Tariff Rate Uncertainty and the Structure of Supply Chains*

Sebastian Heise † Justin R. Pierce ‡ Georg Schaur § Peter K. Schott ¶

September 17, 2021

First Draft: July, 2015

Abstract

We develop a model of international procurement and show that an increase in trade policy uncertainty inhibits the adoption of "Japanese" importing, whereby domestic buyers use long-term relationships to ensure the provision of high-quality inputs from foriegn sellers. Using transaction-level trade data, we demonstrate that a change in U.S. policy that substantially reduced the possibility of a trade war with China induced a relative increase in "Japanese" trading relationships among U.S. importers and Chinese exporters. Quantitative simulations suggest that this increase in policy certainty raised U.S. welfare substantially. (JEL Codes: F13, F14, F15, F23) (Keywords: Supply Chain, Uncertainty, Trade War, Procurement)

*Schott and Heise thank the National Science Foundation (SES-1427027) for research support. We thank seminar participants at Indiana, the Inter-American Development Bank, the NBER Japan Conference, Oregon, Toronto, and the U.S. International Trade Commission for helpful comments. 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 Bank of New York, the Board of Governors, or its research staff. All results have been reviewed to ensure that no confidential information is disclosed.

†Federal Reserve Bank of New York; sebastian.heise@ny.frb.org.

‡Board of Governors of the Federal Reserve System; justin.r.pierce@frb.gov

§University of Tennessee; gschaur@utk.edu

¶Yale School of Management & CEPR & NBER; peter.schott@yale.edu

1 Introduction

Between 1990 and 2008 the rapid expansion of global value chains promoted a substantial increase in world trade and unprecedented convergence in rich and poor country incomes (WorldBank, 2019). Since the financial crisis, however, increasing restrictions have impeded trade and dampened firms' enthusiasm for expanding international supply chains, threatening export-led growth in the developing world and reducing welfare in a variety of ways, as documented in Amiti, Redding, and Weinstein (2019), Fajgelbaum, Goldberg, Kennedy, and Khandelwal (2019), Flaaen and Pierce (2019), and Flaaen, Hortaçsu, and Tintelnot (2020). In this paper, we focus on a subtler impact of these shifts in policy, demonstrating that mere increases in the probability that trade might be restricted can re-shuffle supply chains and reduce welfare by raising firms' international procurement costs.

In the first part of the paper, we develop a model of global sourcing that builds upon the closed-economy supply-chain framework introduced by Taylor and Wiggins (1997), where buyers choose between two procurement systems to minimize the cost of guaranteeing high-quality inputs from sellers. Under the "Japanese" system, buyers motivate a single seller to maintain high input quality by committing to smaller, more frequent purchases at a price above cost over a long-term relationship. In the opposing "American" system, buyers choose larger, less frequent purchases at lower prices from a parade of different sellers in the spot market, and then use costly inspection and contracting to deter cheaters from shipping low quality.[1] "American" buyers are therefore distinguished by the larger number of sellers with which they engage, and lower inspection and contracting costs favor the "American" system.

We extend Taylor and Wiggins (1997) to international procurement by linking domestic importers' ability to purchase inputs from foreign sellers to changes in the probability of a trade war. In equilibrium, each buyer procures and distributes its product using the system that minimizes costs. We show that increases in the probability of a trade war reduce the likelihood that buyers choose "Japanese" procurement because it shortens the expected length of buyer-seller relationships, thereby raising the premia buyers must pay sellers to incentivize high quality.

In the second part of the paper, we employ confidential transaction-level U.S. import data from the U.S. Census Bureau to characterize the prevalence of "Japanese" sourcing among U.S. importers and examine the implications of the model. These data are well-suited for our inquiry as they record both the number of foreign exporters with which U.S. importers trade, as well as quantities and prices associated with each shipment. In our first exercise, we examine the prevalence of "Japanese" importing by source country. Guided by our model, we classify U.S. importing firms as either "Japanese" or "American" based upon the number of foreign suppliers from which they purchase a particular product from each trading partner over the sample period we observe, 1992 to 2016. We find that "Japanese" importing during the late 1990s is most prevalent from Japan and Mexico, and that it increases most substantially in the early 2000s for Brazil, China and Mexico.

Examining patterns of shipments, we then find, consistent with the model, that buyers with fewer

[1] A global movement towards "Japanese" procurement is documented in a series of studies (O'Neal, 1989; Heide and John, 1990; Lyons, Krachenberg, and Henke Jr., 1990; Dyer and Ouchi, 1993; Han, Wilson, and Dant, 1993; Helper and Sako, 1995; Liker and Choi, 2004).

sellers per shipment – i.e., those that are more "Japanese" – do indeed receive smaller, more frequent shipments at a higher price than "American" buyers of the same product. Moreover, we demonstrate that goods with low versus high inspection costs, proxied by their status as commodities versus differentiated goods (Rauch, 1999), are shipped more frequently. To our knowledge, these results provide the first systematic empirical evidence supporting the existence of different procurement systems between buyers and sellers as motivated by Taylor and Wiggins (1997).

In the third part of the paper we compare U.S. importing relationships before and after the 2001 US extension of Permanent Normal Trade Relations (PNTR) towards China, a change in policy that effectively eliminated – at least until the recent period – the annual threat of substantial tariff increases similar to a trade war. Exploiting the fact that the implications of this reduction in uncertainty varied substantially across products, we employ a triple difference-in-differences specification that asks whether U.S. importers' procurement patterns change after the policy is implemented (first difference), for imports from China relative to other countries (second difference), and for products with greater relative exposure to PNTR (third difference).[2] Consistent with the model, we find that shipments of more-exposed products from China do indeed become relatively smaller, more frequent, and more expensive after 2001, consistent with a switch towards "Japanese" procurement. Coefficient estimates suggest that a one standard deviation increase in exposure to PNTR is associated with a relative decline in average shipment quantity and increase in frequency of 4.5 percent, and a relative increase in average shipment price of 2.1 percent.

In the final section of the paper, we embed our model of procurement in an extension of Eaton and Kortum (2002)'s multi-country, general equilibrium model of international trade and analyze the change in trade flows and the welfare costs associated with an increase in trade conflict. In the resulting framework, trading patterns are governed by bilateral probabilities of trade peace in addition to standard cross-country differences in productivity, as they alter the relative attractiveness of the "Japanese" system across trading partners. Quantitative simulations of the model incorporating trends highlighted in our empirical analysis reveal that eliminating "Japanese"-style procurement increases domestic sourcing at the expense of imports. The associated increase in final goods prices lowers U.S. welfare by more than 3 percentage points. The relevance of this scenario is underscored by recent, sharp increases in policy uncertainty in U.S.-China trade and in trade between the United Kingdom and the European Union.

Our analysis makes contributions to several literatures. First, as discussed, we highlight a new channel —changes in firms' procurement patterns—through which changes, or even *potential* changes, in tariffs can affect trade patterns and welfare. Our finding of the relevance of the threat of tariff increases for firms' procurement decisions, therefore, contributes to a literature focusing on the implications of trade policy uncertainty, including (Handley, 2014; Handley and Limão, 2017; Alessandria, Khan, and Khederlarian, 2019). Similarly, both our theoretical and empirical results suggest a novel channel for the welfare implications of trade wars (Ossa, 2014). Lastly, we note that the changes in firms' supply chains associated with switching international procurement systems represent a new

[2]In our model, buyer and seller trade a single product, so the probability of a trade war and the probability the buyer-seller relationship ends are the same. Our empirical analysis, on the other hand, examines firms trading a wide range of products subject to varying potential increases in tariffs prior to the granting of PNTR to China.

channel affecting the expansion and deepening of global value chains (Antràs, Fort, and Tintelnot, 2017; Antràs and Chor, 2018; Alfaro, Antràs, Chor, and Conconi, 2019).

Our work also relates to the much larger literature exploring the effects of changes in trade policy or trade flows on country, industry and worker outcomes.[3] Within this literature, evidence linking trade liberalization to prices is particularly rare (Goldberg and Pavcnik, 2016), despite its importance for understanding the evolution of quality, markups, and overall welfare (De Loecker and Goldberg, 2014; Antoniades, 2015; De Loecker, Goldberg, Khandelwal, and Pavcnik, 2016; Manova and Yu, 2017; Kugler and Verhoogen, 2012; Verhoogen, 2008; Khandelwal, 2010). Here, we find that a reduction in probability of a trade war leads to an increase in unit values as firms switch from "American" to "Japanese" procurement systems, highlighting an additional potential price response to trade liberalization beyond those previously studied.

Third, our results contribute to the literature on incomplete contracts, imperfect contract enforcement and information asymmetries that primarily focuses on buyer-seller relationships within countries.[4] In international economics, papers tend to concentrate on cross-border merger and acquisition as a solution to these impediments (Antràs, 2003, 2005; Antràs and Helpman, 2008; Feenstra and Hanson, 2005; Fisman and Wang, 2010; Grossman and Helpman, 2004; Spencer, 2005). Recent research outside this tradition by Defever, Fischer, and Suedekum (2016) and Kukharskyy (2016) investigates the use of relational contracting, i.e., repeated transactions between buyers and sellers under an informal agreement.[5] In these efforts, the pattern of trade between buyers and sellers is governed by idiosyncratic time preferences. Here, by contrast, discount rates are common and firms choose between procurement systems based on inspection costs and the probability of a trade war. As a result, our model links shipment patterns to trade policies in a manner amenable to empirical inquiry using transaction-level trade data.[6]

Fourth, our analysis adds to recent empirical studies using shipment-level trade data to study trade frictions and inventory costs. Existing research shows that high fixed per-shipment trade costs reduce shipping frequency, thereby raising inventories in a manner that can have substantial impact on how firms adjust to trade shocks (Kropf and Sauré, 2014; Hornok and Koren, 2015b,a; Békés, Fontagné, Muraközy, and Vicard, 2017; Alessandria, Kaboski, and Midrigan, 2010, 2011). Here, we show that trade policy uncertainty can be an important barrier to firms' efforts to reduce inventory costs. An interesting question for future research is whether minimizing such costs provides an additional motivation for lobbying for (more dependable) free trade.

Finally, we contribute to the burgeoning literature using general equilibrium models to evaluate changes in trade policy by extending Eaton and Kortum (2002) in two directions. First, as mentioned

[3] See Goldberg and Pavcnik (2016) for a survey of these studies.

[4] Recent research on domestic procurement includes Tadelis and Zettelmeyer (2015), Cicala (2015), Bajari, Houghton, and Tadelis (2014), Macchiavello and Morjaria (2020), and Kukharskyy and Pflüger (2010).

[5] Macchiavello and Morjaria (2015) examine the value of reputation for reliability among Kenyan rose exporters. While that market relies on relational contracting, the authors do not investigate transaction patterns.

[6] In contrast to heterogeneous-firms models of trade and investment (Melitz, 2003; Bustos, 2011), the fixed and variable costs of trade in our setting are endogenous to firms' choice of procurement system. Our results also related to a broader set of empirical papers analyzing trade flows among importer-exporter relationships (Monarch and Schmidt-Eisenlohr, 2017; Huneeus, 2018; Kikkawa, Magerman, and Dhyne, 2019; Heise, 2019; Alviarez, Fioretti, Kikkawa, and Morlacco, 2021), though not in the context of procurement system choice.

above, we have product prices depending on the probability of a trade war as well as supplier productivity. Second, we allow for increasing returns to scale among sellers as, in our setting, sellers incur fixed logistics costs in addition to variable costs in producing each shipment. We account for the resulting declining average cost curves by assuming sellers compete in a contestable market (Baumol, Panzar, and Willig, 1982), a natural extension to price competition in the presence of economies of scale (Tirole, 1988). We use the model to quantify how a reduction in buyers' ability to form "Japanese" relationships reduces welfare by raising prices, similar to a negative productivity shock. While we develop the model in the context of changes in the probability of trade peace, the underlying manner by which trade regimes affect productivity applies more broadly to any policies that might undermine sellers' beliefs about the viability of establishing long-term relationships with buyers, e.g., uncertainty over the arrival of shipments due to corruption, or weather, at ports.

The remainder of this paper proceeds as follows. Section 2 outlines our theoretical model. Section 3 describes the data, and Section 4 presents empirical evidence for Japanese style procurement. Section 5 analyzes the effect of a change in trade policy uncertainty. Section 6 contains our quantitative simulations. Section 7 concludes. An online appendix contains additional results.

2 Theoretical Model

As noted in the introduction, incomplete contracts, information asymmetries and contract enforcement are common problems associated with firms' purchases of intermediate inputs. While firm integration is one means of addressing these issues (Antràs, 2003, 2005; Antràs and Helpman, 2008), it may not be a viable solution if cross-border integration involves substantial costs.[7] As a result, we focus on an alternate, arm's-length solution to the quality control problem pioneered by Japanese firms such as Toyota.

Our starting point is the setup introduced by Taylor and Wiggins (1997), in which firms overcome contractual and informational frictions using one of two procurement systems. Under the "American" system, buyers use competitive bidding to select the lowest-cost supplier for each shipment of inputs, and use the threat of inspection to deter provision of low-quality goods. Under the "Japanese" system, by contrast, buyers offer sellers incentive premia over long-term relationships to ensure inputs of sufficient quality. We extend Taylor and Wiggins (1997) to an international trade context by linking these incentive premia to the probability of trade peace, and by endogenizing final demand. These extensions allow us analyze the effect of a change in the probability of trade peace on a host of outcomes, including trade flows and welfare.

[7] China, for example, requires foreign ventures to include a domestic partner, while the United States (and other developed countries) mandate national security reviews of foreign acquisitions in certain sectors.

2.1 The Contracting Problem

2.1.1 The Seller's Problem

There is a single country populated by a continuum of homogeneous sellers able to produce the same good.[8] To complete a production run (i.e., produce one shipment) sellers hire labor l at wage $w = 1$ to produce and deliver output $x = \frac{\Upsilon}{\theta} l$, where Υ is a seller's productivity and θ represents the seller's product's level of quality. The unit input requirement, $\frac{\theta}{\Upsilon}$, allows for variation in quality, giving rise to a "quality control" problem.[9] Sellers choose between discrete quality levels, $\theta \in \left\{\underline{\theta}, \overline{\theta}\right\}$, where lower quality is less costly to produce. To complete the shipment, the seller absorbs f units of labor for per-shipment specific logistics services, including transport costs.[10] The sellers' total cost for each production and delivery cycle are therefore $x\frac{\theta}{\Upsilon} + f$.

2.1.2 The Buyer's Procurement Choices

There exist multiple homogeneous buyers that are willing to procure the seller's output and distribute it downstream in the consumer market. These buyers compete in a contestable market, described in greater detail below. Conditional on desired quality, $\overline{\theta}$, consumer demand arrives continuously. Let t denote continuous time and consider time periods $\Delta t = \int_0^1 1dt = 1$, e.g., 1 year. To supply the consumer market over one time period, a buyer procures total quantity, q, in a series of discrete, equally sized, symmetric shipments of size x. We take q as fixed for now, but solve for it in equilibrium in Section 2.3. Consequently, there are q/x shipments during each period. Figure 1 summarizes the shipment and consumption pattern visually. If quality is less than desirable, no downstream consumer demand arrives for the product and buyers must dispose of the obsolete shipment without recompense. To avoid these losses, the buyer seeks to ensure the provision of high-quality inputs using either an "American" (A) or a "Japanese" (J) procurement system.

In the "American" (A) system, buyers pay fixed cost m_A to inspect each shipment's quality before delivery. We assume inspections reveal quality with certainty.[11] Therefore, if buyers inspect, sellers cannot gain by cheating on product quality. Consequently, buyers know that inspections guarantee product quality. We assume that buyers have all the bargaining power. As a result, given an order of size x_A placed with a seller, the buyer sets the per shipment price $v_A(x_A, \overline{\theta})/x_A$ to allow the seller to break even and participate, where

$$v_A(x_A, \overline{\theta}) = f + \frac{\overline{\theta}}{\Upsilon} x_A. \tag{1}$$

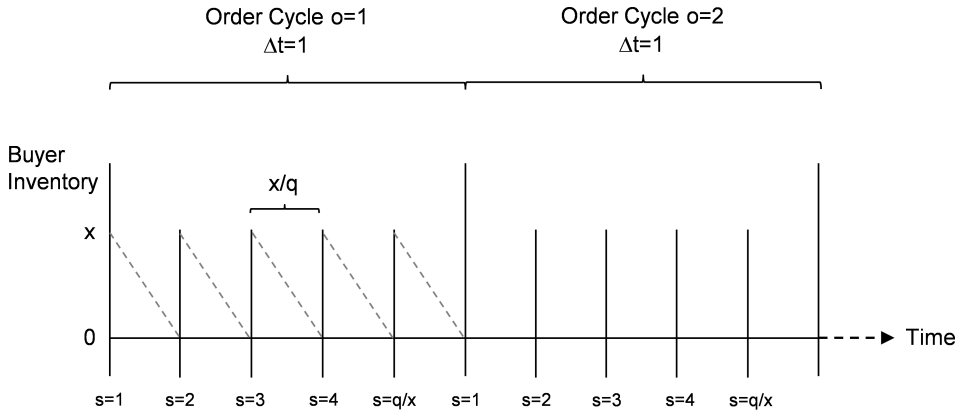
[8] We extend the model to multiple products and sellers in multiple countries in Section 6 below.

[9] See, for example, "Poorly Made," The Economist, May 14th, 2009.

[10] While a large literature within international trade models the cost of shipping between origin and destination as "iceberg", i.e., a fractional loss of the shipped good which rises with distance, recent evidence supports per-unit and per-shipment specific delivery costs (Hummels and Skiba, 2004; Martin, 2012; Kropf and Sauré, 2014; Hornok and Koren, 2015b,a).

[11] Taylor and Wiggins (1997) allow for a more general inspection pattern and show that optimal inspection frequency under the American system is a function of shipment size and quality. This generalization does not affect conclusions below related to per-shipment costs. Our simplifying assumption allows us to derive the properties of the model via implicit function techniques, as opposed to relying on the near-linearity of the problem when r/q is small as in Taylor and Wiggins (1997).

Figure 1: Timing



Notes: The total quantity shipped over an order cycle is q. Order cycles repeat indefinitely and are indexed by $o = \{1, 2, ...\}$. There are $s = \{1, 2, ..., q/x\}$ shipments during an order cycle, each arriving every x/q units of time apart.

Due to the fixed cost, the buyers' average procurement costs are decreasing in order size, and therefore each buyer optimally places each order with a single seller. Since the sellers are homogeneous and all willing to supply at the same price, we assume that for a given buyer the winning seller is chosen randomly for each order. Inclusive of inspection costs, the buyer's total procurement expense equals $v_A(x_A, \theta) + m_A$. The form of this procurement cost is similar to those appearing elsewhere in the literature (Kropf and Sauré, 2014; Hornok and Koren, 2015b,a), where exogenous per-shipment fees such as f and m_A capture administrative barriers.

"Japanese" (J) procurement motivates the production of high quality via payments of a premium over a seller's cost over the course of a long-term relationship. A seller chooses to ship high quality if a long-run relationship with the buyer is of sufficient value, and a contribution of this paper is to show how this value depends upon the stability of trade policy between seller and buyer countries. We assume trade policy shocks that break buyer-seller relationships, e.g., an escalation of the tariff on the product to a prohibitive level, arrive at a constant rate, ρ.[12] In that case, relationships break before time t with probability $F(t) = 1 - e^{-\rho t}$, implying survival over a shipment cycle with probability $e^{-\frac{\rho x}{q}}$.[13] We note that while our focus is on trade policy, there are many other shocks which might have similar effects. For example, natural disasters may affect suppliers' ability to reliably fulfill their implicit agreements with buyers in foreign countries (Boehm, Flaaen, and Pandalai-Nayar, 2019).

For our purposes, $e^{-\frac{\rho x}{q}} < 1$ implies that firms are uncertain about whether future trade policy facilitates relationships: for a given shipping cycle $\frac{x_J}{q}$, a greater arrival rate of trade wars, ρ, increases

[12] Ossa (2014), for example, estimates that the optimal tariffs countries might set in the event of a trade war are substantial, averaging 63 percent worldwide.

[13] For proof, see Wooldridge (2002), page 688. Note that the model considers trade in a single product. An alternate interpretation of ρ related to our data analysis below is that it reflects the overall probability of a buyer-seller relationship becoming unprofitable due to a change in trade policy. That is, ρ reflects both the probability of a trade war (which is the same for all products) and the magnitude of the subsequent rise in tariffs (which might vary across products). The probability of relationship breakup is rising in both of these factors.

the separation probability. Let r be the per-period interest rate and $v_J(x_J, \theta)$ be the payment the buyer sets under the "Japanese" system for each shipment. With continuous compounding, the expected discounted value of the relationship over all future shipment cycles is then $\frac{v_J(x_J,\bar{\theta})}{1-e^{-(r+\rho)x_s/q}}$.[14] Note that here, in contrast to the "American" system, a buyer procures each order from the same seller.

To guarantee desired quality, the buyer must set a per-shipment payment such that the seller's net present value of the continued relationship exceeds the one-time profit from cheating by supplying inferior quality,

$$\frac{v_J(x_J, \bar{\theta}) - f - \frac{\bar{\theta}}{\Upsilon} x_J}{1 - e^{-(r+\rho)x_s/q}} \geq v_J(x_J, \bar{\theta}) - f - \frac{\theta}{\Upsilon} x_J. \tag{2}$$

In this expression, we assume that if the seller provides low quality, the buyer does not find out about it until after the shipment is received and the payment is made, and that the seller delivering low quality is excluded from the market forever. Here, too, we assume buyers own all the bargaining power. Solving (2), buyers under the "Japanese" system set the per-shipment payment to be

$$v_J(x_J, \bar{\theta}) = f + \bar{\theta}\frac{1}{\Upsilon} x_J + \left[e^{(r+\rho)x_J/q} - 1\right](\bar{\theta} - \underline{\theta})\frac{1}{\Upsilon} x_J. \tag{3}$$

In the "Japanese" system buyers pay the per-unit premium $\left[e^{(r+\rho)x_J/q} - 1\right](\bar{\theta} - \underline{\theta})\frac{1}{\Upsilon}$ to induce the seller to provide high quality.

A key feature of the "Japanese" system, therefore, is that more stable trade relationships (i.e., lower separation rates) and smaller shipments sent more frequently (which increase the present discounted value of payments) reduce the premium necessary to guarantee desired quality.[15] As a result, compared to a setting without incentive problems (e.g., an integrated firm), "Japanese" procurement has higher variable costs while American procurement system has higher fixed costs.

Buyers choose between the "American" and "Japanese" system by comparing long-term expected revenues and costs. We assume that a trade war causes buyers to exit irrespective of the system they choose, as they lose access to the suppliers of their good.[16] At a given market price p, long-term expected profits in the two procurement systems are given by

$$\pi_s^b = \left[\int_0^{x_s/q} e^{-rt} pq\, dt - v_s(x_s, \bar{\theta}) - m_s\right] / \left[1 - e^{-(r+\rho)x_s/q}\right] \quad \text{s}\in\text{J,A} \tag{4}$$

where discounted revenues at the beginning of each shipment cycle are $\int_0^{x_s/q} e^{-rt} pq\, dt$ and $m_J = 0$.

[14] The discount rate over a shipping cycle with associated continuous discount factor is $\lim_{N\to\infty} \left(\frac{1}{1+\frac{rx}{q}/N}\right)^N = e^{-\frac{rx}{q}}$.

[15] An alternative approach to incorporating trade policy uncertainty would be to multiply the discount factor by an exogenous probability of trade peace $(1-\rho)$. However, a drawback of that approach is that the probability of relationship separation over a given time period is dependent of the number of shipments made. In our formulation, the likelihood of separation is independent of shipment frequency.

