Tariff Rate Uncertainty and the Structure of Supply Chains*

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

July 19, 2023

Abstract

We show that reductions in trade policy uncertainty promote long-term importer-exporter relationships that ensure provision of high-quality inputs via incentive premia. Empirically, we introduce a method for distinguishing between these "Japanese" versus spot-market (i.e., "American") relationships in customs data, show that their use varies intuitively across trading partners and industries, and find that "Japanese" importing from China increases after a reduction in the possibility of a trade war. Extending the standard general equilibrium trade model to encompass potential trade wars and relational contracts, we estimate that eliminating "Japanese" procurement reduces welfare about a third as much as moving to autarky. (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 George Alessandria, Davin Chor, Kerem Cosar, Teresa Fort, Virgiliu Midrigan, Dan Trefler, and seminar participants at the Bank of Canada, the Barcelona Summer Forum, ERWIT, Indiana, the Inter-American Development Bank, Mannheim, the NBER Fall ITI Meeting, the NBER Japan Conference, Oregon, Toronto, the U.S. International Trade Commission, and the Yale Cowles Conference 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

Since the early 1990s, the rapid expansion of global value chains has promoted a substantial increase in world trade, boosted aggregate productivity, and supported an unprecedented convergence in rich and poor country incomes (Johnson, 2018; World Bank, 2020; Antras and Chor, 2022). For much of this time, the spread of production networks across national borders was bolstered by a global trade policy environment of stable or even declining trade barriers. In recent years, however, increased protectionism has dampened firms' enthusiasm for extending supply chains internationally, threatening the welfare gains of previous decades.[1] In this paper, we demonstrate that mere increases in the *probability* of trade restrictions can reduce welfare by limiting the formation of long-term importer-exporter relationships, thereby raising firms' procurement costs.

In the first part of the paper, we develop a model of global sourcing that builds upon the partial-equilibrium framework for domestic supply chains introduced by Taylor and Wiggins (1997) and inspired by the vast literature on repeated games (e.g., Green and Porter (1984); Abreu, Pearce, and Stacchetti (1990)). In this framework, buyers seek to procure an input from suppliers without observing its quality, and choose an optimal order pattern, payments, and inspections to motivate high-quality shipments from sellers. Buyers' cost-minimizing strategy is one of two simple systems. 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 from a parade of lowest-cost sellers in the spot market. In this system, costly inspections and enforceable contracts deter sellers from cheating. Lower inspection costs favor the "American" system, while policy environments that promote the formation of long-term relationships favor the "Japanese" system.[2]

We extend Taylor and Wiggins (1997) to *international* procurement by linking domestic importers' ability to maintain long-term relationships with 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 under the "Japanese" system.

In the second part of the paper, we use confidential transaction-level U.S. import data from the U.S. Census Bureau to investigate the prevalence of "Japanese" sourcing among U.S. importers. 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 the values and quantities and therefore unit values (prices) associated with each shipment. Guided by our model, we classify U.S. importing firms as using either "Japanese" or "American" procurement based on the number of foreign suppliers from which they purchase a particular product from each country over the sample period, 1992 to 2016. We find that

[1] See, for example, Amiti, Redding, and Weinstein (2019), Fajgelbaum, Goldberg, Kennedy, and Khandelwal (2019), Flaaen and Pierce (2019), Bown, Conconi, Erbahar, and Trimarchi (2020), and Flaaen, Hortaçsu, and Tintelnot (2020).

[2] The global movement towards "Japanese" procurement is documented in a series of studies in the management and economic literatures, e.g., 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), and Liker and Choi (2004).

"Japanese" importing is most prevalent for the source countries of Japan and Mexico, for products classified as transportation and machinery, and for imports obtained by manufacturing firms versus other types of firms. "Japanese" importing increases most substantially in the early 2000s for Brazil, China and Mexico. We then show, consistent with the model, that buyers with a lower ratio of the number of suppliers to the number of shipments – i.e., those buyers that are more "Japanese" – do indeed receive smaller, more frequent shipments at a higher price than "American" buyers of the same product. "Japanese" buyers also tend to be larger, pay higher wages, and have lower inventory to sales ratios. These results provide the first systematic empirical evidence of the "Japanese" and "American" procurement patterns as highlighted by Taylor and Wiggins (1997).

In the third part of the paper, we examine the implications of the model by comparing U.S. importing relationships before and after the 2001 US extension of Permanent Normal Trade Relations (PNTR) towards China, which dramatically reduced the possibility of a trade war between the two countries. Exploiting the fact that the impact of this change in policy 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), in products with greater relative exposure to PNTR (third difference).[3] We find support for the model along two dimensions. First, we show that U.S. importers of more-exposed products from China exhibit a relative reduction in sellers per shipment after the change in policy. Second, examining procurement patterns, we find that these importers' imports become relatively smaller and more frequent, and exhibit higher unit values. Coefficient estimates suggest that a one standard deviation increase in exposure to PNTR is associated with a relative decline in average shipment quantity of 4.5 percent and a relative increase in average shipment frequency and average shipment price of 3.9 percent and 2.1 percent, respectively. Both results indicate a shift from "American" to "Japanese" procurement among the products that benefit most from the elimination of future tariff threats.[4]

In the final part of the paper, we provide the first analysis of the potential consequences of greater trade policy uncertainty on relational contracting, motivated by recent increases in trade policy uncertainty associated with Brexit and trade tensions between the U.S. and its trade partners, in particular China. We assess the importance of "Japanese" sourcing on U.S. welfare by embedding our framework for procurement in the Eaton and Kortum (2002) multi-product, multi-country general equilibrium model. In this setting, sourcing is governed by bilateral trade war arrival rates in addition to standard cross-country differences in productivity, as they affect the relative costs of procurement under the two systems. Quantitative simulations of the model reveal that an increase in the probability of trade war that is sufficient to eliminate "Japanese"-style procurement reduces U.S. welfare about one third as much as placing the U.S. in autarky. Our results show that in the context of quality control, the procurement system is important to understand the consequences of trade policy uncertainty.

[3] 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.

[4] Of course, we only observe the import price paid by the importer to the exporter. The consumer price should still decline due to PNTR since importers switching systems will save the inspection cost under the "American" system.

"American"-style contracts can mitigate the negative effects of trade policy uncertainty since they do not depend on long-term relationships, highlighting the importance of investing into institutions that support such contracts.

Our analysis makes contributions to several literatures. First, we add to the growing body of research on trade wars and trade policy uncertainty (Ossa, 2014; Handley, 2014; Handley and Limão, 2017; Alessandria, Khan, and Khederlarian, 2019) by identifying a new channel – procurement systems – through which changes in the probability of trade war can influence trade patterns and welfare.[5] Our results with respect to changes in unit values as firms switch between procurement systems also highlight a novel source of price variation in response to changes in trade policy that goes beyond the quality premiums and markups studied in the existing literature (Schott, 2004; Verhoogen, 2008; Khandelwal, 2010; Hallak and Schott, 2011; Kugler and Verhoogen, 2012; De Loecker and Goldberg, 2014; Antoniades, 2015; Goldberg and Pavcnik, 2016; De Loecker, Goldberg, Khandelwal, and Pavcnik, 2016; Manova and Yu, 2017).

Second, we contribute to greater understanding of the organization and structure of global value chains (Antràs, Fort, and Tintelnot, 2017; Antràs and Chor, 2018; Antras and Chor, 2022), as well as a larger literature on incomplete contracts, imperfect contract enforcement, and information asymmetries that primarily focuses on buyer-seller relationships within countries (Antràs, 2003, 2005; Grossman and Helpman, 2004; Spencer, 2005; Feenstra and Hanson, 2005; Antràs and Helpman, 2008; Fisman and Wang, 2010).[6] In contrast to much of the research in this area, we consider "Japanese" sourcing rather than firm integration as a solution to firms' quality-control problem. This path is particularly relevant for understanding sourcing in settings where integration is not possible, for example, in China, where foreign firms face numerous formal and informal restrictions regarding ownership of domestic assets. Our work builds on the literature on relational contracting, i.e., repeated transactions between buyers and sellers under an informal agreement, as an alternative to integration (Defever, Fischer, and Suedekum, 2016; Kukharskyy, 2016).[7] In that literature, the pattern of trade between buyers and sellers is usually governed by idiosyncratic time preferences. In our model, 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.[8]

Third, our analysis complements 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

[5] See Handley and Limão (2022) for a review of this literature.

[6] This work also touches on the related literature on trade and networks as summarized in Bernard and Moxnes (2018). Recent research on *domestic* procurement includes Kukharskyy and Pflüger (2010), Bajari, Houghton, and Tadelis (2014), Tadelis and Zettelmeyer (2015), Cicala (2015), and Macchiavello and Morjaria (2020).

[7] See also Macchiavello and Morjaria (2015), who examine the value of reputation for reliability among Kenyan rose exporters.

[8] 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 relate 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; Flaaen, Haberkorn, Lewis, Monken, Pierce, Rhodes, and Yi, 2019; Alviarez, Fioretti, Kikkawa, and Morlacco, 2021), though not in the context of procurement system choice.

on how firms adjust to trade shocks (Alessandria, Kaboski, and Midrigan, 2010, 2011; Kropf and Sauré, 2014; Hornok and Koren, 2015a,b; Békés, Fontagné, Muraközy, and Vicard, 2017). Here, we document that firms that source under the "American" system have higher inventories, and show that trade policy uncertainty can be an important barrier to firms' efforts to reduce inventory costs, as it promotes use of the "American" system. We estimate inspection costs of 0.4 percent of the transaction value for the average import transaction, about one tenth of the average fixed per shipment cost. An interesting question for future research is whether minimizing these inspection costs provides an additional motivation to lobby 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 above, in our setup, product prices depend 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 of 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 a trade war, the underlying manner by which trade regimes affect productivity applies more broadly to any factor 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, pandemics, bad weather, or disruptions at ports.

The remainder of this paper proceeds as follows. Section 2 outlines our partial equilibrium model of international procurement. Section 3 describes the data. Sections 4 through 5 examine "Japanese" sourcing in U.S. trade data. Sections 6 through 8 embed our partial equilibrium model into general equilibrium and examine counterfactuals. Section 9 concludes. An online appendix contains additional explanatory detail and results.

2 Theoretical Model

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.[9] 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 a buyer seeks to obtain inputs from a supplier whose effort is unobservable. This type of problem falls into the class of repeated games with incomplete information (Kandori, 2002; Mailath and Samuelson, 2006). The framework by Taylor and Wiggins (1997) is particularly suitable to our context because the solution

[9]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.

to their repeated game is one of two optimal contracts, corresponding to the two main dimensions by which firms' procurement strategies are often described to differ (Helper and Sako, 1995). 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 deter cheating.

We extend Taylor and Wiggins (1997) to an international trade context by linking incentive premia to the arrival rate of a trade war (this section), and by endogenizing final demand (Section 6). These extensions allow us to analyze the effect of a change in the probability of a trade war on a host of outcomes, including trade flows and welfare.

2.1 The Procurement 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.[10] To complete a production run (i.e., produce one shipment) a seller hires 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 her product's level of quality. The unit input requirement, $\frac{\theta}{\Upsilon}$, allows for variation in quality, giving rise to a "quality control" problem.[11] 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 logistics services, including transport costs.[12] The seller's total costs for each production and delivery cycle are therefore $x\frac{\theta}{\Upsilon} + f$.

2.1.2 The Buyer's Procurement Choices

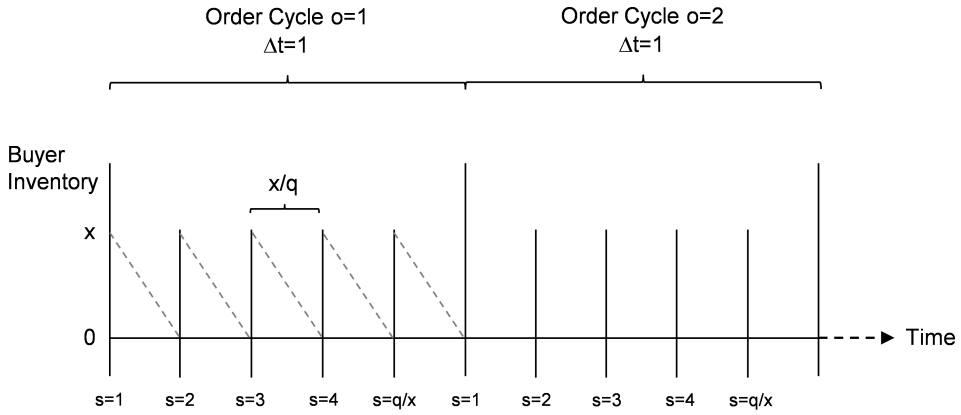
There exist multiple homogeneous buyers that are willing to procure a seller's output and distribute it downstream in the consumer market. These buyers compete in a contestable market, described in more detail below. Conditional on desired quality, $\overline{\theta}$, consumer demand arrives continuously. 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 in this section, but solve for it in equilibrium in Section 6. 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. We take the availability of

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

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

[12] While a large international trade literature 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, 2015a,b). We note that our theory could incorporate an iceberg cost in addition to fixed costs without changing the conclusions; however, it is important that at least some component of costs is fixed to generate discrete shipments.

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.

these two systems as exogenously given, building on the solution to the contracting problem derived by Taylor and Wiggins (1997), and focus on their implications for firms' optimal procurement.

In the "American" (A) system, buyers pay fixed cost m_A to inspect each shipment's quality before delivery.[13] We assume that the buyer always inspects under the "American" system, in contrast to Taylor and Wiggins (1997) who allow for probabilistic inspections.[14] Inspections reveal quality with certainty, and this feature of inspections is known to buyers, so that sellers cannot gain by cheating on product quality. We assume that buyers hold all 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

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

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, 2015a,b), where exogenous per-shipment fees such as f and m_A capture administrative barriers.

[13] These fixed costs could, for example, represent the expense of renting a warehouse where inspections take place, where "[I]t costs the same to have 20 pallets inspected as it does just one." See "What a Year of Brexit Brought U.K. Companies: Higher Costs and Endless Forms," The New York Times, December 29th, 2021.

[14] Generalizing our model to allow for probabilistic inspections would not affect our conclusions below related to per-shipment costs, but would permit us to only obtain numerical solutions. We cannot obtain tractability as Taylor and Wiggins (1997) do by exploiting the fact that the problem is near-linear when the discount factor is small because in our model discounting is additionally affected by the arrival rate of trade wars. We therefore use a different solution approach and derive the properties of the model via implicit function techniques, which necessitates simplifying the model with respect to inspections.

"Japanese" (J) procurement motivates the production of high quality via payments of a premium over a seller's cost during the course of a long-term relationship. A seller chooses to ship high quality if a 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.[15] Thus, over a time interval of length t relationships break with probability $F(t) = 1 - e^{-\rho t}$, implying survival over a shipment cycle with probability $e^{-\frac{\rho x}{q}}$.[16] 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), as would pandemics and associated policy-driven shutdowns.

For our purposes, $e^{-\frac{\rho x}{q}} < 1$ implies that firms are uncertain about whether future trade policy will sustain relationships: for a given shipping cycle $\frac{x_J}{q}$, a greater arrival rate of trade wars, ρ, increases the separation probability. This notion of uncertainty differs from Handley and Limão (2017), where trade policy uncertainty means a greater chance that tariffs either go up or down. In our case, there is no chance that tariffs will decrease, and hence higher uncertainty also raises the expected tariff going forward. 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 $\frac{v_J(x_J,\bar{\theta})}{1-e^{-(r+\rho)x_s/q}}$.[17]

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,

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

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 then excluded from the market forever. Here, too, we assume buyers have 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.

[15] 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.

[16] 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.

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

From (3), a key feature of the "Japanese" system is that more stable trade relationships (i.e., lower separation rates ρ) and smaller shipments, x_J, sent more frequently (which increases the present discounted value of payments) reduce the premium necessary to guarantee desired quality.[18] A long-term relationship with a single seller is optimal under the "Japanese" system since longer streams of payments lower the premium required to incentivize high quality.

Buyers choose between the "American" and "Japanese" system by comparing long-term expected revenues and costs. 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. 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.[19] 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 $m_J = 0$.

2.2 Market Equilibrium and Optimal Procurement Choice

We now study the optimal procurement pattern and system for a given fixed order quantity q. We assume in this section that buyers' profits are zero, and show in Section 6 that our contestable markets setup delivers zero profits as an equilibrium outcome with endogenous q. Setting profits equal to zero implies that the market price must equal average procurement and distribution costs, $AC_s(x_s, q)$, and thus from (4) we obtain the following zero-profit conditions for each system s,

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

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 the optimal 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

[18] 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.

[19] 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.

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, $AC_s\left(x_s^*, q\right)$, 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 with $m_A = 0$, 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, $\frac{\partial AC_J\left(x_J^*, q\right)}{\partial \underline{\theta}} \leq 0$ and $\frac{\partial AC_J\left(x_J^*, q\right)}{\partial \rho} \geq 0$.[20] 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 $AC_A\left(x_A^*, q\right) \to \infty$ because average costs grow without bound, $\frac{\partial AC_A\left(x_A^*, q\right)}{\partial m} = \frac{1}{1 - e^{-\frac{rx_A^*}{q}}} > 1$. 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 "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, which alter the average procurement cost under the "Japanese" system, 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 and reduces the threshold inspection cost m^* at which procurement costs under both system are the same, causing the buyer to switch from the "American" to the "Japanese" system if m^* falls below m. This setting contrasts with existing studies of relational contracts in trade, where exogenous heterogeneity in discount rates determines relationship-based transactions (Kamal and Tang, 2015; Defever, Fischer, and Suedekum, 2016; Kukharskyy, 2016).[21] In our framework, buyers endogenously determine the effective discount rate of rx_s/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., $0 < \frac{rx_s}{q} < 1$.

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

[21] 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.

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 lowers the unit value per shipment and raises the size of shipments (i.e, reduces shipment frequency) in the "American" system.*

Proof. See Appendix A.1.4. □

Under the "Japanese" procurement system an increase in the separation rate ρ 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 (i.e., raising shipping frequency). Given that fixed per-shipment costs are spread over smaller shipment sizes, the increase in ρ causes unit values to increase. Procurement patterns under the "American" system are unaffected by an increase in the separation rate. In contrast, an increase in the inspection cost m raises fixed per-shipment costs under the "American" system, and buyers re-optimize by increasing per-shipment quantities (i.e., decreasing shipping frequency). As a corollary, the unit value paid to the seller must go down in the "American" system since the fixed cost f is spread over more units.

