derive the elements of the value distribution. That is, we weigh the number of firms in a cohort with trade per firm:

$$tradeshare_{k|K} = (R_k alive_k) / \sum_{s=0}^K (R_s alive_s). \quad (12)$$

²⁰Later in the counterfactual section, we reintroduce the time dimension to study impulse responses from shocks to the age distribution.

Plugging in the expression for revenues (8), we obtain

$$tradeshare_{k|K} = \frac{\left[(1 - \alpha)c_1 + \left(\frac{\alpha c_1 + c_2}{\tilde{\theta}_k} \right) \right]^{1-\sigma} alive_k}{\sum_{s=0}^K \left[(1 - \alpha)c_1 + \left(\frac{\alpha c_1 + c_2}{\tilde{\theta}_s} \right) \right]^{1-\sigma} alive_s}. \quad (13)$$

Like expression (11), this expression is also a function of λ , $\hat{\theta}$, and δ , as well as the prepayment share, α , the elasticity of substitution, σ , and the costs, c_1 and c_2 .

4.5 The Inherent Value of a Relationship

How can the value of a relationship be measured? We propose the following calculation: compare the sum of expected profits from a new relationship with the sum of expected profits from a well-established relationship.²¹ In other words, how much would an importer entering into a new relationship pay to have the accumulated knowledge about a supplier inherent in a long-term relationship?

From Equations (6) and (8), expected profits in our framework are:

$$E[\pi_k] = \frac{\tilde{\theta}_k}{\sigma} R_k(\tilde{\theta}_k) = \frac{\tilde{\theta}_k}{\sigma} A \left(\frac{\sigma}{\sigma - 1} \right)^{1-\sigma} \left[(1 - \alpha)c_1 + \frac{\alpha c_1 + c_2}{\tilde{\theta}_k} \right]^{1-\sigma}. \quad (14)$$

Here, we see that the delivery probability, $\tilde{\theta}_k$, affects expected profits via two channels. First, there is a direct effect, as a sale is more likely to succeed. Second, there is an indirect effect: a higher success probability leads to a larger order and hence larger revenues and profits if the trade is successful.

Expected Profit Streams First, consider the standard case, where the importer knows perfectly that the supplier is patient (for example after many interactions). Then the importer knows that $\tilde{\theta}_k = 1$ for all periods, and the probability of a relationship being alive after k periods is simply $(1 - \delta)^k$. Ex-

²¹Note that the sum is not discounted as we assume that importers have a discount factor of one. None of our results would change if we introduced an additional discount factor to the model.

pected future profits are given by

$$\begin{aligned} E[\Pi^o] &= \sum_{k=0}^{\infty} (1-\delta)^k E[\pi_k] \\ &= \frac{A}{\delta\sigma} \left(\frac{\sigma}{\sigma-1}\right)^{1-\sigma} [c_1 + c_2]^{1-\sigma}. \end{aligned} \quad (15)$$

This corresponds to the standard sum of profits under CES in a frictionless world with exogenous death probability δ .

Compare this to the sum of expected profits from a new relationship:

$$\begin{aligned} E[\Pi^n] &= \sum_{k=0}^{\infty} \text{alive}_k E[\pi_k] \\ &= \frac{A}{\sigma} \left(\frac{\sigma}{\sigma-1}\right)^{1-\sigma} \sum_{k=0}^{\infty} (1-\delta)^k \left(\lambda^k (1-\hat{\theta}) + \hat{\theta}\right) \tilde{\theta}_k \left[(1-\alpha)c_1 + \frac{\alpha c_1 + c_2}{\tilde{\theta}_k}\right]^{1-\sigma}. \end{aligned} \quad (16)$$

Here, an importer is matched with a random supplier at $k=0$ that is patient with probability $\hat{\theta}$. The importer then slowly learns the type of the supplier over time.

The sum of profits when buying from a well-known, patient partner is always larger for two reasons. First, at every point in time, the relationship survival probability is higher with a patient supplier. Second, at every point in time, the importer orders more and generates larger expected revenues and profits.

How these two discounted profit streams depend on the parameters λ , $\hat{\theta}$, and δ is illustrated in Figure 9. Profits from a new relationship (the dashed lines) increase in λ (Panel A) and in $\hat{\theta}$ (Panel C) but decline in δ (Panel E). Importantly, within the range of λ and $\hat{\theta}$ that we estimate from the model, the boost to profits from an increase in $\hat{\theta}$ is significantly larger than the boost to profits from an increase of the same magnitude in λ . That is, differences in profits from new relationships are more attributable to differences in $\hat{\theta}$. This is different for profits from a known, reliable partner. When a firm has learned that its supplier is patient, the baseline fraction of patient firms, $\hat{\theta}$, and the enforcement probability, λ , no longer matter for the value of the relationship (that is, the solid lines in Panels A and C are flat). Profits are, however, still

subject to the exogenous dissolution shock, and so the sum of expected profits continues to decrease in δ .

Inherent Value of a Relationship Define the inherent value of a relationship IV to be the ratio of the sum of expected profits from a known, patient supplier over those from a new, random supplier:

$$IV = \frac{[c_1 + c_2]^{1-\sigma}}{\delta} \left(\sum_{k=0}^{\infty} (1-\delta)^k \left(\lambda^k (1-\hat{\theta}) + \hat{\theta} \right) \tilde{\theta}_k \left[(1-\alpha)c_1 + \frac{\alpha c_1 + c_2}{\tilde{\theta}_k} \right]^{1-\sigma} \right)^{-1}. \quad (17)$$

This expression allows us to evaluate the value of long-term relationships across source countries. Although this object is a complex mix of the underlying parameters, we can derive the following result.

Proposition 3 *The inherent value of a relationship from a patient, known supplier and a random, new supplier decreases in $\hat{\theta}$.*

Proof. Note that the partial derivatives of $\left(\lambda^k (1-\hat{\theta}) + \hat{\theta} \right)$ and $\tilde{\theta}_k$ and $\left[(1-\alpha)c_1 + \frac{\alpha c_1 + c_2}{\tilde{\theta}_k} \right]$ with respect to $\hat{\theta}$ are all positive. Then, $\partial IV / \partial \hat{\theta} < 0$, as each element of the infinite sum increases in $\hat{\theta}$. ■

The higher the share of reliable suppliers in a source country, the lower the boost to profits from establishing a long-term relationship. The intuition is that with more reliable suppliers in a country, the initial delivery probability—and, therefore, the profits from a new relationship—are higher. In contrast, the value of old relationships is independent of $\hat{\theta}$.

5 Calibration and Counterfactuals

In this section, we calibrate the model and present counterfactuals. The section has four subsections. The first subsection simulates the model and shows how changes to parameters affect the age distributions of relationships and trade. The second presents our calibration approach and its outcomes. In the third,

we calculate the value of relationships across countries. The fourth presents three counterfactuals.

5.1 Model Simulations

In the next subsection, we will calibrate the parameters δ , λ , and $\hat{\theta}$ by matching the model moments to the count and value distributions from the data. To build intuition on how identification of our set of parameters works, we run simulations that show how each affects the model-generated count and value distributions. For the simulations, we take the following baseline values: $\delta = \lambda = \hat{\theta} = 0.25$.²²

The baseline count and value distributions are shown in the top two panels of Figure 10. The next six panels show how distributions change when each of the three parameters is varied in turn while holding all other parameters constant. In each, the parameter of interest is set to a high value of 0.4 and a low value of 0.1. For the following discussion, it may be helpful for the reader to review Equations (11) and (13).