[16] In the model with multiple sellers discussed below, an alternate assumption is that buyers switch to a seller from another country in the event of a trade war. Given that buyer profits are zero in equilibrium, however, this assumption is equivalent to assuming that the buyer exits.

2.2 Market Equilibrium and Optimal Procurement Choice

We now study the optimal procurement pattern and system for a given fixed order quantity q. Under the assumption that profits are zero in equilibrium, the market price must equal average procurement and distribution costs, $AC_s(x_s, q)$. We will endogenize order quantity and describe a market protocol that delivers zero profits in Section 2.3. Setting profits equal to zero and solving for prices, we obtain zero-profit conditions conditional on the procurement system,

$$p_s = AC_s(x_s, q) = \left(\frac{r}{q}\right) \frac{v_s(x_s, \bar{\theta}) + m_s}{\left[1 - e^{-rx_s/q}\right]} \quad \text{s}\in\text{J,A}, \tag{5}$$

where $m_J = 0$. Buyers choose a shipment size to minimize average procurement costs within each procurement system. Taking first order conditions (FOC_s) for each system and setting them to zero we obtain,

$$\frac{v_s'(x_s, \bar{\theta})}{1 - e^{-rx_s/q}} = \frac{\left[v_s(x_s, \bar{\theta}) + m_s\right] \frac{r}{q} e^{-rx_s/q}}{\left(1 - e^{-rx_s/q}\right)^2} \quad \text{s}\in\text{J,A}. \tag{6}$$

This expression implicitly determines shipment size, x_s^*. The left hand side represents the discounted value of higher costs associated with a small increase in order size. The right hand side measures the savings from an increased discount factor due to spacing these larger orders further apart in time. Trading off these costs and benefits, the firm optimally procures x_s^* to minimize average expected purchasing costs. We present all proofs in Appendix A, and show in Appendix A.1.1 that an interior solution to the first order condition is a unique cost minimizer for $0 < rx/q < 1$ under both procurement systems.

Conditional on procurement quantity q and parameter values, the buyer compares average procurement costs evaluated at the optimum, $AC_s(x_s^*, q)$, across systems to determine the cost-minimizing procurement system. Implicit function techniques provide intuition for this comparison in the presence of non-linearity. With no variation in quality, i.e., for $\bar{\theta} - \underline{\theta} = 0$ and $m_A = 0$, there is no incentive problem and costs in both systems are identical. Compared to this benchmark case, differentiating equation (5) under the "Japanese" system with respect to $\underline{\theta}$ and ρ, respectively, and noting that, by the envelope theorem, the indirect effect coming from the resulting change in x_J is zero, we find that average procurement costs in the "Japanese" system increase with the arrival rate of trade wars, ρ, and with the range of potential qualities, $\bar{\theta} - \underline{\theta}$, due to the greater incentive premia they necessitate, $\frac{\partial AC(x_J^*)}{\partial \underline{\theta}} \leq 0$ and $\frac{\partial AC(x_J^*)}{\partial \rho} \geq 0$.[17] In the "American" system, differentiating (5) with respect to m shows that average costs increase with inspection costs m. Importantly, as $m \to \infty$, we have $AC_A(x_A^*, q) \to \infty$ because average costs grow without bound, $\frac{\partial AC_A(x_A^*)}{\partial m} = \frac{1}{1 - e^{-\frac{rx_A^*}{q}}} > 1$. This result implies the following proposition.

Proposition 2.1. *For $\bar{\theta} - \underline{\theta} > 0$ and $\rho > 0$, there is always a threshold value $m^* \in (0, \infty)$ for inspection costs such that average procurement costs in both systems are the same. This point is the cut-off at which the buyer switches systems: the American system is chosen for $m < m^*$, and the*

[17] See Appendix Section A.1.2 for the proof.

"Japanese" system is chosen for $m > m^*$.

Proof. See Appendix A.1.3. □

This proposition highlights that changes in the arrival rate of trade wars may endogenously affect the choice of procurement system. Starting at a level of inspection cost m slightly below m^*, a reduction in ρ lowers average costs under the "Japanese" system – while also lowering m^* – and may cause the buyer to switch from the American to the "Japanese" system. This setting contrasts with existing studies of relational contracts in trade, where exogenous heterogeneity in discount rates determine relationship-based transactions (Kamal and Tang, 2015; Defever, Fischer, and Suedekum, 2016; Kukharskyy, 2016).[18] In our framework, buyers endogenously determine the effective discount rate of rx_x/q by choosing the optimal procurement system and order size in response to inspection costs and the probability of a future trade conflict.

To map the choice of procurement system into observable trade flows, we examine how order size, frequency, and unit values differ across the two systems as a function of the trade war arrival rate ρ and the inspection cost m for a given q. We restrict our attention to a setting where buyers make a purchase at least once per period, $x^* \leq q$, and where discount rates are small, i.e., $0 < \frac{rx_s}{q} < 1$.

Proposition 2.2. *An increase in the probability of a trade war, which increases ρ, raises the unit value per shipment and reduces the size of shipments (i.e., raises shipment frequency) in the "Japanese" system. An increase in the inspection cost m raises shipment sizes, but reduces shipping frequencies and unit values in the American system.*

Proof. See Appendix A.1.4. □

While an increase in the separation rate ρ does not affect procurement patterns under the American system, under the "Japanese" procurement system it causes sellers to demand a greater premium to maintain quality. As a result, variable procurement costs increase and buyers re-optimize by lowering shipment sizes and raising shipping frequency. Given that fixed per-shipment costs are spread over smaller shipment sizes, the increase in ρ causes unit values to increase. On the other hand, an increase in the inspection cost m raises fixed per-shipment costs, and buyers under the American system re-optimize by decreasing shipping frequency and increasing per-shipment quantities. As a corollary, unit values must go down in the American system since fixed costs are spread over more units.

We can use these results to rank shipping frequencies and unit values across the two systems. If the quality range, $\bar{\theta} - \underline{\theta}$, and inspection costs, m_A, are both zero, then the "American" and "Japanese" procurement systems are identical. As $\bar{\theta} - \underline{\theta}$ rises, variable shipment costs increase under the "Japanese" system due to the larger incentive premium, which raises the marginal benefit of increasing the shipping frequency and lowering shipment size. Under the "American" system, by Proposition 2.2, shipment size increases, and hence shipping frequency decreases, with inspection

[18] Defever, Fischer, and Suedekum (2016), for example, study a setup in which a buyer and a seller interact repeatedly and choose how much to invest into their relationship. They analyze different values of firms' discount rates and show that a cooperative equilibrium can only exist if both firms are sufficiently patient.

costs. For the case of $\bar{\theta} - \underline{\theta} > 0$ and $m \geq 0$, it must therefore be true that shipping frequencies are greater in the "Japanese" system and shipping sizes are greater in the American system. Furthermore, unit values are greater in the "Japanese" procurement system compared to American procurement system. This reasoning forms the basis of our second proposition.

Proposition 2.3. *Batch sizes in the American system are greater than in the "Japanese" system, $x_A^* > x_J^*$, and unit values in the "Japanese" system are greater than in the American system, $v_J(x_J, \bar{\theta})/x_J > v_A(x_A, \bar{\theta})/x_A$.*

Proof. See Appendix A.1.5. □

Proposition 2.2 and Proposition 2.3 state that within the "Japanese" system a decrease in the probability of a future trade conflict, ρ, lowers unit values, raises shipment size and lowers shipping frequency. These patterns should therefore be observed in U.S. trade data if prior to a decline in the probability of a trade war most U.S. trade relationships were already "Japanese". In contrast, if "American" procurement systems were prevalent in U.S. trade, then we would expect that a decrease in the probability of a future trade conflict would lead some relationships to switch from American to "Japanese" procurement, which would lead to higher unit values, greater shipping frequencies, and smaller shipment sizes. In that case, it would be important to distinguish between procurement systems to explain how import patterns adjust to trade policy uncertainty. In Section 5 we show that "Japanese" importing from China was relatively rare prior to a change in trade policy in 2001 that reduced uncertainty with respect to shipments from China, and that the trade patterns after the policy change suggest that it led to the switching of systems rather than adjustment within the "Japanese" system.

We now examine the comparative statics with respect to q as a precursor to solving for q in the next section. We find that an increase in q raises shipment size and shipping frequency in both procurement systems. Intuitively, for a given fixed shipping frequency, firms must increase the batch size x in both systems to meet an increase in q. But by the first-order condition, equation (6), we know that firms trade-off variable procurement costs against fixed per-shipment costs. Therefore, as variable future expected procurement costs increase, buyers respond by spreading the larger quantities over more shipments. As a result, larger quantities purchased lead to greater shipment sizes, but also greater order frequencies and shorter shipment cycles. It follows that unit values in both procurement systems decrease, since fixed per-shipment costs are spread over greater per-shipment quantities. In addition, under the "Japanese" system, an increase in the shipping frequency implies a lower premium to motivate desired quality. We summarize these results with the following proposition.

Proposition 2.4. *An increase in the procurement target q raises batch sizes x_s^* and the shipping frequency q/x_s^* in both systems, and, as a corollary, lowers unit values in both systems.*

Proof. See Appendix A.1.6. □

2.3 Endogenous q Under a Contestable Market

We now consider the determination of q in equilibrium. As we show below, our setup features increasing returns to scale as a result of the fixed cost component in seller's costs. Consequently, our setup generates a natural monopoly, and we cannot assume perfect competition as in many standard trade models which assume constant marginal costs. We therefore assume a contestable market in which homogeneous buyers can freely enter and exit. This assumption is a natural extension of Bertrand competition when firms' (i.e., potential buyers') costs exhibit economies of scale (Baumol, Panzar, and Willig, 1982; Tirole, 1988). Two key features of a contestable market are (i) firms set prices optimally equal to average costs; and (ii) market prices are disciplined by entry if profits are positive, i.e., setting a higher price would lead firms to lose the market to an entering competitor. Our use of these features here is analogous to the assumption of perfect competition in benchmark models of international trade, e.g., Eaton and Kortum (2002).[19]

The trade-off between shipment size and frequency described in Proposition 2.4 has implications for the characteristics of the procurement technologies. As firms order larger batches in each shipping cycle (i.e., q rises), fixed per-shipment costs are spread over more units. In Lemma 2.5, we show that these scale effects dominate under both procurement systems, so that total average costs are downward sloping in q:

Lemma 2.5. *At the optimal order size x_s^*, both procurement systems provide economies of scale, i.e., $\frac{\partial AC(x_s^*,q)}{\partial q} < 0$. As a result, in equilibrium a single buyer procures the product acting as the sole distributor. The second derivative of the average cost with respect to q is positive, $\frac{\partial^2 AC(x_s^*(q),q)}{\partial q^2} > 0$, and the average cost in both systems reaches a positive and finite limit as $q \to \infty$.*

Proof. See Appendix A.1.7. □

Lemma 2.5 implies that average cost curves are convex and that a sustainable and feasible single-buyer equilibrium is characterized by zero profits, no incentive for entry or exit, and cleared markets. A demand curve that uniquely intersects the single buyer's optimized average cost curve from above determines a unique sustainable and feasible equilibrium, q^*. The buyer prices and supplies the market along its average cost curve. Therefore, no firm can undercut the incumbent. If the buyer prices off the average cost curve, then either entrants contest the positive profits or the buyer realizes negative profits. Because consumers are willing to pay prices greater than average costs for $q < q^*$, potential entry forces the incumbent to lower prices and increase quantity such that the market clears where supply equals demand. The buyer is not willing to procure a greater quantity because she would incur losses.

Standard linear or CES demand systems may intersect downward sloping average costs multiple times. In that case, the intersection that determines the greatest equilibrium quantity must cut average costs from above. Otherwise, there is no equilibrium. If demand were to cut from below, then the buyer would procure an infinite quantity because average costs lie below demand forever. Therefore, under appropriate assumptions on the demand system, the market equilibrium is a corollary of Lemma 2.5.

[19] Here, as there, it also facilitates analysis of multiple countries and products, which we explore in Section 6.

Corollary 2.5.1. *If markets are contestable and demand intersects average costs from above at q^* and remains below average costs as $q^* < q \to \infty$, then a single buyer procures the product from the seller and distributes it on the consumer market using the buyer's cost minimizing procurement system at optimal shipping frequencies.*

3 Data

In this section, we test the implications of our theory using confidential data from the U.S. Census Bureau's Longitudinal Foreign Trade Transaction Database (LFTTD). These data track every U.S. import shipment from 1992 to 2016 and include: the dates the shipment left the exporting country and arrived in the United States; identifiers for the U.S. and foreign firm conducting the trade; shipments' value and quantity; a ten-digit Harmonized System (HS10) code classifying the product traded; the country of origin of the exporter; and the mode of transport.[20]

We refine the raw data as follows. First, we drop all transactions that are warehouse entries, so that our dataset represents imports for consumption. Second, we remove all transactions that do not include an importer identifier, an exporter identifier, an HS code, a value, a quantity or a valid transaction date. We also drop observations with invalid exporter identifiers, e.g., those that do not begin with a letter (it should start with the country name) or that have fewer than the requisite number of characters. Third, given our focus on arm's-length trade, we exclude from our analysis all related-party transactions.[21] We choose a conservative approach and exclude all relationships in which the two parties ever report being related, and all observations for which the related-party identifier is missing. Fourth, we use the concordance developed by Pierce and Schott (2012) to create time-consistent HS codes so that purchases of goods can be tracked over time. Fifth, we deflate transaction values using the quarterly GDP deflator of the Bureau of Economic Analysis, so that all values are in 2012 real dollars.[22] Finally, since shipments of the same product between the same buyer and seller spread over multiple containers are recorded as separate transactions, we aggregate the dataset to the weekly level. Thus, multiple transactions in the same week will be summed to one. We perform this aggregation to ensure that each observation in our data reflects a genuinely new transaction rather than being part of a larger shipment.

Table 1 provides an overview of the data. From 1992 to 2016, the U.S. imported 5.68 trillion dollars worth of goods at arm's length, the majority of which arrived by water (vessel). These imports span 360 thousand unique US importers and just over 5 million unique foreign exporters. Our analysis below will focus on "buyer quadruples" that group shipments of a ten-digit HS product (h) imported by a U.S. importer (m) from origin country (c) shipped via mode of transportation (z).[23] As indicated

[20] The four main modes of transportation are vessel, rail, road, and air. We drop the small fraction of transactions that are transported by other means, e.g., hand-carried by passengers. For further information on the LFTTD, see Bernard, Jensen, and Schott (2009) and for detailed information on the foreign firm identifier, see Kamal and Monarch (2018). Though foreign identifiers in the LFTTD include information about country of origin, we continue to mention this dimension of the data for clarity.

[21] The Census Bureau defines parties as related if either party owns, controls or holds voting power equivalent to 6 percent of the outstanding voting stock or shares of the other organization.

[22] https://fred.stlouisfed.org/series/GDPDEF

[23] We include mode of transportation in defining these bins to mitigate the influence of spurious sources of variation

in the table, there are almost 3 million such $mhcz$ quadruples between 1992 and 2016. The final row of the table reports the number of "buyer-seller relationships" associated with these bins, i.e., the number of $mxhcz$ quintuples, where x denotes the exporter.[24] There are nearly 22 million of these relationships within the 3 million buyer quadruples, or an average of about 7 sellers per $mhcz$ cell.[25]

Table 1: U.S. Import Transaction Summary Statistics

Total Imports ($\$Bill$)	5,680
Vessel Imports ($\$Bill$)	4,030
Air Imports ($\$Bill$)	988
Unique Importers	360,000
Unique Exporters	5,037,000
Unique Importer-Product-Country-Mode Quadruples	2,966,000
Unigue Exporter-Importer-Product-Country-Mode Relationship	21,700,000

Source: LFTTD and authors' calculations. Table summarizes U.S. arm's-length imports from 1992 to 2016. Import values are in billions of real 2012 dollars. Vessel imports refer to imports arriving over water. Final four rows of tables provide counts of unique importers, exporters, buyer quadruples, i.e., U.S. importer by HS product by origin country by mode of transport cells, and buyer-seller relationships, i.e., U.S. importer by foreign exporter by HS product by origin country by mode of transport cells

Table 2 summarizes the $mhcz$ quadruples in greater detail. The first four rows of the table reveal that from 1992 to 2016, the average $mhcz$ bin traded 1.9 million dollars, lasted for 304 weeks and encompassed 39 shipments across 7 sellers. Rows 5 through 7 highlight "procurement patterns," showing that value per shipment (VPS_{mhcz}), weeks between shipments (WBS_{mhcz}), and buyer-seller relationship length ($length_{mhcz}$) averaged 36 thousand dollars, 24 weeks and 181 weeks, respectively.[26] In each case, the large standard deviations compared to the means indicate that buyer quadruples exhibit substantial heterogeneity.

– e.g., product quality – that might differ across product varieties shipped using different methods.

[24] We realize that referring to "$mhcz$ quadruples" and "$mxhcz$ quintuples" is awkward but have retained this language for its precision.

[25] In the data, a given seller (i.e., exporter) may supply a particular HS code to multiple buyers (i.e., importers). To match theory and data, we interpret this behavior as sellers producing different varieties within HS codes for each buyer without any costs to the buyer or seller beyond those described in Section 2.1. Under our contestable market assumption, buyers trying to sell a particular variety above average cost in the final goods market face competition from other buyers able to offer the same variety using an input from either the same or a different foreign seller. We assume that American buyers can procure their variety from different sellers over time, and that different buyers procuring the same product from the same seller might use different procurement systems because inspection costs can vary by variety within a product.

[26] While below we also analyze quantity per shipment (QPS_{mhcz}) and unit value per shipment (UV_{mhcz}), we do not summarize these here due to differences in the unit of quantity across products reported in U.S. import data. We note that the relationship lengths can be subject to both left and right censoring around the beginning and end of our period of analysis.

Table 2: Attributes of $mhcz$ Quadruplets

	Mean	*Standard Deviation*
Total Value Traded ($)	1,914,000	36,300,000
Total Shipments	38.6	157.9
Length Between Buyer's First and Last Shipment (Weeks)	304.3	266
Number of Sellers (x)	7.3	25.5
Value per Shipment (VPS), ($)	35,910	386,100
Weeks Between Shipments (WBS)	23.5	28.5
Average Relationship Length in Weeks ($length$)	180.8	154.7
Sellers per Shipment (SPS)	0.334	0.241

Source: LFTTD and authors' calculations. Table reports the mean and standard deviation across importer (m) by country (c) by ten-digit Harmonized System category (h) by mode of transport (z) quadruplets during our 1992 to 2016 sample period. Observations are restricted to quadruplets with more than one transaction.

As discussed in Section 2, a key characteristic of "Japanese" sourcing in the model is the presence of long-term relationships between buyer and seller, in which the buyer incentivizes the seller to provide goods of a sufficient quality by paying a premium over the seller's costs. As a result, buyers engaged in "Japanese" sourcing are expected to have fewer sellers for a particular good than those sourcing under the "American" system.[27] We build on this insight and take the number of sellers used at the level of the importer-product-country-mode ($mhcz$) quadruple as an observable measure of the extent of "Japanese" sourcing. We normalize this measure by the number of shipments to account for the fact that some buyers receive many shipments, and hence had many opportunities to source from a different seller, while others receive only one or two. Hence, we compute sellers per shipment (SPS_{mhcz}) as

$$SPS_{mhcz} = \frac{Sellers_{mhcz}}{Shipments_{mhcz}}. \tag{7}$$

Lower SPS_{mhcz} indicates "Japanese" sourcing, while higher ratios suggest "American" sourcing.

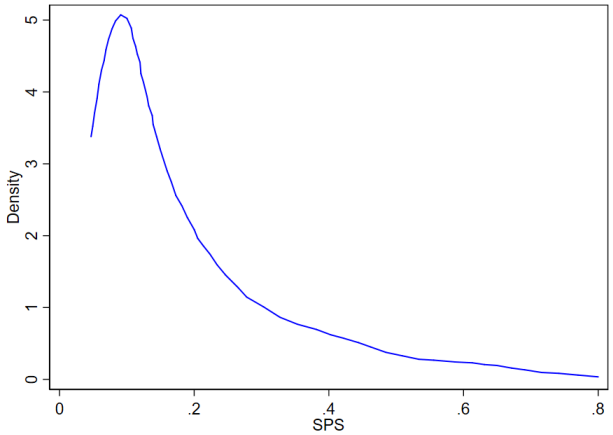
The final row of Table 2 reveals that the mean sellers per shipment across buyer quadruples is 0.33, with a standard deviation of 0.24. The kernal density reported in Figure 2 provides a more complete description of the distribution of this attribute. As indicated in the figure, most buyer quadruples have a relatively small ratio of sellers to shipments. However, observations in the right tail approach a value of 1, i.e., a different seller for each shipment.

4 American vs "Japanese" Buyers

In this section, we rationalize our use of sellers per shipment, SPS_{mhcz}, to identify procurement system. Specifically, we show that U.S. importers with lower SPS_{mhcz} – i.e., those more likely to be

[27] In the model, "Japanese" buyers have one seller for a given product. In practice, "Japanese" buyers might use more than one seller if, for example, they experiment before settling with a long-term partner. Similarly, "American" buyers might be observed to have only one seller if that seller repeatedly offers the lowest price. Nevertheless, we show below that sellers per shipment is correlated with shipment patterns in a manner consistent with the model.

Figure 2: Distribution of Sellers Per Shipment (SPS) Across Relationships, 1992 to 2016



Source: LFTTD and author's calculations. Figure displays the distribution of sellers per shipment (SPS_{mhcz}) across buyer quadruples, i.e., $mhcz$ bins, from 1992 to 2016. The figure was created according to Census Bureau guidelines and omits observations below the 5th percentile and above the 95th percentile.

"Japanese" – do indeed order smaller quantities per shipment and pay a larger average price (unit value) per shipment, than "American" buyers with higher SPS_{mhcz}.

4.1 A vs J Among Buyer Quadruples

We evaluate the appropriateness of using SPS_{mhcz} as a measure of the extent of "Japanese" sourcing by running an $mhcz$-level OLS regression,

$$ln(\overline{Y}_{mhcz}) = \beta_0 + \beta_1 ln(SPS_{mhcz}) + \beta_2 \ln(QPW_{mhcz}) + \beta_3 beg_{mhcz} + \beta_4 end_{mhcz} + \lambda_{hcz} + \epsilon_{mhcz}. \quad (8)$$

The dependent variable, $\overline{Y}_{mhcz}$, represents averages of quantity per shipment (QPS_{mhcz}), weeks between shipments (WBS_{mhcz}), and price (UV_{mhcz}) across all transactions of a given $mhcz$ quadruple in our dataset. Consistent with Proposition 2.3, we condition on the total quantity ordered over a given time horizon by controlling for QPW_{mhcz}, the total quantity traded by the buyer quadruple over its entire lifetime divided by its overall length, in weeks.[28] We also control for quadruples' first and last weeks of trade, beg_{mhcz} and end_{mhcz}, to capture potential time and duration effects. Given these controls, QPS_{mhcz} and WBS_{mhcz} both provide information on the shipping frequency, x/q in the theory, and UV_{mhcz} provides information on the price per shipment, $v_s(x_s, \bar{\theta})/x_s$. Our regression also includes product by country by mode of transportation fixed effects (λ_{hcz}).[29] These fixed effects capture time-invariant characteristics of trade along these dimensions such as distance, transit time, or level of transportation infrastructure. The sample period is 1992 to 2016, and standard errors are

[28] We normalize the total quantity traded by the number of weeks since this is straightforward to implement in our weekly dataset. An alternative would be to use the annual quantity traded, which is simply $52 \cdot QPW_{mhcz}$.