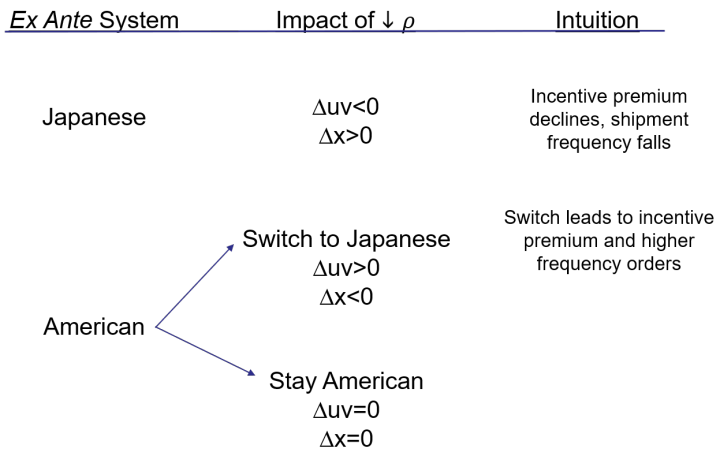
We can use these results to rank shipping frequencies and unit values across the two systems. Starting from the case $\bar{\theta} - \underline{\theta} = 0$ and $m_A = 0$, where the "American" and "Japanese" procurement systems are identical, an increase in $\bar{\theta} - \underline{\theta}$ raises variable shipment costs under the "Japanese" system due to the larger incentive premium. Buyers re-optimize by increasing the shipping frequency and lowering the shipment size, which raises unit values since fixed costs are spread over fewer units. Under the "American" system, by Proposition 2.2, an increase in inspection costs leads to a rise in shipment size, and hence shipping frequency and unit values decrease. For the case of $\bar{\theta} - \underline{\theta} > 0$ and $m \geq 0$, it must therefore be true that shipping sizes are greater in the "American" system. Furthermore, unit values are greater in the "Japanese" procurement system compared to the "American" procurement system. This reasoning forms the basis of our third proposition.

Proposition 2.3. *Batch sizes in the "American" system are greater than in the "Japanese" system, $x_A^* > x_J^*$, and therefore time between shipments is greater under the "American" system, $x_A^*/q > x_J^*/q$. 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, and Figure 2, illustrate that the effect of a decrease in the probability of trade conflict, ρ, depends on whether the adjustment takes place within the "Japanese" system or via a switch from the "American" to the "Japanese" system. If adjustment is within the "Japanese" system, unit values will fall and shipment size will increase (i.e., shipping frequency will decline). On the other hand, if a relationship was "American" before the change in ρ, a decrease in the probability of trade conflict will have no impact on prices and quantities if it does not trigger a change in procurement, or lead to a reduction in quantities and an increase in price if the relationship switches to "Japanese".

Figure 2: Impact of A Decline in the Probability of Trade Conflict (ρ)



Notes: Figure illustrates the impact of a change in the arrival rate of a trade war, ρ on shipment unit values (uv) and quantities (x) under both systems where, e.g., $\Delta uv < 0$ indicates a decline in unit value.

In Section 4, we show that U.S. "Japanese" importing from China is relatively rare in the first part of our sample period, but that it rises over time. Consistent with this finding, we show in Section 5 that a plausibly exogenous reduction in ρ *vis à vis* China primarily results in "American" to "Japanese" switching.

3 U.S. Trade Transaction Data

In this section we describe the confidential data from the U.S. Census Bureau's Longitudinal Foreign Trade Transaction Database (LFTTD) that we use to identify "Japanese" exporters and examine the predictions of the model introduced above. The LFTTD tracks every U.S. import transaction from 1992 to 2016 and includes: 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; the shipment's 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.[22]

We refine the raw data using conventional screens, e.g., by removing all observations that do not include or have an invalid importer identifier, an exporter identifier, an HS code, a value, a quantity, or a valid transaction date. Given our focus on arm's-length trade, we drop all related-party transactions. 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, and deflate all values using the quarterly GDP deflator of the Bureau of Economic Analysis. Since shipments of the same product between the same

[22]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.

buyer and seller spread over multiple containers are recorded as separate transactions, we aggregate the dataset to the weekly level, so that multiple transactions in the same week appear as a single transaction. Finally, to remove unit value outliers, we follow Hallak and Schott (2011) in dropping observations where the unit value is below the 1st or above the 99th percentile within HS10 by country by mode of transportation by quarter cells. For more detail on our data preparation, see Appendix Section B.1.

Our analysis below focuses 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] Since our theory requires that we observe repeated shipments to learn about the procurement system, we exclude buyer quadruples with fewer than five shipments in our analysis. Our sample represents more than 80 percent of all arm's length trade. Table 1 provides an overview.[24] From 1992 to 2016, the importers in our sample purchased 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. The penultimate row shows that our sample contains almost 3 million $mhcz$ 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. 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]

[23] We include mode of transportation in defining these bins to mitigate the influence of spurious sources of variation – e.g., product quality – that might differ across product varieties shipped using different methods.

[24] Quadruples with fewer than five shipments might also represent importers trying out a new product or other idiosyncrasies. In Appendix Section B.2, we provide some statistics comparing our sample against the broader sample of all arm's-length quadruples with at least two transactions. We need at least two transactions to be able to compute some of our variables, such as weeks between shipments (WBS_{mhcz}). As expected, the excluded quadruples with fewer than five transactions tend to be relatively small and trade more rarely.

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

Table 1: U.S. Import Transaction Summary Statistics

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

Source: LFTTD and authors' calculations. Table summarizes U.S. arm's-length imports from 1992 to 2016. Observations are restricted to quadruples with at least five transactions. Import values are in billions of real 2009 dollars. Vessel imports refer to imports arriving over water. The final four rows of the table 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. Observation counts are rounded to the nearest thousand per U.S. Census Bureau disclosure guidelines.

Table 2 summarizes the $mhcz$ quadruples, which are the focus of our study in the next section, 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 (in 2009 units), lasted for 304 weeks and encompassed 39 shipments across 7 sellers. Rows 5 through 7 highlight "procurement patterns," showing that average value per shipment (VPS_{mhcz}), weeks between shipments (WBS_{mhcz}), and buyer-seller relationship length across the relationships within a quadruple ($length_{mhcz}$) 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.

Table 2: Attributes of $mhcz$ Quadruples

	Mean	*Standard Deviation*
Total Value Traded ($)	1,914,000	36,300,000
Length Between Buyer's First and Last Shipment (Weeks)	304.3	266
Total Shipments	38.6	157.9
Number of Sellers (x)	7.3	25.5
Value per Shipment (VPS), ($)	35,910	386,100
Weeks Between Shipments (WBS)	23.5	28.5
Average Relationship Length in Weeks ($length$)	180.8	154.7
Ratio of Sellers to Shipments (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) quadruples during our 1992 to 2016 sample period. Observations are restricted to quadruples with at least five transactions. Observation counts are rounded to the nearest thousand per U.S. Census Bureau disclosure guidelines.

[26] Appendix C provides more details on how all variables are constructed. While below we also analyze quantity per shipment (QPS_{mhcz}) and unit value per shipment (UV_{mhcz}), they are not summarized here due to differences in the quantity units across products. We note that relationship lengths can be subject to both left and right censoring at the beginning and end of our 1992 to 2016 sample period.

4 Identifying "Japanese" Importers in Trade Transaction Data

4.1 Sellers per Shipment (SPS_{mhcz})

A key characteristic of "Japanese" buyers in the model developed in Section 2 is that they trade with just one seller. Guided by this insight, we use the ratio of the number of sellers to the number of shipments (SPS) within importer-product-country-mode ($mhcz$) quadruples,

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

as an observable metric of "Japanese" sourcing. This variable has an upper bound of one, i.e., a different supplier for every shipment, and approaches a lower bound of zero in the case of many transactions sourced from a single seller. Buyers that use fewer sellers relative to the number of shipments (i.e., those with lower values of SPS_{mhcz}) are more likely to be engaged in repeated transactions, and hence in "Japanese" procurement. While "American" buyers might in theory also transact with few sellers if they repeatedly offer the lowest price, introducing noise into our measure, we find below that SPS_{mhcz} is indeed correlated with procurement patterns in a manner consistent with the model.

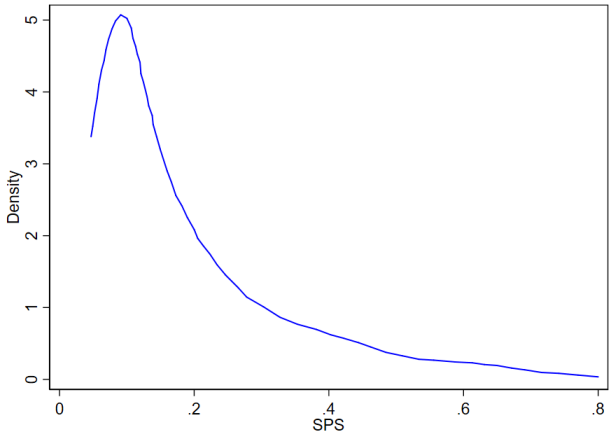
The distribution of SPS_{mhcz} across all buyer quadruples with at least five transactions from 1992 to 2016 is displayed in the kernel density reported in Figure 3. As indicated in the figure, most buyer quadruples have a relatively small ratio of sellers to shipments. Observations in the right tail of the distribution approach a value of 1, i.e., a different seller for each shipment. As reported in the final row of Table 2, the mean ratio of sellers to shipments across buyer quadruples is 0.33, with a standard deviation of 0.24.

Most of the variation in SPS_{mhcz} is driven by importers. We run a series of regressions of SPS_{mhcz} separately on importer, product, country, importer industry, and mode of transportation fixed effects, and examine the R-squared from these regressions as a measure of how much of the variation is explained.[27] We find that importer, product, industry, country, and mode fixed effects individually explain 35%, 12%, 10%, 8%, and 7% of the variation in SPS_{mhcz}, respectively. The large heterogeneity in SPS_{mhcz} across importers is consistent with different firms choosing different procurement strategies.

To provide more detail on the variation of SPS_{mhcz}, we next present several tables that provide summary statistics for different source countries, products, and industries. The first two columns of Table 3 report the weighted average of SPS_{mhcz} for all buyer quadruples with at least five transactions trading with the noted countries, using the total imports of the quadruples as weights. These means are reported for the two five-year time periods used in our regression analysis below, 1995 to 2000 and 2002 to 2007. For the first time period, in column (1), we find that the average SPS_{mhcz} is lowest for U.S. imports from Mexico and Japan, consistent with the prevalence of "Japanese" sourcing in the automobile industry—a key industry in U.S. trade with these countries—including among large

[27] For industry, we use 6-digit NAICS fixed effects. We define the importer's main industry in each year as the one with the largest share of employment, and then take the modal main industry across the years in which the quadruple is active.

Figure 3: 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 all buyer quadruples with at least five transactions between 1992 and 2016. The figure was created according to Census Bureau guidelines and omits observations below the 5th percentile and above the 95th percentile.

Japanese multinationals like Toyota (Boehm, Flaaen, and Pandalai-Nayar, 2020). It is also relatively low for Taiwan and Canada, and relatively high for China, Brazil, and Germany, all of which have values above the rest of the world. Interestingly, there is no obvious effect of distance on SPS_{mhcz}, consistent with the low share of variation explained by the country fixed effect above.

Table 3: "Japanese" Relationships by Country

	Mean SPS		$J^{hz}_{mhcz} = 1$ Share of Import Value	
	(1)	(2)	(3)	(4)
Country	1995-2000	2002-2007	1995-2000	2002-2007
Mexico	0.095	0.068	0.750	0.869
Japan	0.107	0.123	0.756	0.725
Taiwan	0.132	0.114	0.711	0.743
Canada	0.141	0.120	0.602	0.667
United Kingdom	0.146	0.225	0.717	0.519
South Korea	0.156	0.135	0.656	0.724
France	0.177	0.158	0.627	0.667
Rest of the World	*0.180*	*0.156*	*0.625*	*0.678*
Germany	0.184	0.163	0.582	0.606
China	0.185	0.147	0.582	0.693
Brazil	0.190	0.151	0.576	0.706

Source: LFTTD and authors' calculations. Columns (1) and (2) report the weighted average sellers per shipment (SPS_{mhcz}) across buyer quadruples with at least five transactions by country and period, where import values are used as weights. Columns (3) and (4) report the share of the value of U.S. imports accounted for by quadruples with SPS_{mhcz} in the first quartile of the distribution of SPS_{mhcz} within product-mode in the first period. Rows of the table are sorted by column (1).

Results in the second column reveal that, over time, average SPS_{mhcz} generally falls, with imports from Japan and the UK being exceptions. The largest decreases exhibited, both in levels and percent growth, are for Mexico, China and Brazil. The relatively large drop for Mexico may be related to increasingly close supply-chain integration with U.S. producers as a result of NAFTA. In Section 5, we examine whether the decline in SPS_{mhcz} for China is related to the U.S. granting Permanent Normal Trade Relations (PNTR) in 2001.

For our further analysis we make use of a set of dummy variables that indicate plausibly "Japanese" importers, e.g., J^{k}_{mhcz}. In this notation, the subscript defines the quadruple being classified and the superscript k defines the set of quadruples to which it is compared, e.g., the set of all quadruples with the same product-mode, $k = hz$. We construct these dummies as follows. First, we set $J^{k}_{mhcz} = 1$ for buyer quadruples whose SPS_{mhcz} falls within the first quartile of the distribution of SPS in the *first* time period, 1995-2000. For our cross-country comparison, we compute the SPS distribution within product-mode bins, but across countries ($k = hz$). This choice implies that the share of "Japanese" imports can vary between countries even though worldwide 25 percent of quadruples fall into the first quartile by construction. Second, we define analogous dummies for the later time period, also with respect to the distribution of SPS_{mhcz} in the *first* time period, to capture changes with respect to the initial distribution. The final two columns of Table 3 report the share of imports from each country in each time period accounted for by buyer quadruples for which $J^{hz}_{mhcz} = 1$. While the 25^{th} percentile cutoff used in this procedure is arbitrary, it provides a rough indication of variation in "Japanese"

importing across source countries.

As indicated by the high share of U.S. imports accounted for by quadruples for which $J^{hz}_{mhcz} = 1$, "Japanese" quadruples are composed of disproportionately large importers, particularly for Japan and Mexico, where they represent about three quarters of all imports. As we show below, the large share of "Japanese" trade we find here is consistent with the fact that auto trade is the most likely to be "Japanese", and that large firms are more likely to use the "Japanese" system. As with the average SPS_{mhcz} values reported in the first two columns of the table, "Japanese" import value shares increase over time, overall, and most strongly for Brazil, China, and Mexico.

Table 4: "Japanese" Relationships by HS Category

	Mean SPS		$J^{cz}_{mhcz} = 1$ Share of Import Value	
	(1)	(2)	(3)	(4)
Product code (HS chapter)	1995-2000	2002-2007	1995-2000	2002-2007
Transportation (86-89)	0.107	0.081	0.783	0.880
Machinery (84-85)	0.130	0.133	0.754	0.763
Plastics (39-40)	0.130	0.096	0.727	0.820
Optical products (90-92)	0.137	0.127	0.726	0.768
Footwear (64-67)	0.142	0.117	0.750	0.827
Other products (93-99)	*0.151*	*0.124*	*0.697*	*0.808*
Metals (72-83)	0.154	0.128	0.600	0.737
Food (16-24)	0.155	0.120	0.601	0.747
Chemicals (28-38)	0.156	0.121	0.600	0.736
Stones & Jewelry (68-71)	0.159	0.141	0.658	0.674
Animal products & vegetables (01-15)	0.166	0.132	0.511	0.608
Minerals (25-27)	0.182	0.203	0.570	0.500
Leather and wood products (41-49)	0.188	0.153	0.556	0.688
Textiles (50-63)	0.224	0.177	0.463	0.604

Source: LFTTD and authors' calculations. The first two columns report the weighted average sellers per shipment (SPS_{mhcz}) across buyer quadruples with at least five transactions by HS category and period, where import values are used as weights. Numbers in parentheses refer to the Harmonized System chapter of the product. The second two columns report the share of the value of U.S. imports accounted for by quadruples with SPS_{mhcz} in the first quartile of the distribution of SPS_{mhcz} within country-mode in the first period. Rows of the table are sorted by column (1).

Table 4 presents similar results for different groups of ten-digit HS codes. For columns (3) and (4), we define the dummies $J^{k}_{mhcz} = 1$ with $k = cz$, and hence use the distribution of SPS_{mhcz} within country-mode (cz) bins across products in the first time period to retain variation across products. We find that "Japanese" sourcing is most prevalent for transportation equipment, machinery, plastics, and optical products. Textiles and commodities such as minerals and animal products are the least likely to use the "Japanese" system. Table 5 shows results by main 6-digit NAICS industry of the importer where we again use $k = cz$ to identify plausibly "Japanese" sourcing. Manufacturers are the most likely to use "Japanese" sourcing, consistent with these firms obtaining relatively customized inputs for their production processes.

Table 5: "Japanese" Relationships by Main Industry of the Importer

	Mean SPS		$J^{cz}_{mhcz}=1$ Share of Import Value	
	(1)	(2)	(3)	(4)
Industry code (NAICS)	1995-2000	2002-2007	1995-2000	2002-2007
Manufacturing (31-33)	0.119	0.113	0.739	0.778
Agriculture (11)	0.123	0.106	0.584	0.630
Wholesale (42-43)	0.158	0.128	0.623	0.729
Other services	0.160	0.130	0.655	0.713
Professional services (54-55)	0.177	0.220	0.586	0.415
Mining, utilities and construction (21-23)	0.182	0.131	0.561	0.684
Finance and insurance (52-53)	0.187	0.213	0.516	0.514
Retail (44-45)	0.208	0.157	0.532	0.688
Information (51)	0.211	0.182	0.553	0.566
Admin support & waste mgmt (56)	0.213	0.195	0.312	0.423
Transportation and Warehousing (48-49)	0.216	0.210	0.487	0.511

Source: LFTTD and authors' calculations. The first two columns report the weighted average sellers per shipment (SPS_{mhcz}) across buyer quadruples with at least five transactions by main 6-digit NAICS industry-period. To obtain the main NAICS, we find in each year the industry with the importer's largest share of employment, and then take the modal main industry across the years in which the quadruple is active. We aggregate SPS_{mhcz} across quadruples using import values as weights. The second two columns report the share of the value of U.S. imports accounted for by quadruples with SPS_{mhcz} in the first quartile of the distribution of SPS_{mhcz} within country-mode in the first period. Rows of the table are sorted by the column (1).

In the next two sub-sections, we further validate SPS_{mhcz} as measure of "Japanese" sourcing by examining its association with procurement patterns and relationship length.

4.2 SPS_{mhcz} and Procurement Attributes

We evaluate the link between SPS_{mhcz} and procurement patterns via an $mhcz$-level OLS regression,

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

Guided by our theory, the dependent variable, $\overline{Y}_{mhcz}$, represents the key dimensions by which the "American" and the "Japanese" system differ: average quantity per shipment (QPS_{mhcz}), weeks between shipments (WBS_{mhcz}), and unit value (UV_{mhcz}) across all transactions within the $mhcz$ quadruple. In line with our assumption holding quantity fixed in Section 2, we condition on buyers' total order "flow" by controlling for the total quantity imported by the buyer quadruple over its entire lifetime divided by its overall length, in weeks, QPW_{mhcz}.[28] We control for quadruples' first

[28]This variable also controls for the possibility that overall order flow could lead to variation in average shipment sizes or unit values for reasons unrelated to the procurement system. We normalize the total quantity traded by the number of weeks since it is straightforward to implement in our weekly dataset. An alternative would be to use the annual quantity traded, which is simply $52 \cdot QPW_{mhcz}$.

and last weeks of trade, beg_{mhcz} and end_{mhcz}, to capture potential effects of trading in a specific time period—such as a particular stage in the business cycle—and duration effects.[29] Our regression also includes product by country by mode of transportation fixed effects (λ_{hcz}), which 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 two-way clustered at the country and product level.[30]

Results for specification (8) are reported in Table 6. In the first three columns, we find that quadruples with higher SPS_{mhcz}, i.e., those that are more "American", receive shipments for a given total order flow that are larger, less frequent, and lower in price, consistent with Proposition 2.3. Furthermore, the coefficient estimates for our quantity control, QPW_{mhcz}, are in line with Proposition 6.1, discussed below, where an increase in the total quantity procured leads to an increase in shipment size and reductions in the number of weeks between shipments and unit value. Coefficient estimates for SPS_{mhcz} indicate that increasing sellers per shipment by one standard deviation from its mean (from 0.33 to 0.58) is associated with a 23 log point rise in quantity per shipment, a 25 log point increase in weeks between shipments, and a 7 log point decline in price.[31]

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

[30] As before, we only use quadruples with at least five shipments over the entire sample period. In Appendix D.1, we show that results are qualitatively identical for a cutoff of 10 shipments. We describe the construction of all variables in detail in Appendix C.