The Count Distribution Begin by noting that anything that leads to a lower probability of relationship survival, $surv_k = (1 - \delta)\tilde{\theta}_k$, implies higher shares of young relationships with that source country, particularly new relationships—the point where the count distribution meets the y -axis. Thus, a lower share of reliable suppliers ($\hat{\theta}$ in Panel C) and a higher death rate (δ in Panel G) tilt the count distribution towards younger relationships.

Recall from the theory section and Figure 7 that λ has competing effects on survival, hence the pattern for differing λ is more nuanced. Better enforcement implies both a higher initial probability of delivery from unreliable suppliers (which leads to higher survival) as well as slower learning (which leads to lower survival). Since the direct effect of enforcement on ensuring delivery dominates for young relationships, a higher λ implies more survival and less mass in low

²²These values are roughly in the middle of the range of parameters estimated later on. The other model parameters remain as before: $\sigma = 3$, $\alpha = 0.5$, $c_1 = 1$, and $c_2 = 4$.

ages. However, as k increases, the effect of slower learning from a higher λ starts to outweigh the direct effect. High- λ countries have lower survival of older relationships, generating less mass at older ages.

The Value Distribution Identification of our parameters also comes from the value distribution. The intuition is now more challenging, as trade shares may vary both because of differences in survival and because of differences in the quantities traded within relationships.

Since values within relationships increase over time, the value distribution (Panel B) is skewed to the right compared with the count distribution. How skewed the distribution is depends on each of the parameters. With low $\hat{\theta}$, a successful transaction delivers fast learning, and quantities increase substantially within surviving relationships over time. Therefore, with a lower $\hat{\theta}$, the value distribution (Panel D) is much more skewed to the right than the count distribution (Panel C). For high $\hat{\theta}$, the amount traded does not increase much over time as, from the beginning, the importer believes that the supplier is likely to be reliable, meaning the difference between the count and value distributions is much smaller. The same logic applies for high δ (where the value distribution is much more skewed than the count distribution) versus low δ .

As above, λ affects the distribution in two ways. First, lower λ implies a lower initial level of trade. Second, the lower λ is, the more often a supplier has a chance to deviate and the faster the learning by the importer is. Hence, a lower λ implies large within-relationship trade increases. Through this second effect, a lower λ pushes the value distribution toward later years. Combined with the low starting level, this can give rise to a hump-shape like that in panel F.

Taken together, these three parameters can capture a wide range of age distributions for relationship counts and trade values. It should be clear from the graphs as well as from Equations (11) and (13) that identification is based on the full model dynamics, as summarized by the two distributions.

5.2 Matching with the data

We employ an algorithm that searches for the parameter vector $(\delta, \hat{\theta}, \lambda)$ that minimizes the sum of squared differences between the moments in the data and those predicted by the model. As moments, we use the elements of the count and value distributions. We solve the following problem:

$$\arg \min_{\delta, \hat{\theta}, \lambda} \beta \sum_{k=0}^K (\text{ageshare}_k - \widehat{\text{ageshare}}_k)^2 + (1 - \beta) \sum_{k=0}^K (\text{tradeshare}_k - \widehat{\text{tradeshare}}_k)^2,$$

where *ageshare* and *tradeshare* are the values predicted by the model, and $\widehat{\text{ageshare}}$ and $\widehat{\text{tradeshare}}$ are taken from the data. We set $\beta = 0.5$. In addition we set the other parameters to $\sigma = 3$, $\alpha = 0.5$, and $c_1 = c_2 = 1$.²³

We run this procedure for the 20 trading partners with the highest exports to the U.S.²⁴ For each country, we use information on eight cohorts— that is, on relationships that are one to eight years old ($K = 8$).²⁵ The three parameters for each country are therefore selected by matching 16 moments in the model with their empirical counterparts. The parameter estimates are found in the left panel of Table 7.

How well do the model-implied distributions represent their empirical counterparts? Figure 11 shows that the calibrated model matches the underlying data well. China’s trade is, for example, tilted toward new relationships, while Japan has more trade in longer-term relationships. As should be expected, the model matches the data less well when there are outliers. A case in point is Mexico, where in 2011 a large fraction of trade was in relationships that were exactly seven years old. The calibration adjusts parameters to accommodate the outlier at the cost of an inferior fit to the other moments of the value distribution. Similarly, outliers in middle-age relationships in Japan lead the model

²³Changes to either of these parameters have minimal effect on the explanatory power of the model, the ordering of individual countries, or the counterfactuals. They affect the levels of the estimated parameters.

²⁴These are China, Hong Kong, Taiwan, Italy, Germany, United Kingdom, Canada, India, France, Japan, South Korea, Mexico, Thailand, Indonesia, Chile, Spain, Brazil, Netherlands, Philippines, and Venezuela.

²⁵This is the largest number for which we have information on all countries. We could add additional cohorts at the price of dropping some countries.

to over-predict the fraction of trade in long-term relationships in that country.

Checking for external validity, we compare the country-specific institutional parameter λ that we obtain with the World Bank Rule of Law measure. Figure 12 plots the two series against each other. The high positive correlation implies that our model backs out sensible estimates of contract enforcement from the import data.

5.3 The Inherent Value - Quantification

We can now quantify the value of relationships. Columns (4) to (6) of Table 7 present the results. Column (4) reports the sum of expected profits from a new relationship; column (5) reports the sum of expected profits from a relationship with a known, patient supplier; and column (6) reports IV from Equation (17), the ratio between the values in columns (4) and (5).

Consider some examples from Table 7. New relationships with Spain (0.2) and the Netherlands (0.21) have the lowest value, mostly because they rank lowest in $\hat{\theta}$. A new relationship is the most valuable with Taiwan (0.88), which has relatively high estimates for λ and $\hat{\theta}$ combined with the second-lowest estimate for δ . The value of an old relationship maps directly into δ , as it does not depend on λ and $\hat{\theta}$. The expected sum of profits from an established relationship are therefore lowest for China (1.79) and the Philippines (1.84) and highest for Mexico (5.36) and Taiwan (4.22).

Inherent Values According to our estimates, relationships are most valuable in Spain (where an established relationship has a sum of expected profits 13.8 times higher than a new relationship) and the Netherlands (12.4). Both countries feature low $\hat{\theta}$ that depress the value of new relationships. We find the lowest inherent values for the Philippines (3.04) and China (3.88). Both countries have high $\hat{\theta}$, implying high profits from new relationships, and high δ , implying low profits from old relationships. Across all countries, the average value of a long-term relationship is 7.2 times that of a new relationship. The large size of these estimates should not be surprising given our empiri-

cal findings that long-term relationships foster higher trade as well as greater survival.

Importantly, countries with higher inherent values are not simply those countries with fewer long-term relationships or less trade: both Spain and China have nearly identical shares of long-term relationships (≈ 12 percent) and long-term trade (33 percent), but they differ greatly in the inherent value of a relationship (3.9 versus 13.8). Based on our calibration, China has mostly short-term relationships because of a high δ , while Spain has many short-term relationships because of a low $\hat{\theta}$. The high δ for China implies that long-term relationships are not very valuable, as they tend to die quickly for exogenous reasons. The low $\hat{\theta}$ for Spain leads to a very low value of new relationships. At the same time, Spain's δ is not that high, making long-term relationships more valuable. Put together, these parameter differences imply that Spain has a much higher inherent value of relationships than China.

The reverse holds as well: countries with similar values of old relationships need not have similar age distributions. Column (5) of the table shows that old relationships with Germany and Spain are similarly valuable, even though in the data, their shares of old relationships are very different. This is because the model generates a high German share of old relationships via both a fairly high $\hat{\theta}$ (0.23) and a low δ , while the low Spanish share of old relationships comes from a very low $\hat{\theta}$ (0.09) and low δ . Since the δ terms for these countries align (0.215 and 0.210), the value of an old relationship is basically identical.