[29] beg_{mhcz} and end_{mhcz} are continuous variables indicating the week numbers that the relationship commences and ceases.

clustered at the country level. In the baseline regressions presented in this section, we exclude buyer quadruples with fewer than 5 shipments as these cells might represent importers trying out a new product, or other idiosyncrasies. [30]

Results for specification (8) are reported in Table 3. Consistent with Proposition 2.3, we find that quadruples with lower SPS_{mhcz} receive shipments that are smaller, more frequent, and higher in price. The coefficient estimate for quantity per shipment, for example, indicates that a one standard deviation decrease in sellers per shipment (0.24) is associated with a 0.10 log point decline in quantity per shipment, a 0.11 log point decline in weeks between shipment, and a 0.03 log point increase in price. Moreover, the coefficients on QPW_{mhcz} are consistent with Proposition 2.4. In both procurement systems, an increase in the total procured quantity increases shipment size, lowers the number of weeks between shipments and hence raises shipment frequency, and lowers unit import values.

Table 3: A vs J Classification Regression Across $mhcz$ Quadruples

Dep. var.	(1)	(2)	(3)
	$\log(QPS_{mhcz})$	$\log(WBS_{mhcz})$	$\log(UV_{mhcz})$
$\log(SPS_{mhcz})$	0.418^{***}	0.452^{***}	-0.123^{***}
	0.014	0.014	0.019
$\log(QPW_{mhcz})$	0.701^{***}	-0.308^{***}	-0.287^{***}
	0.010	0.010	0.009
Observations	2,966,000	2,966,000	2,966,000
Fixed effects	hcz	hcz	hcz
R-squared	0.947	0.674	0.845
Controls	beg, end	beg, end	beg, end

Source: LFTTD and authors' calculations. Table reports the results of regressing noted attribute of importer by product by country by mode of transport ($mhcz$) bins on bins' sellers per shipment (SPS_{mhcz}) and total quantity shipped per week (QPW_{mhcz}). (QPS_{mhcz}), (WBS_{mhcz}), and (P_{mhcz}) are average quantity per shipment, average weeks between shipment, and average unit value. All regressions include product by country by mode of transport (hcz) fixed effects, control for the beginning and end week of the quadruplet, and exclude quadruplets with less than 5 shipments. Standard errors, adjusted for clustering by country (c) are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

In Table 4, we consider a related specification that relaxes the restriction of a linear relationship between shipment characteristics and SPS_{mhcz} and instead uses a non-parametric specification,

[30] The cutoff for number of shipments in a quadruple is not important for our results. In appendix Section B, we show that results are qualitatively identical for a cutoff of 10 shipments. We also find similar relationships when we compute SPS at the level of broader bins, e.g., for buyer triples (mhc) and buyer doubles (mh).

$$\overline{Y}_{mhczq} = \beta_0 + \sum_{n\in\{2,3,4\}} \alpha_n 1\{Q = n\} + \beta_2 \ln(QPW_{mhczq}) + \beta_3 beg_{mhczq} + \beta_4 end_{mhczq} + \lambda_{hcz} + \epsilon_{mhczq}, \tag{9}$$

where $1\{Q = n\}$ are dummies for the quartile of the SPS_{mhcz} distribution. We compute the quartile cutoffs by looking across SPS_{mhcz} within hcz bins, and assign assign each buyer quadruple to its respective quartile. The first quartile – the most "Japanese" – is the excluded category. $\overline{Y}_{mhcz}$ continues to represent average procurement patterns, and the remaining controls are the same as in equation (8). As before, the sample period is 1992 to 2016, standard errors are clustered at the c level, and observations are restricted to $mhcz$ bins with at least 5 transactions.

Results continue to justify use of SPS_{mhcz} as an indicator of procurement system. As illustrated in the first and second columns of Table 4, we find that the SPS_{mhcz} point estimates are positive and statistically significant at conventional levels for both quantity per shipment and weeks between shipments for all quartiles. Moreover, the coefficients rise in magnitude from quartile 1 to quartile 4, and are statistically different from one another, consistent with the quadruples in those quartiles being increasingly "American". Coefficient estimates indicate that average quantity per shipment and weeks per shipment for U.S. importers in the fourth quartile are roughly 0.79 and 0.86 log points higher than for those in the first quartile. Coefficient estimates for SPS_{mhcz} in column 3 of Table 4 provide similar support for the model, indicating that the average price declines with seller per shipment quartile. Here, the average average price for the fourth quartile of the SPS_{mhcz} distribution is roughly 0.23 log points higher, respectively, than for those in the first quartile. Similar to the earlier results, the coefficients on QPW_{mhcz} are also consistent with Proposition 2.4.

Table 4: A vs J Quartile Classification Regression Across $mhcz$ Quadruples

Dep. var.	(1)	(2)	(3)
	$\log(QPS_{mhcz})$	$\log(WBS_{mhcz})$	$\log(UV_{mhcz})$
$(SPS_{mhcz} = Q2)$	0.328^{***}	0.350^{***}	-0.117^{***}
	0.003	0.003	0.004
$(SPS_{mhcz} = Q3)$	0.552^{***}	0.591^{***}	-0.179^{***}
	0.004	0.004	0.005
$(SPS_{mhcz} = Q4)$	0.792^{***}	0.856^{***}	-0.226^{***}
	0.006	0.006	0.007
$\log(QPW_{mhcz})$	0.687^{***}	-0.323^{***}	-0.282^{***}
	0.003	0.003	0.005
Observations	2,966,000	2,966,000	2,966,000
R-squared	0.945	0.661	0.845
Fixed effects	hcz	hcz	hcz
Controls	beg, end	beg, end	beg, end

Source: LFTTD and authors' calculations. Table reports the results of regressing noted attribute of importer by product by country by mode of transport ($mhcz$) bins on bins' sellers per shipment (SPS_{mhcz}) and total quantity shipped per week (QPW_{mhcz}). (QPS_{mhcz}), (WBS_{mhcz}),and (P_{mhcz}) are average quantity per shipment, average weeks between shipment, and average unit value. All regressions include product by country by mode of transport (hcz) fixed effects, control for the beginning and end week of the quadruplet, and exclude quadruplets with less than 5 shipments. Standard errors, adjusted for clustering by country (c) bin are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

In the next three subsections, we present additional evidence of "Japanese" procurement. First, we analyze procurement patterns within $mxhcz$ buyer-seller relationships. Second, we consider relationship length as an additional indicator of "Japanese" procurement. Finally, we examine procurement patterns for differentiated products versus commodities.

4.2 A vs J Within Sellers

In this sub-section, we examine whether $mhcz$ buyer quadruples' sellers per shipment, SPS_{mhcz}, predicts theory-consistent procurement patterns *within* each of their exporter relationships. In principle, a buyer quadruple could appear "Japanese" in aggregate even if it were not with respect to each of its sellers. For example, a buyer quadruple might obtain frequent shipments from a few sellers, thus appearing to be "Japanese", but shipments within each seller might be dispersed if the buyer alternates among them. We use the following $mxhcz$-level OLS regression,

$$Y_{mxhcz} = \beta_0 + \beta_1 SPS_{mhcz} + \beta_2 \ln(QPW_{mxhcz}) + \beta_3 beg_{mxhcz} + \beta_4 end_{mxhcz} + \lambda_{xhcz} + \epsilon_{mxhcz}. \quad (10)$$

In this specification, Y_{mxhcz} represents procurement attributes at the buyer-seller relationship ($mxhcz$)

level, and the right-hand-side variables are defined at this level, as well, with the exception of SPS_{mhcz} which continues to be at the $mhcz$ level. We also include exporter by product by country by mode fixed effects (λ_{xhcz}) to compare buyer procurement patterns within sellers who may be heterogeneous in a number attributes, including production costs. Standard errors are clustered at the c level.

Results, reported in Table 5, are similar to those in the previous section, providing further support for Proposition 2.3 and 2.4, as well as the use of SPS_{mhcz}. Across U.S. buyer quadruples within foreign exporters, we find that a one standard deviation decrease in SPS_{mhcz} is associated with QPS_{mxhcz} and WBS_{mxhcz} declines of 0.02 and 0.17 log points. For UV_{mxhcz}, this decline coincides with an increase of 0.01 log points. Furthermore, an increase in the total procurement quantity raises shipment size and frequency, but lowers unit import values.

Table 5: A vs J Classification Regression Across $mxhcz$ Quintuples

Dep. var.	(1)	(2)	(3)
	$\log(QPS_{mxhcz})$	$\log(WBS_{mxhcz})$	$\log(UV_{mxhcz})$
$\log(SPS_{mhcz})$	0.100^{***}	0.696^{***}	-0.062^{***}
	0.014	0.040	0.004
$\log(QPW_{mxhcz})$	0.511^{***}	-0.171^{***}	-0.130^{***}
	0.007	0.009	0.007
Observations	4,783,000	4,783,000	4,783,000
R-squared	0.966	0.621	0.953
Fixed effects	$xhcz$	$xhcz$	$xhcz$
Controls	beg, end	beg, end	beg, end

Source: LFTTD and authors' calculations. Table reports the results of regressing noted attribute of US importer by foreign exporter by product by country by mode of transport ($mxhcz$) bins on bins' sellers per shipment (SPS_{mhcz}) and total quantity shipped per week (QPW_{mxhcz}). (QPS_{mxhcz}), (WBS_{mxhcz}), and (P_{mxhcz}) are average quantity per shipment, average weeks between shipment, and average unit value (i.e. value divided by quantity). All regressions include exporter by product by country by mode of transport ($xhcz$) fixed effects, control for the beginning and end week of the quintuple, and exclude buyer quadruplets with less than 5 shipments. Standard errors, adjusted for clustering by country (c) bins are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

4.3 Relationship Length

Since buyers under the "Japanese" system rely on repeat purchases from the same seller to provide incentives, an implication of our theory is that buyers using "Japanese" procurement should have on average longer relationships with their suppliers. We test this prediction by constructing a new variable, average relationship length ($length_{mhcz}$), which captures the average age of each $mxhcz$ buyer-seller relationship across its transactions. We compute this average in two steps. First, for each $mxhcz$ transaction, we compute the number of weeks passed since the first transaction of any good using any mode between the buyer and seller. For each $mhcz$ buyer quadruple, we then take

the average of these lengths across the $mxhcz$ transactions within it.[31] This average allows for the possibility that buyers already sourcing one product from a given supplier, or already using a different mode of transportation with that seller, may structure relational contracts to encompass all products.[32]

We then run the same regressions as in the previous subsections using $length_{mhcz}$ as dependent variable. The results, reported in Table 6, provide support for the hypothesis that buyer quadruples with a lower number of suppliers per shipment are on average in longer relationships. The result in Column 1, which runs specification (8) shows that a one standard deviation decrease in sellers per shipment is associated with a 0.14 log point increase in average relationship length. Column 2, based on specification (9), illustrates that the average relationship length for U.S. importers in the fourth quartile is roughly 0.29 log points lower than in the first quartile. Finally, Column 3 re-runs specification (10), which is at the buyer-seller relationship level. The coefficients indicate that the result also holds within an average buyer-seller relationship.

[31] Using this approach, $length_{mhcz}$ is zero if an mx buyer-seller pair trades only a single good once, and increases if that good is traded more often.

[32] We note that we find similar results – available on request – if we instead define relationship length for $mxhcz$ quintuples.

Table 6: A vs J Classification Regressions, Using Relationship Length

Dep. var.	(1)	(2)	(3)
	$\log(length_{mhcz})$	$\log(length_{mhcz})$	$\log(length_{mxhcz})$
$\log(SPS_{mhcz})$	-0.594^{***}		
	0.013		
$(SPS_{mhcz} = Q2)$		-0.369^{***}	
		0.002	
$(SPS_{mhcz} = Q3)$		0.679^{***}	
		0.003	
$(SPS_{mhcz} = Q4)$		-1.204^{***}	
		0.005	
$\log(SPS_{mxhcz})$			-0.331^{***}
			0.011
$\log(QPW_{mhcz})$	-0.117^{***}	-0.101^{***}	-0.269^{***}
	0.003	0.001	0.007
Observations	2,966,000	2,966,000	4,783,000
R-squared	0.451	0.439	0.773
Fixed effects	hcz	hcz	$xhcz$
Controls	beg, end	beg, end	beg, end

Source: LFTTD and authors' calculations. Table reports the results of regressing the average relationship length ($length_{mhcz}$) of US importer by product by country by mode of transport ($mhcz$) bins on bins' sellers per shipment (SPS_{mhcz}) and total quantity shipped per week (QPW_{mhcz}). The regressions in columns (1) and (2) include product by country by mode of transport (hcz) fixed effects. The regression in column (3) includes exporter by product by country by mode of transport ($xhcz$) fixed effects. All regressions control for the beginning and end week of the quadruplet, and exclude quadruplets with less than 5 shipments. Standard errors, adjusted for clustering by country (c) bin are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

4.4 Differentiated Products Versus Commodities

Finally, we examine whether buyers are more likely to use "Japanese" procurement for some products than for others. If more differentiated or customized products have higher inspection costs, then by Proposition 2.1 buyers are more likely to use "Japanese" procurement for them. By Proposition 2.3, products sourced under the "Japanese" system should be sourced with smaller shipment size, greater frequency, and higher unit import values than products sourced under the "American" system. Moreover, "Japanese" sourcing should be associated with fewer suppliers and longer relationships. While explicit measures of product-level inspection costs do not exist, we examine this feature of the model using the commonly cited measure of product-differentiation from Rauch (1999) using the following $mhcz$-level OLS specification,

$$\overline{Y}_{mhcz} = \beta_0 + \beta_1 Diff_h + \beta_2 \ln(VPW_{mhcz}) + \beta_3 beg_{mhcz} + \beta_4 end_{mhcz} + \lambda_{cz} + \epsilon_{mhcz}. \quad (11)$$

We consider four dependent variables. The first is the average number of weeks between shipments WBS_{mhcz} introduced above. We do not consider quantity per shipment or unit value here since the regression compares shipping systems across products, which are recorded in different units.[33] Instead, we use as our second dependent variable the average transaction value per shipment, VPS_{mhcz}, as a measure of average transaction size. The third variable is the average relationship length $length_{mhcz}$. Finally, the fourth variable is a measure of the buyer's procurement type, SPS_{mhcz}. On the right-hand side, $Diff_h$ is a dummy variable indicating that product h is either differentiated or has a reference price, as opposed to being a commodity, according to the product categorization scheme proposed by Rauch (1999).[34] Because the right-hand-side variable of interest varies only at the product level, we are unable to include product fixed effects, so comparisons are made within country-mode bins by including fixed effects at that level (λ_{cz}). Since we cannot standardize quantities to be consistent across products, we control for potential scale effects using value per week (VPW_{mhcz}), rather than quantity per week, which was used above. As before, the sample period is 1992 to 2016, we include only buyer quadruples with at least 5 transactions, and standard errors are clustered at the country level.

Results, reported in Table 7, are consistent with the model's predictions regarding inspection costs, while providing further support for the use of procurement patterns and sellers per shipment to identify buyer types. As indicated in the first three columns of the table, we find that differentiated products are more Japanese: they are shipped with fewer weeks between shipments, the relationships in which they are traded tend to last longer, and the average transaction size is smaller. Results in the final two column provide further support for this view, as buyer quadruples encompassing differentiated goods tend to have lower sellers per shipment.

[33] For example, we cannot really compare the price of one barrel of oil to the price of one shoe.

[34] Rauch (1999) provides both a liberal and a conservative definition of differentiated goods. We use the liberal definition for the results reported in the main text, but note that these results are robust to using the conservative definition.

Table 7: A vs J Classification Regression for Differentiated Goods

Dep. var.	(1)	(2)	(3)	(4)
	$\log(WBS_{mhcz})$	$\log(VPS_{mhcz})$	$\log(Length_{mhcz})$	$\log(SPS_{mhcz})$
$Differentiated_h$	-0.234***	-0.227***	0.072*	-0.081**
	0.071	0.067	0.040	0.040
$\log(Value/Week_{mhcz})$	-0.463***	0.554***	-0.047***	-0.202***
	0.008	0.008	0.004	0.007
Observations	2,589,000	2,589,000	2,589,000	2,589,000
R-squared	0.613	0.728	0.193	0.276
Fixed effects	cz	cz	cz	cz
Controls	beg, end	beg, end	beg, end	beg, end

Source: LFTTD and authors' calculations. Table reports the results of regressing noted attribute of US importer by product by country by mode of transport ($mhcz$) bins on a dummy for whether the bin's product code is differentiated or reference priced according to the liberal classification by Rauch (1999) and on value shipped per week (VPW_{mhcz}). (WBS_{mhcz}), (VPS_{mhcz}), ($length_{mhcz}$), and SPS_{mhcz}) are average weeks between shipment, average value per shipment, average relationship length, and sellers per shipment. All regressions include country by mode of transport (cz) fixed effects, control for the beginning and end week of the quadruplet, and exclude quadruplets with less than 5 shipments. Standard errors, adjusted for clustering by country, are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

The preceding results demonstrated that key implications of the model are confirmed in the data and that the ratio of sellers per shipment is an appropriate measure for the extent of Japanese-style sourcing. In particular, we showed that buyers with lower SPS_{mhcz} exhibit lower quantity per shipment and weeks between shipments, as well as higher average prices, consistent with Proposition 2.3. Moreover, a higher total quantity traded increases average batch size and reduces weeks between shipments and average price, consistent with Proposition 2.4. We next showed that the relationships between SPS_{mhcz} and shipment characteristics hold even within the average relationship, and that a lower SPS_{mhcz} is associated with longer relationships, as implied by the theory. Lastly, we showed that shipments of differentiated products – which likely face higher inspection costs than homogeneous products – exhibit more "Japanese" characteristics, including lower SPS_{mhcz}, consistent with Proposition 2.1 and 2.3.

5 The Effect of PNTR on the Choice of Procurement System

One of the fundamental contributions of the model presented in Section 2 is that the share of "American" and "Japanese" procurement relationships in an economy can vary with trade policy. In particular, the model suggests that a decrease in the possibility that a trade war will sever relationships can induce buyers and sellers to switch away from "American" to "Japanese" procurement. In this section we examine this implication using a plausibly exogenous change in U.S. trade policy: the U.S. granting of permanent normal trade relations (PNTR) to China in October 2000, which substantially reduced the possibility of a trade war between the two countries. We first provide some background on the policy before connecting the policy change to a shift towards "Japanese" sourcing.

5.1 Description of PNTR

U.S. imports from non-market economies such as China are generally subject to relatively high non-Normal Trade Relations (NTR) tariff rates originally set under the Smoot-Hawley Tariff Act of 1930, as oppposed to the generally low NTR tariff rates the U.S. offers to trading partners that are members of the World Trade Organization (WTO). A provision of U.S. trade law, however, allows imports from non-market economies to enter the United States under NTR tariffs subject to *annual* approval by both the President and Congress. Chinese imports first began entering the United States under this provision in 1980, after the warming of relations between the United States and China.

Annual approval became controversial and less certain after the Tiananmen Square incident in 1989, however, and this uncertainty continued throughout the 1990s. During this time, firms engaged in or considering U.S.-China trade faced the possibility, each year, of substantial tariff increases—32 percentage points for the average product—if China's NTR status was not re-approved. Pierce and Schott (2016) document extensive discussion by U.S. firms of the trade-dampening effects of this uncertainty in the 1990s, and Alessandria, Khan, and Khederlarian (2019) show that it reduced U.S. imports from China, while also driving seasonal patterns in imports within the year.

When the U.S. granted PNTR to China in 2001, it eliminated the need for these annual renewals, along with the potential for tariff increases and associated uncertainty. Pierce and Schott (2016) show that this elimination of uncertainty was associated with reductions in U.S. manufacturing employment as imports from China increased, and Handley and Limão (2017) estimate that the elimination of uncertainty was equivalent to a 13 percentage point decline in tariff rates. Note that this policy change did not affect actual tariffs, and would therefore have no effect on trade flows in models such as in Eaton and Kortum (2002), while in our framework it affects ρ and hence the procurement costs under the "Japanese" system.

This policy environment provides an ideal setting for empirically examining the model's implications of the effect of a change in the probability of trade peace on importing firms' sourcing behavior. Specifically, as described by Proposition 2.1 in Section 2, the model implies that a decrease in the probability of a trade war should lead to the adoption of more "Japanese" sourcing behavior, as the premium over marginal cost paid to foreign suppliers to ensure sufficient quality declines.

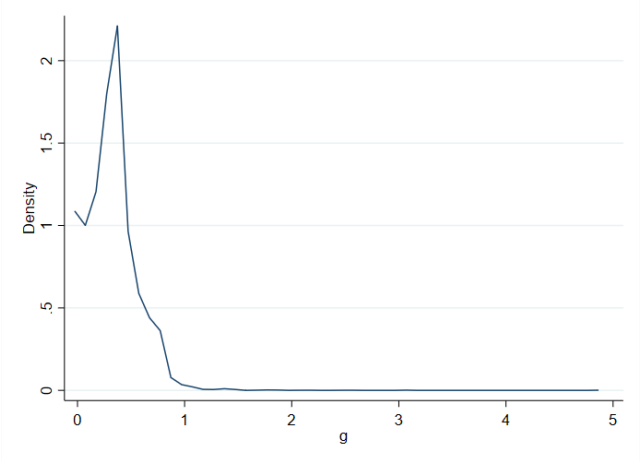
The policy environment also suggests a straightforward way to assess a firm's exposure to PNTR. Specifically, we assess an imported product's exposure to PNTR in terms of the NTR gap from Pierce and Schott (2016), i.e., the difference between the non-NTR rate to which tariffs could have risen before PNTR and the lower NTR rate that was locked in by the 2000 policy change,

$$NTR\,Gap_h = Non\,NTR\,Rate_h - NTR\,Rate_h. \tag{12}$$

We compute these gaps as of 1999, the year before the change in policy, using *ad valorem* equivalent tariff rates provided by Feenstra, Romalis, and Schott (2002).[35] As indicated in Figure 3, these gaps vary widely across products, and have a mean and standard deviation of 0.32 and 0.23.

[35]While U.S. tariffs are set at the level of eight-digit HS products, we observe trade at the ten-digit HS level. We therefore match each ten-digit HS product with the tariff associated with its first eight digits.

Figure 3: Distribution of the *NTR Gap*



Source: Feenstra, Romalis, and Schott (2002) and author's calculations. Figure displays the distribution of the $NTR\ Gap_h$, the difference between the relatively low NTR tariff rate that was locked in by PNTR and the higher rate to which U.S. tariffs on Chinese goods might have risen absent the change in policy.

Our identification strategy exploits this variation in the NTR gap to determine whether U.S.-China procurement patterns change relative to procurement patterns with exporters from other source countries (first difference) after the change in U.S. policy is implemented (second difference) in industries with higher NTR gaps (third difference). The last difference captures the fact that industries with larger NTR gaps experience a larger increase in the relationship continuation probability than industries with smaller gaps.[36] We expect the largest shifts toward Japanese-style procurement after PNTR to occur in U.S. imports of high-NTR gap products from China.

5.2 "American" versus "Japanese" U.S. Imports

We next use the empirical procedure for identifying "Japanese" buyers introduced in Section 4.1 to explore the extent to which "Japanese" procurement varies across U.S. trading partners around the time that PNTR was introduced. We begin with descriptive information, and then proceed to a differences-in-differences analysis.

We define two five-year periods straddling the granting of PNTR in 2001: 1995 to 2000 and 2002 to 2007. We then compute the ratio of suppliers to shipments for each buyer quadruple in each of the two periods p, SPS_{mhczp}.[37] Finally, we classify buyer quadruples as "Japanese" in a given period if their SPS_{mhcz} falls within the first quartile of the sellers per shipment distribution for their product-mode (hz) in the *first* period. Thus, the share of buyer quadruples classified as "Japanese" within these hz bins in the first period will be 25 percent *worldwide* by construction, but can vary across countries in that period. In the second period, the share of such buyer quadruples worldwide

[36] As described above, this relationship continuation probability depends on the change in the probability of China's NTR status – which is identical for all products – and the change tariff rates that could have occurred before PNTR, which varies by product.