[31] Our analysis computes 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 Appendix D.2, we re-run specification (8) where we compute SPS using the ratio of sellers to transactions within buyer-product-country triples (SPS_{mhc}) and buyer-product doubles (SPS_{mh}), and find similar results.

Table 6: SPS_{mhcz} and Procurement Attributes

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. var.	$\ln(QPS_{mhcz})$	$\ln(WBS_{mhcz})$	$\ln(UV_{mhcz})$	$\ln(QPS_{mhcz})$	$\ln(WBS_{mhcz})$	$\ln(UV_{mhcz})$
$\ln(SPS_{mhcz})$	0.418^{***}	0.452^{***}	-0.123^{***}			
	0.017	0.017	0.021			
$1\{SPS_{mhcz} = Q2\}$				0.328^{***}	0.350^{***}	-0.117^{***}
				0.014	0.015	0.014
$1\{SPS_{mhcz} = Q3\}$				0.552^{***}	0.591^{***}	-0.179^{***}
				0.024	0.024	0.023
$1\{SPS_{mhcz} = Q4\}$				0.792^{***}	0.856^{***}	-0.226^{***}
				0.034	0.035	0.038
$\log(QPW_{mhcz})$	0.701^{***}	-0.308^{***}	-0.287^{***}	0.687^{***}	-0.323^{***}	-0.282^{***}
	0.014	0.014	0.020	0.013	0.014	0.019
Observations	2,966,000	2,966,000	2,966,000	2,966,000	2,966,000	2,966,000
Fixed effects	hcz	hcz	hcz	hcz	hcz	hcz
R-squared	0.947	0.674	0.845	0.945	0.661	0.845
Controls	beg, end	beg, end	beg, end	beg, end	beg, end	beg, end

Source: LFTTD and authors' calculations. Table reports the results of regressing noted attributes of importer by product by country by mode of transport ($mhcz$) bins on bins' sellers per shipment (SPS_{mhcz}), sellers per shipment quartile dummies, and total quantity shipped per week (QPW_{mhcz}). (QPS_{mhcz}), (WBS_{mhcz}), and (UV_{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 quadruple, and exclude quadruples with less than five shipments. Standard errors, adjusted for clustering by country (c) and product (h) are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

In the final three columns of Table 6, we consider a related specification that relaxes the restriction of a linear relationship between procurement attributes and sellers per shipment by replacing SPS_{mhcz} with a series of dummy variables indicating the quartile into which buyer quadruples fall. We compute these quartiles separately for each hcz bucket using the entire sample period. The first quartile, $1\{SPS_{mhcz} = Q1\}$, is the left-out category. This specification further justifies the use of SPS_{mhcz} as a metric of "Japanese" sourcing, as coefficient estimates for SPS_{mhcz} are of the expected sign, statistically significant at conventional levels, and rise or fall monotonically from quartile 1 to quartile 4 in a manner consistent with the quartiles representing increasingly "American" quadruples. Coefficient estimates in columns 4 and 5 indicate that average quantity per shipment and weeks per shipment for U.S. importers in the fourth quartile are roughly 79 and 86 log points higher than for those in the first quartile. The point estimate for unit value indicates that the average price for the fourth quartile is 23 log points lower than that for the first quartile.[32]

[32]In Appendix Sections D.3 and D.4 we show that the relationships displayed above are robust to analyzing procurement patterns separately by mode of transport, i.e., vessel versus air, and for goods that are differentiated versus commodities. In Appendix Section D.5, we find similar results when we examine procurement patterns within $mxhcz$ buyer-seller relationships. In Appendix Section D.6, we show that the results hold separately within each sector, such as manufacturing or retail.

4.3 SPS_{mhcz} and Relationship Length

Buyers under the "Japanese" system rely on repeat purchases from the same seller, while buyers under the "American" system choose potentially different lowest-cost suppliers for each transaction. An implication of these choices is that "Japanese" buyers have longer relationships with their suppliers. We investigate this prediction using the variable $length_{mhcz}$, which tracks the average length of the mx buyer-seller relationships associated with $mhcz$ buyer quadruples. This variable is constructed in two steps. First, for each $mxhcz$ quintuple, we compute the total number of weeks passed between the first and the last transaction of any product by any mode between the buyer m and seller x, i.e., their total relationship length. Second, for each $mhcz$ buyer quadruple, we take the average of these numbers of weeks across all $mxhcz$ quintuples within the quadruple. 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, add products over time.[33]

Table 7: SPS_{mhcz} and Relationship Length

	(1)	(2)
Dep. var.	$\ln(length_{mhcz})$	$\ln(length_{mhcz})$
$\ln(SPS_{mhcz})$	-0.576^{***}	
	0.015	
$(SPS_{mhcz} = Q2)$		-0.383^{***}
		0.015
$(SPS_{mhcz} = Q3)$		-0.683^{***}
		0.027
$(SPS_{mhcz} = Q4)$		-1.139^{***}
		0.047
$\ln(QPW_{mhcz})$	-0.147^{***}	-0.130^{***}
	0.006	0.005
Observations	2,966,000	2,966,000
R-squared	0.431	0.413
Fixed effects	hcz	hcz
Controls	beg, end	beg, end

Source: LFTTD and authors' calculations. Table reports the results of regressing the average buyer-seller quadruple relationship length ($length_{mhcz}$) on quadruples' sellers per shipment (SPS_{mhcz}), sellers per shipment quartile dummies and total quantity shipped per week (QPW_{mhcz}). The regressions include product by country by mode of transport (hcz) fixed effects. All regressions control for the beginning and end week of the quadruple, and exclude quadruples with less than 5 shipments. Standard errors, adjusted for clustering by country (c) and product (h) bin are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

We use the same specification outlined in equation (8) but using $length_{mhcz}$ as the dependent

[33] See Appendix Section C for more details on the variable construction.

variable. The results, reported in Table 7, show that $mhcz$ buyer quadruples with lower ratios of suppliers per shipments tend to have longer relationships. In column (1), we find that a one standard deviation increase of sellers per shipment from its mean is associated with a 31 log point decrease in average relationship length. In column (2), the average relationship length for quadruples in the fourth quartile (most "American") is about 114 log points lower than that in the first quartile ("most Japanese").[34]

4.4 SPS_{mhcz} and Buyer Characteristics

As shown above, the importer dimension is the most important for explaining variation in SPS_{mhcz}. We therefore next investigate how various firm-level attributes are related to import sourcing strategy, measured by SPS_{mhcz}. We aggregate the quadruples across products, countries, and modes to the importer-level and run

$$\ln(\overline{Y}_m) = \beta_1 \ln(SPS_m) + \epsilon_{m\cdot}, \tag{9}$$

where $\overline{Y}_m$ is one of importer m's total sales, total payroll, average wage, or the firm's inventory-to-sales ratio, and SPS_m is an average of SPS_{mhcz} across all quadruples of the importer. We obtain sales, payroll, and wages at the firm-level from the Longitudinal Business Database (LBD), where the average wage is constructed as the firm's total payroll divided by the number of employees. We obtain beginning-of-year inventories for all manufacturing firms from the Census of Manufactures (CMF). We use for each firm attribute the earliest non-missing observation available for the firm.[35] Table 8 shows the regression results.

Table 8: SPS_m and Firm Characteristics

	(1)	(2)	(3)	(4)
Dep. var.	$\ln(sales_m)$	$\ln(pay_m)$	$\ln(wage_m)$	$(inv/sales)_m$
$\ln(SPS_m)$	-0.291^{***}	-0.350^{***}	-0.056^{***}	0.015^{***}
	0.005	0.006	0.002	0.001
Observations	184,000	184,000	184,000	48,500
R-squared	0.015	0.018	0.003	0.006

Source: LFTTD and authors' calculations. Table reports the results of regressing importer characteristics in the year of the importer's first transaction on sellers per shipment (SPS_{mhcz}) averaged across all quadruples involving the importer. All regressions exclude quadruples with less than five shipments. $(sales_m)$, (pay_m), $(wage_m)$, and $((inv/sales)_m)$ are total sales, total payroll, average wage (i.e., payroll divided by number of employees), and total inventory at the beginning of the year divided by total sales, respectively. Robust standard errors are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

We find that firms that on average rely on more "American" procurement practices tend to be

[34] In Appendix sections D.1 to D.6, we show that all our robustness checks also go through for the length variable. In Appendix D.7, we consider an alternative definition of relationship length where we treat each quintuple as a separate relationship, rather than using the overall importer-supplier pair, and show that our results still hold.

[35] Appendix D.8 shows that the results are robust to using an average across all years in which the firm is active.

smaller, pay lower wages, and hold higher inventories. An increase in average sellers per shipment by one standard deviation from its mean is associated with a 16 log point decline in sales, 19 log point decline in payroll, and a 3 log point decline in the average wage.[36] A one standard deviation increase in SPS_{mhcz} from its mean raises the inventory-to-sales ratio by 0.8 log points, consistent with "American" procurement leading to larger inventories.[37]

Finally, consistent with a firm-wide sourcing strategy, we find that importers' procurement system is correlated across products. Using all importers with at least two products in a given country-mode bin, we randomly draw two of these products for each importer. We then use the "Japanese" dummy J^{hcz}_{mhcz}, computed using the distribution of SPS_{mhcz} within hcz bins for the entire sample period, and regress $J^{hcz}_{mhcz,1}$ of the first product on $J^{hcz}_{mhcz,2}$ of the second product. We re-run this regression 1000 times, where we re-draw the two products on every run. Our estimated average coefficient on $J^{hcz}_{mhcz,2}$ is 0.234 (bootstrap s.e. $= 0.001$) with a constant of 0.214 (s.e. $= 0.001$), indicating that the probability of the second product being "Japanese" approximately doubles when the first one is. Thus, firms' sourcing strategy is not independent across products. Instead, firms that rely on "Japanese" sourcing for some products are more likely to use "Japanese" procurement for other products as well.

5 The Effect of PNTR on the Choice of Procurement System

A key insight from the model presented in Section 2 is that trade policy can alter buyers' choice of procurement system by affecting the probability of trade wars. In this section, we examine the prediction that a decrease in the probability of a trade war can induce buyers to shift from "American" to "Japanese" procurement using a plausibly exogenous change in U.S. trade policy, the U.S. granting of Permanent Normal Trade Relations (PNTR) to China in 2001. We first provide a brief description of PNTR and how exposure to it can be measured, and then examine how exposure affects procurement patterns using a difference-in-differences framework.

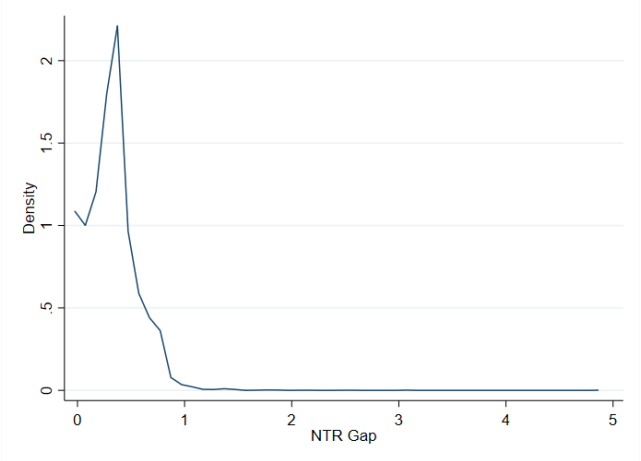
5.1 Description of PNTR

U.S. imports from non-market economies such as China are generally subject to relatively high "column two" tariff rates originally set under the Smoot-Hawley Tariff Act of 1930, as opposed to the generally low Normal Trade Relations (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 bilateral relations. Annual approval became controversial and less certain after the Tiananmen Square incident in 1989, and this uncertainty continued throughout the 1990s. During this time, firms engaged in or considering U.S.-China trade faced the

[36] As we will show in Section 7 below, these findings are qualitatively consistent with our model, where we find that larger importers are more likely to use the "Japanese" system.

[37] Note that we only observe the overall inventories of the firm, across all products and including domestic purchases. Our results suggest that variation in international procurement is associated with tangible differences in firms' overall inventories.

Figure 4: Distribution of the *NTR Gap*



Source: Feenstra, Romalis, and Schott (2002) and authors' 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.

possibility, each year, of substantial tariff increases if China's NTR status was not re-approved. The magnitude of these potential tariff increases—32 percentage points for the average product—make clear that some buyer-seller relationships that were profitable under NTR tariff rates would not be profitable under a shift to "column two" tariffs. Indeed, Pierce and Schott (2016) document extensive discussion by U.S. firms of the trade-dampening effects of this uncertainty in the 1990s, and Handley and Limão (2017) provide a theoretical basis for these effects that operates via suppressed entry by Chinese exporters.[38] Alessandria, Khan, and Khederlarian (2019) show that uncertainty regarding the annual renewal of China's NTR status each summer reduced U.S. imports from China, while also driving intra-year seasonal patterns in imports. When the United States granted PNTR to China in 2001, it locked in NTR rates, eliminating the need for annual renewals and the potential for relationship-severing tariff increases. This plausibly exogenous policy change provides a useful opportunity for testing Proposition 2.1, i.e., whether a decrease in the probability of a trade war leads to the adoption of more "Japanese" sourcing.[39] We follow Pierce and Schott (2016) in defining a product's exposure to PNTR as the difference between the non-NTR rate to which its tariff could have risen before PNTR and the lower NTR rate that was locked in by the policy change,

$$NTR\,Gap_h = Non\,NTR\,Rate_h - NTR\,Rate_h. \tag{10}$$

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). As indicated in Figure 4, these gaps vary widely across products, and have a mean and standard deviation of 0.32 and 0.23, respectively.

[38] Handley and Limão (2017) also estimate that the reduction in uncertainty associated with PNTR's ultimate approval was equivalent to a 13 percentage point permanent reduction in tariff rates.

[39] See also Blanchard, Bown, and Johnson (2016), who examine how the presence of global value chains can affect the longer-term endogenous determination of tariff rates as part of multilateral trade negotiations.

5.2 PNTR and Procurement Attributes Among Continuing $mxhcz$ Quintuples

Our first approach to testing whether PNTR influences procurement is to examine its impact on the procurement attributes examined in Section 4: quantity per shipment, weeks between shipments and unit value. These attributes are observed at the buyer-seller $mxhcz$ quintuple level. We therefore analyze procurement attributes among *continuing* quintuples, which trade in both the pre- and the post-PNTR period, in this subsection, and for new quintuples in the next subsection.

Our OLS triple difference-in-differences (DID) identification strategy examines the relationship between PNTR and the procurement attributes before versus after the change in policy (first difference), for imports from China versus other source countries (second difference), for products with higher versus lower NTR gaps (third difference),

$$\begin{aligned}\ln(Y_{mxhczt}) =& \beta_1 1\{t = Post\} * 1\{c = China\} * NTR\,Gap_h + \beta_2 ln(QPW_{mxhczt}) + \\ & \beta_3 \chi_{mxhczt} + \lambda_{mxhcz} + \lambda_t + \epsilon_{mxhczt}.\end{aligned} \tag{11}$$

The last difference captures the fact that products with larger NTR gaps experience a greater decline in the relationship termination probability, which is a function of the change in China's NTR status (identical for all products) and the increase in tariff rates that could have occurred before PNTR, which varies by product. We expect the largest shift toward "Japanese"procurement after PNTR to occur in U.S. imports of high-NTR-gap products from China.

The variable Y_{mxhczt} on the left-hand side of specification (11) represents one of the three procurement attributes: quantity per shipment (QPS_{mxhczt}), weeks between shipments (WBS_{mxhczt}), and unit value (UV_{mxhczt}).[40] The first term on the right-hand side is the triple difference-in-differences (DID) term of interest, an interaction of an indicator for the post period, $1\{t = Post\}$, a dummy for imports from China, $1\{c = China\}$, and the $NTR\ Gap_h$. The variable χ_{mxhczt} 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_{mxhczt}) as well as buyer-seller quintuple (λ_{mxhcz}) and period (λ_t) fixed effects. Our two five-year periods (t), 1995 to 2000 and 2002 to 2007, are chosen to straddle the change in policy in 2001 and end before the Great Recession.[41] Standard errors are two-way clustered at the country and product level.

Conducted at the $mxhcz$ level, equation (11) is restricted to continuing buyer-seller relationships via the $mxhcz$ quintuple fixed effect. We restrict the sample to quintuples that transact at least twice both before PNTR and after the policy change so that weeks between shipments (WBS_{mxhczt}) can be computed.

Results, reported in Table 9, indicate that higher exposure to PNTR is associated with changes in shipping attributes that are consistent with a move toward Japanese-style procurement within existing buyer-seller quintuples. Coefficient estimates in the first two columns show that a one standard deviation increase in the $NTR\ Gap$ (0.23) induces a relative decline in quantity per shipment and

[40] Appendix C.2 provides more details on how the variables in this section are constructed.

[41] In Appendix Section E.1, we demonstrate that all results in this section are robust to using a different post-PNTR period, 2004-2009.

weeks between shipments of 4.5 log points and 3.9 log points respectively. Moreover, results in column 3 reveal that a one standard deviation increase in exposure to PNTR is associated with a relative increase in unit value of 2.1 log points. In each case, the findings in Table 9 are consistent with the predictions of Propositions 2.1 and 2.3, indicating a switch from "American" to "Japanese" procurement, as opposed to an adjustment within the "Japanese" system.[42]

Table 9: Baseline Within *mxhcz* Quintuple PNTR DID Regression

	(1)	(2)	(3)
Dep. var.	$\ln(QPS_{mxhczt})$	$\ln(WBS_{mxhczt})$	$\ln(UV_{mxhczt})$
$Post_t * China_c * NTRGap_h$	-0.197***	-0.168***	0.092***
	0.009	0.009	0.023
$ln(QPW_{mxhczt})$	0.368***	-0.632***	-0.124***
	0.009	0.008	0.013
Observations	439,000	439,000	439,000
R-squared	0.982	0.894	0.985
Fixed effects	$mxhcz, t$	$mxhcz, t$	$mxhcz, t$
Controls	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_{mxhczt}, WBS_{mxhczt}, and UV_{mxhczt} are average quantity per shipment, average weeks between shipment, and average unit value (i.e. value divided by quantity) in period t. All regressions include $mxhcz$ and period t fixed effects, control for the beginning and end week of the quadruple as well as all variables needed to identify the DID term of interest. Standard errors, adjusted for clustering by country (c) and product (h), are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

5.3 PNTR and Procurement Attributes Among New *mxhcz* Quintuples

In this section, we compare the procurement attributes of *new* buyer-seller *mxhcz* quintuples formed in the post-PNTR period to relationships that were new in the pre-PNTR period. For both periods, we define new quintuples as those involving buyer-seller *mx* pairs that had not yet appeared before the beginning of the period, i.e., from 1992 to 1994 for the first period and from 1992 to 2001 for the second period.