We next illustrate how the expected profits from a relationship change with relationship age. Figure 13 shows the expected sum of profits as a function of relationship age for three economies— China, Germany, and Spain— based on their estimated parameters: λ , δ , and $\hat{\theta}$. The starting point of the lines are given by Equation (16)— the value of a new relationship— while the end points can be approximated by Equation (15)— the value of an old relationship. First, note that because of learning, the sum of expected profits from a relationship strictly increases in the relationship's age.

Consider China and Germany. A new relationship has about the same

expected profits in both countries. However, a first successful year of trading raises the value of a Chinese relationship by more. This is because in China, one successful transaction is a stronger indicator for a reliable supplier: after one year, the belief about being matched with a reliable suppliers increases from 36 percent to 90 percent. By comparison, a delivery from Germany increases the belief about being matched with a patient supplier only from 23 to 55 percent. However, since δ is higher for China than for Germany, expected profits in a long-term relationship plateau at a lower level for China than for Germany.

Now consider Spain and Germany. Spain's lower $\hat{\theta}$ than Germany's implies a much lower value of profits from a new relationship. In the long run, however, all relationships that are still alive are with patient suppliers. So after six years, $\hat{\theta}$ and λ hardly have any effect on the relationship value, and the key determinant becomes δ . Since Spain and Germany have about the same δ , their relationship values converge to roughly the same number.

Would it suffice to only consider the exogenous death shock and abstract from learning? We check this by running our calibration exercise on a restricted model where only δ is active and λ is set to one (a frictionless world). This delivers an extremely poor fit to the data as reported in Table 8. The table also compares the fully estimated model to versions where either $\hat{\theta}$ or λ is fixed across countries (and chosen optimally)- the sum of squared errors from Equation (18) is lowest by allowing all three parameters to vary across countries. Figure 14 plots the goodness of fit (the sum of squared errors) of the whole model against one where only δ matters, and shows that only relying on exogenous relationship death to explain the data is insufficient. Hence, while δ is a central parameter for the value of long-term relationships, it can only be calibrated credibly when learning is active.

5.4 Counterfactuals

This section presents counterfactuals of the learning model. Counterfactuals start with the relationship distribution, entry, and trade flows in a steady state. Then, there is an initial (or multi-period) shock to entry or the relationship dis-

tribution. After the shock, entry returns to its steady-state level or to a value that keeps the number of relationships constant, while the relationship distribution and trade flows move slowly back to steady state. While entry is thus fixed exogenously, the CES price level is allowed to adjust. Since dynamics we are interested in mostly unfold in the short to medium term (within five years), we close the model by restricting entry responses while allowing adjustments at the intensive margin.

The experiments are solved country by country. We solve the dynamics of the partial equilibrium model using the parameters of each country and then compare the outcomes. For comparability across countries, the main outcome variable— the CES aggregate for trade flows, Q_{t-} is normalized by its country-specific steady-state value. To keep exposition straightforward, each counterfactual presents outcomes for both Japan (a country with mostly older relationships) and China (a country with mostly young relationships) and discusses how results are driven by differences in the count and value distributions between the two countries.

Reset of All Trade Relationships The first counterfactual addresses the value of long-term relationships from an alternative angle: how much trade would be lost if all long-term relationships were “reset” to new relationships? In the experiment, the economy starts in a steady state. A shock then wipes out all established relationships and replaces them with the same number of new relationships. After the initial shock, entry is set such that the number of relationships is held constant. That is, entry exactly corresponds to the number of relationships that get dissolved in a period.

Outcomes are presented in Figure 15. In both countries, trade drops substantially and takes several years to recover. The paths for Japan and China differ. Long-term relationships die less quickly in Japan than in China and are therefore more relevant for trade. Consequently, a larger fraction of trade is lost when the shock hits. At the same time, rebuilding the relationships takes longer, so recovery is less quick in Japan than in China. All told, Japan

loses exactly 2.0 times its annual steady-state level of trade over the transition, while China loses 0.67 times its one-year trade level. Across all countries, the average loss in the transition back to a steady state is about 1.5 times the annual steady-state level of trade, with the Philippines losing the least (0.17) and Mexico losing the most (2.45).

To build intuition, consider the similarity between capital and the stock of relationships. Like relationships, capital builds up slowly over time. When capital gets destroyed, the country with the higher steady-state capital loses more output. As with relationships, that country also needs more time to rebuild its capital stock and return to a steady state.

One period of increased entry The second counterfactual studies the effect of a one-off increase in the number of relationships. Again, the counterfactual begins in steady state. Then, for a single period, entry increases by 20 percent. Thereafter, entry returns to its steady-state level. Results are shown in Figure 16. The initial increase is captured by the dotted line. Thereafter, trade slowly moves back to a steady state as the additional relationships die out.

There is a notable heterogeneity across countries: initially, imports from China surge relative to imports from Japan. This is because imports from China are more tilted towards new relationships and so increased entry has a stronger effect on aggregate trade. However, after three years, the two lines intersect and the shock implies higher trade gains for Japan. Because relationships die less quickly in Japan, the one-off higher entry benefits Japan more in the long run; trade actually continues to increase after the shock for a few years.

Global Trade Collapse The final counterfactual focuses on the Great Trade Collapse of 2008-09. The counterfactual feeds the average reduction in relationship formation and survival in 2008 and 2009 into the model. Results are shown in Figure 17. Now, China is hurt more than Japan. This is sensible given the importance of young relationships for imports from China. The same trait implies that China recovers faster than Japan. Even many years later, Japan's

trade is depressed because of the missing relationships not created in the crisis years.

Our counterfactuals show that learning and relationships can have first-order effects on the level and the evolution of trade, giving a separate methodology to evaluate the value of long-run relationships. The variation in the count and value distributions across countries imply very different dynamics for trade flows from those countries- those countries with trade concentrated in long-run relationships are less affected by shocks to relationship formation in the short run, but they also recover more slowly from these shocks.

6 Conclusions

Relationships between importers and exporters are the core of international trade transactions. This paper presented new facts on these relationships and developed a model, building on Araujo et al. (2012), to quantify the role of learning in international trade. The model fits our empirical findings well: trade increases within successful relationships over time, relationship survival increases over time, and countries differ systematically in their count and value distributions in ways our model can capture.

Our counterfactuals generate substantial cross-country heterogeneity in trade dynamics. In the short run, shocks to entry have the strongest effects on countries like China that are tilted toward relationships of short duration. In the medium run, however, entry shocks affect trade more with countries that rely more on longer-term relationships, like Japan. This has implications, for example, for the study of trade policy. Our model suggests that trade dynamics around FTAs should depend on country institutions and on how trade is spread across relationships of different ages. More generally, our analysis suggests that to capture trade dynamics in the aggregate, we need to take dynamics at the relationship-level seriously.

Finally, our paper provides a way to estimate the quality of country institutions from transaction-level data in a transparent way. The obtained estimates

can complement measures made available by the World Bank and other institutions.