[37] Some quadruples extend through both periods, others show up in just one or the other. Our baseline analysis does not make use of longer periods so that the post-PNTR interval ends before the onset of the Great Recession in 2007.

need not be 25 percent.

The first two columns of Table 8 reports the share of U.S. imports accounted for by U.S. importer quadruples classified as "Japanese" for the top ten U.S. trading partners from 1995 to 2000 and 2002 to 2007. Since the first-quartile cutoff for classifying relationships as "Japanese" is arbitrary, in columns 3 and 4 we report a value-weighted average of SPS_{mhczp} by country and period. The rows of the table are sorted by the first column, the J buyer share in the first period.

Table 8: Japanese Relationships by Country

	J Buyers' Import Share		Mean SPS	
Country	1995-2000	2002-2007	1995-2000	2002-2007
Japan	0.756	0.725	0.118	0.131
Mexico	0.750	0.869	0.093	0.069
United Kingdom	0.717	0.519	0.157	0.229
Taiwan	0.711	0.743	0.138	0.122
France	0.627	0.667	0.192	0.168
Canada	0.602	0.667	0.147	0.130
South Korea	0.656	0.724	0.166	0.141
Rest of the World	*0.625*	*0.678*	*0.187*	*0.160*
Germany	0.582	0.606	0.203	0.179
China	0.582	0.693	0.192	0.150
Brazil	0.576	0.706	0.197	0.158

Source: LFTTD and authors' calculations. First two columns report the share of U.S. imports accounted for by importers classified as using the "Japanese" system, by country and time period. The second two columns report the weighted average sellers per shipment (SPS_{mhczp}) across buyer quadruples by country and period, where import values are used as weights. Rows of the table are sorted by the first column.

As indicated in the first column of the table, we find that the value-weighted share of buyer quadruples classified as "Japanese" is indeed highest for U.S. imports from Japan, perhaps due to influence of large "Japanese" multinationals like Toyota producing in the United States (Boehm, Flaaen, and Pandalai-Nayar, 2020). It is also relatively high for Mexico, the UK, and Taiwan, and relatively low for Germany, China and Brazil, all of which have values below the rest of the world.

Results in the second column reveal that, over time, the share of J buyer quadruples generally rises, with Japan and the UK being exceptions. The largest increases exhibited, both in levels and percent growth, are for Brazil, China and Mexico.[38]

5.3 Response to PNTR Among Continuing Buyer-Seller Relationships

In this section we examine whether the rise in China's "Japanese" trade is related to PNTR. Our analysis is guided by Propositions 2.1, 2.2 and 2.3, which allow us to distinguish empirically between

[38]The relatively large growth of "Japanese" imports from Mexico may be related to increasingly close supply-chain integration with the United States after NAFTA, including maquiladora trade.

adjustments within the "Japanese" system and a switch between systems. Propositions 2.1 and 2.3 predict that if an increase in the probability of trade peace is large enough, buyers will switch from the "American" to the "Japanese" system, inducing shipment size to fall, shipment frequency to increase, and unit values to rise for a given total quantity ordered. Proposition 2.2 implies that if most procurement already takes place under the "Japanese" system prior to PNTR, shipment size will increase, shipment frequency fall, and unit values will decline for a given quantity due to an adjustment within the "Japanese" system. As noted in Table 8, "Japanese" U.S. imports from China are relatively low from 1995 to 2000.

We start our analysis by analyzing procurement patterns in a balanced panel of *continuing* buyer-seller relationships, i.e., relationships that exist before PNTR and persist afterwards. This initial focus is motivated by the fact that PNTR led to a significant increase in Chinese firms exporting to the United States (Pierce and Schott, 2016; Amiti, Dai, Feenstra, and Romalis, 2020). If this initial entry of new Chinese suppliers is associated with experimentation and supplier switching by U.S. buyers, then initially *mhcz* quadruples could increase their number of suppliers and PNTR could appear to promote the "American" procurement system, even though each relationship became more "Japanese".

Our first specification examines the relationship between PNTR and *existing* buyer-seller relationships using the *mxhczp*-level OLS regression, where period $p \in \{Pre, Post\}$ represents the periods before and after the change in policy,

$$\begin{aligned} ln(Y_{mxhczp}) = &\beta_1 1\{p = Post\} * 1\{c = China\} * NTR\,Gap_h + \beta_2 ln(QPW_{mxhczp}) + \\ &\beta_3 \chi_{mxhczp} + \beta_4 beg_{mxhczp} + \beta_5 end_{mxhczp} + \lambda_{mxhcz} + \lambda_p + \epsilon_{mxhczp}. \end{aligned} \tag{13}$$

In this regression, Y_{mxhczp} represents buyer-seller relationship-level procurement patterns computed separately for each period – i.e., $\ln(QPS_{mxhczp})$, $\ln(WBS_{mxhczp})$, and $\ln(UV_{mxhczp})$ – which now also includes $length_{mxhczp}$ .Similar to before, the latter is computed for every *mxhcz* quintuple as the number of weeks between the first transaction of the overall importer-exporter (*mx*) relationship in period *Pre* and its last transaction in period *p*. The triple difference-in-differences (DID) term of interest is an interaction of a dummy for the post period ($1\{p = Post\}$), a dummy for imports from China ($1\{c = China\}$), and the *NTR Gap*. The variable χ_{mxhczp} represents the full set of interactions of those variables required to identify β_1. The remaining terms on the right-hand side control for the average quantity traded per week in each of the two periods (QPW_{mxhczp}), the first and last week of the relationship within the given period, as well as buyer-seller relationship (λ_{mxhcz}) and period (λ_p) fixed effects. As in Table 8, the pre- and post periods are 5-year intervals straddling 2001: 1995 to 2000 and 2002 to 2007.[39] Standard errors are clustered at the country level.

Results, reported in Table 9, indicate that higher exposure to PNTR is associated with a shift toward Japanese-style procurement within existing buyer-seller quintuples. Coefficient estimates in

[39] As noted earlier, we choose these periods to have symmetric length around the 2001 change in policy, and so that the latter window ends before the Great Recession. In Appendix Section C.1, we show that we find substantially similar results using different time periods.

the first two columns show that a one standard deviation increase in the $NTR\ Gap$ induces a relative decline in quantity per shipment and weeks between shipments of 4.5 percent.[40] Results in columns 3 and 4, by contrast, reveal that a one standard deviation increase in exposure to PNTR is associated with relative increases in price and relationship length of 2.1 percent and 2.9 percent, respectively. In each case, the findings in Table 9 are consistent with the predictions of Proposition 2.3. Moreover, consistent with Proposition 2.4, the coefficient estimates for $ln(QPW_{mxhczp})$ indicate that an increase in the procurement quantity increases the size of shipments, raises shipping frequency, and reduces unit values.

Table 9: Baseline Within-Exporter PNTR Regression: 2002-2007 vs 1995-2000

Dep. var.	(1)	(2)	(3)	(4)
	$\ln(QPS_{mxhczp})$	$\ln(WBS_{mxhczp})$	$\ln(UV_{mxhczp})$	$\ln(length_{mxhczp})$
$Post_p * China_c * NTR Gap_p$	-0.1970***	-0.1970***	0.0922***	0.1278***
	0.0156	0.0156	0.0167	0.0265
$ln(QPW_{mxhczp})$	0.3680***	-0.6320***	-0.1238***	-0.1449***
	0.0077	0.0077	0.0113	0.0036
Observations	439,000	439,000	439,000	439,000
R-squared	0.98	0.88	0.98	0.85
Fixed effects	$mxhcz, p$	$mxhcz, p$	$mxhcz, p$	$mxhcz, p$
Controls	Yes	Yes	Yes	Yes

Source: LFTTD and authors' calculations. Table reports the results of regressing noted attribute of US importer by exporter by product by country by mode of transport ($mxhcz$) bins on the difference-in-differences term of interest and quantity shipped per week. Pre-and post periods are 1995 to 2000 and 2002 to 2007. (QPS_{mxhczp}), (WBS_{mxhczp}), (P_{mxhczp}), and ($length_{mxhczp}$) are average quantity per shipment, average weeks between shipment, average unit value (i.e. value divided by quantity), and average relationship length in period p. All regressions include $mxhcz$ and period p fixed effects, control for the beginning and end week of the quadruplet as well as all variables needed to identify the DID term of interest, and exclude quadruplets with less than 5 shipments. Standard errors, adjusted for clustering by country, are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

5.4 Response to PNTR Among New Buyer-Seller Relationships

While the preceding analysis offers a tight examination of the implications of PNTR on procurement patterns by focusing on changes *within* relationships, it excludes, by design, implications of the policy change on the formation of new relationships. As described above, PNTR led to the creation of new trade relationships as China increased its footprint in the U.S. market. Our second estimation, therefore, explores whether new buyer-seller relationships formed after PNTR are more likely to adhere to the "Japanese" system than those that were newly formed before PNTR. We continue to use equation (13), but in this case restrict the observations in each period to new relationships.

In the pre-PNTR period, we define a new relationship as an mx pair that appears for the first time in 1995 to 2000, i.e., it is not present in any prior years for which data are available (from 1992 to 1994). Likewise, a new relationship in the 2002 to 2007 period is an mx pair not observed previously,

[40] The DID point estimates and standard errors in the first two columns of Table 9 are identical because we also control for QPW, end, and beg. That is, regressing either x/q or x on the DID term yields identical coefficients given that we control for q. We note that the sign and significance of the DID terms of interest are virtually the same if we do not control for QPW_{mxhcz}.

from 1992 to 2001. As above, the estimation is performed at the $mxchzp$ level, and standard errors are clustered at the country level. Here, however, because we are comparing different buyer-seller relationships in each period, we do not include buyer-seller relationship (i.e. $mxhcz$) fixed effects. Rather, we look across these relationships within buyer quadruples ($mhcz$) and include separate exporter (x) and period (p) fixed effects. In this sense, our regression focuses on the new relationships formed by existing buyers and sellers. A second difference with respect to the estimation above is that examining relationship length ($length_{mxhcz}$) is not very meaningful because all relationships are new, by construction.

Results, reported in Table 10, are also consistent with the model's predictions. As indicated in the table, we find that buyer-seller relationships formed after PNTR trading goods with greater exposure to the policy change exhibit relatively small and more frequent shipments, at relatively high prices. The coefficient estimates reported in the table indicate that a one standard deviation increase in exposure is associated with a 2.7 percent decline in shipment size and frequency, and a 2.1 percent rise in price.

Table 10: New Relationship PNTR Regression: 2002-2007 vs 1995-2000

Dep. var.	(1)	(2)	(3)(4)
	$\ln(QPS_{mxhczp})$	$\ln(WBS_{mxhcz})$	$\ln(UV_{mxhcz})$
$Post_p * China_c * NTRGap_p$	-0.1164***	-0.1164***	0.0895***
	0.0247	0.0247	0.0316
$ln(QPW_{mxhcz})$	0.4087***	-0.5913***	-0.1291***
	0.0117	0.0117	0.0163
Observations	3,184,000	3,184,000	3,184,000
R-squared	0.97	0.83	0.97
Fixed effects	$mhcz, x, p$	$mhcz, x, p$	$mhcz, x, p$
Controls	Yes	Yes	Yes

Source: LFTTD and authors' calculations. Table reports the results of comparing new buyer-seller relationships in the pre-versus post-PNTR period. Pre-and post periods are 1995 to 2000 and 2002 to 2007. New relationships are defined as mx pairs appear for the first time in each period. (QPS_{mxhczp}), (WBS_{mxhczp}), (P_{mxhczp}), and ($length_{mxhczp}$) are average quantity per shipment, average weeks between shipment, and average unit value (i.e. value divided by quantity) in period p. All regressions include $mhcz$, x and period p fixed effects, control for the beginning and end week of the quadruplet as well as all variables needed to identify the DID term of interest, and exclude quadruplets with less than 5 shipments. Standard errors, adjusted for clustering by country, are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

5.5 Response to PNTR at the Buyer-Level

We now analyze the response of procurement patterns at the more aggregated buyer quadruple ($mhcz$) level. These regressions illustrate how sourcing patterns change for the average importer across all relationships, combining new and continuing buyer-seller relationships. Intuitively, changes to procurement patters are weaker at this more aggregate level, though still present for suppliers per shipment.

We examine the extent to which buyer quadruples become more "Japanese" after PNTR versus before using the following $mhczp$-level DID regression,

$$ln(Y_{mhczp}) = \beta_1 1\{p = Post\} * 1\{c = China\} * NTR\,Gap_h + \beta_2 ln(QPW)_{mhczp} + \beta_3 \chi_{mhczp} + \beta_4 beg_{mhcz} + \beta_5 end_{mhcz} + \lambda_{mhcz} + \lambda_p + \epsilon_{mhczp}. \quad (14)$$

In this specification, Y_{mhczp} represents the four shipment attributes introduced above, but now computed for buyer quadruples in each period rather than buyer-seller quintuples. We include separate $mhcz$ and p fixed effects. As a result of this specification, the regression now also includes exporters x that appear in only one of the two periods, and the results are identified from buyer quadruples that trade in both periods. As above, the pre-PNTR period is 1995 to 2000 and the post period is 2002 to 2007.

Results, displayed in Table 11, show a significant decline in the average shipping size and weeks between shipments, consistent with a shift towards "Japanese" procurement. However, the increases in unit values and in relationship length, while positive, are insignificant, possibly because some importers switched to new suppliers which charged low prices after PNTR.

Table 11: Within-Importer PNTR Regression: 2002-2007 vs 1995-2000

Dep. var.	(1)	(2)	(3)	(4)
	$\ln(QPS_{mhczp})$	$\ln(WBS_{mhczp})$	$\ln(UV_{mhczp})$	$\ln(length_{mhczp})$
$Post_p * China_c * NTR\,Gap_p$	-0.043***	-0.043***	0.018	0.061
	0.014	0.014	0.024	0.047
$ln(QPW_{mhczp})$	0.436***	-0.564***	-0.207***	-0.015**
	0.017	0.017	0.024	0.006
Observations	738,000	738,000	738,000	738,000
R-squared	0.98	0.88	0.97	0.71
Fixed effects	$mhcz, p$	$mhcz, p$	$mhcz, p$	$mhcz, p$
Controls	Yes	Yes	Yes	Yes

Source: LFTTD and authors' calculations. Table reports the results of regressing noted attribute of US importer by product by country by mode of transport ($mhcz$) bins on the difference-in-differences term of interest and quantity shipped per week. Pre-and post periods are 1995 to 2000 and 2002 to 2007. (QPS_{mhczp}), (WBS_{mhczp}), (P_{mhczp}), and ($length_{mhczp}$) are average quantity per shipment, average weeks between shipment, average unit value (i.e. value divided by quantity), and average relationship length in period p. All regressions include $mhcz$ and period p fixed effects, control for the beginning and end week of the quadruplet as well as all variables needed to identify the DID term of interest, and exclude quadruplets with less than 5 shipments. Standard errors, adjusted for clustering by country, are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

At the level of $mhcz$ quadruples we can consider two additional variables introduced earlier that are correlated with Japanese procurement: the number of suppliers per shipment used by the quadruple, SPS_{mhczp}, and the dummy variable indicating that suppliers per shipment are within the first quartile of the hz distribution computed between 1995 and 2000, J_{mhzp}.[41] We examine the relationship between the DID term of interest and these variables in Table 12.

In columns 1 and 2, we find that buyer quadruples importing from China *vis à vis* other countries exhibit lower sellers per shipment and are more likely to be classified as J after PNTR versus before,

[41] These variables are defined at the buyer quadruple level rather than at the level of quintuples, and hence were not included in the earlier tables.

though the continuous measure is close to zero and not significant at conventional levels. To explore whether this weakness might be driven by initial buyer experimentation during the years immediately after PNTR, we also consider, in columns 3 and 4, similarly constructed outcomes but for a slightly later 2004 to 2009 time period of the same length. Coefficient estimates for this alternate post period have the same sign patterns but are larger in absolute magnitude and statistically significant, providing some support for the idea for a buyer-adjustment period after PNTR.[42]

In Appendix C.2, we analyze similar regression specifications where instead of aggregating to the *mhcz* quadruple level we aggregate to the *hcz* level. The results are relatively similar to our findings for the buyer quadruple level regressions considered in this section.

Table 12: Within-Importer PNTR Regression, Buyer Characteristics

Dep. var.	(1)	(2)	(3)	(4)
	$\ln(SPS_{mhcz})$	J_{mhcz}	$\ln(SPS_{mhcz})$	J_{mhcz}
$Post_p * China_c * NTR\,Gap_p$	-0.0064	0.0414*	-0.0755*	0.0757**
	0.0309	0.0218	0.0383	0.031
$ln(QPW_{mhcz})$	0.4361***	-0.5639***	-0.2068***	-0.0150*
	0.0172	0.0172	0.0242	0.0062
Observations	738,000	291,000	556,000	225,000
R-squared	0.7727	0.6752	0.7573	0.66
Fixed effects	$mhcz, p$	$mhcz, p$	$mhcz, p$	$mhcz, p$
Controls	Yes	Yes	Yes	Yes
Post Period	2002-7	2002-7	2004-9	2004-9

Source: LFTTD and authors' calculations. Table reports the results of regressing noted attribute of US importer by product by country by mode of transport (*mhcz*) bins on the difference-in-differences term of interest and quantity shipped per week. Pre-PNTR period is 1995 to 2000. First and second two columns summarize results when the post period is 2002 to 2007, and 2004 to 2009, respectively. All regressions include *mhcz* and period *p* fixed effects, control for the beginning and end week of the quadruplet as well as all variables needed to identify the DID term of interest, and exclude quadruplets with less than 5 shipments. Standard errors, adjusted for clustering by country, are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

Overall, the results in this section provide evidence consistent with one of the key implications of the model, namely that a change in trade policy that changes the probability of relationship continuation can induce changes in firms' sourcing behavior. Across existing and new relationships and defined at different levels of aggregation, we consistently find that the granting of PNTR to China in 2001, which reduced the probability of tariff increases, was associated with a switch toward "Japanese" sourcing. In addition to providing empirical support for the model's framework, the results also suggest important implications of more recent U.S. trade policy. In particular, the results suggest that the more uncertain trade policy environment since the imposition of sizable tariffs by the U.S. and China in 2018 and 2019 is likely already suppressing the formation of long-term "Japanese" relationships. In the following section, we quantify the extent to which trade policy's encouragement

[42] One caveat associated with this later time period is that it encompasses part of the Great Recession. In Appendix C.1, we re-run all regressions reported in this section using this later post period. Results are similarly consistent with our motivating propositions.

or suppression of particular types of relationships affects trade and aggregate welfare via a general equilibrium framework.

6 Quantitative Analysis

6.1 Multi-Country Model

Environment

We embed our model into the multi-country framework of Eaton and Kortum (2002), generalizing their setup to include the possibility of a trade war and downward-sloping average costs.

The economy consists of three countries indexed by n. In the estimation below, we interpret these countries to be the United States, China, and the Rest of the World.[43] Each country is populated by L_n consumers, who purchase a continuous flow of goods from a manufacturing sector and a non-manufacturing sector to maximize a Cobb-Douglas utility of the form $Q_n^{\alpha} Z_n^{1-\alpha}$, where Q_n is the quantity of a composite manufactured good and Z_n is the quantity of a homogeneous, non-manufactured good. The manufactured good is a CES aggregate of a continuum of differentiated products indexed by $\omega \in [0, 1]$,

$$Q_n = \left(\int_0^1 q_n(\omega)^{(\sigma-1)/\sigma} d\omega \right)^{\sigma/(\sigma-1)}, \tag{15}$$

where $\sigma > 0$ is the elasticity of substitution and $q_n(\omega)$ is the quantity bought of product ω. Consumers purchase each manufactured good ω from homogeneous buyer firms in their country, which offer the goods at prices $p_n(\omega)$. Country n's price index for manufactured goods is thus

$$p_n = \left(\int p_n(\omega)^{1-\sigma} d\omega \right)^{1/(1-\sigma)}. \tag{16}$$

We assume that each consumer supplies one unit of labor.

Manufactured good ω can be produced by country i with the linear production function $q = \frac{\Upsilon}{\theta} l$, where l is the quantity of labor used. Labor is paid the wage rate w_i in country i. The productivity $\Upsilon_i(\omega)$ is specific to each origin country-product combination and, as in Eaton and Kortum (2002), drawn from a Frechet distribution according to

$$F_i(\Upsilon) = e^{-T_i \Upsilon^{-\zeta}}, \tag{17}$$

where the country-specific parameter T_i scales the mean of the distribution and ζ scales the variation. The productivity draws are independent across products within each country. We assume that a country's firms are owned by the household.

We depart from Eaton and Kortum (2002)'s assumption of perfect competition and assume that within each product-destination country pair buyer firms compete in a contestable market. In each such market, trade between buyers and manufacturers takes place in the way described in Section

[43] Our model generalizes to an arbitrary number of countries. However, for our purposes three are sufficient.

2. Buyers compete to procure a given quantity $q_n(\omega)$ by choosing the origin country i, procurement system $s \in \{A, J\}$, and shipment size $x^*_{ni,s}(\omega)$ such that long-run average costs (5) are minimized. A buyer that prices above this minimum average cost curve loses his market and is replaced by an entrant.

Manufacturing in country i incurs fixed costs f_i in units of labor for each destination country that is supplied. Buyers choosing the "American" system need to use an additional $m(\omega)$ labor units, drawn from a distribution $G(m)$, to inspect the quality of product ω. Sellers obtain profits per order cycle that are equal to the incentive premium under the "Japanese" system and zero under the "American" system. We denote by $\pi^s_{ni,s}(\omega)$ the continuous flow of profits to country i from sales to country n of variety ω under system s.

Non-manufacturing output is produced according to $Z_n = a_n L_n^{NM}$, where a_n is country n's productivity in non-manufacturing and L_n^{NM} is labor used. The non-manufacturing good can be costlessly traded across countries, and is therefore used as numeraire. As in Eaton and Kortum (2002) and Antras, Fort, and Tintelnot (2017), we assume that labor is perfectly mobile between manufacturing and non-manufacturing, and that the non-manufacturing sector is sufficiently large so that at least some of its output is produced. The wage rate is then exogenously pinned down by the productivity of the non-manufacturing sector as $w_n = a_n$.

Equilibrium

Equilibrium requires that buyer firms minimize costs such that the contestable market equilibrium is feasible and sustainable, the household maximizes the CES objective, and the goods and labor markets clear.

Buyer firms in country n compare the average costs $AC_{ni,s}(x^*_{ni,s}, q_n(\omega))$ of purchasing a quantity $q_n(\omega)$ across all systems s and origin countries i to minimize overall average costs

$$AC_n(q_n(\omega))^* = \min\left\{\min\left\{AC_{ni,A}(x^*_{ni,A}(\omega), q_n(\omega)), AC_{ni,J}(x^*_{ni,J}(\omega), q_n(\omega))\right\};\ i = 1, ..., N\right\}, \quad (18)$$

where $x^*_{ni,s}(\omega)$ is the optimal batch size for each country-system combination determined by the first-order condition (6). Since marginal costs are downward sloping in q, in equilibrium there will only be one buyer in each market as shown in Section 2, and this buyer procures under one system from exactly one seller country, the one with the lowest average costs. The contestable market price is $p_n(\omega) = AC_n(q_n(\omega))^*$.