As in the previous section, the regression is performed at the *mxhcz* level and standard errors are two-way clustered at the country and product level. Instead of *mxhcz* quintuple fixed effects, however, we include separate buyer quadruple ($mhcz$), exporter (x), and period (t) fixed effects, thereby focusing on buyers and sellers that exist in both time periods (with at least one trading partner), but who form new relationships across the time periods.[43]

[42] Consistent with Proposition 6.1, 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. We show in Appendix E.2 that our conclusions are qualitatively unchanged, though the coefficient on WBS_{mxhcz} is not statistically significant, when we do not include QPW_{mxhczt} as a covariate. In Appendix E.3, we analyze the effect of PNTR on the three procurement attributes at the *mhcz* quadruple level and find similar results.

[43] As noted in Appendix E.4, results are robust to including both continuing and new *mxhcz* buyer-seller quintuples simultaneously in one regression.

Results, reported in Table 10, are consistent with relatively greater entry of "Japanese" relationships after PNTR: buyer-seller quintuples trading goods with greater exposure to the change in policy formed after it was implemented exhibit relatively smaller and more frequent shipments, at relatively higher prices, than quintuples formed before PNTR. Point estimates indicate that a one standard deviation increase in exposure is associated with a 2.7 log points and 2.2 log point decline in shipment size and frequency, respectively, and a 2.1 log point rise in price.

Table 10: New $mxhcz$ Quintuple PNTR DID Regression

	(1)	(2)	(3)
Dep. var.	$\ln(QPS_{mxhczt})$	$\ln(WBS_{mxhczt})$	$\ln(UV_{mxhczt})$
$Post_t * China_c * NTRGap_h$	-0.116***	-0.097***	0.090**
	0.023	0.023	0.038
$ln(QPW_{mxhczt})$	0.409***	-0.594***	-0.129***
	0.013	0.012	0.018
Observations	3,184,000	3,184,000	3,184,000
R-squared	0.966	0.842	0.972
Fixed effects	$mhcz, x, t$	$mhcz, x, t$	$mhcz, x, t$
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_{mxhczt}), (WBS_{mxhczt}), and (UV_{mxhczt}) are average quantity per shipment, average weeks between shipment, and average unit value (i.e. value divided by quantity) in period t. All regressions include $mhcz$, x and period t fixed effects, control for the beginning and end week of the quadruple as well as all variables needed to identify the DID term of interest. Standard errors, adjusted for clustering by country (c) and product (h), are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

5.4 PNTR and Sellers per Shipment (SPS)

The previous two sections demonstrate that higher exposure to PNTR is associated with relatively more "Japanese" procurement attributes after the policy change. In this section, we focus on the impact of PNTR on buyers' sellers per shipment, the metric for identifying "Japanese" relationships introduced in Section 4. We consider both the continuous measure SPS_{mhcz} as well as an indicator for whether this ratio falls into the first quartile of pre-PNTR period distribution within product by country by mode bins, $J^{hcz}_{mhcz} = 1$.[44]

Our triple DID specification is similar to equation (11), but takes place at the buyer $mhcz$ quadruple level,

$$ln(Y_{mhczt}) = \beta_1 1\{t = Post\} * 1\{c = China\} * NTRGap_h + \beta_2 ln(QPW_{mhczt}) + \beta_3 \chi_{mhczt} + \lambda_{mhcz} + \lambda_t + \epsilon_{mhczt}. \quad (12)$$

The triple DID term of interest is the same as above, an interaction of post-period and China-import

[44] As before, we define the cutoffs of the distribution in the first time period and keep them fixed for the second period.

dummies with the NTR gap, and the variable χ_{mhczt} 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_{mhczt}) as well as buyer quadruple (λ_{mhcz}) and period (λ_t) fixed effects. Once again, standard errors are two-way clustered at the country and product level. Conducted at the *mhcz* level, equation (12) is restricted to continuing importers—i.e. those active before and after granting of PNTR—via the *mhcz* buyer quadruple fixed effect.

While our model requires repeated interactions between buyers and sellers, it does not mandate relationships be long-established. Moreover, existing research finds substantial relative growth in U.S.-importer-Chinese-exporter relationships after PNTR (Pierce and Schott, 2016). As a result, we also estimate equation (12) aggregated at the more aggregated *hcz* level, which broadens the analysis to include entering importers. For this regression, SPS_{hcz} is defined as the average SPS_{mhcz} within *hcz* cells. The "Japanese" indicator J^{hcz}_{hcz} is defined analogously.

Results in Table 11 indicate that PNTR induced a shift towards more "Japanese" importing, and that this effect is influenced by the entry of new importers. As shown in the first two columns, continuing buyer quadruples more exposed to the change in policy exhibit relative increases in the probability of being in the first SPS quartile ($J^{hcz}_{mhczt} = 1$) after PNTR, with no relationship for SPS_{mhczt}. When we re-estimate equation (12) at the *hcz* level—which accounts for the role of new importers that enter in the post-PNTR period—we find that higher exposure to PNTR is associated not only with a higher probability of ($J^{hcz}_{mhczt} = 1$), but also with a precisely estimated reduction in SPS_{hczt}, as expected. The estimates in columns 3 and 4 indicate that a one standard deviation increase in exposure to PNTR is associated with a 0.5 percent relative decline in SPS and a 1.3 percent increase in falling within the first quartile of the SPS distribution.[45]

[45]To analyze the influence of initial buyer experimentation during the years immediately after PNTR on our results, we also consider, in Appendix E.1, similarly constructed outcomes but for a slightly later — 2004 to 2009 — post-PNTR time period. Coefficient estimates for this alternate post period have the same sign patterns, but are larger in absolute magnitude and are more precisely estimated, suggesting adjustment to PNTR may have occurred gradually.

Table 11: Within-Importer PNTR Regression, Buyer Characteristics

	(1)	(2)	(3)	(4)
Dep. var.	$\ln(SPS_{mhczt})$	$1\{J^{hcz}_{mhczt}=1\}$	$\ln(SPS_{hczt})$	$1\{J^{hcz}_{hczt}=1\}$
$Post_t * China_c * NTRGap_h$	-0.006	0.041*	-0.021**	0.034*
	0.031	0.022	0.009	0.019
$ln(QPW_{mhczt})$	-0.171***	0.124***	-0.062***	0.032***
	0.006	0.005	0.002	0.003
Observations	738,000	291,000	368,000	28,500
R-squared	0.772	0.675	0.695	0.547
Fixed effects	$mhcz, t$	$mhcz, t$	hcz, t	hcz, t
Controls	Yes	Yes	Yes	Yes

Source: LFTTD and authors' calculations. First two columns report 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. Second two columns are analogous but at the hcz level of aggregation. Pre- and post-PNTR periods are 1995 to 2000 and 2002 to 2007. All regressions include period t fixed effects, control for the beginning and end week of the quadruple as well as all variables needed to identify the DID term of interest. Columns 2 and 4 exclude quadruples with less than five shipments in both periods. Standard errors, adjusted for clustering by country (c) and product (h), are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

Overall, the results in this section provide support for the model's prediction that a lower likelihood of a trade war can bring buyers to switch from "American" to "Japanese" procurement, in terms of the procurement attributes with importers' partners, the formation of new "Japanese" relationships, and the number of seller partners.

6 Multi-Country Framework with Endogenous Demand

Our examination of PNTR indicates that changes in the probability of a trade war can cause firms to switch procurement strategies even if tariff levels remain unchanged. Such reactions are of particular relevance given recent efforts to reverse globalization, such as "Brexit" and the Trump trade wars, that have increased trade policy uncertainty worldwide. In this section, we embed the partial equilibrium model introduced in Section 2 within the multi-country, multi-product general equilibrium Eaton and Kortum (2002) model of trade. We use numerical solutions of the resulting framework to assess the potential welfare implications of shutting down "Japanese" procurement in Section 8.

Our main point of departure from the standard Eaton and Kortum (2002) framework is the introduction of homogeneous *buyer* firms in each country, which purchase manufacturing goods from *seller* firms and distribute these goods to consumers. Buyer and seller firms are subject to the procurement problem described in Section 2. Our model generates variation in procurement system within country-products via variation in sellers' productivity, inspection costs, and trade war probabilities.

6.1 Environment

Households. Our modeling here is standard. There are N countries, which we index by n and i. Each country is populated by L_n consumers, who purchase a continuous flow of a manufactured

composite good and a homogeneous "outside" good 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 good. The composite good is a CES aggregate of a continuum of differentiated products indexed by $\omega \in [0,1]$,

$$Q_n = \left(\int_0^1 q_n(\omega)^{(\sigma-1)/\sigma} d\omega \right)^{\sigma/(\sigma-1)}, \tag{13}$$

where $\sigma > 0$ is the elasticity of substitution and $q_n(\omega)$ is quantity. This aggregator implies a price index given by

$$P_n = \left(\int p_n(\omega)^{1-\sigma} d\omega \right)^{1/(1-\sigma)}. \tag{14}$$

We assume that each consumer supplies one unit of labor.

Homogeneous good. The homogeneous "outside" good in country n is produced by a representative firm according to $Z_n = a_n L_n^O$, where a_n is country n's productivity in the homogeneous good sector and L_n^O is the aggregate labor used in the production of the good. Labor is paid the wage rate w_n in country n. The homogeneous good is directly sold to households and can be costlessly traded across countries. We therefore set its price as the numeraire and normalize it to one. Labor is perfectly mobile between manufacturing and the homogeneous good sector, as in Eaton and Kortum (2002), Handley and Limão (2017), and Antràs, Fort, and Tintelnot (2017), and the homogeneous good sector is sufficiently large so that at least some of its output is produced.

Manufacturing sellers. Manufactured good ω can be produced by homogeneous seller firms in country n with the linear production function $q = \frac{\Upsilon}{\theta} l$, where l is the quantity of labor used by the seller and $\theta \in \{\underline{\theta}, \bar{\theta}\}$ is quality. As in the standard Eaton-Kortum model, sellers are perfect competitors, taking prices as given. The productivity $\Upsilon_n(\omega)$ is specific to each origin country-product combination and drawn from a distribution $F_n(\Upsilon)$. Sellers in country n also incur fixed logistic and transportation costs f_n in units of seller country labor for each destination that is supplied.[46] We assume that a country's firms are owned by their households.

Manufacturing buyers. We add a continuum of homogeneous buyer firms in each country into the standard framework. Buyer firms purchase manufactured goods from sellers domestically or abroad, and offer them to the households in their country at prices $p_n(\omega)$. The transactions between buyers and sellers take place as described in Section 2: given household demand $q_n(\omega)$, buyers in country n choose the lowest-cost sourcing country i for product ω, whether to source the good using "American" or "Japanese" procurement, and the optimal order size. Buyers using the "American" system need to use an additional $m_n(\omega)$ labor units to inspect the quality of the good, where $m_n(\omega)$ is drawn from a destination country-specific distribution $G_n(m)$. Buyers choosing the "Japanese" system pay an incentive premium to ensure quality. The incentive payments imply that sellers obtain positive

[46] Hence, the fixed costs are $f_n w_n$ for shipments from country n.

profits under the "Japanese" system.[47] As discussed in Section 2, due to the fixed procurement cost each buyer optimally places each order with a single seller.

6.2 Partial Equilibrium with Endogenous q in a Contestable Market

In this section, we describe how the quantity demanded $q_n(\omega)$ is determined in equilibrium in our model, assuming that the buyer has already chosen the source country and procurement system. In the next section, we discuss how the source country and the procurement system are chosen in general equilibrium. As we are focusing on a single market in this section, we omit country and product subscripts.

Our first step to determine the equilibrium, Proposition 6.1, establishes that batch size and shipping frequency increase with quantity ordered, q:

Proposition 6.1. *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. □

Intuitively, for a given fixed shipping frequency, buyers must increase the batch size x_s in both systems to meet an increase in q. But by the first-order condition, equation (6), we know that buyers trade-off variable procurement costs against fixed per-shipment costs. Therefore, as variable procurement costs increase, buyers respond by spreading the larger quantities over more shipments, raising shipping frequency. As a result, larger quantities purchased lead to both greater shipment sizes, x_s, but also greater order frequencies. It follows that unit values in both procurement systems decrease, since fixed per-shipment costs are spread over greater per-shipment quantities. Additionally, in the "Japanese" system, an increase in the shipping frequency implies a lower premium to motivate desired quality, which lowers the unit value further.

The comparative statics with respect to q are supported by the empirical estimates in Sections 4 and 5. As indicated in Tables 6, 9, and 10, we find that shipment size is positively related to the quantity shipped per week (QPW), and that both shipping frequency and unit values are negatively related to QPW.

We next show that greater shipment sizes x_s^* at larger total order flow q imply that buyers' average cost curves are downward sloping:

Lemma 6.2. *At the optimal order size x_s^*, both procurement systems provide economies of scale, i.e., $\frac{\partial AC(x_s^*,q)}{\partial q} < 0$. Moreover, 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. □

Downward sloping average cost curves are a key departure of our model from standard trade models based on Eaton and Kortum (2002), which generally assume constant marginal cost. In

[47] These profits are not competed away under perfect competition since sellers offering a lower price would violate the incentive constraint and would therefore not be able to supply the good under the "Japanese" system.

our model, sellers face the standard constant marginal costs and perfect competition, but the fixed logistic and transport costs generate a natural monopoly for buyers in the downstream market. We therefore need to choose an appropriate market structure. Our assumption is that buyers compete in a "contestable" market for consumers. Introduced by Baumol, Panzar, and Willig (1982) and Tirole (1988), contestable markets are a natural extension of Bertrand competition when firms' costs exhibit economies of scale. In a contestable market, there exist several homogeneous competitors whose entry is costless. In our setup, due to economies of scale it must therefore be the case in equilibrium that a single buyer serves the entire consumer market for each product.

Lemma 6.2 indicates that average cost curves are convex, and therefore a demand curve that uniquely intersects the single buyer's optimized average cost curve from above determines a unique, sustainable and feasible equilibrium in the product market, q^*. The buyer prices and supplies the final consumers along its average cost curve. If the buyer were to price above average costs, entrants would contest the positive profits, pushing the buyer out and taking over the market. If the buyer were to price below average costs, she would realize negative profits. Since for any $q < q^*$ consumers are willing to pay prices greater than average costs, potential entry forces an incumbent offering q to lower its prices and to increase quantity to the level q^* where the market clears and supply equals demand. The buyer is not willing to procure a greater quantity because she would incur losses.

Under appropriate assumptions on the demand system, the market equilibrium is a corollary of Lemma 6.2.[48]

Corollary 6.2.1. *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.*

6.3 General Equilibrium with Endogenous q

We now embed the product market equilibrium into the equilibrium of the overall economy. Equilibrium requires that buyer firms minimize costs such that the contestable market equilibrium is feasible and sustainable in each product-destination country market, the household maximizes the CES objective, and the goods and labor markets clear.

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

$$AC_n(q_n(\omega))^* = \min\left\{\min\left\{AC_{ni,A}(x^*_{ni,A}(\omega), q_n(\omega)), AC_{ni,J}(x^*_{ni,J}(\omega), q_n(\omega))\right\};\ i = 1, ..., N\right\}, \quad (15)$$

where $x^*_{ni,s}(\omega)$ is the optimal batch size for shipments from i to n under procurement system s

[48] In principle our CES demand system may intersect the downward sloping average cost curve multiple times. For equilibrium to exist in that case, the demand curve must cut the average cost curve from above at the intersection that determines the greatest equilibrium quantity, q^*_{high}. Intuitively, if the demand curve were to cut from below, it would be above the average cost curve for all $q^*_{high} < q \to \infty$, implying that consumers are willing to buy an infinite quantity of the good when the buyer sets price equal to average costs.

determined by the first-order condition (6). As described in the preceding section, since average costs are downward sloping in q and the market is contestable, in equilibrium there is only one buyer firm serving each market. This firm procures under one system from the seller country 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 total income on manufactured goods. Consumption of each manufactured good is chosen to maximize (13) subject to the budget constraint

$$\int_0^1 p_n(\omega)q_n(\omega)\,d\omega \leq \alpha\left(w_n L_n + \sum_i \sum_s \int \pi^s_{in,s}(\omega) I_{in,s}(\omega) d\omega\right), \tag{16}$$

where $I_{in,s}(\omega)$ is an indicator function that is equal to one if the buyer in country i procures product ω from country n under system s, and zero otherwise. The right-hand side of the equation is the share of country n's total income, W_n, spent on manufacturing. Given our assumption of perfect labor mobility between sectors, the wage rate is pinned down by the productivity of the homogeneous good sector as $w_n = a_n$. The second term on the right-hand side, which is new relative to the standard framework, represents the incentive premia collected from shipments to countries i under $s = J$, where $\pi^s_{ni,s}(\omega)$ is the continuous flow of profits to sellers in country i from sales to country n of product ω under system s. This term is zero if shipments are under the "American" system. Thus, all else equal, households' income in country n rises with the number of products that are shipped under the "Japanese" system.

Consumption of the homogeneous good satisfies $Z_n = (1-\alpha)W_n$. Markets clear for each manufactured good ω and for the homogeneous good, and the labor market clears in each country. We provide these market clearing conditions in Appendix F.[49]

7 Quantitative Analysis

In this section, we estimate the model quantitatively before using it in Section 8 to analyze the effects of changes in trade policy on trade flows and aggregate welfare. This analysis, therefore, underscores the relevance and implications of the model in considering highly topical changes and potential changes in the international trading environment.

We parametrize the model using a combination of external calibration and within-model moment matching. Since our model does not admit an analytical solution due to the non-linearity of the buyer's problem, we use an iterative procedure to find the equilibrium for a given set of parameters. 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

[49] In the quantitative simulations, we verify that a positive amount of labor is allocated to both manufacturing and the production of the homogeneous good in each country in equilibrium.

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 G provides further details on our equilibrium solution algorithm.

7.1 Parametrization

We set $N = 3$ countries and interpret these countries to be the United States, China, and the Rest of the World (RoW).[50] As in Eaton and Kortum (2002), the productivity $\Upsilon_n(\omega)$ is drawn from a Fréchet distribution according to

$$F_n(\Upsilon) = e^{-\Lambda_n \Upsilon^{-\zeta}}, \tag{17}$$

where the country-specific parameter Λ_n scales the mean of the distribution and ζ scales the variation. The productivity draws are independent across products within each country.

We assume inspection costs for domestic procurement to be zero, implying that all domestic sourcing takes place under the "American" system.[51] For imports, we assume that the distribution of inspection costs, $G_n(m)$, is Pareto and given by

$$G_n(m) = 1 - \left(\frac{\underline{m}}{m}\right)^{\gamma_n}, \tag{18}$$

where $\underline{m}$ is the lower bound of inspection costs, and γ_n is a parameter to be estimated.[52] We set the lower bound of the inspection cost distribution to $\underline{m} = 0.001$, as inspection of certain commodities is virtually costless. We set each time period to one quarter.