References

- Albornoz, Facundo, Héctor F Calvo Pardo, Gregory Corcos, and Emanuel Ornelas**, “Sequential Exporting,” *Journal of International Economics*, 2012, *88* (1), 17–31.
- Alessandria, George and Horag Choi**, “Establishment heterogeneity, exporter dynamics, and the effects of trade liberalization,” *Journal of International Economics*, 2014, *94* (2), 207–223.
- , – , and **Kim J Ruhl**, “Trade adjustment dynamics and the welfare gains from trade,” Technical Report, National Bureau of Economic Research 2014.
- Alvarez, Fernando and Robert E Lucas**, “General equilibrium analysis of the Eaton–Kortum model of international trade,” *Journal of Monetary Economics*, 2007, *54* (6), 1726–1768.
- Antras, Pol and C Fritz Foley**, “Poultry in Motion: A Study of International Trade Finance Practices,” *Journal of Political Economy*, 2015, *123* (4).
- Araujo, Luis F, Giordano Mion, and Emanuel Ornelas**, “Institutions and Export Dynamics,” Working Paper 220, National Bank of Belgium 2012.
- Arkolakis, Costas, Theodore Papageorgiou, and Olga Timoshenko**, “Firm Learning and Growth,” Mimeo 2015.
- Atkeson, Andrew and Ariel Tomas Burstein**, “Innovation, Firm Dynamics, and international Trade,” *Journal of political economy*, 2010, *118* (3), 433–484.
- Bastos, Paulo, Daniel Dias, and Olga Timoshenko**, “Learning, Prices and Firm Dynamics,” Mimeo 2015.

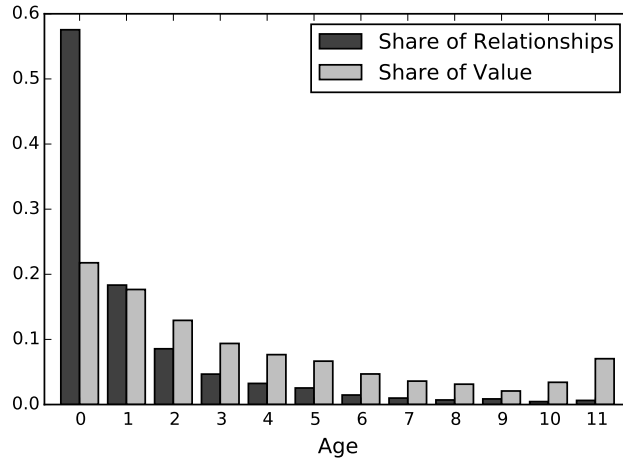
- Bernard, Andrew B, Andreas Moxnes, and Karen Helene Ulltveit-Moe**, “Two-sided Heterogeneity and Trade,” Working Paper 20136, National Bureau of Economic Research 2014.
- , **J Bradford Jensen, and Peter K Schott**, “Importers, Exporters and Multinationals: A Portrait of Firms in the US that Trade Goods,” in “Producer dynamics: New evidence from micro data,” University of Chicago Press, 2009, pp. 513–552.
- Besedeš, Tibor**, “A Search Cost Perspective on Formation and Duration of Trade,” *Review of International Economics*, 2008, 16 (5), 835–849.
- and **Thomas J Prusa**, “Product Differentiation and the Duration of U.S. Import Trade,” *Journal of International Economics*, 2006, 70 (2), 339–358.
- 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.
- Eaton, Jonathan and Samuel Kortum**, “Technology, geography, and trade,” *Econometrica*, 2002, 70 (5), 1741–1779.
- , **Marcela Eslava, Cornell J Krizan, Maurice Kugler, and James Tybout**, “A Search and Learning Model of Export Dynamics,” 2014. Mimeo.
- , **Samuel Kortum, Brent Neiman, and John Romalis**, “Trade and the global recession,” Technical Report, National Bureau of Economic Research 2011.
- Egan, Mary Lou and Ashoka Mody**, “Buyer-Seller Links in Export Development,” *World Development*, 1992, 20 (3), 321–334.

- Heise, Sebastian**, “Firm-to-Firm Relationships and Price Rigidity: Theory and Evidence,” Mimeo 2015.
- Impullitti, Giammario, Alfonso A Irarrazabal, and Luca David Opro-molla**, “A Theory of Entry into and Exit from Export Markets,” *Journal of International Economics*, 2013, 90 (1), 75–90.
- Jovanovic, Boyan**, “Selection and the Evolution of Industry,” *Econometrica*, 1982, 50 (3), 649–670.
- Kamal, Fariha and Asha Sundaram**, “Buyer-Seller Relationships in International Trade: Do Your Neighbors Matter?,” 2013. Mimeo.
- , **C.J. Krizan, and Ryan Monarch**, “Identifying Foreign Suppliers in U.S. Merchandise Import Transactions,” Working Paper 15-10, U.S. Census Center for Economic Studies 2015.
- Kaufmann, Daniel, Aart Kraay, and Massimo Mastruzzi**, “The World-wide Governance Indicators: Methodology and Analytical Issues,” Policy Research Working Paper Series 4978, The World Bank 2010.
- Krautheim, Sebastian**, “Heterogeneous firms, exporter networks and the effect of distance on international trade,” *Journal of International Economics*, 2012, 87 (1), 27–35.
- 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.
- Macchiavello, Rocco and Ameet Morjaria**, “The value of relationships: evidence from a supply shock to Kenyan rose exports,” *The American Economic Review*, 2015, 105 (9), 2911–2945.
- 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**, “Business and Social Networks in International Trade,” *Journal of Economic Literature*, 2001, pp. 1177–1203.
- **and Joel Watson**, “Starting Small in an Unfamiliar Environment,” *International Journal of Industrial Organization*, 2003, 21 (7), 1021–1042.
- **and –**, “Network Intermediaries in International Trade,” *Journal of Economics & Management Strategy*, 2004, 13 (1), 69–93.
- Ruhl, Kim and Jonathan Willis**, “New Exporter Dynamics,” Mimeo 2014.
- Ruhl, Kim J**, “The international elasticity puzzle,” Mimeo 2008.
- Schmidt-Eisenlohr, Tim**, “Towards a theory of trade finance,” *Journal of International Economics*, 2013, 91 (1), 96 – 112.
- Timoshenko, Olga A**, “Learning versus sunk costs explanations of export persistence,” *European Economic Review*, 2015, 79, 113–128.
- , “Product switching in a model of learning,” *Journal of International Economics*, 2015, 95 (2), 233–249.

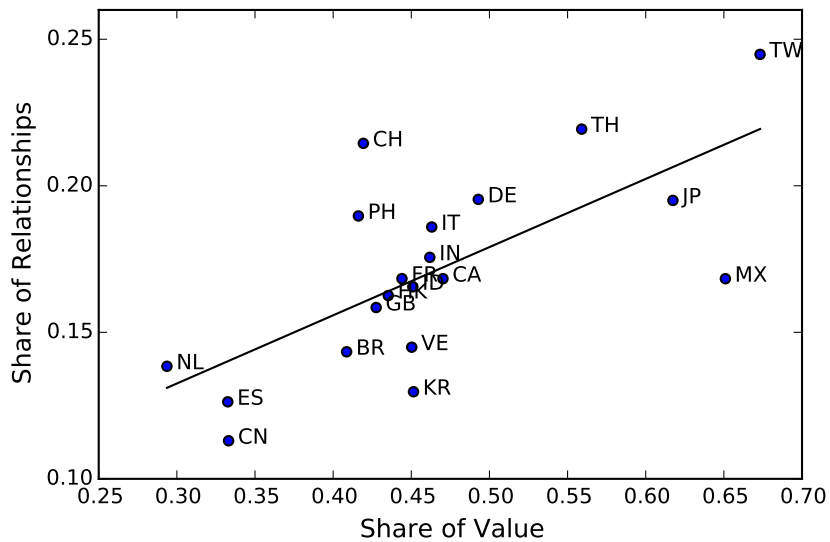
A Figures

Figure 1: Count and Value Distributions, 2011




Note: This figure plots the count distribution and the value distribution for all U.S. imports in 2011. The count distribution is the distribution of relationship counts over different ages. The value distribution is the distribution of trade values across relationships of different ages. Age 11 includes all relationships that are 11 years or older.