From the properties of the Cobb-Douglas utility function, households spend a fraction α of their income on manufactured goods. Consumption of each manufactured good is chosen to maximize (15) subject to the budget constraint

$$\int_0^1 p_n(\omega) q_n(\omega)\, d\omega \leq \alpha \left(w_n L_n + \sum_{n'} \sum_s \int \pi^s_{n'n,s}(\omega) I_{n'n,s}(\omega) d\omega \right), \quad (19)$$

where $I_{ni,s}(\omega)$ is an indicator function that is equal to one if the buyer in country n procures product

ω from country i under system s, and zero otherwise. The right-hand side is country n's total income, W_n. The second term on the right-hand side, which is new relative to Eaton and Kortum (2002), represents the incentive premia collected from shipments to countries n' under $s = J$. Consumption of the non-manufactured good satisfies $Z_n = (1-\alpha)W_n$.

Goods market clearing implies for each manufactured good ω

$$\sum_n \sum_i \sum_s I_{ni,s}(\omega) x^*_{ni,s}(\omega) = \sum_n \sum_i \sum_s I_{ni,s}(\omega) \int_0^{x^*_{ni,s}(\omega)/q_n(\omega)} q_n(\omega) dt \quad \forall \omega, \tag{20}$$

where the left-hand side of the equation indicates worldwide production of product ω in one order cycle and the right-hand side is worldwide consumption over the same cycle. The market for non-manufactured goods clears as well,

$$\sum_n Z_n = \sum_n a_n L_n^{NM}. \tag{21}$$

Finally, labor market clearing in each country requires that

$$\begin{aligned} L_n = &\sum_{n'} \sum_s I_{n'n,s}(\omega) \int_0^1 \frac{\bar{\theta}}{\Upsilon_n(\omega)} q_{n'}(\omega) d\omega + f_n \sum_{n'} \sum_s I_{n'n,s}(\omega) \int_0^1 \frac{q_{n'}(\omega)}{x^*_{n'n,s}(\omega)} d\omega \\ &+ \sum_i I_{ni,A} \int_0^1 (\omega) m(\omega) \frac{q_{n'}(\omega)}{x^*_{n'n,s}(\omega)} d\omega + L_n^{NM} \quad \forall n \in N, \end{aligned} \tag{22}$$

where the left-hand side is total labor supply in country n, and on the right-hand side we have labor used in manufacturing production, labor used for fixed costs, labor used for inspections, and non-manufacturing labor, respectively. Since the fixed costs and the inspection costs are paid for each shipment, we scale these costs by the number of shipments per period.

Due to our departure from competitive markets, our problem does not admit an analytical solution. We solve for the equilibrium using an iteration procedure. First, we trace out the AC curves in each market by solving the cost minimization problem for different values of $q_n(\omega)$. We then guess a price index for manufactured goods p_n and a total income W_n and compute the demand curve for each market from household utility maximization under these choices. Using the intersection of supply and demand, we obtain a candidate equilibrium price vector $p_n(\omega)$, quantity vector $q_n(\omega)$, and a new price index p_n, which determine trade flows, and hence the amount of labor used in each sector and thus total income W_n. We then use the new values of p_n and W_n to obtain new demand curves and iterate to convergence. Appendix D provides further details on our solution algorithm.

6.2 Parametrization and Calibrated Parameters

We now describe how we estimate the model.

Parametrization We assume that the inspection costs for domestic procurement are zero. This assumption implies that all domestic sourcing will take place under the "American" system. For

imports, we assume that the distribution of inspection costs, $G(m)$, is Frechet

$$G(m) = e^{-\Theta m^{-\gamma}}, \tag{23}$$

with the parameters to be estimated. We choose this distribution to be symmetric with productivity, since their are now empirical estimates of the inspection cost distribution.

Calibrated Parameters We set each period to one quarter, and follow Caliendo, Dvorkin, and Parro (2019) by setting $r = 0.01$. We choose an elasticity of substitution of $\sigma = 3.85$ as estimated in Antras, Fort, and Tintelnot (2017), and normalize $\underline{\theta} = 0$. We choose the cost of high quality $\bar{\theta} = 1$. This parameter only appears via the ratio $\bar{theta}/\Upsilon_i(\omega)$ and we will estimate the mean of the distribution of Υ below to match the data. We set the share of consumption accounted for by manufactured goods to $\alpha = 0.5$ to match the share spent on manufactured goods and food in developed countries from Duarte (2020). We follow Eaton and Kortum (2002) and set the dispersion of productivities to be the same in each country and equal to $\zeta = 3.6$.

The wages of each country, w_n, are exogenous in the model and pinned down by the productivity of the non-manufacturing sector. We normalize the U.S. to $w_{US} = 1$. Since comparable wage data is difficult to obtain for different countries, we set each country's wages using data for 2016 from the World Bank World Development Indicators. Specifically, we assume that two thirds of GDP go to labor in each country and divide the resulting figure by the size of each country's labor force to obtain each country's labor income per capita. For the Rest of the World, we take a weighted average across countries using each country's exports to the U.S. in 2016 as weight. This yields $w_{CN} = 0.12$ and $w_{RoW} = 0.47$. We normalize the U.S. labor force $L_{US} = 1$, and set labor for China so that we match the size of the labor force in 2016 from the World Bank World Development Indicators, yielding $L_{CN} = 5$. We set the labor force for the rest of the world to match the total workforce in the ten main trading partner countries of the U.S. from Table 8, which yields $L_{RoW} = 2.5$. The size of the labor force mainly affects total income and hence U.S. exports to these countries.

We assume that the trade war shocks are symmetric between country pairs, $\rho_{ni} = \rho_{in}$, and assume that a trade war between the U.S. and the Rest of the World is unlikely in steady state by setting $\rho_{US,RoW} = 0$. Nearly all U.S. imports from the Rest of the World from 1993 to 2016 were from WTO members, which are subject to a formal dispute settlement system to avoid a trade war, and hence a trade war with these countries would have been seen as very unlikely. To calibrate the probability of a trade war between the U.S. and China, we proceed as follows. First, we take all buyer-seller ($mxhcz$) quintuples in the data that we estimate to be associated with "Japanese" procurement as described in Section 5.2, i.e., our dummy $J_{mhcz} = 1$. For these relationships, we compute the separation hazard rate, that is, the probability that a quintuple trades for the last time at age t quarters conditional on having lasted that long, separately for U.S.-China and for U.S.-rest of the world trade.[44]. We then fit the exponential decay function $e^{-\psi_{ni}t}$ to minimize the squared deviation between this function and the empirically observed hazard rate for each country pair. We fit this function from quarter two of the relationship onwards since a large share of quintuples break up after only one transaction for

[44]We exclude 2016 from this computation due to the censoring of our data

exogenous reasons. We thus obtain $\psi_{US,RoW} = 0.0873$ and $\psi_{US,CN} = 0.1137$. Given our assumption of a zero probability of trade wars between the U.S. and the rest of the world, we interpret the hazard rate of breakups with the rest of the world as normal churn due to firm exits, product obsolescence, and so on. Given this baseline separation rate, we take the excess probability of break-ups with Chinese suppliers as reflecting the added probability of a trade war, and hence set $\rho_{US,CN} = 0.0264$. We assume that a trade war between the Rest of the World and China is unlikely and set $\rho_{CN,RoW} = 0$. Table 13 summarizes the calibrated parameters.

Table 13: Calibrated Parameters

Parameter	Value	Source
Interest rate (r)	0.01	Caliendo, Dvorkin, and Parro (2019)
Elasticity of substitution (σ)	3.85	Antras, Fort, and Tintelnot (2017)
Cost of low quality ($\underline{\theta}$)	0	Normalization
Cost of high quality ($\bar{\theta}$)	1	Normalization
Consumption share of manufactured goods (α)	0.5	Duarte (2020)
Dispersion of productivities (ζ)	3.6	Eaton and Kortum (2002)
Wage (w_i)		
- U.S.	1	Normalization
- China	0.12	World Bank, authors' calculations
- RoW	0.47	World Bank, authors' calculations
Labor Force (L_i)		
- U.S.	1	Normalization
- China	5	World Bank, authors' calculations
- RoW	2.5	World Bank, authors' calculations
Rate of trade wars, U.S.-China ($\rho_{US,CN}$)	0.0264	Census Bureau (LFTTD)

6.3 Targeted Moments and Estimation

Three sets of parameters remain to be estimated: the productivity scales T_n, the country-specific fixed costs f_n, and the inspection cost parameters Θ and γ. We estimate these parameters via a simulated method of moments procedure using moments observed in the LFTTD data and in external data. While the parameters are jointly estimated, we proceed to describe the empirical moments targeted and the underlying identification assumptions for each parameter in turn. The column labeled "Data" in Table 14 summarizes the targeted moments and their values in the data. We provide more details on how each moment is constructed in Appendix E.

The first set of parameters to be estimated are the productivity scales T_n. We normalize $T_{US} = 1$, and estimate the other two parameters using the U.S.' import penetration from China and from the rest of the world in 2016 (rows 1 and 2 of Table 14). A lower value of T_n increases country n's productivity, which raises that country's share in U.S. consumption and hence imports.

Table 14: Estimated Parameters and Targeted Moments

	Parameter	Value	Moment	Data	Model
(1)	Productivity China (T_{CN})	14.498	Share of Chinese imports in consumption	0.074	0.098
(2)	Productivity RoW (T_{RoW})	2.550	Share of RoW imports in consumption	0.270	0.299
(3)	Fixed costs, China (f_{CN})	0.018	$WBS_{A,CN}$	39.07	35.124
(4)	Fixed costs, RoW (f_{RoW})	0.012	$WBS_{A,RoW}$	42.57	46.821
(5)	Scale of inspection costs (Θ)	9.298	WBS difference Q4-Q1, China	0.871	1.466
(6)			WBS difference Q4-Q1, RoW	0.822	0.810
(7)	Dispersion of inspection costs (γ)	0.122	Sd of WBS, China	0.390	0.323
(8)			Sd of WBS, RoW	0.398	0.291

The country-specific fixed costs f_n are estimated using the frequency of shipments from China and the rest of the world to the United States. Proposition 2.2 implies that higher fixed costs lead to shipments that are more spaced out within a given system. Since we do not observe the procurement system with certainty in the data, we focus on buyer quadruples that fall into the fourth quartile of the distribution of suppliers per shipment (SPS_{mhcz}) within a given region (China or rest of the world) by product by mode bin to isolate quadruples that most likely use the "American" system. We then compute the average weeks between shipments (WBS_{mhcz}) across these quadruples (rows 3 and 4), and construct the same moments in the model.

The scale parameter of inspection costs, Θ governs the frequency of shipments under the "American" system relative to shipments under the "Japanese" system. A lower Θ raises the average inspection cost, making shipments relatively less frequent under "American" procurement by Proposition 2.2. We estimate this moment by running the classification regression (8) with weeks between shipments (WBS_{mhcz}) as dependent variable for China and for the rest of the world,

$$ln(\overline{WBS}_{mhcz}) = \beta_0 + \beta_1 \mathbb{I}_{SPS_{mhcz}=Q4} + \beta_2 \ln(QPW_{mhcz}) + \beta_3 beg_{mhcz} + \beta_4 end_{mhcz} + \lambda_{hcz} + \epsilon_{mhcz}, \tag{24}$$

where we replace SPS_{mhcz} on the right-hand side with a dummy, $\mathbb{I}_{SPS=Q4}$, indicating whether SPS_{mhcz} falls into the fourth quartile of its distribution within a given region by product by mode bin. We keep only the first and the fourth quartile of the SPS distribution for the regression to maximize the likelihood of distinguishing between the two systems. A higher estimated coefficient on the dummy indicates that shipments are relatively more dispersed under the "American" system than under "Japanese" sourcing. Rows 5 and 6 present the estimated coefficients.

Finally, the parameter governing the dispersion of inspection costs, γ, is estimated from the dispersion in shipping times across $mhcz$ quadruples. When gamma is low, the inspection cost draws are more dispersed, leading to a higher variance of the shipping frequencies. We construct this moment by running regression (24) with country-mode fixed effects to preserve variation across products. We then keep only $mhcz$ quadruples that fall into the fourth quartile of the SPS distribution, and compute for these observations the standard deviation of the regression residuals across product-country-mode

bins. Rows 7 and 8 show the estimated moments.

Our estimation algorithm is standard: we solve for a vector of parameters satisfying

$$\phi^* = \arg\min_{\phi \in \mathbb{F}} \sum_x T(m_x(\phi), \hat{m}_x) \tag{25}$$

where $T(\cdot)$ is the percentage difference between the model, $m_x(\phi)$, and data, $\hat{m}_x$, moments. Appendix F provides more details on the estimation algorithm and outcomes.

6.4 Discussion

Our estimation relies on several key choices. First, we assume that inspection costs vary across products. This assumption generates dispersion in the relative costs of American- and Japanese-style procurement and thus variation in the procurement system used across goods coming from the same country, since Table 8 suggests that both systems are used by each exporter country. While the dispersion of productivity across products could in principle generate some variability in shipping systems by affecting the relative demand for each product, and hence the relative importance of fixed and variable costs, in practice this mechanism has only a minor effect. Our assumption implies that we need to parametrize the shape of the distribution of inspection costs. This distribution determines the relative costs of "Japanese" versus "American" procurement and hence the welfare costs of changing the probability of a trade war. Importantly, we need to know not only the inspection costs of products that end up being sourced under the "American" system, but also of products that are not sourced via the "American" system.[45] We choose a Frechet distribution to mirror the assumptions on productivity, and fix the shape of the distribution using the dispersion in shipping times of those products that are most likely procured under the "American" system. This distributional form of products observed under the "American" system then pins down the inspection costs for all products.

Second, an alternative assumption to induce variation in shipping systems would be to allow for heterogeneity in the probability of a trade war across products. While both variation in inspection costs and trade war probabilities across products are likely to be relevant in practice, it is difficult to identify both separately without additional information about which shipping system is used. We therefore load all the heterogeneity onto the inspection costs and assume that the trade war probability varies only across countries. We opt for a parsimonious specification and assume that inspection costs are only product-specific, but not also country-specific.

Third, we need to take a stand on how we estimate moments that are specific to one of the two systems, since we do not observe the actual contract choice. Guided by our empirical findings, we interpret buyer quadruples in the first quartile of the SPS_{mhcz} distribution as more "Japanese" and cells in the fourth quartile as more "American", and construct moments based on these two quartiles to maximize the likelihood of picking up the two distinct systems. While in the model we *do* know the actual procurement system used, we construct the moments in the model in the same way. We

[45] As an example, nuclear reactors may never be sourced via American-style procurement, but for our counterfactuals we need to know what the costs under the "American" system would have been to compute the welfare costs of disallowing "Japanese" sourcing.

then use the estimated model to infer structurally the shares of each procurement system.

Fourth, the model necessarily does not capture all the heterogeneity present in the data, such as different modes of transportation, heterogeneity across goods, etc. In constructing our moments we try to remove this heterogeneity by controlling for product-country-mode fixed effects in the regression. Additionally, we control for the quantity traded per week in order to account for general equilibrium effects across estimation runs, and use the relationships between parameters and moments holding quantity fixed to estimate the parameter values.

6.5 Model Fit and Results

The column labeled "Value" in Table 14 presents the estimated values of the parameters, and the "Model" column shows the values of the targeted moments. The model provides a good fit along several dimensions. First, it matches the trade shares and generates shipping frequencies consistent with the data. It also generates substantial variation in shipping frequencies across goods, although our model slightly undershoots the empirically observed heterogeneity. This is not surprising: in the data, there are other reasons for differences in shipping times across products than inspection costs, such as seasonality of the good, characteristics of buyers and sellers involved, and the ports most frequently used for a given good. However, overall, our model-generated moments are relatively close to their empirical analogues.

The first column of Table 15 presents some summary statistics of U.S. imports in the estimated equilibrium. The first four rows show the share of manufactured goods consumption that is imported from China and the rest of the world, as well as the share of the imports that are obtained under the "Japanese" system. Our estimates imply that 69 percent of imports from China are under the "Japanese" system, while 76 percent of imports from the rest of the world take place under "Japanese" procurement. The higher share for the rest of the world reflects the higher trade war probability with China, which discourages trade under the "Japanese" system. The structurally estimated "Japanese" shares are in the ballpark of the empirical estimates we obtained using shipments in the first quartile of the SPS_{mhcz} distribution in Table 8, providing an ex-post check of our estimates. Overall, our model suggests that "Japanese" procurement accounts for the majority of imports by value.

Row 5 presents the price index in manufacturing in the United States, $p_{U.S.}$, and row 6 shows the utility $Q_{U.S.}^{\alpha} Z_{U.S.}^{1-\alpha}$. Both of these are normalized to 1 for ease of interpretation. Row 7 shows that the share of labor in the manufacturing sector is 81%, with the rest of labor used to make the non-manufactured good.

Our estimated inspection cost parameters imply that the average product imported by the U.S. is subject to an inspection cost of 0.13 percent of the import value. We estimate fixed costs of 1.93 percent of the import value. These figures are in line with estimates by Kropf and Sauré (2014), who estimate that Swiss exporters face total fixed shipment costs of 0.8 percent to 5.4 percent of the value imported, providing an external validity check. Our method of computing the fixed costs of shipments is similar to theirs, using the shipping frequency and a structural model.

Table 15: Summary Statistics

		Baseline	**No J**
(1)	Value imported from China (%)	9.7%	6.5%
(2)	- of which, "Japanese"	68.9%	0.0%
(3)	Value imported from ROW (%)	29.9%	16.1%
(4)	- of which, "Japanese"	76.1%	0.0%
(5)	Manufacturing price index	1.000	1.065
(6)	Utility	1.000	0.966
(7)	Labor in Manufacturing	0.809	0.787
(8)	Avg. inspection costs	0.13%	2.15%
(9)	Avg. fixed costs	1.93%	0.74%

6.6 Counterfactuals

We next analyze the effect of the arrival rate of trade wars on a country's choice of procurement system and source country. The choice of country and system is shaped by three factors: (i) productivity, (ii) the product-specific inspection costs, (iii) and the country-pair-specific probability of trade peace. Products in which the domestic country has a higher productivity than foreign countries are sourced domestically, since domestic sourcing avoids any cost of inspection. Products in which the domestic country is not the most productive and inspection costs are low tend to be sourced under the "American" system. Products in which the domestic country is not the most productive and inspection costs are high are sourced under the "Japanese" system if the trade war arrival rate is sufficiently low. A higher trade war arrival rate reduces the share of goods that are imported by raising the incentive premium for importing, making domestic sourcing more attractive. As the arrival rate of trade wars goes to infinity, no goods are imported under the "Japanese" system since the incentive premium becomes prohibitive.

To illustrate these mechanisms, Figure 4a shows the value of manufactured goods purchased by origin and system at the estimated parameters as $\rho_{US,RoW}$ increases from zero to infinity while keeping the arrival rate of a trade war with China at $\rho_{US,RoW}+0.0264$.[46] As the trade war arrival rate rises, the incentive premium of "Japanese" sourcing increases, reducing imports under the "Japanese" system. Buyers switch to sourcing under the "American" system and to domestic sourcing, leading to an increase in expenditures for these two sources. In general, buyers switch to "American" sourcing for goods where the productivity advantage of importing is sufficiently great to outweigh the increase in sourcing costs due to having to pay the inspection costs. If the productivity advantage is small or inspection costs are high, buyers switch to domestic sourcing.

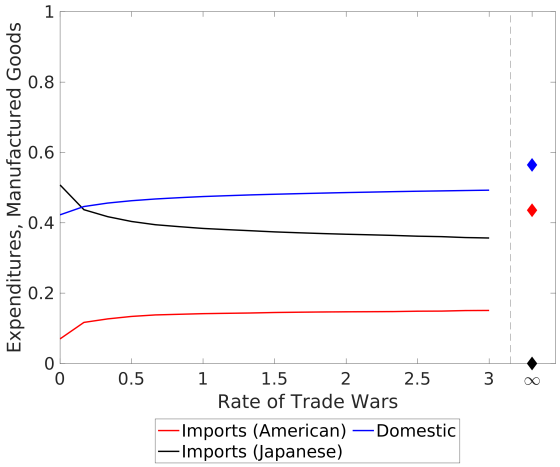
The increase in the trade war arrival rate increases sourcing costs, which raises prices and lowers consumption. Figure 4b shows the total utility of U.S. consumers. As the arrival rate of trade wars rises, overall consumption declines and consumer welfare falls.

The second column of Table 15 presents some statistics for our estimated economy when "Japanese"

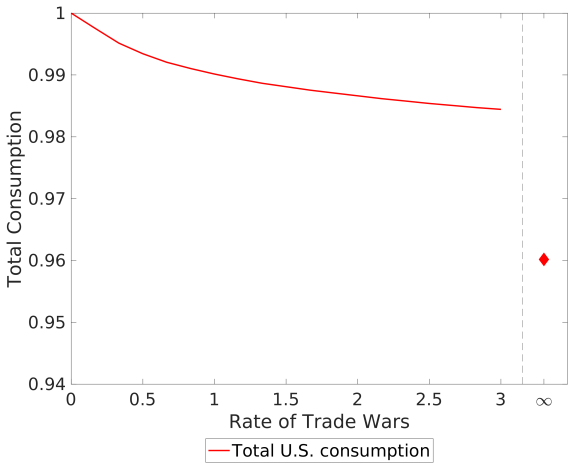
[46] We normalize the total expenditures of manufactured goods at a trade war arrival rate of zero to one.

Figure 4: Effect of Trade War Arrival Rate on Sourcing and Consumption

(a) Manufacturing Expenditures



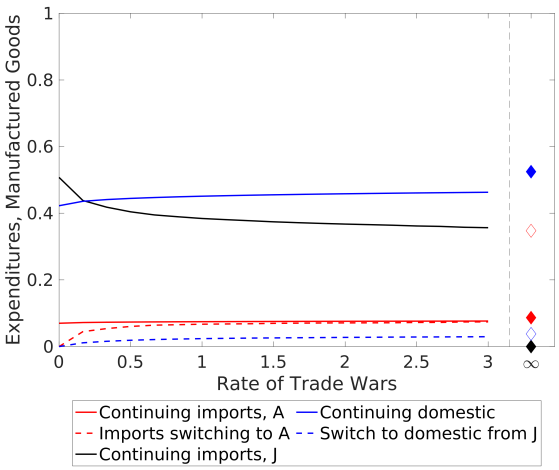
(b) Utility



procurement is not possible, i.e., $\rho_{US,RoW} = \infty$. The share of consumption imported from China falls from 9.7 percent to 6.5 percent (row 1), and sourcing from the rest of the world declines by 14 percent percentage points to 16 percent (row 3). The manufacturing price index p_{US} rises due to the higher sourcing costs by 6.5 percent (row 5), causing a decline in consumer welfare by 3.4 percent (row 6). This result highlights that non-traditional trade policies on their own can generate welfare effects by changing the probability of maintaining long-term relationships, even without the presence of fundamental differences in productivity. While some manufacturing production is brought back to the United States, world production declines due to the less efficient production, and U.S. labor in manufacturing falls by 2 percentage points due to the lower world demand (row 7). Finally, since the products which switch from "Japanese" to "American" imports have relatively high inspection costs, the average inspection cost of imported products rises (row 8). However, the average fixed cost as a share of import value falls because import prices have increased (row 9).

Figure 5 further decomposes the change in manufacturing expenditures relative to the baseline into different adjustment margins. First, the solid red line plots the level of expenditures on imports under the "American" system for goods that were imported under the "American" system at $\rho = 0$. We refer to these imports as imports that continue to be American. The dashed red line shows imports that were not under the "American" system at $rho = 0$ but that are "American" for greater trade war arrival rates, i.e., products that switch towards the "American" system. Both of these lines increase with the rate of trade wars. The majority of the increase in imports under the "American" system comes from products that switch systems. However, there is also a small increase for continuing goods under the "American" system. Consumers increase their expenditures on these goods because they become relatively cheaper compared to products sourced under the "Japanese" system, and hence there is a substitution effect towards them.