7.2 Calibrated Parameters

We follow Caliendo, Dvorkin, and Parro (2019) in setting the quarterly interest rate, r, to be 0.01, and use the estimate of the elasticity of substitution in Antràs, Fort, and Tintelnot (2017), $\sigma = 3.85$. We normalize $\underline{\theta} = 0$ and choose $\bar{\theta} = 1$ for the cost of high quality. This parameter only appears via the ratio $\bar{\theta}/\Upsilon_n(\omega)$, and we estimate the mean of the distribution of Υ_n 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 from Duarte (2020). We follow Eaton and Kortum (2002) in fixing the dispersion of productivity to be the same in each country and equal to $\zeta = 3.6$.

The productivity of the homogeneous good sector in each country, a_n, is directly linked to the country's wage level, since $w_n = a_n$ under perfect labor mobility. We therefore pin down each country's productivity in the homogeneous good sector using its wage level. Since consistent wage

[50] While our model generalizes to an arbitrary number of countries, for our purposes three are sufficient.

[51] This is a normalization. Since we do not have domestic transactions data, we cannot estimate the share of "Japanese" and "American" trade for within-country transactions. Equivalently, we could also refer to "domestic" sourcing as a third type of procurement system that does not face an incentive problem and hence corresponds to the first-best outcome. In our quantitative exercises below, sourcing choices will depend only on the relative cost of foreign versus domestic sourcing.

[52] We perform an alternative estimation below where we assume that inspection costs are distributed according to a Fréchet distribution instead of Pareto. We found that the model fit is better under Pareto, and therefore choose it as our baseline. We also assume that a given destination country has the same distribution of inspection costs for all origins to reduce the number of parameters to be estimated. We show below that the model fits the data quite well despite this restriction.

data are difficult to obtain, we estimate each country's wage to be two-thirds of their GDP divided by the size of their labor force, using data reported by the World Bank World Development Indicators (WDI) in 2016. We compute the average wage of the rest of the world as a weighted average over the wages of the U.S.' top-ten trading partners (listed in Table 3) excluding China, using U.S. imports in 2016 as weights. We set the U.S. productivity of the homogeneous good sector to $a_{US} = 1$ and set the other countries' productivity relative to that.[53] This procedure yields $a_{CN} = 0.12$ and $a_{RoW} = 0.5$.

We set the U.S. labor force to $L_{US} = 1$, and set labor for China so that we match the size of the labor force in 2016 from the WDI, yielding $L_{CN} = 5$. We set the labor force for the rest of the world as the combined workforce of the United States' ten largest trading partners excluding China, which results in $L_{RoW} = 2.5$.

We assume that the trade war shocks are symmetric between country pairs, $\rho_{ni} = \rho_{in}$, and that trade wars between the United States and the RoW and China and RoW are unlikely in steady state by setting $\rho_{US,RoW} = 0$ and $\rho_{CN,RoW} = 0$.[54] To calibrate the probability of a trade war between the United States and China, we compute the difference in the break-up rates of U.S. firms' relationships with Chinese suppliers versus suppliers in RoW. We interpret a higher likelihood of break-ups with Chinese suppliers as arising due to the additional uncertainty of trading with China. Specifically, we compute break-up probabilities as follows. First, we take all buyer-seller ($mxhcz$) quintuples whose associated $mhcz$ is in the first quartile of the within-product-country-mode SPS distribution over the full sample period, i.e. those quintuples that are most likely to be associated with "Japanese" procurement. Next, we compute the separation hazard rate of these quintuples in each quarter, separately for U.S.-China and for U.S.-RoW trade. 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 thus obtain $\psi_{US,RoW} = 0.0873$ and $\psi_{US,CN} = 0.1137$, leading to $\rho_{US,CN} = 0.0264$. 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 our assumption of zero trade wars with these countries. Our hazard rates translate into a 37 percent probability of a relationship break-up with a Chinese supplier over four quarters compared to 29 percent with a supplier from RoW. Appendix H provides more information on how we calibrate the parameters, and Table 12 summarizes their values.

[53] Since both the fixed shipment costs and the inspection costs are in terms of labor, doubling all wages will double all prices in the model. We therefore have one degree of freedom and use it to set the U.S. productivity.

[54] While narrow trade disputes between the United States and RoW—such as safeguards and antidumping duties—occur often, the WTO's formal dispute settlement system was an effective deterrent to full-fledged trade war between the US and RoW during our sample period. Similarly, while trade tensions were also present between RoW and China, a variety of bilateral disagreements between the US and China meant that the risk of RoW-China trade war was substantially lower.

Table 12: Calibrated Parameters

Parameter	Value	Source
Interest rate (r)	0.01	Caliendo, Dvorkin, and Parro (2019)
Elasticity of substitution (σ)	3.85	Antràs, 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)
Homogeneous good sector productivity (a_n)		
- U.S.	1	Normalization
- China	0.12	World Bank, authors' calculations
- RoW	0.50	World Bank, authors' calculations
Labor Force (L_n)		
- 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)
Rate of trade wars, U.S.-RoW ($\rho_{US,RoW}$)	0	Assumption

Notes: Table presents the exogenously fixed parameters. Column (1) displays the parameter value, and column (2) shows its source.

7.3 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 distribution parameters γ_n. We estimate these parameters via a simulated method of moments procedure using moments observed in the LFTTD and in aggregate 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 "Moment in Data" in Table 13 summarizes the targeted moments and their values in the data. We provide more details on how each moment is constructed in Appendix H.

We normalize $T_{US} = 1$, and estimate the other two productivity parameters using the share of imports from China and from the rest of the world in U.S. domestic manufacturing sales in 2016 (rows 1 and 2 of Table 13). A lower value of T_n increases country n's productivity, which raises that country's share in U.S. domestic sales and hence imports.

We estimate the remaining four parameters using the observed shipping patterns from China and the rest of the world to the United States. A corollary of Proposition 2.2 is that, given a total quantity ordered q, higher fixed costs lead to shipments that are more spaced out under both systems. We can therefore estimate the fixed shipment costs f_{CN} and f_{RoW} by running a modified classification regression (8) with average weeks between shipments ($\overline{WBS}_{mhcz}$) as dependent variable, separately

for China and for the rest of the world,

$$ln(\overline{WBS}_{mhcz}) = \beta_0 + \beta_1 1\{WBS_{mhcz} = Q4\} + \beta_2 \ln(QPW_{mhcz}) + \beta_3 beg_{mhcz} + \beta_4 end_{mhcz} + \lambda_{hcz} + \epsilon_{mhcz}. \quad (19)$$

The regression allows us to control for the total quantity per week, QPW, to be consistent with the theory, and for time variation and fixed effects by product by country by mode to remove potentially confounding variation in shipping frequency that is unrelated to fixed costs. To isolate sourcing that is most likely under the "American" and the "Japanese" system, respectively, our regression sample includes only quadruples that fall into the first or the fourth quartile of the SPS_{mhcz} distribution. We then include a dummy, $1\{WBS_{mhcz} = Q4\}$, indicating whether WBS_{mhcz} falls into the fourth quartile, to differentiate "American" from "Japanese" sourcing. We set f_n by targeting the average shipping frequency from each country under the "American" system, $\exp(\hat{\beta}_0 + \hat{\beta}_1)$ (rows 3 and 4).[55] We target the average shipping frequency within this system to remove variation in shipping patterns that is due to different mixes of "American" versus "Japanese" sourcing across countries. Since our target includes the intercept β_0, we will interpret as a higher fixed cost f_n any factor that lengthens the time between shipments and that is common across modes and products for the same country, for example distance. We do not have information on the procurement choice of buyers outside of the United States that source from U.S. firms, and therefore we assume that f_{US} is equal to the fixed cost from the rest of the world, $f_{US} = f_{RoW}$.

Table 13: Estimated Parameters and Targeted Moments

	(1)	(2)	(3)	(4)	(5)
	Parameter	Estimated Value	Moment that Primarily Identifies the Parameter	Moment in Data	Moment in Model
(1)	Productivity China (T_{CN})	15.117	Share of Chinese imports in domestic sales	0.074	0.072
(2)	Productivity RoW (T_{RoW})	2.798	Share of RoW imports in domestic sales	0.270	0.249
(3)	Fixed costs, China (f_{CN})	0.545	$\exp(\hat{\beta}_0 + \hat{\beta}_1)$ from (19) for CN	115.71	115.33
(4)	Fixed costs, RoW (f_{RoW})	0.044	$\exp(\hat{\beta}_0 + \hat{\beta}_1)$ from (19) for RoW	49.83	54.06
(5)	Dispersion of inspection costs, China (γ_{CN})	0.249	$\hat{\beta}_1$ from (19) for China	0.871	0.816
(6)			Sd of $\hat{\epsilon}$ from (19) for China	0.227	0.173
(7)	Dispersion of inspection costs, RoW (γ_{RoW})	0.105	$\hat{\beta}_1$ from (19) for RoW	0.822	0.836
(8)			Sd of $\hat{\epsilon}$ from (19) for RoW	0.219	0.205
(9)	Total objective $T(\cdot)$				0.079

Source: LFTTD and authors' calculations. Column (1) lists the parameters estimated for the model. Column (2) contains the estimated parameter values. Column (3) reports the moment targeted to identify the parameter. Column (4) presents the value of the moment in the data, and Column (5) presents the value of the moment computed in our simulated model. The last row presents the value of the function $T(\cdot)$ from (20).

[55] Note that by using the intercept of the regression we are targeting the shipping frequency at $\ln(QPW_{mhcz}) = 0$.

Regression (19) is also informative about the dispersion of the inspection cost parameters γ_{CN} and γ_{RoW}, which pin down the share of "Japanese" sourcing estimated by the model. Starting from $\gamma_n \to \infty$, at which point all inspection costs are zero and all sourcing is under the "American" system, lowering γ_n increases the number of high inspection cost draws and therefore raises the share of "Japanese" sourcing. We target two sets of moments that we obtain from regression (19) to pin down these parameters. First, a greater dispersion of inspection costs (a smaller γ_n) increases the difference in average shipping times between those quadruples that are more likely "American" and those that are more likely "Japanese". We therefore target the difference in shipping frequencies between the first and the fourth quartile of the WBS_{mhcz} distribution, given by $\hat{\beta}_1$ in specification (19) (rows 5 and 7). Second, we target the dispersion in shipping times across more "American" $mhcz$ quadruples. When γ_n is low, the inspection cost draws are more dispersed, leading to a higher variance of the shipping frequencies. We construct this moment by taking the residuals from (19) for all observations that fall into the fourth quartile of the WBS distribution, and compute the standard deviation of these residuals. Rows 6 and 8 show the estimated moments. Similar to the fixed costs, we assume that $\gamma_{US} = \gamma_{RoW}$.

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

$$\phi^* = \arg\min_{\phi \in \mathbb{F}} \sum_x T(\mathcal{M}_x(\phi), \hat{\mathcal{M}}_x) \tag{20}$$

where $T(\cdot)$ is the percentage difference between the model, $\mathcal{M}_x(\phi)$, and data, $\hat{\mathcal{M}}_x$, moments. Appendix I.1 provides more details on the estimation algorithm and outcomes. We present the estimated values of the parameters in the column labeled "Estimated Value" in Table 13, and the "Moment in Model" column shows the values of the simulated moments with these parameters.

The model provides a good fit along several dimensions. First, the model-generated shares of Chinese and RoW imports in U.S. manufacturing consumption are 7.2% and 24.9%, respectively, compared to 7.4% and 27.0% in the data. Second, the model generates shipping frequencies that are consistent with the data, of around 115 weeks for China and 50 weeks for the rest of the world. We note that these estimation targets are for $\ln(QPW) = 0$: the empirically observed number of weeks between shipments is much lower since shipping frequency increases with quantity shipped.[56] Finally, the model also generates substantial variation in shipping frequencies across goods, similar to the data. Our results slightly underestimate the dispersion of inspection costs for China (rows 5-6). Increasing the dispersion of inspection costs further would raise the average time between shipments beyond its empirical target (row 3), leading the model to strike a balance between matching these two sets of moments. A higher dispersion of inspection costs would tend to increase the share of "Japanese" sourcing we find, and hence our estimated share of "Japanese" sourcing from China should be on the lower side.

Our estimation yields larger fixed costs of production in terms of labor for China than for the rest of the world. However, Chinese wages are only about a quarter of those in the rest of the world, and

[56] In the first quartile of the WBS distribution from China the average number of weeks between shipments is 9 weeks, in the fourth quartile it is 39 weeks.

hence the fixed costs in terms of the numeraire good are only about three times higher. These higher fixed costs are an implication of the lower shipping frequency from China. It is likely that some of these fixed costs capture factors excluded in the model, such as distance, which act as a trade barrier (Hummels and Schaur, 2013). More generally, since we estimate the model using observed trade flows during our sample period, any trade barriers between countries are reflected in our estimated parameters.

7.4 Discussion

Our estimation relies upon several key assumptions. First, we assume that inspection costs vary across products. This assumption generates dispersion in the relative costs of "American" and "Japanese" procurement and thus variation in the procurement system used across goods coming from the same country.[57] The shape of the inspection cost distribution is important since it determines the welfare effects of changing the probability of a trade war. If some products have more extreme inspection costs, then removing the ability to import these products under the "Japanese" system leads to large welfare losses. We fix the shape of the distribution based on the observed variability of shipping times within quadruples after controlling for total quantity, which through the lens of our model is due to differences in inspection costs.

Second, while we interpret quadruples with a higher WBS_{mhcz} as more "American" when constructing the targets, our estimation does not assume a specific share of "American" or "Japanese" trade ex-ante. Instead, we construct the same moments from the first and the fourth quartile of the WBS_{mhcz} distribution in the model as in the data, and find the best fit of these moments. We can then read off the share of "Japanese" relationships as an estimation result. Thus, our model allows us to back out the share of "Japanese" trade using moments related to shipping frequencies. We prefer this approach to the alternative of simply setting the shares of "American" and "Japanese" sourcing exogenously in the estimation. Below, we will perform an alternative estimation with a Fréchet distribution for inspection costs, which generates significantly higher shares of "Japanese" sourcing. This estimation can be used to evaluate the effects of a higher share of "Japanese" sourcing on our counterfactuals.

Third, our model assumes away important sources of heterogeneity present in the data, such as variation in mode of transportation. To the extent possible, we try to account for such heterogeneity in our estimations, e.g., by using product-country-mode fixed effects. We also control for the quadruples' overall quantity traded per week in recognition of general equilibrium effects across estimation runs.

7.5 Model Results

The first column of Table 14 summarizes 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

[57] Without the variability in inspection costs, nearly all products from a given country would be under the same system. While the dispersion of productivity across products could in principle generate some variability in shipping systems by affecting the quantity shipped, in practice this mechanism has only a minor effect.

imply that 9 percent of imports from China are under the "Japanese" system, while 52 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, as well as the higher fixed costs for China. The structurally estimated "Japanese" shares are somewhat smaller than the empirical estimates we obtained using shipments in the first quartile of the SPS_{mhcz} distribution in Table 3 for China, but they are in the ballpark for the rest of the world.

Table 14: Comparison of Base-Model and Counterfactual Equilibria

		(1)	(2)	(3)	(4)
		Baseline Equilibrium	**Equilibrium Without Japanese Sourcing**	**Autarky**	**PNTR**
(1)	Value imported from China (%)	7.2%	7.5%	.	7.3%
(2)	- of which, "Japanese"	8.6%	.	.	14.6%
(3)	Value imported from ROW (%)	24.9%	17.9%	.	24.9%
(4)	- of which, "Japanese"	51.6%	.	.	51.6%
(5)	Value imported from U.S. (%)	67.9%	74.6%	100.0%	67.8%
(6)	Avg. inspection costs	0.4%	1.3%	.	0.3%
(7)	Avg. fixed costs (imports)	4.1%	3.2%	.	4.1%
(8)	Manufacturing price index	1.000	1.025	1.113	1.000
(9)	Utility	1.000	0.984	0.944	1.0002

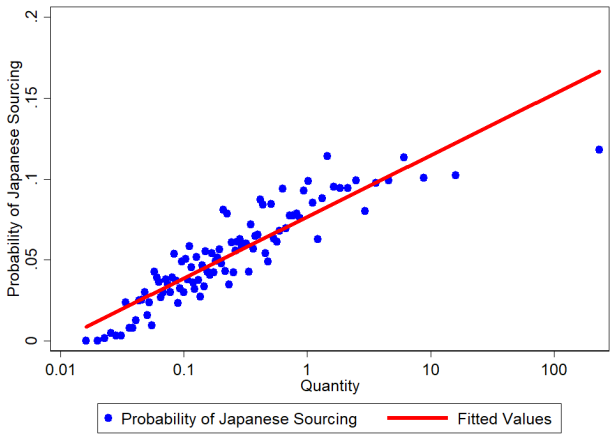
Table shows various statistics of the equilibrium under the assumption of a Pareto distribution for inspection costs. The first column presents the statistics for the baseline equilibrium, using the parameters that minimize the objective function. The second column shows the same statistics for a counterfactual economy in which the formation of "Japanese" relationships is not possible due to $\rho \to \infty$. The third column shows an autarky economy in which trade is not possible. The fourth column shows a counterfactual economy in which we reduce the arrival rate of trade wars from China to zero. Rows 1-5 show the share of U.S. manufacturing sales, $P_{US}Q_{US}$, that is from China, from the rest of the world, and from the U.S., respectively, and the share of these manufacturing sales that is sourced under the "Japanese" system. Row 6 presents the average inspection costs as a share of the import value, computed over all imports, including under the "Japanese" system. Row 7 shows the average fixed costs as a share of the import value. Row 8 shows the manufacturing price index, P_{US}, normalized to one in the baseline. Row 9 shows total utility, $W_{US} = Q_{US}^{\alpha} Z_{US}^{1-\alpha}$, normalized to one in the baseline.

As a check of the model, we verify that larger importers are more likely to use the "Japanese" system, as found in Table 8 above. We generate separately for U.S. imports from China and from the rest of the world 100 buckets corresponding to the percentiles of the quantity distribution of imports. Figure 5 plots the share of "Japanese" importers for each bucket against the average quantity imported, for China and RoW.[58] The figure shows a strong positive relationship, indicating that larger importers are more likely to use the "Japanese" system, as in the data. Intuitively, buyers import more when sellers have a higher productivity $\Upsilon_n(\omega)$, which reduces the variable costs per unit under both systems. Under the "Japanese" system a higher seller productivity additionally lowers the

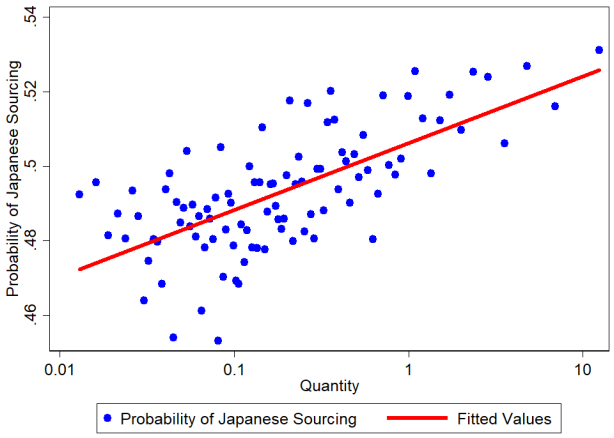
[58] We drop outliers below the 1st and above the 99th percentile of the distribution.