Figure 2: Distribution of Relationship Age and Value Shares, 2011




Note: This figure plots the share of long-term (3 or more years) relationships against the share of trade in long-term relationships for the top 20 U.S. trading partners ($R^2 = 0.47$).

Figure 3: Evolution of Trade Within a Relationship



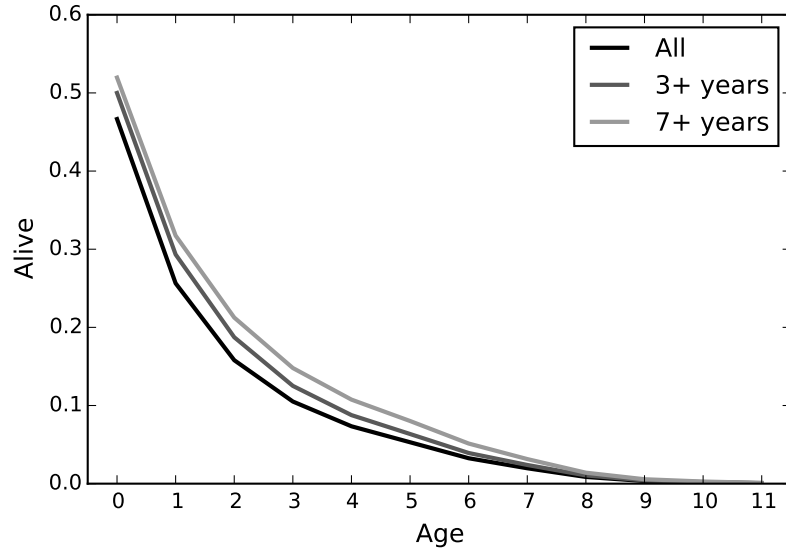
Note: This figure plots coefficients from regressions of traded value on relationship age dummies, run separately for relationships that lasted 5, 7, 10, and (at least) 13 years.

Figure 4: Relationship Survival Probabilities



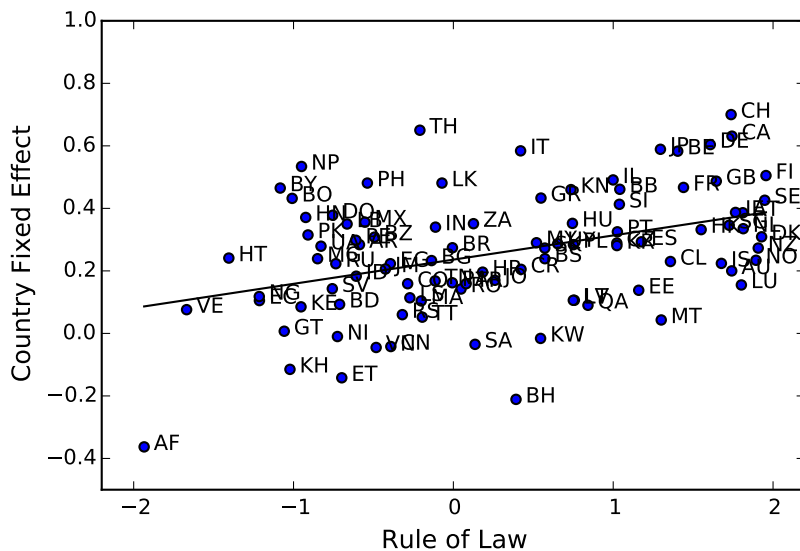
Note: This figure plots the conditional survival probability by age in 2011; that is, the share of relationships of age k that survive to age $k + 1$.

Figure 5: Survival Function




Note: This figure plots the survival function based on estimates from the Cox proportional hazard model for all relationships, firms that are at least 3 years old, and firms that are at least 7 years old.

Figure 6: Relationship Longevity and Source Country Rule-of-Law



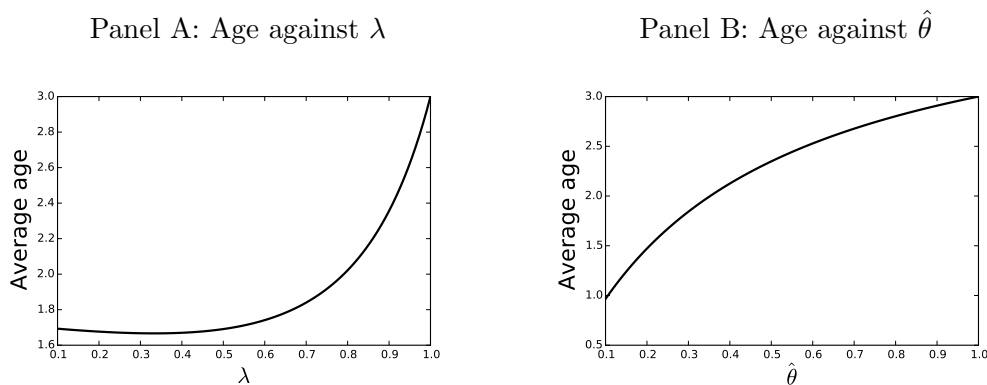
Note: This figure plots the country fixed effects from the regression in Column (6) of Table 6 against the World Bank's Rule of Law measure.

Figure 7: Delivery Probabilities



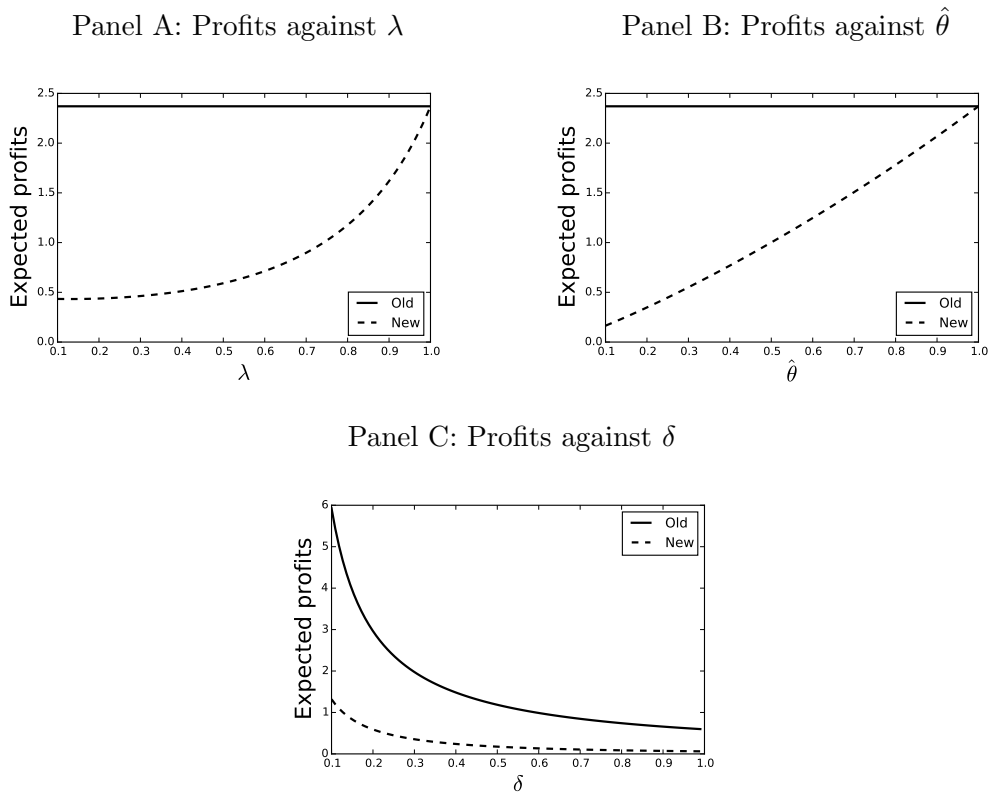
Note: This figure illustrates the role of λ for $\tilde{\theta}_k$, the probability of an additional successful transaction following k periods of trade. λ takes the values 0.1 and 0.4. $\hat{\theta}$ is set at 0.25.