Figure 5: Decomposition of Adjustment Margins



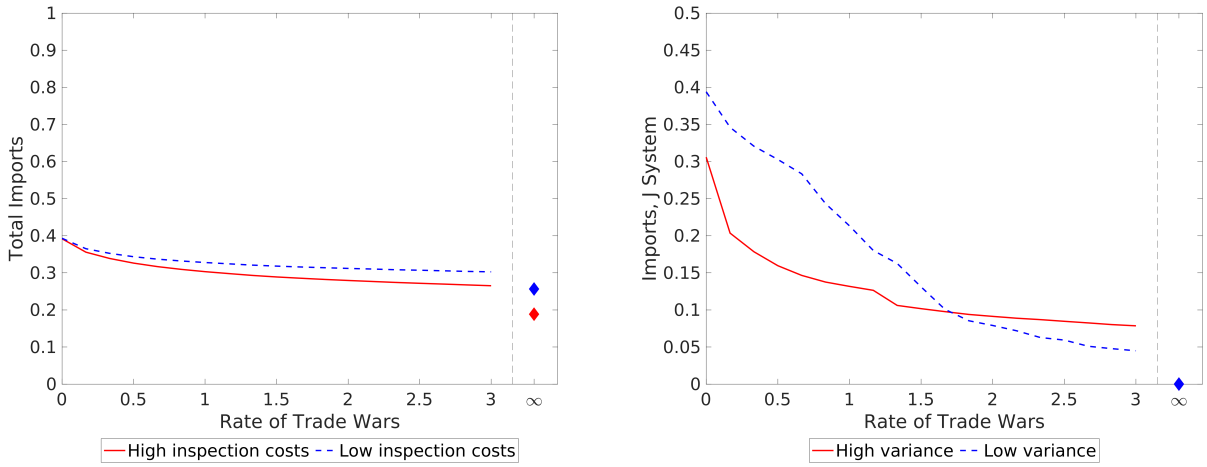
The blue lines decompose the domestic sourcing, for goods that were were already sourced domestically at $\rho = 0$ (solid blue line) and for goods that were sourced via Japanese-style imports at $\rho = 0$ and switch to domestic (blue dashed line). All two lines increase with the arrival rate of trade wars. However, only relatively few goods switch from importing to domestic procurement. Instead, most of the increase in domestic sourcing comes from products that were already obtained domestically at $\rho = 0$, since these products see an increase in demand due to a substitution effect.

The effect of changes in the probability of trade peace on trade depends crucially on the distributions of inspection costs. To illustrate this point, we start from the baseline equilibrium but assume an inspection cost parameter $\Theta = 1$, which increases the average inspection costs relative to the baseline. The red solid line in Figure 6a shows total U.S. imports for different trade war arrival rates under this scenario. For comparison, the blue dashed line presents U.S. imports when average inspection costs are low, $\Theta = 100$. In the high inspection cost scenario, total imports decline relatively quickly with the arrival rate of trade wars. In this case, a worsening in the trade policy environment raises the costs of the "Japanese" system, and since inspections under the "American" system are expensive total imports fall. In contrast, the relationship between imports and trade wars is relatively flat when inspection costs are low. In this case, the "American" system can be used to circumvent the low quality of institutions, which prohibit trade under the "Japanese" system.

The distribution of inspection costs γ across products is similarly important. The red solid line in Figure 6b shows how a change in the trade war arrival rate affects U.S. imports under the "Japanese" system when inspection costs are very dispersed across products, $\zeta = 0.2$. In that case, an increase in the arrival rate of trade wars decreases "Japanese" procurement only slightly because only a few additional products become competitive under the "American" system because their inspection costs are so dispersed. The blue dashed line instead shows the analogous effects when the dispersion of inspection costs is small, $\zeta = 2$. In this case, nearly all products have the same inspection costs. As a result, when the trade war arrival rate is low nearly all products are sourced under the "Japanese"

Figure 6: Effect of Distribution of Inspection Costs on Sourcing

(a) Mean Parameter (Θ) (b) Dispersion Parameter (γ)



system, and when the arrival rate is high all products are sourced under the "American" system, leading to a steep drop in "Japanese" procurement as the trade war arrival rate rises.

7 Conclusion

This paper analyzes the impact of changes in trade policy on procurement patterns along a supply chain using theory, data and quantitative methods. We develop a theoretical model in which importers' ability to solve a quality control problem depends upon exporters' beliefs about the possibility of a trade war breaking out between the firms' countries. When the probability of trade peace is small, buyers choose American-style procurement, characterized by competitive bidding for large, infrequent orders, and costly inspections to ensure the provision of high-quality goods. When the probability of trade peace is high, buyers can induce sellers to provide high quality without inspections by paying them a premium above their costs over a long-term relationship. We show that changes in trade policy that reduce the likelihood of trade wars increase welfare by lowering procurement costs.

We examine the model's key implications using transaction-level U.S. import data. We begin by classifying importer-exporter relationships as American- or Japanese-style and show that these relationships differ along the dimensions – such as shipment size, shipment frequency and shipment size – emphasized in the model. Next we estimate the effect of the U.S. granting of Permanent Normal Trade Relations – which substantially reduced the possibility of a U.S.-China trade war – on the procurement patterns of U.S.-based firms. Using triple difference-in-differences specification, we show that PNTR is associated with a movement toward more Japanese-style procurement among U.S. importers and Chinese exporters along the dimensions highlighted by the model.

Our findings suggest that an important but under-examined aspect of trade agreements in a world with already low tariffs may be their effect on relationship formation. That is, trade agreements promoting institutions that allow firms to develop more stable relationships may give rise to an additional source of welfare gains from trade associated with reducing inventory and monitoring costs.[47] The extent to which such gains are smaller or larger than those that allow firms better access to contract enforcement or dispute resolution is an interesting area for further research.

[47]Indeed, improving the efficiency of trade relationships is a goal of the recent WTO agreement on trade facilitation. See https://www.wto.org/english/thewto_e/minist_e/mc9_e/desci36_e.htm.

References

Alessandria, G., J. P. Kaboski, and V. Midrigan (2010): "Inventories, Lumpy Trade, and Large Devaluations," *American Economic Review*, 100(5), 2304–39.

——— (2011): "US Trade and Inventory Dynamics," *American Economic Review*, 101(3), 303–07.

Alessandria, G. A., S. Y. Khan, and A. Khederlarian (2019): "Taking Stock of Trade Policy Uncertainty: Evidence from China's Pre-WTO Accession," Working Paper 25965, National Bureau of Economic Research.

Alfaro, L., P. Antràs, D. Chor, and P. Conconi (2019): "Internalizing Global Value Chains: A Firm-Level Analysis," *Journal of Political Economy*, 127(2), 508–559.

Alviarez, V., M. Fioretti, K. Kikkawa, and M. Morlacco (2021): "Two-Sided Market Power in Firm-to-Firm Trade," .

Amiti, M., M. Dai, R. C. Feenstra, and J. Romalis (2020): "How Did China's WTO Entry Affect US Prices?," *Journal of International Economics*, 126, 103339.

Amiti, M., S. J. Redding, and D. E. Weinstein (2019): "The Impact of the 2018 Tariffs on Prices and Welfare," *Journal of Economic Perspectives*, 33(4), 187–210.

Antoniades, A. (2015): "Heterogeneous Firms, Quality, and Trade," *Journal of International Economics*, 95(2), 263 – 273.

Antràs, P. (2003): "Firms, Contracts, and Trade Structure," *Quarterly Journal of Economics*, 118(4), 1375–1418.

——— (2005): "Incomplete Contracts and the Product Cycle," *American Economic Review*, 95(4), 1054–1073.

Antràs, P., and D. Chor (2018): "Chapter 5 - On the Measurement of Upstreamness and Downstreamness in Global Value Chains," World Trade Evolution: Growth, Productivity and Employment. Routledge.

Antràs, P., T. C. Fort, and F. Tintelnot (2017): "The Margins of Global Sourcing: Theory and Evidence from U.S. Firms," *American Economic Review*, 107(9), 2514–64.

Antras, P., T. C. Fort, and F. Tintelnot (2017): "The Margins of Global Sourcing: Theory and Evidence from US Firms," *American Economic Review*, 107(9), 2514–64.

Antràs, P., and E. Helpman (2008): "Contractual Frictions and Global Sourcing," in *The Organization of Firms in a Global Economy*, ed. by E. Helpman, T. Verdier, and D. Marin, pp. 9–54. Harvard University Press.

Bajari, P., S. Houghton, and S. Tadelis (2014): "Bidding for Incomplete Contracts: An Empirical Analysis of Adaptation Costs," *American Economic Review*, 104(4), 1288–1319.

Baumol, W., J. Panzar, and R. Willig (1982): *Contestable Markets and the Theory of Industry Structure*. Harcourt Brace Jovanovich, New York.

Békés, G., L. Fontagné, B. Muraközy, and V. Vicard (2017): "Shipment frequency of exporters and demand uncertainty," *Review of World Economics*, 153(4), 1610–2886.

Bernard, A. B., J. B. Jensen, and P. K. Schott (2009): "Importers, Exporters and Multinationals: A Portrait of Firms in the U.S. that Trade Goods," in *Producer Dynamics: New Evidence from Micro Data*, ed. by T. Dunne, J. B. Jensen, and M. J. Roberts. Chicago University Press, Chicago.

Boehm, C. E., A. Flaaen, and N. Pandalai-Nayar (2019): "Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tōhoku Earthquake," *The Review of Economics and Statistics*, 101(1), 60–75.

——— (2020): "Multinationals, Offshoring, and the Decline of U.S. Manufacturing," *Journal of International Economics*, 127, 103391.

Bustos, P. (2011): "Trade Liberalization, Exports and Technology Upgrading: Evidence on the impact of MERCOSUR on Argentinian Firms," *American Economic Review*, 101(1), 304–340.

Caliendo, L., M. Dvorkin, and F. Parro (2019): "Trade and labor market dynamics: General equilibrium analysis of the china trade shock," *Econometrica*, 87(3), 741–835.

Cicala, S. (2015): "When does Regulation Distort Costs? Lessons from Fuel Procurement in U.S. Electricity Generation," *American Economic Review*, 105(1), 411–444.

De Loecker, J., and P. K. Goldberg (2014): "Firm Performance in a Global Market," *Annual Review of Economics*, 6, 201–27.

De Loecker, J., P. K. Goldberg, A. Khandelwal, and N. Pavcnik (2016): "Prices, Markups, and Trade Reform," *Econometrica*, 84(2), 445–510.

Defever, F., C. Fischer, and J. Suedekum (2016): "Relational Contracts and Supplier Turnover in the Global Economy," *Journal of International Economies*, 103, 147–165.

Duarte, M. (2020): "Manufacturing consumption, relative prices, and productivity," *Journal of Macroeconomics*, 65, 103232.

Dyer, J. H., and W. G. Ouchi (1993): "Japanese-Style Partnerships: Giving Companies a Competitive Edge," *Sloan Management Review*, 35(1), 51–63.

Eaton, J., and S. Kortum (2002): "Technology, Geography, and Trade," *Econometrica*, 70(5), 1741–1779.

Fajgelbaum, P. D., P. K. Goldberg, P. J. Kennedy, and A. K. Khandelwal (2019): "The Return to Protectionism*," *The Quarterly Journal of Economics*, 135(1), 1–55.

Feenstra, R. C., and G. H. Hanson (2005): "Ownership and Control in Outsourcing to China: Estimating the Property-Rights Theory of the Firm," *Quarterly Journal of Economics*, 120(2), 729–761.

Feenstra, R. C., J. Romalis, and P. K. Schott (2002): "U.S. Imports, Exports, and Tariff Data, 1989-2001," NBER Working Paper No. 9387.

Fisman, R., and Y. Wang (2010): "Trading Favors within Chinese Business Groups," *American Economic Review: Papers & Proceedings*, 100(2), 429–433.

Flaaen, A., A. Hortaçsu, and F. Tintelnot (2020): "The Production Relocation and Price Effects of US Trade Policy: The Case of Washing Machines," *American Economic Review*, 110(7), 2103–27.

Flaaen, A., and J. Pierce (2019): "Disentangling the Effects of the 2018-2019 Tariffs on a Globally Connected U.S. Manufacturing Sector," *Federal Reserve Finance and Economics Discussion Paper*, 2019-086.

Goldberg, P. K., and N. Pavcnik (2016): "The Effects of Trade Policy," mineo.

Grossman, G. M., and E. Helpman (2004): "Managerial Incentives and International Organization of Production," *Journal of International Economics*, 63(2), 237–262.

Han, S.-L., D. P. Wilson, and P. Dant, Shirish (1993): "Buyer-Supplier Relationships Today," *Industrial Marketing Management*, 22(4), 331–338.

Handley, K. (2014): "Exporting under Trade Policy Uncertainty: Theory and Evidence," *Journal of International Economics*, 94(1), 50–66.

Handley, K., and N. Limão (2017): "Policy Uncertainty, Trade, and Welfare: Theory and Evidence for China and the United States," *American Economic Review*, 107(9), 2731–83.

Heide, J. B., and G. John (1990): "Alliances in Industrial Purchasing: The Determinants of Joint Action in Buyer-Supplier Relationships," *Journal of Marketing Research*, 27(1).

Heise, S. (2019): "Firm-to-firm relationships and the Pass-Through of Shocks: Theory and Evidence," *FRB of New York Staff Report*, (896).

Helper, S. R., and M. Sako (1995): "Supplier Relations in Japan and in the United States: Are They Converging?," *Sloan Management Review*, 36(3), 77–84.

Hornok, C., and M. Koren (2015a): "Administrative Barriers to Trade," *Journal of International Economics*, 96, S110–S122.

——— (2015b): "Per-Shipment Costs and the Lumpiness of International Trade," *Review of Economics and Statistics*, 97(2), 525–530.

Hummels, D., and A. Skiba (2004): "Shipping the Good Apples Out? An Empirical Confirmation of the Alchian-Allen Conjecture," *Journal of Political Economy*, 112(6), 1384–1402.

Huneeus, F. (2018): "Production Network Dynamics and the Propagation of Shocks," *Graduate thesis, Princeton University, Princeton, NJ.*

Jarosch, G. (2016): "Searching for Job Security and the Consequences of Job Loss," University of Chicago.

Kamal, F., and R. Monarch (2018): "Identifying Foreign Suppliers in US Import Data," *Review of International Economics*, 26(1), 117–139.

Kamal, F., and H. Tang (2015): "Relational Trade: Evidence from the United States," mimeograph.

Khandelwal, A. (2010): "The Long and Short (of) Quality Ladders," *Review of Economic Studies*, 77(4), 1450–1476.

Kikkawa, A. K., G. Magerman, and E. Dhyne (2019): "Imperfect Competition in Firm-to-Firm trade," *SSRN Working paper No. 3389836*.

Kropf, A., and P. Sauré (2014): "Fixed Costs Per Shipment," *Journal of International Economics*, 92(1), 166–184.

Kugler, M., and E. Verhoogen (2012): "Prices, Plant Size, and Product Quality," *Review of Economic Studies*, 79(1), 307–339.

Kukharskyy, B. (2016): "Relational contracts and global sourcing," *Journal of International Economics*, 101(C), 123–147.

Kukharskyy, B., and M. Pflüger (2010): "Relational Contracts and the Economic Well-Being of Nations," IZA Discussion Paper No. 5394.

Liker, J. K., and T. Y. Choi (2004): "Building Deep Supplier Relationships," *Harvard Business Review*, 82(12), 104–113.

Lyons, Thomas F., K., A. R. Krachenberg, and J. W. Henke Jr. (1990): "Mixed Motive Marriages: What's Next for Buyer-Supplier Relations?," *Sloan Management Review*, 31(3), 29–36.

Macchiavello, R., and A. Morjaria (2015): "The Value of Relationships: Evidence from a Supply Shock to Kenyan Rose Exports," *American Economic Review*, 105(9), 2911–45.

——— (2020): "Competition and Relational Contracts in the Rwanda Coffee Chain*," *The Quarterly Journal of Economics*, 136(2), 1089–1143.

Manova, K., and Z. Yu (2017): "Multi-product firms and product quality," *Journal of International Economics*, 109, 116 – 137.

Martin, J. (2012): "Markups, Quality, and Transport Costs," *European Economic Review*, 56(4), 777–791.

Melitz, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71, 1695–1725.

Monarch, R., and T. Schmidt-Eisenlohr (2017): "Learning and the Value of Trade Relationships," *FRB International Finance Discussion Paper*, (1218).

O'Neal, C. R. (1989): "JIT Procurement and Relationship Marketing," *Industrial Marketing Management*, 18(1), 55–63.

Ossa, R. (2014): "Trade Wars and Trade Talks with Data," *American Economic Review*, 104(12), 4104–4146.

Pierce, J., and P. K. Schott (2016): "The Surprisingly Swift Decline of U.S. Manufacturing Employment," *American Economic Review*, 106(7), 1632–1662.

Pierce, J. R., and P. K. Schott (2012): "Concording U.S. Harmonized System Categories Over Time," *Journal of Official Statistics*, 28(1), 53–68.

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

Spencer, B. J. (2005): "International Outsourcing and Incomplete Contracts," *Canadian Journal of Economics*, 38(4), 1107–1135.

Tadelis, S., and F. Zettelmeyer (2015): "Information Disclosure as a Matching Mechanism: Theory and Evidence from a Field Experiment," *American Economic Review*, 105(2), 886–905.

Taylor, C. R., and S. N. Wiggins (1997): "Competition or Compensation: Supplier Incentives under the American and Japanese Subcontracting Systems," *American Economic Review*, 87(4), 598–618.

Tirole, J. (1988): *The Theory of Industrial Organization*. MIT Press, Cambridge, MA.

Verhoogen, E. A. (2008): "Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector," *The Quarterly Journal of Economics*, 123(2), 489–530.

WorldBank (2019): "Trading for Development in the Age of Global Value Chains, World Development Report 2020," .

Appendix

A Analytical Results

A.1 Proofs

A.1.1 Second Order Conditions Hold

American System The second derivative of the average cost yields

$$AC''_A(x,q) = \frac{r}{q}\frac{\left(\frac{r}{q}\right)e^{-rx/q}\frac{\bar{\theta}}{\Upsilon}\left[-2\left(1-e^{-rx/q}\right)+\left(\frac{r}{q}\right)\left[1+e^{-rx/q}\right]\left[x+\frac{f+m}{\bar{\theta}/\Upsilon}\right]\right]}{\left[1-e^{-rx/q}\right]^3}.$$

Thus the first order condition is strictly upward sloping, $AC''_A(x,q)>0$, if and only if

$$\left[1+e^{-rx/q}\right]\left[r\frac{x}{q}+\left(\frac{r}{q}\right)\left(\frac{f+m}{\bar{\theta}/\Upsilon}\right)\right]-2\left[1-e^{-rx/q}\right]>0. \tag{A.1}$$

Consider the case when $f+m=0$. If the condition holds for this case, it must also hold for $f+m>0$,because (A.1) is increasing in $f+m$. Define $y\equiv rx/q$. Note that for $y=0$ and $f+m=0$ the left-hand side of equation (A.1) is equal to zero. Taking the derivative of the left-hand side of equation (A.1) with respect to y we obtain

$$1-e^{-y}(1-y),$$

Thus, the left-hand side of (A.1) is strictly increasing in y for $0<y<1$. Therefore, if $0<y<1$, then $AC''_A(x,q)>0$.

Japanese System

$$\begin{aligned}AC''_J(x) = \Bigg[&\frac{\left(\frac{r}{q}\right)^2 e^{-rx/q}\left[f+\underline{\theta}\frac{1}{\Upsilon}x+e^{(r+\rho)x/q}(\bar{\theta}-\underline{\theta})\frac{1}{\Upsilon}x\right]\left[1+e^{-rx/q}\right]}{\left[1-e^{-rx/q}\right]^3}\\ &-\frac{2\left(\frac{r}{q}\right)e^{-rx/q}\left[\underline{\theta}\frac{1}{\Upsilon}+e^{(r+\rho)x/q}(\bar{\theta}-\underline{\theta})\frac{1}{\Upsilon}\left(1+\left(\frac{r+\rho}{q}\right)x\right)\right]\left[1-e^{-rx/q}\right]}{\left[1-e^{-rx/q}\right]^3}\\ &+\frac{\left(\frac{r+\rho}{q}\right)e^{(r+\rho)x/q}(\bar{\theta}-\underline{\theta})\frac{1}{\Upsilon}\left[2+\left(\frac{r+\rho}{q}\right)x\right]\left[1-e^{-rx/q}\right]^2}{\left[1-e^{-rx/q}\right]^3}\Bigg]\frac{r}{q}.\end{aligned}$$

Then $AC''_J(x)>0$, if and only if the numerator is greater than zero. Note that the numerator increases in f. Therefore if the numerator is positive for $f=0$, it must also be positive for $f>0$. Assume $f=0$, and factor the numerator of $AC''_J(x)$ to obtain

$$\begin{aligned}&\left(\frac{r}{q}\right)e^{-rx/q}\left[\underline{\theta}\frac{1}{\Upsilon}+e^{(r+\rho)x/q}(\bar{\theta}-\underline{\theta})\frac{1}{\Upsilon}\right]\left[\left(\frac{r}{q}\right)x\left(1+e^{-rx/q}\right)-2\left(1-e^{-rx/q}\right)\right]\\ &+\left(\frac{r+\rho}{q}\right)e^{(r+\rho)x/q}(\bar{\theta}-\underline{\theta})\frac{1}{\Upsilon}\left[1-e^{-rx/q}\right]\left\{\left[1-e^{-rx/q}\right]\left[2+\left(\frac{r+\rho}{q}\right)x\right]-2\left(\frac{r}{q}\right)xe^{-rx/q}\right\}\end{aligned}$$

Define $y\equiv rx/q$. For the first term note that $\left(1+e^{-y}\right)y-2\left(1-e^{-y}\right)>0$ for $0<y<1$. For the second term to be positive, we require that $\left(\left[1-e^{-y}\right]\left[2+y+\left(\frac{\rho}{q}\right)x\right]-2ye^{-y}\right)>0$. If $\rho=0$, then $(\cdot)>0$ for $0<y<1$.

Because $(\cdot)$increases in ρ, it must be true that $(\cdot) > 0$ for $\rho > 0$ and $0 < y < 1$. Therefore, if $\rho > 0$ and $0 < y < 1$, then $AC_J''(x) > 0$.

A.1.2 Effect of Quality and Arrival Rate of Trade Wars on Average Costs

$$\left.\frac{\partial p_J}{\partial \underline{\theta}}\right|_{\underline{\theta}<\bar{\theta}} = \left[\left(e^{(r+\rho)x_J/q} - 1\right) x_J r\right] \Big/ \left[\Upsilon(e^{-rx/q} - 1)q\right] < 0$$

$$\left.\frac{\partial p_J}{\partial \bar{\theta}}\right|_{\underline{\theta}<\bar{\theta}} = x_J r e^{(r+\rho)x_J/q} \Big/ \left[\Upsilon(e^{-rx_J/q} - 1)q\right] > 0$$

$$\frac{\partial p_J}{\partial \rho} = \left(e^{(r+\rho)x_J/q} x_J^2(\bar{\theta} - \underline{\theta})r\right) \Big/ q^2\Upsilon\left(1 - e^{-\frac{rx}{q}}\right) > 0$$

Finally, comparing procurement costs in both systems note that:

$$\frac{r}{q}\frac{f + \bar{\theta}\frac{1}{\Upsilon}x_J^* + (\bar{\theta} - \underline{\theta})\frac{1}{\Upsilon}x_J^*\left[e^{rx_J^*/q} - 1\right]}{1 - e^{-rx_J^*/q}} > \frac{r}{q}\frac{f + \bar{\theta}\frac{1}{\Upsilon}x_J^*}{1 - e^{-rx_J^*/q}} > \frac{r}{q}\frac{f + \bar{\theta}\frac{1}{\Upsilon}x_A^*}{1 - e^{-rx_A^*/q}}$$

The first inequality holds since $e^{rx_J^*/q} > 1$, the second inequality holds because the batch size that minimizes average costs in the "Japanese" system is strictly less than the batch size that minimizes average costs in the "American" system when $m = 0$,$x_J^* < x_A^*(m = 0)$. Hence, the average procurement cost under the "Japanese" system is strictly greater than under the "American" system for any $\rho \geq 0$ when $m = 0$.