Figure 5: Relationship between Quantity Imported and Share of "Japanese" Importers

(a) China



(b) Rest of the World



Notes: The figure shows for each percentile of the distribution of U.S. imports the average quantity imported against the average share of importers using the "Japanese" system. The left panel presents the results for imports from China, the right panel is for imports from the rest of the world.

incentive premium, which makes "Japanese" sourcing relatively more attractive for high-productivity imports.

Rows 6 to 7 of Table 14 present our estimates for the inspection costs and fixed costs of imports. The average product imported by the U.S. is subject to an inspection cost of 0.4 percent of the import value. We estimate fixed costs of 4.1 percent of the import value. These figures provide a second validity check of the model, since our estimates 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. Our method of computing the fixed costs of shipments is similar to theirs, using the shipping frequency and a structural model.

The final two rows present the price index in manufacturing in the United States, $P_{U.S.}$, and the utility $Q_{U.S.}^{\alpha} Z_{U.S.}^{1-\alpha}$. Both of these are normalized to 1 for ease of interpretation.

In Appendix I.2, we provide a robustness check of our estimates by re-estimating the model with a Fréchet distribution for inspection costs. This estimation matches our targeted moments slightly less well than the baseline, but generates shares of imports that are under the "Japanese" system that are significantly higher. This alternative therefore allows us to examine the sensitivity of our findings to the share of "Japanese" relationships in the economy.

8 Counterfactuals

In this section, we use the calibrated model developed above to estimate U.S. welfare under two counterfactuals. First, to determine the importance of "Japanese" sourcing, we compare welfare under two extreme scenarios: elimination of "Japanese" sourcing *vis à vis* autarky. This counterfactual then provides context for estimating the welfare gain associated with PNTR. Each of these welfare analyses is highly relevant for considering the effects of recent increases in uncertainty in the global trading

system arising either from policy considerations or the possibility of periodic disruptions from Covid or other natural disasters.

No Japanese Sourcing: In our extension of Eaton and Kortum (2002), buyers' choice of country and procurement system is shaped by three factors: (i) seller productivity, (ii) country-product inspection costs, and (iii) the country-pair probabilities of trade peace. Products in which the domestic country has a higher productivity than foreign countries are sourced domestically, since, by assumption, domestic inspection costs are zero. Products not sourced domestically are imported under the "American" system from some partner if country-product inspection costs are low, and under the "Japanese" system if they are high and a trade war is sufficiently unlikely. An increased likelihood of trade war raises incentive premia under the "Japanese" system, rendering domestic production more attractive in particular for high-inspection-cost goods. As the arrival rate of trade wars goes to infinity, no goods are imported under the "Japanese" system since the incentive premium becomes prohibitive.

Our first counterfactual compares that scenario – no "Japanese" importing – to autarky, no trade at all. We shut down "Japanese" importing by setting $\rho_{US,n} = \infty$ for trade between the U.S. and both of its trading partners. This exercise entails a relatively larger increase in trade costs for RoW than for China since in the baseline $\rho_{US,CN} > \rho_{US,RoW}$. As shown in the second column of Table 14, U.S. imports rise slightly from 7.2 to 7.5 percent for China (row 1) due to the switch of trade from China to RoW under the "American" system, while imports fall significantly from 25 to 18 percent from RoW (row 3), leading to an increase in domestic sourcing (row 5). Average inspection costs jump from 0.4 to 1.3 percent (row 6) because the imports which switch from "Japanese" to "American" are precisely those with relatively high inspection costs, for which "American" sourcing was previously not optimal. Similarly, higher import prices drive down the average fixed cost as a share of import value (row 7). The manufacturing price index P_{US} rises 2.5 percent (row 8) due to higher sourcing costs associated with switching high-inspection-cost imported goods from the "Japanese" to the "American" system. Overall, welfare falls 1.6 percent. This drop in welfare is about one third as severe as moving the U.S. to autarky. As illustrated in the third column of Table 14, that counterfactual leads to a 5.6 percent drop in welfare (row 9).

When we re-estimate the model with a Fréchet distribution for inspection costs in Appendix I.2, we find that, consistent with a higher share of "Japanese" relationships in that estimation, welfare costs of removing "Japanese" relationships are significantly larger at 3.5 percent, about one half of the losses from autarky. Thus, our baseline estimates with a relatively low share of "Japanese" sourcing from China likely provide a lower bound on the welfare losses.

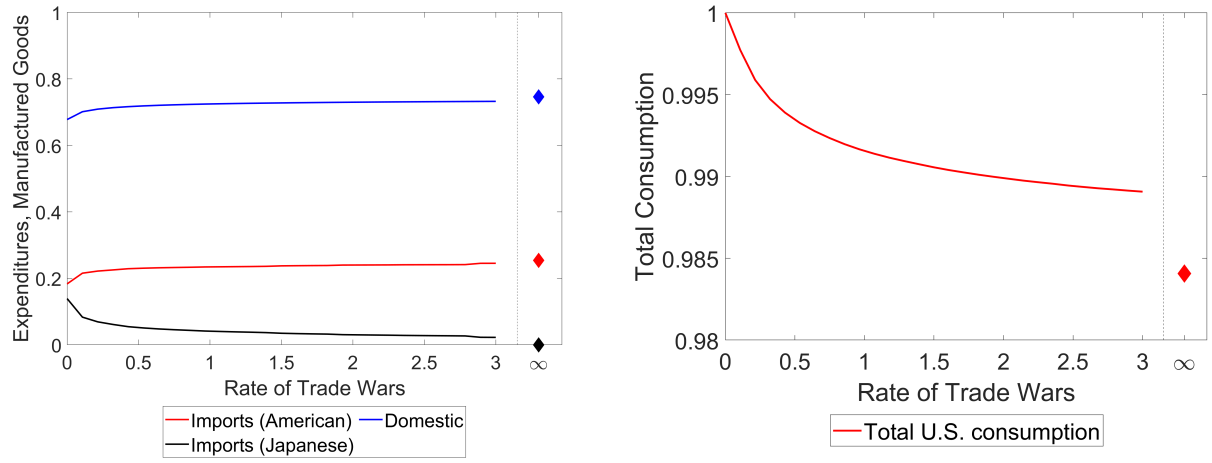
Further intuition for the welfare implications of eliminating "Japanese" sourcing is provided by Figure 6, where the left and right panels display the value of U.S. imported versus domestically manufactured goods and welfare as $\rho_{US,n}$ increases from zero to infinity, respectively.[59] As the trade war arrival rate rises, "Japanese" sourcing (black line) declines as buyers switch to "American"

[59]We set the trade war arrival rate from China and from ROW to be equal in this exercise, $\rho_{US,CN} = \rho_{US,ROW}$, to facilitate the interpretation of the figure. We normalize the total expenditures of manufactured goods at $\rho_{US,n} = 0$ to one.

Figure 6: Effect of Trade War Arrival Rate on Sourcing and Consumption

(a) Manufacturing Expenditures

(b) Utility



Notes: The left panel displays the share of expenditures on manufactured goods by the United States as a function of the arrival rate of trade wars from the rest of the world, where we distinguish imports under the "American" system (red), imports under the "Japanese" system (black), and domestic sourcing (blue). The right panel shows U.S. utility, calculated as $Q_{US}^{\alpha} Z_{US}^{1-\alpha}$, as a function of the trade war arrival rate from the rest of the world. Welfare at an arrival rate of zero is normalized to one.

sourcing (red line) for goods where the foreign productivity advantage is relatively large, and to domestic sourcing (blue line) for goods where it is relatively low. These trade responses, and the associated declines in welfare displayed in the right panel, are most dramatic at initial increases in the arrival rate of trade war. [60]

PNTR: Our second counterfactual, summarized in the fourth column of Table 14, analyzes a hypothetical "PNTR" scenario, which we simulate by reducing $\rho_{US,CN}$ to zero to match the trade war arrival rate of the rest of the world. As indicated in the table, we find that the share of "Japanese" imports from China increases by 6 percentage points, about one half of the increase in "Japanese" relationships for China in moving from column 3 to 4 of Table 3 in the empirical analysis.

The share of "Japanese" trade with China does not rise to the level of the RoW (rows 1 and 2) because of differences in other calibrated parameters, such China's higher fixed costs. The overall price and welfare effects are very minor, though the welfare impact is positive, for two reasons. First, the U.S. share of imports from China is relatively small. Second, products that switch towards "Japanese" procurement as a result of PNTR are those that were already at the margin, i.e., that had very similar costs under both systems.

Our results differ significantly from Handley and Limão (2017), who find large effects of PNTR on consumer income, for several reasons. First, in the Handley-Limão model a reduction of uncertainty leads to the entry of foreign exporters, thus expanding the set of available varieties and driving down

[60] In Appendix J, we analyze the effect of changing the arrival rate of trade wars on consumer income. We find that income is maximized for a *positive* trade war arrival rate, since exporters collect an incentive premium profit under the "Japanese" system and this premium rises with the arrival rate of trade wars. The maximum income trades off the size of the incentive premium collected by domestic exporters against the loss of exporting opportunities as the trade war arrival rate rises.

the price index, as in the Melitz (2003) model. This margin is absent in our model. Similar to Eaton and Kortum (2002), the set of varieties available to consumers is fixed in our framework: a change in trade policy uncertainty only changes the identity of the producer and the procurement system used. Second, since we estimate a relatively small quarterly probability of a trade war, eliminating this probability completely reduces the cost of "Japanese" procurement only little. Therefore, only a few marginal buyers with procurement costs that are relatively similar under the "American" and the "Japanese" system before the policy change switch systems, and the benefits to consumers are small. Generating large welfare gains through the mechanism we study would require a much larger shift in the probability of a trade war, as in the first counterfactual above where "Japanese" procurement is completely eliminated. Finally, Handley and Limão (2017) allow exporters to pay a fixed cost to reduce their marginal cost, and the set of firms that choose to pay this cost rises when uncertainty is lower, which further increases the gains from low uncertainty they find. We view our channel as complementary to the mechanisms described by Handley and Limão (2017).

9 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 war is high, 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 war is low, 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 can induce a switch between procurement systems.

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 unit value – 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.

We embed our framework for procurement in the Eaton and Kortum (2002) multi-product, multi-country general equilibrium model to assess the importance of "Japanese" sourcing in U.S. welfare. Quantitative simulations of the model reveal that an increase in the probability of trade war that is sufficient to eliminate "Japanese"-style procurement reduces U.S. welfare about one third as much as placing the U.S. in autarky.

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.[61] 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.

[61] 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

Abreu, D., D. Pearce, and E. Stacchetti (1990): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica: Journal of the Econometric Society*, pp. 1041–1063.

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.

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

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.

Antras, P., and D. Chor (2022): "Global Value Chains," in *Handbook of International Economics*, vol. 5, pp. 297–376. Elsevier.

Antràs, 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.

Bernard, A. B., and A. Moxnes (2018): "Networks and Trade," *Annual Review of Economics*, 10(1), 65–85.

Blanchard, E. J., C. P. Bown, and R. C. Johnson (2016): "Global Supply Chains and Trade Policy," Working Paper 21883, National Bureau of Economic Research.

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.

Bown, C., P. Conconi, A. Erbahar, and L. Trimarchi (2020): "Trade Protection Along Supply Chains," Discussion paper, Working Paper.

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.

Chernozhukov, V., and H. Hong (2003): "An MCMC Approach to Classical Estimation," *Journal of Econometrics*, 115(2), 293–346.

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.

Engbom, N., and C. Moser (2021): "Earnings Inequality and the Minimum Wage: Evidence from Brazil," Discussion paper, NBER Working Paper No. 28831.

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., F. Haberkorn, L. Lewis, A. Monken, J. Pierce, R. Rhodes, and M. Yi (2019): "Bill of Lading Data in International Trade Research with an Application to the COVID-19 Pandemic," *Federal Reserve Finance and Economics Discussion Paper*, 2019-086.

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.

Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin (2013): *Bayesian Data Analysis*. Chapman and Hall/CRC.

Goldberg, P., and N. Pavcnik (2016): "The Effects of Trade Policy," vol. 1 of *Handbook of Commercial Policy*, pp. 161–206. North-Holland.

Green, E. J., and R. H. Porter (1984): "Noncooperative Collusion Under Imperfect Price Information," *Econometrica*, pp. 87–100.

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

Hallak, J. C., and P. K. Schott (2011): "Estimating Cross-Country Differences in Product Quality," *Quarterly Journal of Economics*, 126(1), 417–474.

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.

——— (2022): "Trade Policy Uncertainty," *Annual Review of Economics*, 14(1), null.

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.

Hummels, D. L., and G. Schaur (2013): "Time as a Trade Barrier," *American Economic Review*, 103(7), 2935–59.

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.

Johnson, R. C. (2018): "Measuring Global Value Chains," *Annual Review of Economics*, 10(1), 207–236.

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.

Kandori, M. (2002): "Introduction to Repeated Games with Private Monitoring," *Journal of Economic Theory*, 102(1), 1–15.

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

Khandelwal, A., P. Schott, and S.-J. Wei (2013): "Trade Liberalization and Embedded Institutional Reform: Evidence from Chinese Exporter," *American Economic Review*, 103(6), 2169–95.

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.

Lise, J., C. Meghir, and J.-M. Robin (2016): "Matching, sorting and wages," *Review of Economic Dynamics*, 19, 63–87, Special Issue in Honor of Dale Mortensen.

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.

Mailath, G. J., and L. Samuelson (2006): *Repeated Games and Reputations: Long-Run Relationships*. Oxford university press.

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.

Schott, P. K. (2004): "Across-Product versus Within-Product Specialization in International Trade," *Quarterly Journal of Economics*, 119(2), 647–678.

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.

World Bank (2020): "Trading for Development in the Age of Global Value Chains," in *World Development Report*. World Bank, Washington, D.C.

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

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

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

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

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 $e^{rx_J^*/q} > 1$, and 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 $m = 0$, i.e., $x_J^* < x_A^*(m = 0)$. 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 $\frac{d}{dy}\left(\frac{y}{2} + 1\right)e^y - y - 1 = -1 + \frac{1}{2}(y + 3)e^y > 0$. Therefore $\frac{\partial FOC_J}{\partial \rho} > 0$. Then by the implicit function theorem

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

where we denote by SOC_J the second-order condition, which is greater than zero as shown in Section A.1.1.

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, $\frac{\partial x_A^*}{\partial m} > 0 \Rightarrow \frac{\partial \frac{v_A(x_A)}{x_A}}{\partial m} < 0$.

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 6.1

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 6.2: 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 Data

B.1 Data Refinement

We refine the raw US trade transactions data from the LFTTD as follows. First, we drop all transactions that are warehouse entries, so that our dataset represents imports for consumption.[62] 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 (identifiers should start with the country name). Third, given our focus on arm's-length trade, we exclude from our analysis all related-party transactions.[63] 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 HS10 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 2009 real dollars.[64] Sixth, 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. Finally, we remove outliers in terms of unit values. We compute the distribution of unit values in each HS10 by country by mode of transportation by quarter cell and drop transactions with a unit value below the 1st or above the 99th percentile.

B.2 Baseline Sample

Our baseline sample restricts our cleaned dataset to buyer quadruples with at least five transactions. In this section, we compare this sample to the arm's-length sample that does not impose this restriction. Since we cannot compute some variables such as weeks between shipments (WBS_{mhcz}) for quadruples that trade only a single time, we focus for consistency on the arm's length sample consisting of quadruples with two or more transactions.

Table A.1 presents an overview of the samples. The first row shows that the baseline sample accounts for slightly more than 80 percent of the broader sample of arm's-length trade by quadruples with at least two transactions. The next row shows that the broader sample contains almost twice as many importers, suggesting that most of these do not have substantial imports. The third row presents the number of unique exporters and the fourth row shows the number of unique quadruples. The latter rises more than twofold in the broader sample. The last row presents the number of unique quintuplets. These do not increase nearly as much in percentage terms as the number of quadruples,

[62] The term "Imports for consumption" denotes the total value of merchandise that physically clears customs or is withdrawn from customs bonded warehouses or U.S. Foreign Trade Zones; it does not refer to whether imports are purchased by consumers or firms.

[63] 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.

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

as most of the quadruples unique to the broader sample have only few suppliers.

Table A.2 compares the $mhcz$ quadruples in the two samples. The first row shows that the average value traded by a quadruple in the broader sample is only about half of the trade value in the baseline sample. Rows two to four show that quadruples in the broader sample are shorter-lived, contain fewer shipments, and source from fewer suppliers on average. However, the average value per shipment is relatively similar to the baseline sample (row 5). Shipments in the broader sample are significantly more spaced out over time (row 6). The last two rows show that the average importer-exporter relationship length associated with a quadruple in the broader sample is shorter than in the baseline sample and that quadruples in the broader sample have a higher ratio of suppliers to shipments. The latter fact suggests that many of the additional quadruples not in the baseline sample conduct their few transactions with different suppliers.

Table A.1: U.S. Import Transaction Summary Statistics

	Baseline $t \geq 5$	Sample $t \geq 2$
Total Imports ($\$Bill$)	5,680	6,990
Unique Importers (m)	360,000	637,000
Unique Exporters (x)	5,037,000	6,531,000
Unique Importer-Product-Country-Mode Quadruples ($mhcz$)	2,966,000	7,615,000
Unigue Exporter-Importer-Product-Country-Mode Quintuples ($mxchz$)	21,700,000	30,600,000

Source: LFTTD and authors' calculations. Table summarizes U.S. arm's-length imports from 1992 to 2016. Observations are based on the cleaned data described in Section B.1. The first column further restricts to our baseline sample of quadruples with at least five transactions ($t \geq 5$), analogous to Table 1. The final column restricts to the broader sample of quadruples with two or more transactions ($t \geq 2$). Import values are in billions of real 2009 dollars. The final four rows of the table 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. Observation counts are rounded to the nearest thousand per U.S. Census Bureau disclosure guidelines.

Table A.2: Attributes of $mhcz$ Quadruples

	Baseline Sample $t \geq 5$		Broader Sample $t \geq 2$	
	Mean	*Standard Deviation*	*Mean*	*Standard Deviation*
Total Value Traded ($)	1,914,000	36,300,000	918,400	24,100,000
Length Between Buyer's First and Last Shipment (Weeks)	304.3	266	187.9	229.8
Total Shipments	38.6	157.9	17.8	100.4
Number of Sellers (x)	7.3	25.5	4.0	16.2
Value per Shipment (VPS), ($)	35,910	386,100	38,090	470,500
Weeks Between Shipments (WBS)	23.5	28.5	44.5	79.8
Average Relationship Length in Weeks ($length$)	180.8	154.7	147.2	156.7
Ratio of Sellers to Shipments (SPS)	0.334	0.241	0.512	0.306

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) quadruples during our 1992 to 2016 sample period. Observations are restricted to the sample described in Section B.1. The first two columns further restrict to our baseline sample of quadruples with at least five transactions, analogous to Table 2. The final two columns restrict to the broader sample of quadruples with two or more transactions. Observation counts are rounded to the nearest thousand per U.S. Census Bureau disclosure guidelines.