Figure 8: Average Age of Relationships



Note: These figures plot the average age of a relationship implied by Equation (6) against λ and $\hat{\theta}$. Average age increases in the share of reliable suppliers $\hat{\theta}$, and increasing in the enforcement probability λ for λ higher than $1 - \frac{\delta}{(1-\delta)\sqrt{\hat{\theta}}}$. The baseline parameters are $\delta = \hat{\theta} = \lambda = 0.25$.


Figure 9: Expected profits from new and old relationships




Note: These figures show the expected sum of future profits from new relationships (dashed lines) and old relationships (solid lines) for different values of λ , $\hat{\theta}$, and δ . The baseline parameters are $\delta = \hat{\theta} = \lambda = 0.25$.

Figure 10: Mechanics of the model


Panel A: Countshare - Baseline




Panel B: Valueshare - Baseline




Panel C: Countshare - $\hat{\theta}$




Panel D: Valueshare - $\hat{\theta}$




Panel E: Countshare - λ




Panel F: Valueshare - λ



Panel G: Countshare - δ

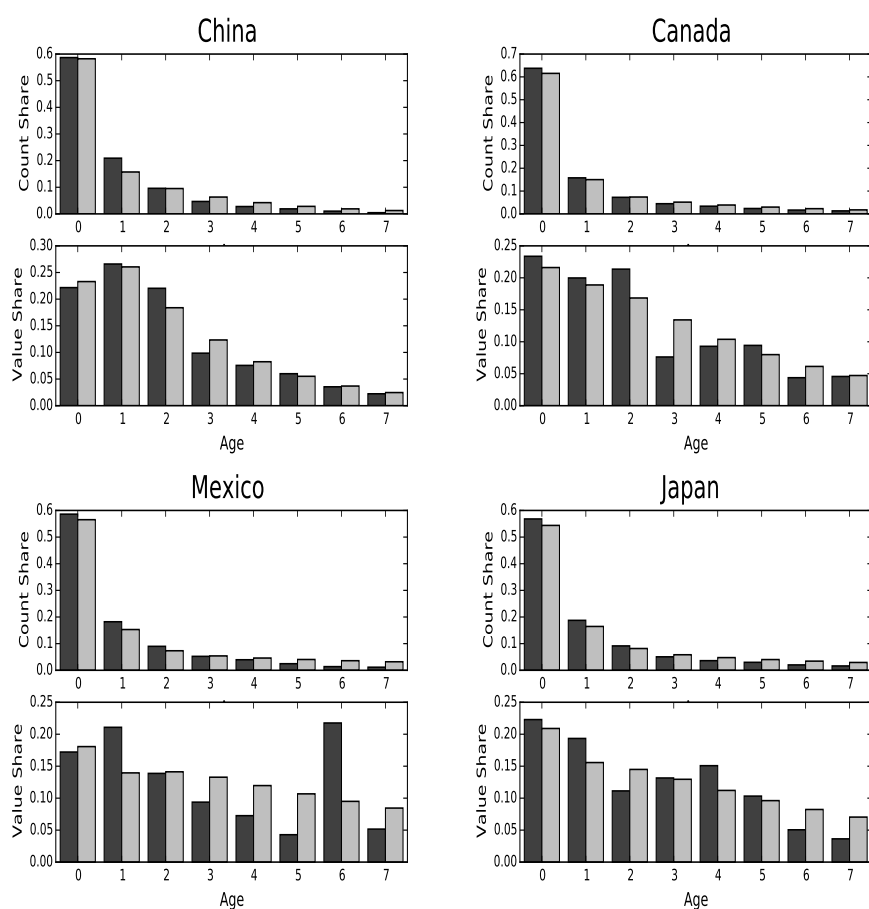


Panel H: Valueshare - δ



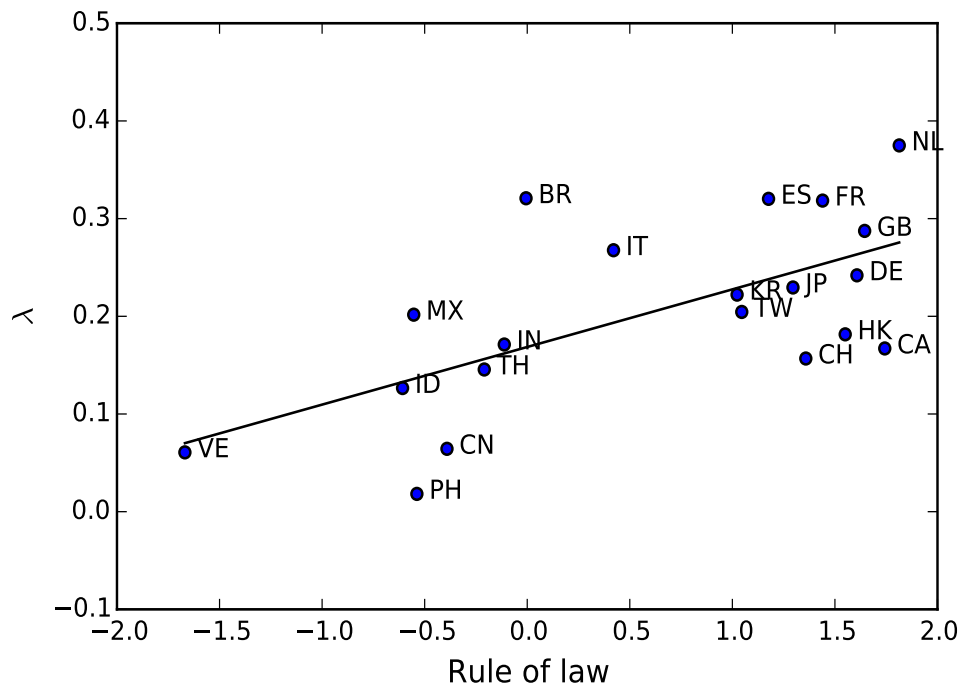
Note: This figure shows the count and value distributions under the baseline parametrization ($\hat{\theta} = \delta = \lambda = 0.25$) and by varying each parameter from 0.1 to 0.4, keeping the others constant.

Figure 11: Count and Value distributions, Model vs. Data




Note: This figure shows the count distributions and the value distributions for four countries. The dark bars represent data moments, while the light bars show moments from the calibrated model.

Figure 12: Institutional Quality: Model vs. Data




Note: This figure plots the country-specific enforcement parameter λ backed out from relationship data against the World Bank rule of law measure for the 20 countries in our sample ($R^2 = 0.37$).

Figure 13: Expected profits and relationship age




Note: This figure plots the expected sum of future profits at ages k for China, Germany, and Spain for their estimated λ, δ , and $\hat{\theta}$. The starting point is given by Equation (15), while the end point can be approximated by Equation (16).