A.1.3 Proof of Proposition 2.1

For $\bar{\theta} - \underline{\theta} > 0$ and $\rho > 0$, when $m_A = 0$ average costs under the "Japanese" system must be higher than under the "American" system by the discussion above Proposition 2.1 and in Appendix A.1.2. Since average costs under the "American" system grow without bound as $m_A \to \infty$, there must be an m^* such that average costs under the systems are equalized.

A.1.4 Proof of Proposition 2.2

Japanese System: We apply the implicit function theorem to the FOC (6). We obtain

$$\frac{\partial FOC_J}{\partial \rho} = \frac{2xe^{\frac{x\rho}{q}}\left(\bar{\theta} - \underline{\theta}\right)}{q^2\Upsilon(e^{-\frac{rx}{q}} - 1)}\left[\frac{x\rho}{2}\left(e^{\frac{rx}{q}} - 1\right) + q\left(\left(\frac{rx}{2q} + 1\right)e^{\frac{rx}{q}} - \frac{rx}{q} - 1\right)\right]$$

Define $y = rx/q$. Note that $\lim_{y \downarrow 0}\left(\frac{y}{2} + 1\right)e^y - y - 1 = 0$ and $\frac{d}{dy}\left(\frac{y}{2} + 1\right)e^y - y - 1 = -1 + \frac{1}{2}(x + 3)e^x > 0$. Therefore $\frac{\partial FOC_J}{\partial \rho} > 0$. Then by the implicit function theorem

$$\frac{\partial x}{\partial \rho} = -\frac{\frac{\partial FOC_J}{\partial \rho}}{SOC_J} < 0.$$

Remember that $v_J(x_J, \rho) = f + \bar{\theta}\frac{1}{\Upsilon}x_J^* + (\bar{\theta} - \underline{\theta})\frac{1}{\Upsilon}x_J^*\left[e^{rx_J^*/q} - 1\right]$. Average costs in the "Japanese" system are then $\frac{r}{q}\frac{v_J(x_J,\rho)}{1 - exp(-\frac{rx_J}{q})}$. Taking the first order condition of these average costs and setting zero we can write.

$$\frac{\partial v(x_J, \rho)}{\partial x_J} = \frac{r}{q}\frac{v(x_J, \rho)exp(-\frac{rx_J}{q})}{1 - exp(-\frac{rx_J}{q})}$$

Now take the derivative of the unit value, $\frac{v_J(x_J,\rho)}{x_J}$,with respect to ρ to obtain

$$\left(\frac{\partial v(x_J,\rho)}{\partial x_J}\frac{\partial x_J}{\partial \rho}x + \frac{\partial v(x_J,\rho)}{\rho}x_J - v(x_J,\rho)\frac{\partial x_J}{\partial \rho}\right)\frac{1}{x_J^2}$$

Substituting for $\frac{\partial}{\partial x}v(x_J,\rho)$ from the equilibrium condition (22) into (23) we can rewrite (23) to obtain

$$\left[\left(\frac{rx_J}{q}\frac{exp(-\frac{rx_J}{q})}{1-exp(-\frac{rx_J}{q})} - 1\right)\frac{\partial x_J}{\partial \rho}v(x_J,\rho) + \frac{\partial v(x_J,\rho)}{\rho}x_J\right]\frac{1}{x_J^2}$$

Note that $\frac{\partial v(x_J,\rho)}{\rho}x_J = \frac{x_J^3(\bar{\theta}-\underline{\theta})}{exp(-\frac{(r+\rho)x_J}{q})q\Upsilon} > 0$. Also note that $\frac{rx_J}{q}\frac{exp(-\frac{rx_J}{q})}{1-exp(-\frac{rx_J}{q})} - 1 < 0$ for $0 < \frac{rx}{q} < 1$. Then because $\frac{\partial x_J}{\partial \rho} < 0$ we have shown that $\frac{\partial}{\partial \rho}\frac{v_J(x_J,\rho)}{x_J} > 0$

American System: We apply the implicit function theorem to show:

$$\frac{\partial x_A^*}{\partial m} = -\frac{\frac{\partial FOC_A}{\partial m}}{SOC_A} = \frac{r^2 e^{-\frac{rx_A}{q}}}{q^2\left(1-e^{-\frac{rx_A}{q}}\right)^2} > 0$$

Note that unit values in the "American" system are simply $\frac{v_A(x_A)}{x_A} = \frac{f}{x_A} + \frac{\bar{\theta}}{\Upsilon}$.Therefore, $\frac{\partial x_A^*}{\partial m} > 0 \Rightarrow \frac{\partial \frac{v_A(x_A)}{x_A}}{\partial x_A} < 0$.

A.1.5 Proof of Proposition 2.3

Part 1: Comparing shipping sizes: $x_J^* < x_A^*$ First note that if $m = 0$ and $\bar{\theta} - \underline{\theta} = 0$, then average costs in the two procurement systems are identical. If $\frac{\partial x_A^*}{\partial m} > 0$ and $\frac{\partial x_J^*}{\partial \underline{\theta}} > 0$, then $x_J^* < x_A^*$all else equal. We apply the implicit function theorem. Let $FOC_A$and FOC_J denote the first order conditions to minimize average procurement costs, and, let $SOC_A > 0$ and $SOC_J > 0$ be the associated second order conditions that are greater than zero by (A.1.1).

American System

$$\frac{\partial x_A^*}{\partial m} = -\frac{\frac{\partial FOC_A}{\partial m}}{SOC_A} = \frac{r^2 e^{-\frac{rx_A}{q}}}{q^2\left(1-e^{-\frac{rx_A}{q}}\right)^2} > 0$$

Japanese System

$$\begin{aligned}\frac{\partial x_J^*}{\partial \underline{\theta}} = -\frac{\frac{\partial FOC_J}{\partial \underline{\theta}}}{SOC_J} &= \left(\frac{r}{q}\right)\frac{1}{\Upsilon}\frac{\left[1-e^{(r+\rho)x_J^*/q}\left[1+\left(\frac{r+\rho}{q}\right)x_J^*\right]\right]\left[1-e^{-rx_J^*/q}\right]}{\left(1-e^{-rx_J^*/q}\right)^2} \\ &- \left(\frac{r}{q}\right)^2\frac{1}{\Upsilon}\frac{x_J^* e^{-rx_J^*/q}\left[1-e^{(r+\rho)x_J^*/q}\right]}{\left(1-e^{-rx_J^*/q}\right)^2}.\end{aligned}$$

For $(r+\rho)x_J^*/q > 0$, this expression is negative if and only if

$$\frac{\left[1-e^{(r+\rho)x_J^*/q}\left[1+\left(\frac{r+\rho}{q}\right)x_J^*\right]\right]}{\left[1-e^{(r+\rho)x_J^*/q}\right]} > \frac{\left(\frac{r}{q}\right)x_J^* e^{-rx_J^*/q}}{\left[1-e^{-rx_J^*/q}\right]}. \tag{A.2}$$

Note that the left-hand side is greater than 1. Hence, we need to show that the right-hand side is less than 1. Define $y \equiv rx_J^*/q$, where $0 < y < 1$. We find for the right-hand side $\lim_{y\to 0}\frac{ye^{-y}}{1-e^{-y}} = \lim_{y\to 0} 1 - y = 1$. Next,

note that $\frac{d}{dy}\frac{ye^{-y}}{1-e^{-y}} = \frac{e^{-y}\left[(1-y)-e^{-y}\right]}{\left[1-e^{-y}\right]^2} < 0$. It follows that the right-hand side of (A.2) is never greater than 1. Therefore, $\partial FOC/\partial\underline{\theta} < 0$ and$\partial x_J^*/\partial\underline{\theta} > 0$.

Part 2: Comparing unit values: $v_A(x_A)/x_A < v_J(x_J)/x_J$

$$v_s(x_s)/x_s = \begin{cases} \frac{f}{x_A^*} + \frac{\bar{\theta}}{\Upsilon} & \text{if } s = A \\ \frac{f}{x_J^*} + \frac{\bar{\theta}}{\Upsilon} + \left(e^{\frac{(r+\rho)x}{q}} - 1\right)(\bar{\theta} - \underline{\theta})\frac{1}{\Upsilon} & \text{if } s = J \end{cases}$$

Comparing the expressions, $x_A^* > x_J^*$(see Part 1) and $\left(e^{\frac{(r+\rho)x}{q}} - 1\right)(\bar{\theta} - \underline{\theta})\frac{1}{\Upsilon} \Rightarrow v_A(x_A)/x_A < v_J(x_J)/x_J$.

A.1.6 Proof of Proposition 2.4

Part 1: Order size and shipping frequency increase in q.

American System We apply the implicit function theorem to the first order condition in the "American" system. From the first order condition and setting to zero we obtain $v'(x) = \frac{r(v(x)+m)e^{-rx/q}}{q\left(1-e^{-rx/q}\right)}$. Substituting this optimality condition into $\frac{\partial FOC_A}{\partial q}$we obtain

$$\frac{\partial x_A}{\partial q} = -\frac{\frac{\partial FOC_A}{q}}{SOC_A} = \frac{\left[1 - \frac{\frac{rx}{q}\mathrm{e}^{-\frac{rx}{q}}}{1-\mathrm{e}^{-\frac{rx}{q}}} - \frac{rx}{q}\right]}{SOC_A}\frac{r^2\left(v\left(x\right)+m\right)\mathrm{e}^{-\frac{rx}{q}}}{q^3\left(1-\mathrm{e}^{-\frac{rx}{q}}\right)^2}$$

Then, $0 < \frac{rx}{q} < 1 \Rightarrow [\cdot] < 0 \Rightarrow \frac{\partial x_A}{\partial q} > 0$ over the relevant parameter range where costs are positive.

For the shipment frequency, $d(x_A^*/q)/dq < 0$, define $\psi_A = x_A^*/q$. Then, simplifying the first-order condition under the "American" system we have

$$FOC(\psi_A) = \bar{\theta}\frac{1}{\Upsilon}\left[1 - e^{-r\psi_A}\right] - \left(\frac{r}{q}\right)e^{-r\psi_A}\left[f + m + \bar{\theta}\frac{1}{\Upsilon}q\psi_A\right] = 0.$$

Applying the implicit function theorem to this expression yields

$$\frac{\partial\psi_A}{\partial q} = -\frac{\frac{\partial FOC(\psi_A)}{\partial q}}{\frac{\partial FOC(\psi_A)}{\partial\psi_J}} = -\frac{[f+m]}{rq\left[f + m + \bar{\theta}\frac{1}{\Upsilon}q\psi_A\right]} < 0,$$

and hence the time between shipments decreases, i.e., shipping frequency increases.

Japanese System We follow the same strategy as in the proof for the American system. From the first order condition, FOC_J, we obtain $\frac{\partial v_J(x_J,q)}{\partial x_J} = \frac{rv_J(x_J,q)e^{-\frac{rx}{q}}}{q\left(1-e^{-\frac{rx}{q}}\right)}$ which we substitute into $\frac{\partial FOC_J}{\partial q}$to obtain:

$$\begin{aligned}\frac{\partial FOC_J}{q} &= \left[1 - \frac{rxe^{-\frac{rx}{q}}}{q\left(1-e^{-\frac{rx}{q}}\right)} - \frac{rx}{q}\right]\left(\frac{r^2v(x,q)e^{-\frac{rx}{q}}}{q^3\left(1-e^{-\frac{rx}{q}}\right)^2}\right) \\ &- \frac{2(r+\rho)(\bar{\theta}-\underline{\theta})xre^{\frac{x\rho}{q}}}{q^4\Upsilon(e^{-\frac{rx}{q}}-1)^2}\left(\frac{x\rho}{2}\left(e^{\frac{rx}{q}}-1\right) + \left[\left(\frac{rx}{2q}+1\right)e^{\frac{rx}{q}} - \frac{rx}{q} - 1\right]q\right)\end{aligned}$$

Note that $0 < \frac{rx}{q} < 1 \Rightarrow \left[1 - \frac{rxe^{-\frac{rx}{q}}}{q\left(1-e^{-\frac{rx}{q}}\right)} - \frac{rx}{q}\right] < 0$ & $\left[\left(\frac{rx}{2q}+1\right)e^{\frac{rx}{q}} - \frac{rx}{q} - 1\right] > 0 \Rightarrow -\frac{\frac{\partial FOC_J}{q}}{SOC_J} > 0 \Rightarrow \frac{\partial x_J^*}{\partial q} > 0$, because all other terms are positive by inspection.

To see that $d(x_J^*/q)/dq < 0$, define $\psi_J = x_J^*/q$. The first-order condition under the "Japanese" system can then be simplified to

$$\begin{aligned} FOC(\psi_J) = &\left[\underline{\theta}\tfrac{1}{\Upsilon} + \left(\bar{\theta}-\underline{\theta}\right)\tfrac{1}{\Upsilon}e^{(r+\rho)\psi_J}\left[1+(r+\rho)\psi_J\right]\right]\left(1-e^{-r\psi_J}\right) \\ &- \left(\frac{r}{q}\right)e^{-r\psi_J}\left[f + \underline{\theta}\tfrac{1}{\Upsilon}\psi_J q + (\bar{\theta}-\underline{\theta})\tfrac{1}{\Upsilon}e^{(r+\rho)\psi_J}\psi_J q\right] = 0. \end{aligned} \tag{A.3}$$

Applying the implicit function theorem to this expression yields

$$\frac{\partial \psi_J}{\partial q} = -\frac{\frac{\partial FOC(\psi_J)}{\partial q}}{\frac{\partial FOC(\psi_J)}{\partial \psi_J}}.$$

For the numerator, we have

$$\frac{\partial FOC(\psi_J)}{\partial q} = \frac{r}{q^2}e^{-r\psi_J}f > 0.$$

For the denominator we find

$$\begin{aligned} \frac{\partial FOC(\psi_J)}{\partial \psi_J} = &(r+\rho)(\bar{\theta}-\underline{\theta})\tfrac{1}{\Upsilon}e^{(r+\rho)\psi_J}\left[2+(r+\rho)\psi_J\right]\left[1-e^{-r\psi_J}\right] \\ &+ \frac{r^2}{q}e^{-r\psi_J}\left[f + \underline{\theta}\tfrac{1}{\Upsilon}\psi_J + (\bar{\theta}-\underline{\theta})\tfrac{1}{\Upsilon}e^{(r+\rho)\psi_J}\psi_J\right] > 0. \end{aligned}$$

Therefore, $\partial FOC(\psi_J)/\partial q > 0$, and thus $d(x_J^*/q)/dq < 0$.

A.1.7 Proof of Lemma 2.5: Average cost curves are downward sloping, convex and reach a limit

Part 1: Average cost curves are downward sloping

American System The average cost function under the "American" system is

$$AC(q) = \frac{\theta\frac{x}{q} + \frac{f}{q} + \frac{m}{q}}{1 - exp(-\frac{rx}{q})}.$$

Taking the first derivative of the expression with respect to q, and fully writing out also the terms that involve x, we get

$$AC'(q) = \frac{-\frac{f+m}{q^2} + \theta\frac{x'(q)}{q} - \theta\frac{x}{q^2}}{1-exp(-\frac{rx}{q})} - \frac{\frac{r}{q}exp(-\frac{rx}{q})\left[\theta\frac{x}{q} + \frac{f}{q} + \frac{m}{q}\right]x'(q)}{\left[1-exp(-\frac{rx}{q})\right]^2} + \frac{\left(\frac{rx}{q^2}\right)exp(-\frac{rx}{q})\left[\theta\frac{x}{q} + \frac{f}{q} + \frac{m}{q}\right]}{\left[1-exp(-\frac{rx}{q})\right]^2}.$$

Re-arranging this expression, we obtain

$$AC'(q) = \frac{-\frac{f+m}{q^2}}{1-exp(-\frac{rx}{q})} + \frac{1}{q}x'(q)\left\{\frac{\theta}{1-exp(-\frac{rx}{q})} - \frac{\frac{r}{q}exp(-\frac{rx}{q})\left[\theta x + f + m\right]}{\left[1-exp(-\frac{rx}{q})\right]^2}\right\}$$
$$-\frac{x}{q^2}\left\{\frac{\theta}{1-exp(-\frac{rx}{q})} - \frac{\frac{r}{q}exp(-\frac{rx}{q})\left[\theta x + f + m\right]}{\left[1-exp(-\frac{rx}{q})\right]^2}\right\}.$$

Note that the two terms in brackets are the first-order condition of the cost function with respect to x, which is equal to zero (this is the "Envelope condition")! This is key: because in the average cost function x and q almost always appear as x/q, we can re-arrange terms to not only cancel the expression containing $x'(q)$, but also the term involving x/q^2. Thus, we get

$$AC'(q) = \frac{-\frac{f+m}{q^2}}{1-exp(-\frac{rx}{q})}.$$

Now, it is much easier to work with this! Also, note that this clearly shows that average cost curves are decreasing.

Japanese System The proof proceeds in the same way as before. Average costs under the "Japanese" system are

$$AC(q) = \frac{\theta\frac{x}{q}exp(\frac{(r+\rho)x}{q}) + \frac{f}{q}}{1-exp(-\frac{rx}{q})}.$$

The first derivative with respect to q is (ignoring the derivative with respect to x here, which we know must be zero)

$$AC'(q) = \frac{-\frac{f}{q^2} - \theta\frac{x}{q^2}exp(\frac{(r+\rho)x}{q}) - \theta(r+\rho)\frac{x^2}{q^3}exp(\frac{(r+\rho)x}{q})}{1-exp(-\frac{rx}{q})} + \frac{\left(\frac{rx}{q^2}\right)exp(-\frac{rx}{q})\left[\theta\frac{x}{q}exp(\frac{(r+\rho)x}{q}) + \frac{f}{q}\right]}{\left[1-exp(-\frac{rx}{q})\right]^2}.$$

Re-arranging yields

$$AC'(q) = \frac{-\frac{f}{q^2}}{1-exp(-\frac{rx}{q})} - \frac{x}{q^2}\left\{\frac{\theta exp(\frac{(r+\rho)x}{q})\left[1+(r+\rho)\frac{x}{q}\right]}{1-exp(-\frac{rx}{q})} - \frac{\frac{r}{q}exp(-\frac{rx}{q})\left[\theta x exp(\frac{(r+\rho)x}{q}) + f\right]}{\left[1-exp(-\frac{rx}{q})\right]^2}\right\}.$$

Similar to before, the term in curly brackets is the first-order condition with respect to x and is equal to zero. Therefore, we have

$$AC'(q) = \frac{-\frac{f}{q^2}}{1-exp(-\frac{rx}{q})}.$$

This function must be convex because the function under the American system was convex for all m, and thus also for $m = 0$.

Part 2: Average cost curves are convex

American System From A.1.6 we obtain the slope of the average cost curve in q :

$$AC'(q) = \frac{-\frac{f+m}{q^2}}{1 - exp(-\frac{rx}{q})}.$$

Taking the second derivative of average costs then yields

$$AC''(q) = \frac{2\frac{f+m}{q^3}}{1 - exp(-\frac{rx}{q})} - \frac{\left(\frac{rx}{q^2}\right) exp(-\frac{rx}{q}) \left(\frac{f+m}{q^2}\right)}{\left[1 - exp(-\frac{rx}{q})\right]^2} + \frac{\left(\frac{rx'(q)}{q}\right) exp(-\frac{rx}{q}) \left(\frac{f+m}{q^2}\right)}{\left[1 - exp(-\frac{rx}{q})\right]^2}.$$

The last term is positive since $x'(q) > 0$. Therefore, to prove that the average cost function is convex, we only need to show that the first two terms together are positive. These terms can be re-written as

$$\frac{2\left[1 - exp(-\frac{rx}{q})\right] \left(\frac{f+m}{q^3}\right) - \left(\frac{rx}{q}\right) exp(-\frac{rx}{q}) \left(\frac{f+m}{q^3}\right)}{\left[1 - exp(-\frac{rx}{q})\right]^2},$$

which is positive if

$$2\left[1 - exp(-\frac{rx}{q})\right] > \left(\frac{rx}{q}\right) exp(-\frac{rx}{q}).$$

This expression holds if

$$2\left[exp(\frac{rx}{q}) - 1\right] > \left(\frac{rx}{q}\right),$$

which is true. Therefore, average costs are convex, for any m and f.

Japanese System From A.1.6 we obtain the slope of the average cost curve in the "Japanese" system.

$$AC'(q) = \frac{-\frac{f}{q^2}}{1 - exp(-\frac{rx}{q})}.$$

By the same arguments as in the "American" system $AC''(q) > 0$.

Asymptote for both systems We first show $(x(q)/q) \to 0$ as $q \to \infty$.

From the Monotone Convergence Theorem, since $(x(q)/q)$ is strictly decreasing and bounded from below by zero, it must converge to a limit. Call this limit $\psi^* \geq 0$. To show that $\psi^* = 0$, assume for contradiction that $\psi^* = K > 0$. Then, it must be the case that there exists no combination of $\psi = x(q)/q < K$ and q that solves the first-order condition of the cost minimization problem. Thus, if we can find a q solving the first-order condition for a $\psi < K$, then K cannot have been the limit since ψ is strictly decreasing.

For the "American" system, pick any $0 \leq \psi_A < K$. The first-order condition of the cost minimization problem under the American system is

$$\bar{\theta}\frac{w_z}{\Upsilon}\left[1 - e^{-r\psi_A}\right] = \left(\frac{r}{q}\right) e^{-r\psi_A}\left[f + mw_b + \bar{\theta}\frac{w_z}{\Upsilon} q\psi_A\right].$$

Re-arranging this expression, we can solve for q as a function of ψ_A and find that

$$q = \frac{[f + mw_b]\, re^{-r\psi_A}}{\bar{\theta}\frac{w_z}{\Upsilon}\left[1 - e^{-r\psi_A}\left[1 + r\psi_A\right]\right]}. \tag{A.4}$$

This expression gives the q that solves the first-order condition for a given pick of $\psi_A = x_A/q$. If we can show

that for any pick $\psi_A \geq 0$ there exists a $q \geq 0$ solving the equation, then it cannot be the case that $K > 0$ is the limit. For this result to hold, we need to show that the denominator is non-negative. To see that it is non-negative, note that

$$\begin{aligned} 1 - e^{-r\psi_A}\left[1 + r\psi_A\right] &\geq 0 \\ \Leftrightarrow \quad e^{r\psi_A} &\geq 1 + r\psi_A, \end{aligned}$$

which holds. Thus, for any $\psi_A \geq 0$ there exists a $q \geq 0$ solving the equation. In particular, such a q exists for any $\psi_A < K$. Therefore, $(x(q)/q)$ must converge to zero. Indeed, from the equation we can see that for $\psi_A = 0$, q must be infinite.