C Construction of the Variables

C.1 Classification Regressions

In this section, we provide more details on the variables used in Sections 3 and 4. As discussed in the main text, we collapse all transactions of the same importer (m) - product (h) - country (c) - mode of transportation (z) quadruple in the same week into one. Therefore, a "transaction" (i) refers to a week in which the quadruple imports.

Quantity per Shipment (QPS_{mhcz}). This variable is constructed as

$$QPS_{mhcz} = \frac{\sum_i Quantity_{mhczi}}{Ntrans_{mhcz}}, \tag{A.6}$$

where $Quantity_{mhczi}$ is the quantity imported by quadruple $mhcz$ at transaction i and $Ntrans_{mhcz}$ is the total number of transactions by the quadruple.

Weeks between Shipments (WBS_{mhcz}). We construct this variable as

$$WBS_{mhcz} = \frac{end_{mhcz} - beg_{mhcz}}{Ntrans_{mhcz} - 1}, \tag{A.7}$$

where end_{mhcz} is the number of the week of the last transaction of the quadruple (see below for the construction of this variable), beg_{mhcz} is the number of the week of the first transaction of the quadruple, and $Ntrans_{mhcz}$ is the total number of transactions by the quadruple. The denominator represents the number of time gaps between subsequent transactions of the quadruple, which is one less than the number of transactions. If $Ntrans_{mhcz} = 1$, the average time gap cannot be computed.

This case never arises in our analysis since we restrict the sample to quadruples with at least five transactions.

Unit Value (UV_{mhcz}). We construct this variable as

$$UV_{mhcz} = \frac{1}{Ntrans_{mhcz}} \sum_i \frac{Value_{mhczi}}{Quantity_{mhczi}}, \tag{A.8}$$

where $Value_{mhczi}$ is the value imported by quadruple $mhcz$ at transaction i, $Quantity_{mhczi}$ is the corresponding quantity, and $Ntrans_{mhcz}$ is the total number of transactions by the quadruple.

Quantity per Week (QPW_{mhcz}). This variable is constructed as

$$QPW_{mhcz} = \frac{\sum_i Quantity_{mhczi}}{end_{mhcz} - beg_{mhcz}}, \tag{A.9}$$

where $Quantity_{mhczi}$ is the quantity imported by quadruple $mhcz$ at transaction i, end_{mhcz} is the number of the week of the last transaction of the quadruple (see below for the construction of this variable), and beg_{mhcz} is the number of the week of the first transaction of the quadruple. In contrast to QPS_{mhcz}, this variable does not divide by the number of transactions but by the "flow" of imports in an average week. We note that since we require at least five transactions in our baseline, the beginning and end week are never the same and therefore the expression is finite.

First week (beg_{mhcz}). This variable is constructed as the first week in which a $mhcz$ quadruple transacts,

$$beg_{mhcz} = \min_i \{Week_{mhczi}\}, \tag{A.10}$$

where $Week_{mhczi}$ is the week number of the transaction, relative to the first week of 1960. Thus, for example the first week of 2016 has week number 2912.

Last week (end_{mhcz}). This variable is constructed as the last week in which a $mhcz$ quadruple transacts,

$$end_{mhcz} = \max_i \{Week_{mhczi}\}, \tag{A.11}$$

where $Week_{mhczi}$ is the week number of the transaction, relative to the first week of 1960. Thus, for example the first week of 2016 has week number 2912.

Quartile of the SPS distribution ($1\{SPS_{mhcz} = Q_n\}$). We construct this variable by taking the distribution of SPS_{mhcz} within each hcz bin over the entire sample period 1992-2016 and assign each buyer quadruple to its respective quartile. The variable is an indicator that is equal to one if SPS_{mhcz} is in quartile n.

Average relationship length ($length_{mhcz}$). This variable is constructed in two steps. First, for every importer (m) - exporter (x) pair, we compute the total relationship length as

$$length_{mx} = \max_i\{Week_{mxi}\} - \min_i\{Week_{mxi}\}, \tag{A.12}$$

where $Week_{mxi}$ is the week number of a transaction i of the buyer-seller pair mx in any good or mode of transportation, relative to the first week of 1960. Thus, for example the first week of 2016 has week number 2912.

We then compute the average relationship length of a $mhcz$ quadruple as the average of these lengths across all exporters with which the buyer quadruple trades:

$$length_{mhcz} = \frac{\sum_x length_{mx}}{Sellers_{mhcz}}, \tag{A.13}$$

where $Sellers_{mhcz}$ is the number of exporters (x) with which the quadruple ($mhcz$) has an $mxhcz$ quintuple relationship.

Value per Shipment (VPS_{mhcz}). This variable is constructed as

$$VPS_{mhcz} = \frac{\sum_i Value_{mhczi}}{Ntrans_{mhcz}}, \tag{A.14}$$

where $Value_{mhczi}$ is the value imported by quadruple $mhcz$ at transaction i and $Ntrans_{mhcz}$ is the total number of transactions by the quadruple.

Dummy for differentiated goods ($Diff_h$). This variable is a dummy that is equal to one if the 6-digit HS code associated with product h is differentiated or has a reference price based on Rauch (1999) under the liberal classification.

Value per Week (VPW_{mhcz}). This variable is constructed as

$$VPW_{mhcz} = \frac{\sum_i Value_{mhczi}}{end_{mhcz} - beg_{mhcz}}, \tag{A.15}$$

where $Value_{mhczi}$ is the value imported by quadruple $mhcz$ at transaction i, end_{mhcz} is the number of the week of the last transaction of the quadruple (see above for the construction of this variable), and beg_{mhcz} is the number of the week of the first transaction of the quadruple. In contrast to VPS_{mhcz}, this variable does not divide by the number of transactions but by the "flow" of imports in an average week. We note that since we require at least five transactions in our baseline, the beginning and end week are never the same and therefore the expression is finite.

C.2 PNTR Regressions

In this section, we provide more details on the variables used in the baseline PNTR regressions in Section 5. As discussed in the main text, our data cleaning process collapses all transactions of the same importer (m) - exporter (x) - product (h) - country (c) - mode of transportation (z) quintuple

in the same week into one. Therefore, a "transaction" (i) refers to a week in which the quintuple imports.

Quantity per Shipment (QPS_{mxhczt}). This variable is constructed as

$$QPS_{mxhczt} = \frac{\sum_i Quantity_{mxhczti}}{Ntrans_{mxhczt}}, \tag{A.16}$$

where $Quantity_{mxhczti}$ is the quantity imported by quintuple $mxhcz$ in period t (either 1995-2000 or 2002-2007) at transaction i and $Ntrans_{mxhczt}$ is the total number of transactions by the quintuple in period t.

Weeks between Shipments (WBS_{mxhczt}). We construct this variable as

$$WBS_{mxhczt} = \frac{end_{mxhczt} - beg_{mxhczt}}{Ntrans_{mxhczt} - 1}, \tag{A.17}$$

where end_{mxhczt} is the number of the week of the last transaction of the quintuple in period t (either 1995-2000 or 2002-2007), beg_{mxhczt} is the number of the week of the first transaction of the quintuple, and $Ntrans_{mxhczt}$ is the total number of transactions by the quintuple in period t. The denominator represents the number of time gaps between subsequent transactions of the quintuple, which is one less than the number of transactions. If $Ntrans_{mxhczt} = 1$, the average time gap cannot be computed. The PNTR regressions therefore require for each quintuple at least two transactions in each period t.

Unit Value (UV_{mxhczt}). We construct this variable as

$$UV_{mxhczt} = \frac{1}{Ntrans_{mxhczt}} \sum_i \frac{Value_{mxhczti}}{Quantity_{mxhczti}}, \tag{A.18}$$

where $Ntrans_{mxhczt}$ is the total number of transactions of quintuple $mxhcz$ in period t (either 1995-2000 or 2002-2007), $Value_{mxhczti}$ is the value imported by quintuple $mxhczt$ at transaction i in period t, and $Quantity_{mxhczti}$ is the corresponding quantity.

Quantity per Week (QPW_{mxhczt}). This variable is constructed as

$$QPW_{mxhczt} = \frac{\sum_i Quantity_{mxhczti}}{end_{mxhczt} - beg_{mxhczt}}, \tag{A.19}$$

where $Quantity_{mxhczti}$ is the quantity imported by quintuple $mxhcz$ at transaction i in period t (either 1995-2000 or 2002-2007), end_{mxhczt} is the number of the week of the last transaction of the quintuple in period t, and beg_{mxhczt} is the number of the week of the first transaction of the quintuple in period t. In contrast to QPS_{mxhczt}, this variable does not divide by the number of transactions but by the "flow" of imports in an average week. As described above for WBS_{mxhczt}, we require for each quintuple at least two transactions in each period t so that this variable can be computed. The week number is relative to the first week of 1960. Thus, for example the first week of 2016 has week number 2912.

D Additional A vs J Classification Regressions

D.1 A versus J Classification for Thicker Relationships

Our baseline regressions in Section 4.2 are restricted to $mhcz$ quadruples with at least five transactions over our sample period. One concern might be that for quadruples that trade only relatively few times, our variable suppliers per shipment (SPS_{mhcz}) is mismeasured because we did not observe a sufficient number of transactions. In Table A.3, we show that our 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

	(1)	(2)	(3)	(4)
Dep. var.	$\log(QPS_{mhcz})$	$\log(WBS_{mhcz})$	$\log(UV_{mhcz})$	$\log(length_{mhcz})$
$\log(SPS_{mhcz})$	0.359^{***}	0.370^{***}	-0.064^{***}	-0.504^{***}
	0.015	0.016	0.020	0.013
$\log(QPW_{mhcz})$	0.700^{***}	-0.306^{***}	-0.273^{***}	-0.134^{***}
	0.014	0.014	0.019	0.005
Observations	1,645,000	1,645,000	1,645,000	1,645,000
R-squared	0.950	0.659	0.855	0.488
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 (SPS_{mhcz}) and total quantity shipped per week (QPW_{mhcz}). (QPS_{mhcz}), (WBS_{mhcz}), (P_{mhcz}), and ($length_{mhcz}$) are average quantity per shipment, average weeks between shipment, average unit value, 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 quadruple, and exclude quadruples with less than 10 shipments. Standard errors, adjusted for clustering by country (c) and product (h) are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

D.2 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 4.2 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.4 and A.5 we show that our results relating suppliers per shipment and procurement patterns are robust to defining SPS at these higher levels of aggregation, where we keep all other variables at the $mhcz$ level of the baseline.

Table A.4: A vs J Classification Regression With SPS at mhc Level

	(1)	(2)	(3)	(4)
Dep. var.	$\log(QPS_{mhcz})$	$\log(WBS_{mhcz})$	$\log(UV_{mhcz})$	$\log(length_{mhcz})$
$\log(SPS_{mhc})$	0.346^{***}	0.376^{***}	-0.083^{***}	-0.578^{***}
	0.014	0.015	0.018	0.013
$\log(QPW_{mhcz})$	0.687^{***}	-0.322^{***}	-0.279^{***}	-0.147^{***}
	0.015	0.015	0.020	0.005
Observations	2,966,000	2,966,000	2,966,000	2,966,000
R-squared	0.944	0.654	0.844	0.442
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}), (P_{mhcz}), and ($length_{mhcz}$) are average quantity per shipment, average weeks between shipment, average unit value, 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 quadruple, and exclude quadruples with less than five shipments. Standard errors, adjusted for clustering by country (c) and product (h) are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

Table A.5: A vs J Classification Regression With SPS at mh Level

	(1)	(2)	(3)	(4)
Dep. var.	$\log(QPS_{mhcz})$	$\log(WBS_{mhcz})$	$\log(UV_{mhcz})$	$\log(Length_{mhcz})$
$\log(SPS_{mh})$	0.285^{***}	0.311^{***}	-0.063^{***}	-0.483^{***}
	0.019	0.020	0.021	0.009
$\log(QPW_{mhcz})$	0.668^{***}	-0.343^{***}	-0.274^{***}	-0.115^{***}
	0.014	0.014	0.020	0.006
Observations	2,966,000	2,966,000	2,966,000	2,966,000
R-squared	0.940	0.631	0.844	0.379
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}), (P_{mhcz}), and ($Length_{mhcz}$) are average quantity per shipment, average weeks between shipment, average unit value, 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 quadruple, and exclude quadruples with less than five shipments. Standard errors, adjusted for clustering by country (c) and product (h) are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

D.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 one of the two main modes of transportation, and in particular not for air shipments. Results are reported in Table A.6, 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.6: A vs J Classification Regression Across Mode of Transport

	(1)	(2)	(3)	(4)
Dep. var.	$\log(QPS_{mhcz})$	$\log(WBS_{mhcz})$	$\log(UV_{mhcz})$	$\log(length_{mhcz})$
	Vessel			
$\log(SPS_{mhcz})$	0.419***	0.451***	-0.172***	-0.570***
	0.015	0.015	0.013	0.018
$\log(QPW_{mhcz})$	0.661***	-0.347***	-0.263***	-0.177***
	0.011	0.011	0.018	0.008
Observations	1,506,000	1,506,000	1,506,000	1,506,000
R-squared	0.924	0.686	0.829	0.434
	Air			
$\log(SPS_{mhcz})$	0.410***	0.443***	-0.058**	-0.609***
	0.022	0.022	0.025	0.018
$\log(QPW_{mhcz})$	0.737***	-0.272***	-0.300***	-0.106***
	0.015	0.015	0.023	0.005
Observations	1,029,000	1,029,000	1,029,000	1,029,000
R-squared	0.933	0.635	0.764	0.416

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}), (UV_{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 quadruple, and exclude quadruples with less than five shipments. Standard errors, adjusted for clustering by country (c) and product (h), are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

D.4 Differentiated Products Versus Commodities

In this section we examine whether buyers are more likely to use "Japanese" procurement for differentiated goods. If differentiated products have higher inspection costs, then by Proposition 2.1 buyers are more likely to use "Japanese" procurement for them, which implies smaller shipment size, greater frequency, and higher unit import values than products sourced under the "American" system (Proposition 2.3). Moreover, as discussed in Section 4.3, this "Japanese" sourcing of differentiated products should be associated with fewer suppliers and longer relationships. We examine these features of the model using the commonly cited measure of product-differentiation from Rauch (1999) in the following $mhcz$-level OLS specification,

$$\overline{Y}_{mhcz} = \beta_0 + \beta_1 Diff_h + \beta_2 \ln(VPW_{mhcz}) + \beta_3 beg_{mhcz} + \beta_4 end_{mhcz} + \lambda_{cz} + \epsilon_{mhcz}. \tag{A.20}$$

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.[65] Instead, we use as our second dependent variable the average transaction value per shipment, VPS_{mhcz}, as a measure of average transaction size. Third, we consider the average relationship length ($length_{mhcz}$) as in Section 4.3. Finally, the fourth variable is a measure of the buyer's procurement type, sellers per shipment (SPS_{mhcz}) introduced in the main text. 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).[66] 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 in the main text. The sample period is 1992 to 2016, we include only buyer quadruples with at least five transactions, and standard errors are clustered at the country-product level.

Results, reported in Table A.7, are consistent with the model's predictions regarding inspection costs, while providing further support for the use of 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 average transaction size is smaller, and the average relationship length is longer. Results in the final column provide further support for this view, as buyer quadruples encompassing differentiated goods tend to have lower sellers per shipment.

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

[66] 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 similar when we use the conservative definition.

Table A.7: A vs J Classification Regression for Differentiated Goods

	(1)	(2)	(3)	(4)
Dep. var.	$\log(WBS_{mhcz})$	$\log(VPS_{mhcz})$	$\log(length_{mhcz})$	$\log(SPS_{mhcz})$
$Diff_h$	-0.234***	-0.225***	0.073**	-0.082***
	0.026	0.025	0.028	0.025
$\log(VPW_{mhcz})$	-0.464***	0.557***	-0.045***	-0.203***
	0.002	0.002	0.001	0.001
Observations	2,589,000	2,589,000	2,589,000	2,589,000
R-squared	0.611	0.730	0.193	0.278
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 quadruple, and exclude quadruples with less than five shipments. Standard errors, adjusted for clustering by country and product, are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

D.5 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,

$$Y_{mxhcz} = \beta_0 + \beta_1 SPS_{mhcz} + \beta_2 \ln(QPW_{mxhcz}) + \beta_3 beg_{mxhcz} + \beta_4 end_{mxhcz} + \lambda_{xhcz} + \epsilon_{mxhcz}. \tag{A.21}$$

In this specification, Y_{mxhcz} represents procurement attributes at the buyer-seller relationship quintuple ($mxhcz$) 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 two-way clustered at the country (c) and product (h) level.

Results, reported in Table A.8, are similar to those in Section 4.2, providing further support for Proposition 2.3, as well as the use of SPS_{mhcz}. Across U.S. buyer quadruples within foreign exporters, we find that a one percent increase in SPS_{mhcz} is associated with QPS_{mxhcz} and WBS_{mxhcz} increases

of 0.1 and 0.7 percent. This rise in SPS_{mhcz} coincides with declines of UV_{mxhcz} and $length_{mhcz}$ of 0.06 and 0.3 percent (for reference, a one standard deviation increase in SPS_{mhcz} evaluated at the mean is equal to an increase of SPS_{mhcz} of 72 percent).

Table A.8: A vs J Classification Regression Across $mxhcz$ Quintuples

	(1)	(2)	(3)	(4)
Dep. var.	$\log(QPS_{mxhcz})$	$\ln(WBS_{mxhcz})$	$\ln(UV_{mxhcz})$	$\ln(length_{mxhcz})$
$\ln(SPS_{mhcz})$	0.100^{***}	0.696^{***}	-0.062^{***}	-0.302^{***}
	0.015	0.041	0.006	0.011
$\ln(QPW_{mxhcz})$	0.511^{***}	-0.171^{***}	-0.130^{***}	-0.241^{***}
	0.010	0.009	0.011	0.008
Observations	4,783,000	4,783,000	4,783,000	4,783,000
R-squared	0.966	0.621	0.953	0.786
Fixed effects	$xhcz$	$xhcz$	$xhcz$	$xhcz$
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 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}), (P_{mxhcz}), and ($length_{mxhcz}$) 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 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 quadruples with less than five shipments. Standard errors, adjusted for clustering by country (c) and product (h) bins are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

D.6 Regressions by Sector

In this sub-section, we re-run our baseline regressions of Section 4.2 separately for different sectors. One concern with our findings could be that the results might only hold in some sectors, such as manufacturing, but not in others. We show in Tables A.9 to A.12 that our results regarding the relationship between SPS_{mhcz} and shipment attributes hold across a broad range of sectors.