Figure 14: Count and Value distributions, Model vs. Data, all countries




Note: These figures plots the data moments against the model moments for all 20 countries. Panel A shows this when calibrating the full model. Panel B shows results when only the dissolution shock δ is active; that is, $\lambda = 1$.

Figure 15: Experiment 1: Ending All Relationships




This figure shows results from a counterfactual where all relationships are “reset” to new for Japan and China.

Figure 16: Experiment 2: One Period of Higher Entry



This figure shows results from a counterfactuals where there is a one-time increase in relationship formation of 20%.

Figure 17: Experiment 3: Entry During the Great Recession



This figure shows results from a counterfactual where relationship formation takes the average values from the Great Trade Collapse.

B Tables

Table 1: Relationship Structure of U.S. Imports, 2011

	New	1-2 Years	3 or More Years
Share of Relationships	57.6	26.9	15.6
Share of Trade	21.8	30.6	47.6

Note: For this table, a relationship is defined as a U.S. importing firm transacting with a non-U.S. exporting firm. A new relationship is one that is not found in any previous year of data.

Table 2: Relationship Structure of U.S. Imports, 2011

	% Counts, 3+ Year Relationships		% Trade, 3+ Year Relationships
China	0.11	Netherlands	0.29
Spain	0.13	Spain	0.33
Korea	0.13	China	0.33
Netherlands	0.14	Brazil	0.41
Brazil	0.14	Philippines	0.42
Venezuela	0.14	Chile	0.42
United Kingdom	0.16	United Kingdom	0.43
Hong Kong	0.16	Hong Kong	0.44
Indonesia	0.17	France	0.44
Mexico	0.17	Venezuela	0.45
Canada	0.17	Indonesia	0.45
France	0.17	Korea	0.45
India	0.18	India	0.46
Italy	0.19	Italy	0.46
Philippines	0.19	Canada	0.47
Japan	0.20	Germany	0.49
Germany	0.20	Thailand	0.56
Chile	0.21	Japan	0.62
Thailand	0.22	Mexico	0.65
Taiwan	0.24	Taiwan	0.67

Note: This table shows for each of the top 20 trading partners of the U.S., what share of relationship counts (left panel) and what share of trade value (right panel) are in relationships that have appeared for four or more consecutive years.

Table 3: Traded Value within a Relationship

Dependent Variable: $Value_{mx,t}$

	(1)	(2)	(3)	(4)
<i>Age=1</i>	0.2923 (0.007)	0.3148 (0.009)	0.4760 (0.018)	0.4499 (0.044)
<i>Age=2</i>	0.3066 (0.007)	0.3634 (0.009)	0.5641 (0.018)	0.5129 (0.045)
<i>Age=3</i>	0.2655 (0.007)	0.3712 (0.009)	0.5920 (0.018)	0.6304 (0.047)
<i>Age=4</i>	0.1616 (0.007)	0.3535 (0.009)	0.6046 (0.019)	0.6374 (0.046)
<i>Age=5</i>	-0.1805 (0.007)	0.3101 (0.009)	0.6117 (0.020)	0.6837 (0.046)
<i>Age=6</i>		0.2319 (0.009)	0.6020 (0.019)	0.6270 (0.046)
<i>Age=7</i>		-0.0163 (0.009)	0.5460 (0.019)	0.6594 (0.047)
<i>Age=8</i>			0.4828 (0.019)	0.6438 (0.047)
<i>Age=9</i>			0.4267 (0.019)	0.6587 (0.048)
<i>Age=10</i>			0.2609 (0.019)	0.6656 (0.048)
<i>Age=11</i>				0.5643 (0.046)
<i>Age=12</i>				0.5581 (0.047)
<i>Age=13</i>				0.4525 (0.048)
Constant	-1.3288 (0.469)	-2.5295 (0.955)	-2.797 (0.177)	-2.078 (0.344)
Total Age (Years)	5	7	10	13
N	192,300	169,800	67,000	13,200
R ²	0.49	0.49	0.45	0.46

Note: This table presents results on the growth of trade values within relationships. The value of trade is regressed on age fixed effects, controlling for relationship fixed effects - see Equation 1. The omitted category is trade in the first year ($Age=0$). Data is from all relationships formed in the years 1998-2006. Trade values are available through 2011. The regressions controls for importer and exporter firm size, number of HS10 products sold/bought, and number of relationships. Standard errors are in parentheses.

Table 4: Relationship Margins during the Trade Collapse

Counts Relative to 2007

Year	2007	2008	2009	2010
Relationship Counts	1	0.97	0.90	0.98
New Relationship Counts	1	0.93	0.85	0.96
Continuing-Relationship Counts	1	1.03	0.97	0.99
Ending-Relationship Counts	1	1.02	1.02	0.88

Note: This table describes relationship formation, dissolution and continuation in the years around the 2008/2009 Great Trade Collapse. Each category is normalized by the baseline year 2007. A new relationship is one not observed in any prior year. A continuing relationship is one that is found in both the current year and the previous year. An ending relationship is one that is found in the current year, but not the next year.

Table 5: Survival Probabilities During the Trade Collapse

Relationship Age	0-2	3-5	6+
Crisis	-0.051	-0.056	-0.007
	(0.008)	(0.014)	(0.015)
CrisisXRule of Law	0.019	0.043	-0.012
	(0.008)	(0.014)	(0.012)
Age Fixed Effects	Yes	Yes	Yes
N	654	963	1605
R ²	0.54	0.08	0.03

Note: This table presents results on relationship survival during the Great Trade Collapse. An observation is age cohort k in year t in country c . The dependent variable is the probability that a relationship of a given age survives one more year. The “Crisis” variable is 1 in 2008-2009 and 0 otherwise. Standard errors clustered by source country are in parentheses.

Table 6: Determinants of Importer-Exporter Relationship Length, 2011

Dependent Variable: Age_{mxc}						
	(1)	(2)	(3)	(4)	(5)	(6)
Importer Controls						
$Size_m$	-0.027 (0.026)	-0.027 (0.026)	-0.027 (0.026)	-0.027 (0.026)	-0.026 (0.025)	-0.026 (0.025)
Age_m	0.044 (0.005)	0.044 (0.005)	0.044 (0.004)	0.044 (0.004)	0.045 (0.004)	0.044 (0.004)
$RelCount_m$	0.064 (0.035)	0.064 (0.035)	0.062 (0.034)	0.062 (0.034)	0.061 (0.034)	0.06 (0.033)
$NoProducts_m$	0.294 (0.053)	0.296 (0.053)	0.297 (0.053)	0.297 (0.053)	0.298 (0.052)	0.298 (0.052)
Exporter Controls						
$Size_x$	0.232 (0.013)	0.232 (0.013)	0.233 (0.013)	0.233 (0.013)	0.234 (0.013)	0.234 (0.013)
$RelCount_x$	-0.375 (0.016)	-0.376 (0.016)	-0.378 (0.016)	-0.379 (0.016)	-0.383 (0.016)	-0.384 (0.015)
Country Controls						
$RuleOfLaw_c$		0.122 (0.012)	0.124 (0.013)	0.119 (0.016)	0.126 (0.012)	
$PC - GDP_c$	0.164 (0.020)	0.073 (0.019)	-0.006 (0.014)		0.031 (0.011)	
$OECD_c$			0.238 (0.021)	0.228 (0.023)	0.142 (0.015)	
RTA_c				0.021 (0.017)	0.009 (0.038)	
$Distance_c$					-0.018 (0.032)	
Constant	-2.585 (0.297)	-1.766 (0.287)	-1.132 (0.255)	-1.19 (0.209)	-2.595 (0.337)	-1.304 (0.218)
N	1,075,100	1,075,100	1,075,100	1,075,100	1,075,100	1,075,100
Fixed Effects	HS2	HS2	HS2	HS2	HS2	c , HS2
R ²	0.11	0.11	0.12	0.12	0.12	0.12

Note: This table presents results on the role of firm and country institutions on the age of relationships. Each regression includes source country and primary-HS2 product effects. Standard errors clustered by source country are in parentheses. The dependent variable is the age of a relationship in 2011, defined as the number of consecutive years a relationship has been found in the data, and is capped at 11.