We can construct a similar proof for the "Japanese" system. The first-order condition under the "Japanese" system is

$$\frac{e^{(r+\rho)\psi_J}\bar{\theta}\frac{w}{\Upsilon}\left[1 + (r+\rho)\psi_J\right]}{1 - e^{-r\psi_J}} = \frac{\left(\frac{r}{q}\right)e^{-r\psi_J}\left[f + e^{(r+\rho)\psi_J}\bar{\theta}\frac{w}{\Upsilon}q\psi_J\right]}{\left[1 - e^{-r\psi_J}\right]^2}.$$

We can re-arrange this expression to solve for q and find that

$$q = \frac{fre^{-r\psi_J}}{\bar{\theta}\frac{w_z}{\Upsilon}e^{(r+\rho)\psi_J}\left[(r+\rho)\psi_J\left[1 - e^{-r\psi_J}\right] + 1 - e^{-r\psi_J}\left[1 + r\psi_J\right]\right]}. \tag{A.5}$$

By the same argument as before, the term in the denominator is non-negative and therefore for any $\psi_J \geq 0$ there exists a $q \geq 0$ solving the equation. Therefore, $(x(q)/q)$ must converge to zero. Indeed, from the equation we can see that for $\psi_J = 0$, q must be infinite.

Convergence in the "American" System Consider average costs $C(x,q)/q$. Under the "American" system, we have that

$$\frac{C(x,q)}{q} = \frac{\theta\frac{x}{q}}{1 - exp(-\frac{rx}{q})} + \frac{\frac{f}{q} + \frac{m}{q}}{1 - exp(-\frac{rx}{q})}.$$

We want to show the limit of this expression goes to a positive number as $q \to \infty$. For the second term we have that

$$\lim_{q\to\infty}\frac{(f+m)\frac{x^*(q)}{q}\frac{1}{x^*(q)}}{1 - exp(-r\frac{x^*(q)}{q})} = \lim_{q\to\infty}\frac{(f+m)\frac{x^*(q)}{q}}{1 - exp(-r\frac{x^*(q)}{q})} \cdot \lim_{q\to\infty}\frac{1}{x^*(q)} = \lim_{\psi_A\to 0}\frac{(f+m)\psi_A}{1 - exp(-r\psi_A)} \cdot 0 = \frac{f+m}{r} \cdot 0,$$

by the multiplication rule of limits, where the first term converges to $(f+m)/r$ by L'Hopital's rule since $\psi_A \to 0$ as $q \to \infty$, and the second term converges to zero because $x^*(q) \to \infty$ as $q \to \infty$. Therefore, the overall term converges to 0.

For the first term we have that

$$\lim_{q\to\infty}\frac{\theta\frac{x}{q}}{1 - exp(-\frac{rx}{q})} = \lim_{\psi_A\to 0}\frac{\theta\psi_A}{1 - exp(-r\psi_A)} = \frac{\theta}{r},$$

where we again applied L'Hopital's rule. Therefore, overall, the average cost function under the "American" system converges to (θ/r), which is positive.

Convergence in the "Japanese" System Next consider the "Japanese" system. We have that average costs are

$$\frac{C(x,q)}{q} = \frac{\theta e^{(r+\rho)(x/q)\frac{x}{q}}}{1 - exp(-\frac{rx}{q})} + \frac{\frac{f}{q}}{1 - exp(-\frac{rx}{q})}.$$

The second term converges to zero by the same argument as before. For the first term we find

$$\lim_{\psi_J \to 0} \frac{\theta e^{(r+\rho)\psi_J}\psi_J}{1 - exp(-r\psi_J)} = \lim_{\psi_J \to 0} e^{(r+\rho)\psi_J} \cdot \lim_{\psi_J \to 0} \frac{\theta\psi_J}{1 - exp(-r\psi_J)} = 1 \cdot \frac{\theta}{r},$$

and hence average costs under the "Japanese" system asymptote to exactly the same positive limit as under the "American" system.

B Additional A vs J Classification Regressions

B.1 A vs J Classification with More Aggregated Suppliers per Shipment

Our baseline regressions for examining whether procurement patterns are related to sellers per shipment in Section 4.1 compute SPS at the level of buyer quadruples ($mhcz$). One concern with this definition might be that buyers obtain shipments across multiple modes of transportation, and therefore procurement systems – and hence SPS – should be better defined at the mhc level. Analogously, SPS could be defined at at an even more aggregated mh level.In Tables A.1 and A.2 we show that our results relating suppliers per shipment and procurement patterns are robust to defining SPS at these higher levels of aggregation.

Table A.1: A vs J Classification Regression With SPS at mhc Level

Dep. var.	(1)	(2)	(3)	(4)
	$\log(QPS_{mhcz})$	$\log(WBS_{mhcz})$	$\log(UV_{mhcz})$	$\log(Length_{mhcz})$
$\log(SPS_{mhc})$	0.346^{***}	0.376^{***}	-0.083^{***}	-0.588^{***}
	0.011	0.011	0.015	0.012
$\log(QPW_{mhcz})$	0.687^{***}	-0.322^{***}	-0.279^{***}	-0.115^{***}
	0.010	0.010	0.009	0.003
Observations	2,966,000	2,966,000	2,966,000	2,966,000
R-squared	0.944	0.654	0.844	0.458
Fixed effects	hcz	hcz	hcz	hcz
Controls	beg, end	beg, end	beg, end	beg, end

Source: LFTTD and authors' calculations. Table reports the results of regressing noted attribute of importer by product by country by mode of transport ($mhcz$) bins on sellers per shipment defined for broader mhc bins (SPS_{mhc}) and total quantity shipped per week (QPW_{mhcz}). (QPS_{mhcz}), (WBS_{mhcz}), and (P_{mhcz}) are average quantity per shipment, average weeks between shipment, and average unit value. All regressions include product by country by mode of transport (hcz) fixed effects, control for the beginning and end week of the quadruplet, and exclude quadruplets with less than 5 shipments. Standard errors, adjusted for clustering by country (c) are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

Table A.2: A vs J Classification Regression With SPS at mh Level

Dep. var.	(1)	(2)	(3)	(4)
	$\log(QPS_{mhcz})$	$\log(WBS_{mhcz})$	$\log(UV_{mhcz})$	$\log(Length_{mhcz})$
$\log(SPS_{mh})$	0.285^{***}	0.311^{***}	-0.063^{***}	-0.489^{***}
	0.017	0.018	0.020	0.012
$\log(QPW_{mhcz})$	0.668^{***}	-0.343^{***}	-0.274^{***}	-0.083^{***}
	0.009	0.009	0.009	0.005
Observations	2,966,000	2,966,000	2,966,000	2,966,000
R-squared	0.940	0.631	0.844	0.391
Fixed effects	hcz	hcz	hcz	hcz
Controls	beg, end	beg, end	beg, end	beg, end

Source: LFTTD and authors' calculations. Table reports the results of regressing noted attribute of importer by product by country by mode of transport ($mhcz$) bins on sellers per shipment defined for broader mh bins (SPS_{mh}) and total quantity shipped per week (QPW_{mhcz}). (QPS_{mhcz}), (WBS_{mhcz}), and (P_{mhcz}) are average quantity per shipment, average weeks between shipment, and average unit value. All regressions include product by country by mode of transport (hcz) fixed effects, control for the beginning and end week of the quadruplet, and exclude quadruplets with less than 5 shipments. Standard errors, adjusted for clustering by country (c) are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

B.2 A versus J Classification for Thicker Relationships

Our baseline regressions for examining whether procurement patterns are related to buyer quadruples' sellers per shipment (SPS_{mhcz} in Section 4.1 are restricted to quadruples with at least five transactions. One concern might be that for quadruples that trade only few times, our variable suppliers per shipment is mismeasured because we did not observe a sufficient number of transactions. In Table A.3, we show that these results are robust to restricting observations to quadruples with at least 10 transactions.

Table A.3: A vs J Classification Regression With Buyer Quadruples With At Least 10 Transactions

Dep. var.	(1)	(2)	(3)	(4)
	$\log(QPS_{mhcz})$	$\log(WBS_{mhcz})$	$\log(UV_{mhcz})$	$\log(Length_{mhcz})$
$\log(SPS_{mh})$	0.285^{***}	0.311^{***}	-0.063^{***}	-0.489^{***}
	0.017	0.018	0.020	0.012
$\log(QPW_{mhcz})$	0.668^{***}	-0.343^{***}	-0.274^{***}	-0.083^{***}
	0.009	0.009	0.009	0.005
Observations	2,966,000	2,966,000	2,966,000	2,966,000
R-squared	0.940	0.631	0.844	0.391
Fixed effects	hcz	hcz	hcz	hcz
Controls	beg, end	beg, end	beg, end	beg, end

Source: LFTTD and authors' calculations. Table reports the results of regressing noted attribute of importer by product by country by mode of transport ($mhcz$) bins on sellers per shipment defined for broader mh bins (SPS_{mh}) and total quantity shipped per week (QPW_{mhcz}). (QPS_{mhcz}), (WBS_{mhcz}), and (P_{mhcz}) are average quantity per shipment, average weeks between shipment, and average unit value. All regressions include product by country by mode of transport (hcz) fixed effects, control for the beginning and end week of the quadruplet, and exclude quadruplets with less than 5 shipments. Standard errors, adjusted for clustering by country (c) are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

B.3 A versus J Classification for Different Modes of Transportation

In this section we examine whether the relationship between buyer quadruples' sellers per shipment and their procurement patterns vary by mode of transport. One concern with the results in the paper might be that the relationship between suppliers per shipment and procurement patterns holds for only some modes of transportation, and in particular not for air shipments. Results are reported in Table A.4, with the top panel focusing on vessel and the bottom on air. As indicated in the table, we find similar results for both forms of transport.

Table A.4: A vs J Classification Regression Across Mode of Transport

Dep. var.	(1)	(2)	(3)	(4)
	$\log(QPS_{mhc})$	$\log(WBS_{mhc})$	$\log(UV_{mhc})$	$\log(Length_{mhc})$
	Vessel			
$\log(SPS_{mhcz})$	0.4185***	0.4512***	-0.1722***	-0.5855***
	0.0131	0.0128	0.0094	0.0142
$\log(Quantity/Week_{mhcz})$	0.6608***	-0.3474***	-0.2631***	-0.1411***
	0.0077	0.0077	0.0089	0.0041
Observations	1,506,000	1,506,000	1,506,000	1,506,000
R-squared	0.9236	0.6855	0.8294	0.4632
	Air			
$\log(SPS_{mhcz})$	0.4104***	0.4434***	-0.0579**	-0.6528***
	0.0185	0.0195	0.0224	0.0087
$\log(Quantity/Week_{mhcz})$	0.7372***	-0.2723***	-0.3001***	-0.0830***
	0.0105	0.0108	0.0155	0.0037
Observations	1,029,000	1,029,000	1,029,000	1,029,000
R-squared	0.9328	0.6348	0.7641	0.4308

Source: LFTTD and authors' calculations. Table reports the results of regressing noted attribute of importer by product by country by mode of transport ($mhcz$) bins on bins' sellers per shipment (SPS_{mhcz}) and total quantity shipped per week. (QPS_{mhcz}), (WBS_{mhcz}), (P_{mhcz}), and ($length_{mhcz}$) are average quantity per shipment, average weeks between shipment, average unit value (i.e. value divided by quantity, and average relationship length. All regressions include product by country by mode of transport (hcz) fixed effects, control for the beginning and end week of the quadruplet, and exclude quadruplets with less than 5 shipments. Standard errors, adjusted for clustering by country, are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

C Additional DID Regressions

C.1 Alternate Time Periods

In the baseline DID regression results presented in the main text we compare relationships from 2002 to 2007 to those from 1995 to 2000. In this section, we show that these results are robust to use of alternate time periods.

In Table A.5...

Table A.5: PNTR Regression, 2004-2009 vs 1995-2000

Dep. var.	(1)	(2)	(3)	(4)	(5)	(6)
	$\ln(QPS_{mxhcz})$	$\ln(WBS_{mxhcz})$	$\ln(UV_{mxhcz})$	$\ln(Length_{mxhcz})$	$\ln(VPW_{mxhcz})$	(QPW_{mxhcz})
$Post_p * China_c * NTRGap_p$	-0.1986***	-0.1986***	0.1492***	0.0701**	-0.0415	-0.2200***
	0.027	0.027	0.024	0.0294	0.0453	0.0504
$ln(QPW_{mxhcz})$	0.4033***	-0.5967***	-0.1333***	-0.1253***		
	0.0073	0.0073	0.0121	0.0034		
Observations	221,000	221,000	221,000	221,000	221,000	221,000
R-squared	0.9797	0.8732	0.9818	0.8476	0.8267	0.9222
Fixed effects	$mxhcz, t$	$mxhcz, t$	$mxhcz, t$	$mxhcz, t$	$mxhcz, t$	$mxhcz, t$
Controls	Yes	Yes	Yes	Yes	Yes	Yes

Source: LFTTD and authors' calculations. Table reports the results of regressing noted attribute of US importer by exporter by product by country by mode of transport ($mxhcz$) bins on the difference-in-differences term of interest and quantity shipped per week. Pre-and post periods are 1995 to 2000 and 2004 to 2009. (QPS_{mxhcz}), (WBS_{mxhcz}), (P_{mxhcz}), ($length_{mxhcz}$), VPW_{mxhcz}), and QPW_{mxhcz} are average quantity per shipment, average weeks between shipment, average unit value (i.e. value divided by quantity, average relationship length, value per week and quantity per week. All regressions include $mxhcz$ and period p fixed effects, control for the beginning and end week of the quadruplet as well as all variables needed to identify the DID term of interest, and exclude quadruplets with less than 5 shipments. Standard errors, adjusted for clustering by country, are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

C.2 Alternate Time Periods

Our final specification ignores both importer and exporter identity and analyzes shipments within products across periods,

$$\begin{aligned} ln(\overline{Y}_{hczp}) =& \beta_0 + \beta_1 1\{p = Post\} * 1\{c = China\} * NTRGap_h + \gamma\chi_{hczp} \\ &+ \beta_2 ln(TotQuantity)_{hczp} + \delta_{hcz} + \delta_p + \epsilon_{hczp} \end{aligned} \tag{A.6}$$

As above, we require at least one shipment within each $hczp$ bin, and the data are collapsed to the $hczp$ level after the procurement attributes are computed.

Results

Table A.6: Across Product PNTR Regression: 2002-2007 vs 1995-2000

Dep. var.	(1)	(2)	(3)	(4)	(5)	(6)
	$\ln(QPS_{hcz})$	$\ln(WBS_{hcz})$	$\ln(UV_{hcz})$	$\ln(Length_{hcz})$	$\ln(VPW_{hcz})$	(QPW_{hcz})
$Post_p * China_c * NTRGap_p$	-0.3082***	-0.3082***	0.021	-0.1946***	0.6110***	0.8760***
	0.0282	0.0282	0.033	0.0332	0.0793	0.0717
$ln(QPW_{hcz})$	0.5985***	-0.4015***	-0.3265***	0.0947***		
	0.0142	0.0142	0.0155	0.0045		
Observations	355,000	355,000	355,000	355,000	355,000	355,000
R-squared	0.9746	0.9026	0.9594	0.6671	0.8403	0.9063
Fixed effects	hcz, p	hcz, p	hcz, p	hcz, p	hcz, p	
Controls	Yes	Yes	Yes	Yes	Yes	

Source: LFTTD and authors' calculations. Table reports the results of regressing noted attribute of US importer by exporter by product by country by mode of transport (hcz) bins on the difference-in-differences term of interest and quantity shipped per week. Pre-and post periods are 1995 to 2000 and 2002 to 2007. (QPS_{mhcz}), (WBS_{mhcz}), (P_{mhcz}), $(length_{mhcz})$, $VPW_{mxhcz})$, and QPW_{mxhcz} are average quantity per shipment, average weeks between shipment, average unit value (i.e. value divided by quantity, average relationship length, value per week and quantity per week. All regressions include $mxhcz$ and period p fixed effects, control for the beginning and end week of the quadruplet as well as all variables needed to identify the DID term of interest, and exclude quadruplets with less than 5 shipments. Standard errors, adjusted for clustering by country, are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

D Equilibrium Solution Algorithm

We solve for the equilibrium of the model in several steps. We discretize the product space to $\Omega = 5,000$ products for the estimation algorithm.

1. Define a four-dimensional grid with $(K_1 \times K_2 \times K_3 \times Q)$ grid points. Let $\mathbf{k} \equiv (k_1, k_2, k_3, q_k)$ denote a given grid point. Solve numerically for the average costs $AC(\mathbf{k})$ at each grid point under each system, using equation (5), i.e.

$$AC_A(\mathbf{k}) = \min_x \left(\frac{r}{q_k}\right) \frac{k_1 + k_2 x}{\left[1 - e^{-rx/q_k}\right]} \tag{A.7}$$

and

$$AC_J(\mathbf{k}) = \min_x \left(\frac{r}{q_k}\right) \frac{k_1 + e^{(r+k_3)x/q_k} k_2 x}{\left[1 - e^{-rx/q_k}\right]}. \tag{A.8}$$

We denote by $x_A(\mathbf{k})$ and $x_J(\mathbf{k})$ the cost-minimizing shipment sizes under each system at grid point $\mathbf{k}$.

2. Initiate the model by drawing an inspection cost $m(\omega)$ for each product ω from $G(m)$ and by drawing a productivity $\Upsilon_i(\omega)$ from $F_i(\Upsilon)$. Also set the trade war arrival rates ρ_{ni} for each country pair.

3. Compute the average cost for each origin country (i)-destination country (n)-product (ω) triplet under each system s, $AC_{ni,s}(x^*_{ni,s}, q_k)$, at each of the Q quantity grid points. Specifically, we map each $ni\omega$'s draw $(m(\omega), \Upsilon_i(\omega), \rho_{ni})$ to an estimated average cost using linear interpolation on the grid of average costs computed in Step 1, where under the "American" system we use $k_1 = f_i w_i + m(\omega) w_n$, $k_2 = \frac{\bar{\theta}}{\Upsilon_i(\omega)} w_i$ and under the "Japanese" system we use $k_1 = f_i w_i$,

$k_2 = \frac{\bar{\theta}}{\Upsilon_i(\omega)} w_i$, and $k_3 = \rho_{ni}$. Similarly, obtain the shipment sizes from linear interpolation on the grid of shipment sizes computed in Step 1.

4. Determine the cost minimizing system and origin country at each quantity q_k for each destination-product market $n\omega$, using equation (18). This traces out the average cost curve of each tuple.

5. Guess an initial manufacturing price index in each destination country, $p_n(0)$, and an initial total income, $W_n(0)$.

6. Begin iteration, t=0.

   (a) Compute each destination-product market $n\omega$'s demand curve, using utility maximization, by computing for each q_k

$$p_n(\omega; q_k, t) = \left(\frac{\alpha W_n(t)}{q_k} \right)^{\frac{1}{\sigma}} p_n(t)^{\frac{\sigma-1}{\sigma}}. \tag{A.9}$$

   (b) Find the intersection between supply and demand curve in each market, using linear interpolation between grid points, to obtain the equilibrium $(p_n^*(\omega), q_n^*(\omega))$. If there are several intersections, find the last intersection at which the demand curve intersects the supply curve from above.

   (c) Use the equilibrium prices in each market to compute a new price index, $p_n(t+1)$.

   (d) Determine the labor used for production, fixed costs, and inspection costs in each period, and compute the labor used in non-manufacturing L_n^{NM} from the labor market clearing condition (22). Verify that this labor is non-negative in all markets.

   (e) Compute the total income of each country n, which is equal to labor income $w_n L_n$ plus profits under the "Japanese" system, see equation (19).

   (f) Repeat the iteration steps with the new values of $p_n(t+1)$ and $W_n(t+1)$ and iterate to convergence.

We choose to solve for the average costs on a grid for efficiency reasons. While in principle it would be possible to solve numerically for the average costs of each $ni\omega$ tuple directly, thus combining steps 1-3, in practice the numerical solution of the average cost problem is quite time consuming. Since we have to solve the model many times using our MCMC algorithm, this would lengthen the estimation time considerably. By separating out the steps, we can in a first step numerically solve the cost minimization problem for a large number of combinations of parameter values and then save the results prior to the actual estimation procedure. During the estimation we then only use linear interpolation on this estimated cost grid without doing any further numerical solution, which is much faster.

E Construction of Empirical Moments

In this section, we describe how each of the moments used in the estimation is constructed.

1. Share of Chinese imports in domestic sales

We target the U.S. import penetration from China in 2016. We obtain U.S. imports of goods from China in 2016 from the Census Bureau[48], and obtain gross output in the manufacturing sector in 2016 from the Bureau of Economic Analysis (BEA).[49] We also obtain total U.S. imports and exports of goods in 2016 from the Census Bureau. We then construct the target

$$IP_{CN} = \frac{\text{Imports}_{CN}}{\text{Domestic output} + \text{Total imports} - \text{Total exports}}, \tag{A.10}$$

where Imports_{CN} are U.S. imports from China, Domestic output denotes gross output in the manufacturing sector, and Total imports and Total exports are total U.S. imports and exports of goods.

2. Share of rest of the world imports in domestic sales

We target the U.S. import penetration from the rest of the world in 2016. We construct this target analogously to import penetration from China. We obtain U.S. imports of goods from China in 2016 from the Census Bureau, and obtain gross output in the manufacturing sector in 2016 from the Bureau of Economic Analysis (BEA). We also obtain total U.S. imports and exports of goods in 2016 from the Census Bureau. We then construct the target

$$IP_{RoW} = \frac{\text{Imports}_{RoW}}{\text{Domestic output} + \text{Total imports} - \text{Total exports}}, \tag{A.11}$$

where Imports_{RoW} are U.S. imports from the rest of the world, Domestic output denotes gross output in the manufacturing sector, and Total imports and Total exports are total U.S. imports and exports of goods.

F Estimation Algorithm and Outcomes

The objective is to find a parameter vector ϕ^* that solves

$$\arg\min_{\phi \in \mathbb{F}} \sum_x T(m_x(\phi), \hat{m}_x) \tag{A.12}$$

where $T(\cdot)$ is the percentage difference between the model, $m_x(\phi)$, and data, $\hat{m}_x$, moments, and $\mathbb{F}$ is the set of admissible parameter vectors, which is bounded to be strictly positive and finite. In the choice of the function $T(m_x(\phi), \hat{m}_x)$ we follow Jarosch (2016) and **?** and minimize the sum of the percentage deviations between model-generated and empirical moments.

The minimization algorithm that we use to solve the problem combines the approaches of Jarosch (2016) and **?**, and **?**, adapted to our needs.

We simulate, using Markov Chain Monte Carlo for classical estimators as introduced in **?**, 200 strings of length 10,000 (+ 1,000 initial scratch periods used only to calculate posterior variances)

[48] https://www.census.gov/foreign-trade/balance/c5700.html.
[49] https://apps.bea.gov/iTable/iTable.cfm?reqid = 150&step = 2&isuri = 1&categories = gdpxind.

starting from 200 different guesses for the vector of parameters ϕ_0. In the first run, we choose the initial guesses to span a large space of possible parameter vectors. In updating the parameter vector along the MCMC simulation, we pick the variance of the shocks to target an average rejection rate of 0.7, as suggested by **?**. The average parameter values across the 200 strings for the last 1,000 iterations provide a first estimate of the vector of parameters. We then repeat the same MCMC procedure, but we start each string from the parameter estimates of the first step. We pick our final estimates as the average across the parameter vectors, picked from all strings, that are associated with the 100 smallest values of the likelihood functions.

Figure XX illustrates our approach. The black dotted line shows the density function of the last 1,000 iterations across all strings. We pick the optimal parameter following **?** and select the vector of parameters that minimizes the objective function among all our draws.[50] Our estimates are shown with red dotted lines in the figure. These correspond to the estimates reported in Table 14 in the main text. Finally, the blue density functions shows the density, across all strings, of the 10 best outcomes within each string. This density provides a visual representation of the tightness of our estimates, which are, in general, quite good. It is also relevant to notice that all the densities are single-peaked, which suggests that the model is, at least locally, tightly identified.

Figure YY illustrates how each parameter is identified.

[50] More precisely, we take the average across the 100 best outcomes across all the draws.