Table 7: Parameters and the Value of Relationships

	(1)	(2)	(3)	(4)	(5)	(6)
Country	λ	$\hat{\theta}$	δ	Value of a New Relationship	Value of an Old Relationship	Inherent Value of a Relationship
Brazil	0.32	0.11	0.21	0.25	2.77	10.92
Canada	0.17	0.18	0.23	0.33	2.56	7.73
Chile	0.16	0.33	0.27	0.54	2.16	4.00
China	0.06	0.36	0.33	0.46	1.79	3.88
Germany	0.24	0.23	0.21	0.49	2.76	5.69
Spain	0.32	0.09	0.21	0.20	2.82	13.83
France	0.32	0.10	0.17	0.31	3.54	11.56
United Kingdom	0.29	0.13	0.22	0.28	2.73	9.87
Hong Kong	0.18	0.26	0.26	0.43	2.30	5.29
Indonesia	0.13	0.29	0.24	0.52	2.43	4.66
India	0.17	0.27	0.24	0.51	2.52	4.91
Italy	0.27	0.18	0.20	0.41	2.91	7.01
Japan	0.23	0.16	0.14	0.54	4.09	7.61
Korea	0.22	0.18	0.23	0.35	2.61	7.54
Mexico	0.20	0.13	0.11	0.58	5.36	9.26
Netherlands	0.37	0.09	0.23	0.21	2.60	12.44
Philippines	0.02	0.44	0.32	0.61	1.84	3.04
Thailand	0.15	0.24	0.17	0.65	3.45	5.3
Taiwan	0.20	0.25	0.14	0.88	4.22	4.78
Venezuela	0.06	0.32	0.25	0.57	2.33	4.09

Note: This table presents parameter estimates for the 20 countries from which the U.S. imported the most in 2011. In columns (4) and (5), it also presents the expected sum of profits from a new and an old relationship, respectively. Column (6) reports the ratio of the expected sum of profits from an old relationship over those from a new relationship.

Table 8: Alternative models

Model	Error	λ		$\hat{\theta}$		δ	
		min	max	min	max	min	max
Full	0.92	0.02	0.37	0.09	0.44	0.11	0.33
δ Only	7.57		1		1	0.29	0.55
Fixed $\hat{\theta}$	0.96	0.11	0.37		0.25	0.17	0.38
Fixed λ	1.09		0.19	0.16	0.54	0.17	0.49

Note: This figure presents results from alternative models. Error is the sum of squared errors as defined in Equation (18).

C Proof of Proposition 1

In our model, the average age is given by:

$$a\bar{g}e = \frac{\sum_{k=0}^{\infty} \text{alive}_k k}{\sum_{k=0}^{\infty} \text{alive}_k} = \frac{\sum_{k=0}^{\infty} (1-\delta)^k \left(\lambda^k (1-\hat{\theta}) + \hat{\theta} \right) k}{\sum_{k=0}^{\infty} (1-\delta)^k \left(\lambda^k (1-\hat{\theta}) + \hat{\theta} \right)} \quad (\text{C.1})$$

Proposition 1 *The average age of a relationship increases in the share of reliable suppliers in a country $\hat{\theta}$, and in enforcement λ , for $\lambda > 1 - \frac{\delta}{(1-\delta)\sqrt{\hat{\theta}}}$.*

Proof. Since $\sum_{k=0}^{\infty} x^k = \frac{1}{1-x}$ and $\sum_{k=0}^{\infty} kx^k = \frac{x}{(x-1)^2}$, we can write the right side of Equation C.1 as:

$$\frac{\frac{(1-\hat{\theta})(1-\delta)\lambda}{((1-\delta)\lambda-1)^2} + \frac{\hat{\theta}(1-\delta)}{((1-\delta)-1)^2}}{\frac{1-\hat{\theta}}{1-(1-\delta)\lambda} + \frac{\hat{\theta}}{(1-(1-\delta))}} \quad (\text{C.2})$$

Taking the derivative of this expression with respect to $\hat{\theta}$, we obtain:

$$\frac{\partial a\bar{g}e}{\partial \hat{\theta}} = - \frac{(1-\delta)(\lambda-1)}{\left((1-\delta)\hat{\theta}\lambda + (1-\delta) - 1 \right)^2}$$

which is always greater than zero, since λ, δ , and $\hat{\theta}$ are in $(0,1)$ by construction.

Taking the derivative of the expression with respect to λ , we obtain:

$$\frac{\partial a\bar{g}e}{\partial \lambda} = \frac{(\hat{\theta}-1)(1-\delta) \left[(1-\delta)^2 \left(\hat{\theta}(\lambda-1)^2 - 1 \right) + 2(1-\delta) - 1 \right]}{\left((1-\delta)\lambda - 1 \right)^2 \left((1-\delta) \left(\hat{\theta}(\lambda-1) + 1 \right) - 1 \right)^2}$$

The denominator is always positive, while the numerator is positive if and only if the expression in brackets is negative, i.e.

$$(1-\delta)^2 \left(\hat{\theta}(\lambda-1)^2 - 1 \right) + 2(1-\delta) - 1 < 0 \iff \quad (\text{C.3})$$

$$(1-\delta)^2 \hat{\theta} \lambda^2 - (1-\delta)^2 \hat{\theta} 2\lambda + (1-\delta)^2 \hat{\theta} - (1-\delta)^2 + 2(1-\delta) - 1 < 0 \quad (\text{C.4})$$

Using the quadratic formula to solve the left side for λ yields:

$$\begin{aligned}
\lambda &= \frac{2\hat{\theta}(1-\delta)^2 \pm \sqrt{4\hat{\theta}^2(1-\delta)^4 - 4\hat{\theta}(1-\delta)\left[(1-\delta)^2(\hat{\theta}-1) + 2(1-\delta) - 1\right]}}{2\hat{\theta}(1-\delta)^2} \\
&= 1 \pm \frac{\sqrt{4\hat{\theta}^2(1-\delta)^4 \left[1 - \frac{\hat{\theta}-1}{\hat{\theta}} - 2\frac{1}{\hat{\theta}(1-\delta)} + \frac{1}{\hat{\theta}(1-\delta)^2}\right]}}{2\hat{\theta}(1-\delta)^2} \\
&= 1 \pm \sqrt{\frac{\delta^2}{\hat{\theta}(1-\delta)^2}} = 1 \pm \frac{\delta}{(1-\delta)\sqrt{\hat{\theta}}}
\end{aligned}$$

This means that we can rewrite the Inequality C.4 as

$$\left(\lambda - \left[1 + \frac{\delta}{(1-\delta)\sqrt{\hat{\theta}}}\right]\right) \left(\lambda - \left[1 - \frac{\delta}{(1-\delta)\sqrt{\hat{\theta}}}\right]\right) < 0$$

The left term is always negative for $\lambda, \delta, \hat{\theta} \in (0, 1)$, so the condition for average age to be increasing in λ is that the right term is positive:

$$\lambda > 1 - \frac{\delta}{(1-\delta)\sqrt{\hat{\theta}}}$$

■