Table A.9: SPS_{mhcz} and Procurement Attributes - Mining and Utilities

	(1)	(2)	(3)	(4)
Dep. var.	$\ln(QPS_{mhcz})$	$\ln(WBS_{mhcz})$	$\ln(UV_{mhcz})$	$\ln(length_{mhcz})$
$\ln(SPS_{mhcz})$	0.413^{***}	0.455^{***}	-0.106^{**}	-0.692^{***}
	0.021	0.022	0.041	0.017
$\log(QPW_{mhcz})$	0.704^{***}	-0.305^{***}	-0.283^{***}	-0.190^{***}
	0.031	0.032	0.019	0.014
Observations	25,500	25,500	25,500	25,500
Fixed effects	hcz	hcz	hcz	hcz
R-squared	0.972	0.756	0.925	0.562
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 bins' sellers per shipment (SPS_{mhcz}) and total quantity shipped per week (QPW_{mhcz}). Industries are assigned using the main 6-digit NAICS industry of the importer based on total employment. (QPS_{mhcz}), (WBS_{mhcz}), (UV_{mhcz}), and ($length_{mhcz}$) are average quantity per shipment, average weeks between shipment, average unit value, 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 quadruple, and exclude quadruples with less than five shipments. Standard errors, adjusted for clustering by country (c) and product (h) are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

Table A.10: SPS_{mhcz} and Procurement Attributes - Manufacturing

	(1)	(2)	(3)	(4)
Dep. var.	$\ln(QPS_{mhcz})$	$\ln(WBS_{mhcz})$	$\ln(UV_{mhcz})$	$\ln(length_{mhcz})$
$\ln(SPS_{mhcz})$	0.500^{***}	0.538^{***}	-0.181^{***}	-0.540^{***}
	0.014	0.014	0.022	0.012
$\log(QPW_{mhcz})$	0.769^{***}	-0.238^{***}	-0.367^{***}	-0.131^{***}
	0.018	0.018	0.022	0.008
Observations	560,000	560,000	560,000	560,000
Fixed effects	hcz	hcz	hcz	hcz
R-squared	0.950	0.712	0.816	0.434
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 bins' sellers per shipment (SPS_{mhcz}) and total quantity shipped per week (QPW_{mhcz}). Industries are assigned using the main 6-digit NAICS industry of the importer based on total employment. (QPS_{mhcz}), (WBS_{mhcz}), (UV_{mhcz}), and ($length_{mhcz}$) are average quantity per shipment, average weeks between shipment, average unit value, 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 quadruple, and exclude quadruples with less than five shipments. Standard errors, adjusted for clustering by country (c) and product (h) are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

Table A.11: SPS_{mhcz} and Procurement Attributes - Wholesale

	(1)	(2)	(3)	(4)
Dep. var.	$\ln(QPS_{mhcz})$	$\ln(WBS_{mhcz})$	$\ln(UV_{mhcz})$	$\ln(length_{mhcz})$
$\ln(SPS_{mhcz})$	0.443^{***}	0.475^{***}	-0.181^{***}	-0.571^{***}
	0.015	0.015	0.013	0.020
$\log(QPW_{mhcz})$	0.682^{***}	-0.328^{***}	-0.281^{***}	-0.167^{***}
	0.012	0.012	0.017	0.007
Observations	1,215,000	1,215,000	1,215,000	1,215,000
Fixed effects	hcz	hcz	hcz	hcz
R-squared	0.945	0.708	0.856	0.469
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 bins' sellers per shipment (SPS_{mhcz}) and total quantity shipped per week (QPW_{mhcz}). Industries are assigned using the main 6-digit NAICS industry of the importer based on total employment. (QPS_{mhcz}), (WBS_{mhcz}), (UV_{mhcz}), and ($length_{mhcz}$) are average quantity per shipment, average weeks between shipment, average unit value, 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 quadruple, and exclude quadruples with less than five shipments. Standard errors, adjusted for clustering by country (c) and product (h) are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

Table A.12: SPS_{mhcz} and Procurement Attributes - Retail

	(1)	(2)	(3)	(4)
Dep. var.	$\ln(QPS_{mhcz})$	$\ln(WBS_{mhcz})$	$\ln(UV_{mhcz})$	$\ln(length_{mhcz})$
$\ln(SPS_{mhcz})$	0.424^{***}	0.458^{***}	-0.120^{***}	-0.556^{***}
	0.030	0.031	0.023	0.022
$\log(QPW_{mhcz})$	0.643^{***}	-0.366^{***}	-0.195^{***}	-0.115^{***}
	0.007	0.007	0.012	0.008
Observations	525,000	525,000	525,000	525,000
Fixed effects	hcz	hcz	hcz	hcz
R-squared	0.945	0.708	0.856	0.955
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 bins' sellers per shipment (SPS_{mhcz}) and total quantity shipped per week (QPW_{mhcz}). Industries are assigned using the main 6-digit NAICS industry of the importer based on total employment. (QPS_{mhcz}), (WBS_{mhcz}), (UV_{mhcz}), and ($length_{mhcz}$) are average quantity per shipment, average weeks between shipment, average unit value, 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 quadruple, and exclude quadruples with less than five shipments. Standard errors, adjusted for clustering by country (c) and product (h) are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

D.7 Alternative Definition of Relationship Length

In this sub-section, we examine whether our results with respect to relationship length depend on our definition of that variable. If firms treat relationships with the same supplier across different products or modes as different relationships, then relationship length should not be defined using the time passed since the first ever transaction with the supplier but instead the duration of the quintuple. We therefore construct an alternative relationship duration variable. First, for each $mxhcz$ quintuple, we compute the total number of weeks passed between the first and the last transaction. Second, for each $mhcz$ buyer quadruple, we take the average over the length of the $mxhcz$ quintuples within it. We refer to this variable as $Qlength_{mhcz}$ to indicate that it is based on the duration of the quintuple, rather than the overall length of the relationship between the importer and the exporter.

We run the same specification outlined in equation (8) using $Qlength_{mhcz}$ as the dependent variable. The results, reported in Table A.13, are similar to those in the main text, with coefficients that are about twice as large. The first column of the table shows that a one percent increase in SPS_{mhcz} is associated with a 1.1 percent decrease in average relationship length. The second column shows that the average relationship length for quadruples in the fourth quartile is about 235 log points lower than the average relationship length for quadruples in the first quartile.

Table A.13: SPS_{mhcz} and Alternative Relationship Length

	(1)	(2)
Dep. var.	$\log(Qlength_{mhcz})$	$\log(Qlength_{mhcz})$
$\log(SPS_{mhcz})$	-1.126^{***}	
	0.039	
$(SPS_{mhcz} = Q2)$		-0.653^{***}
		0.013
$(SPS_{mhcz} = Q3)$		-1.230^{***}
		0.024
$(SPS_{mhcz} = Q4)$		-2.348^{***}
		0.046
$\log(QPW_{mhcz})$	-0.164^{***}	-0.137^{***}
	0.008	0.006
Observations	2,966,000	2,966,000
R-squared	0.619	0.613
Fixed effects	hcz	hcz
Controls	beg, end	beg, end

Source: LFTTD and authors' calculations. Table reports the results of regressing the average quintuple relationship length within each quadruple ($Qlength_{mhcz}$) quadruples' sellers per shipment (SPS_{mhcz}), sellers per shipment quartile dummies and total quantity shipped per week (QPW_{mhcz}). The regressions include product by country by mode of transport (hcz) fixed effects. All regressions control for the beginning and end week of the quadruple, and exclude quadruples with less than five shipments. Standard errors, adjusted for clustering by country (c) and product (h) bins are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

D.8 Average Firm Attributes

In this sub-section, we examine whether firm-level attributes are affected by firms' import sourcing strategy, measured by SPS_{mhcz}. In contrast to the main text, where we use the firm-level attribute in the year of the first import transaction, here we compute for each buyer quadruple an average across all years in which the quadruple is active. The two specifications could generate different results if the firm's attributes change significantly over time, for example because the firm is growing. We average across quadruples to the firm-level. Similar to the main text, we run

$$ln(\overline{\hat{Y}}_m) = \beta_1 ln(SPS_m) + \epsilon_{m\cdot}, \tag{A.22}$$

where $\overline{\hat{Y}}_m$ is one of importer m's total sales, total payroll, average wage, or the firm's inventory-to-sales ratio averaged across years and quadruples.

The results, shown in Table A.14, are relatively similar to those in the main text.

Table A.14: SPS_m and Firm Characteristics

	(1)	(2)	(3)	(4)
Dep. var.	$\log(sales_m)$	$\log(pay_m)$	$\log(wage_m)$	$(inv/sales)_m$
$\log(SPS_m)$	-0.255^{***}	-0.313^{***}	-0.066^{***}	0.016^{***}
	0.005	0.006	0.002	0.001
Observations	184,000	184,000	184,000	48,500
R-squared	0.012	0.014	0.004	0.007

Source: LFTTD and authors' calculations. Table reports the results of regressing importer characteristics averaged across all years in which the importer is active on sellers per shipment (SPS_{mhcz}) averaged across all quadruples involving the importer. All regressions exclude quadruples with less than five shipments. $(sales_m)$, (pay_m), $(wage_m)$, and $((inv/sales)_m)$ are total sales, total payroll, average wage (i.e., payroll divided by number of employees), and total inventory at the beginning of the year divided by total sales, respectively. Robust standard errors are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

E Additional DID Regressions

E.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 also hold if we use a different post-PNTR period from 2004 to 2009. Table A.15 presents the results from the continuing relationship PNTR regression (11) using these alternative time periods, and Table A.16 shows the results for the regression with only new relationships. All results retain their expected sign and remain significant. Table A.17 presents the results from the within-importer regression, equation (12), both at the $mhcz$ level and at the hcz level. On average, we find that the results from the main text become stronger for this later post-period, possibly because the shift of systems takes time.

Table A.15: Within $mxhcz$ Quintuple PNTR DID Regression: 2004-2009 vs 1995-2000

	(1)	(2)	(3)
Dep. var.	$\ln(QPS_{mxhczt})$	$\ln(WBS_{mxhczt})$	$\ln(UV_{mxhczt})$
$Post_t * China_c * NTRGap_h$	-0.199***	-0.163***	0.149***
	0.017	0.021	0.031
$ln(QPW_{mxhczt})$	0.403***	-0.606***	-0.133***
	0.009	0.008	0.014
Observations	221,000	221,000	221,000
R-squared	0.980	0.883	0.982
Fixed effects	$mxhcz, t$	$mxhcz, t$	$mxhcz, t$
Controls	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_{mxhczt}), (WBS_{mxhczt}), and (UV_{mxhczt}) are average quantity per shipment, average weeks between shipments, and average unit value (i.e. value divided by quantity) in period t. All regressions include $mxhcz$ and period t fixed effects, control for the beginning and end week of the quintuple as well as all variables needed to identify the DID term of interest. Standard errors, adjusted for clustering by country (c) and product (h), are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

Table A.16: New $mxhcz$ Quintuple PNTR DID Regression: 2004-2009 vs 1995-2000

	(1)	(2)	(3)
Dep. var.	$\ln(QPS_{mxhczt})$	$\ln(WBS_{mxhczt})$	$\ln(UV_{mxhczt})$
$Post_t * China_c * NTRGap_h$	-0.087**	-0.067*	0.075*
	0.036	0.035	0.045
$\ln(QPW_{mxhczt})$	0.414***	-0.590***	-0.127***
	0.012	0.011	0.017
Observations	3,158,000	3,158,000	3,158,000
R-squared	0.968	0.845	0.973
Fixed effects	$mhcz, x, t$	$mhcz, x, t$	$mhcz, x, t$
Controls	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_{mxhczt}), (WBS_{mxhczt}), and (UV_{mxhczt}) are average quantity per shipment, average weeks between shipments, and average unit value (i.e. value divided by quantity) in period t. All regressions include $mxhcz$ and period t fixed effects, control for the beginning and end week of the quintuple as well as all variables needed to identify the DID term of interest. Standard errors, adjusted for clustering by country (c) and product (h), are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

Table A.17: Within-Importer PNTR Regression: 2004-2009 vs 1995-2000

	(1)	(2)	(3)	(4)
Dep. var.	$\ln(SPS_{mhczt})$	$1\{J_{mhczt}=1\}$	$\ln(SPS_{hczt})$	$1\{J_{hczt}=1\}$
$Post_t * China_c * NTRGap_h$	-0.076**	0.076**	-0.027**	0.042
	0.037	0.029	0.011	0.027
$ln(QPW_{mhczt})$	-0.186***	0.125***	-0.059***	0.031***
	0.005	0.005	0.002	0.004
Observations	556,000	225,000	355,000	28,000
R-squared	0.757	0.660	0.687	0.550
Fixed effects	$mhcz, t$	$mhcz, t$	hcz, t	hcz, t
Controls	Yes	Yes	Yes	Yes

Source: LFTTD and authors' calculations. First two columns report 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. Second two columns are analogous but at the hcz level of aggregation. Pre- and post-PNTR periods are 1995 to 2000 and 2004 to 2009. All regressions include period t fixed effects, and control for the beginning and end week of the quadruple as well as all variables needed to identify the DID term of interest. Regressions in columns two and four are restricted to quadruples with at least five transactions in both periods. Standard errors, adjusted for clustering by country (c) and product (h), are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

E.2 PNTR Regression Without Quantity Control

In this sub-section, we examine whether our baseline PNTR results from Section 5.2 also hold when we do not control for quantity per week, QPW_{mxhczt}. One concern with our analysis could be that by conditioning on quantity, we do not take into account that PNTR also affects the quantity traded, which could in turn affect the procurement system. Our results in Table A.18 show that we still find a decline in the quantity per shipment and an increase in the unit value even without the quantity control. The effect on weeks between shipments is qualitatively consistent with the theory, though not significant at conventional levels.

Table A.18: Baseline Within $mxhcz$ Quintuple PNTR DID Regression Without Quantity: 2002-2007 vs 1995-2000

	(1)	(2)	(3)
Dep. var.	$\ln(QPS_{mxhczt})$	$\ln(WBS_{mxhczt})$	$\ln(UV_{mxhczt})$
$Post_t * China_c * NTRGap_h$	-0.2753***	-0.0339	0.1186***
	0.0076	0.0318	0.0191
Observations	439,000	439,000	439,000
R-squared	0.97	0.69	0.98
Fixed effects	$mxhcz, t$	$mxhcz, t$	$mxhcz, t$
Controls	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_{mxhczt}), (WBS_{mxhczt}), and (UV_{mxhczt}) are average quantity per shipment, average weeks between shipment, and average unit value (i.e. value divided by quantity) in period t. All regressions include $mxhcz$ and period t fixed effects, control for the beginning and end week of the quadruple as well as all variables needed to identify the DID term of interest. Standard errors, adjusted for clustering by country (c) and product (h), are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

E.3 Response to PNTR Among Continuing $mhcz$ Quadruples

We examine in this section whether the shift from "American" to "Japanese" procurement in response to PNTR also affected shipping patterns at the $mhcz$ quadruple level. Compared to the regressions of continuing relationships at the $mxhcz$ level, this regression aggregates across the supplier dimension, and computes shipping attributes of the quadruple using transactions with all suppliers. It also allows for an additional margin of extensive margin adjustment, namely the formation of relationships with new suppliers that did not sell to the United States prior to PNTR. We use the following $mhczt$-level DID regression,

$$\ln(Y_{mhczt}) = \beta_1 1\{t = Post\} * 1\{c = China\} * NTRGap_h + \beta_2 ln(QPW)_{mhczt} + \beta_3 \chi_{mhczt} + \lambda_{mhcz} + \lambda_t + \epsilon_{mhczt}. \quad \text{(A.23)}$$

As before, Y_{mhczt} represents one of the three procurement attributes. The pre-PNTR period is 1995 to 2000 and the post period is 2002 to 2007.

Results, displayed in Table A.19, show a significant decline in the average shipping size and weeks between shipments, consistent with a shift towards "Japanese" procurement. The increase in unit values, while positive, is statistically insignificant at conventional levels. One potential explanation for this outcome is the entry of new Chinese exporters during this period (Pierce and Schott, 2016; Amiti, Dai, Feenstra, and Romalis, 2020), including privately owned firms that tend to have lower prices than state-owned incumbents (Khandelwal, Schott, and Wei, 2013). New suppliers might also charge low, introductory prices to gain market share, further dampening unit values.

Table A.19: Within $mhcz$ Quadruple PNTR DID Regression

	(1)	(2)	(3)
Dep. var.	$\ln(QPS_{mhczt})$	$\ln(WBS_{mhczt})$	$\ln(UV_{mhczt})$
$Post_t * China_c * NTRGap_h$	-0.043***	-0.058***	0.018
	0.014	0.013	0.024
$ln(QPW_{mhczt})$	0.436***	-0.584***	-0.207***
	0.018	0.018	0.026
Observations	738,000	738,000	738,000
R-squared	0.978	0.887	0.974
Fixed effects	$mhcz, t$	$mhcz, t$	$mhcz, t$
Controls	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_{mhczt}), (WBS_{mhczt}), and (UV_{mhczt}) are average quantity per shipment, average weeks between shipment, and average unit value (i.e. value divided by quantity) in period t. All regressions include $mhcz$ and period t fixed effects, control for the beginning and end week of the quadruple as well as all variables needed to identify the DID term of interest. Standard errors, adjusted for clustering by country (c) and product (h), are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

E.4 PNTR Regression For All Relationships

In this sub-section, we re-run our relationship-level PNTR regression (11) using both continuing and new relationships simultaneously for all buyer quadruples and sellers that appear in both. Specifically, we run a modified version of the regression,

$$\begin{aligned}\ln(Y_{mxhczt}) =& \beta_1 1\{t = Post\} * 1\{c = China\} * NTRGap_h + \beta_2 ln(QPW_{mxhczt}) + \\ & \beta_3 \chi_{mxhczt} + \lambda_{mhcz} + \lambda_x + \lambda_t + \epsilon_{mxhczt},\end{aligned} \tag{A.24}$$

where we use importer-product-country-mode of transportation ($mhcz$) fixed effects, exporter (x) fixed effects, and period (t) fixed effects and include all $mxhcz$ quintuples in the sample. Our results in Table A.20 indicate that PNTR leads to a decline in the quantity per shipment and the number of weeks between shipments, and an increase in the unit value for this set of relationships.

Table A.20: Within $mxhcz$ Quintuple PNTR DID Regression Using All Relationships: 2002-2007 vs 1995-2000

	(1)	(2)	(3)
Dep. var.	$\ln(QPS_{mxhczt})$	$\ln(WBS_{mxhczt})$	$\ln(UV_{mxhczt})$
$Post_t * China_c * NTRGap_h$	-0.131***	-0.115**	0.078***
	0.012	0.012	0.027
$ln(QPW_{mxhczt})$	0.407***	-0.597***	-0.130***
	0.013	0.012	0.018
Observations	4,023,000	4,023,000	4,023,000
R-squared	0.966	0.838	0.971
Fixed effects	$mhcz, x, t$	$mhcz, x, t$	$mhcz, x, t$
Controls	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_{mxhczt}), (WBS_{mxhczt}), and (UV_{mxhczt}) are average quantity per shipment, average weeks between shipment, and average unit value (i.e. value divided by quantity) in period t. All regressions include $mhcz$, exporter x, and period t fixed effects, and control for the beginning and end week of the quadruple as well as all variables needed to identify the DID term of interest. Standard errors, adjusted for clustering by country (c) and product (h), are reported below coefficient estimates. ***, **, and * represent statistical significance at the 1, 5 and 10 percent levels.

F Market Clearing Conditions

This section provides the market clearing conditions of the multi-country model described in Section 6.

Goods market clearing implies for each manufactured good ω

$$\sum_n \sum_i \sum_s I_{ni,s}(\omega) x^*_{ni,s}(\omega) = \sum_n \sum_i \sum_s I_{ni,s}(\omega) \int_0^{x^*_{ni,s}(\omega)/q_n(\omega)} q_n(\omega) dt \quad \forall \omega, \tag{A.25}$$

where $I_{in,s}(\omega)$ is an indicator function that is equal to one if the buyer in country i procures product ω from country n under system s, and zero otherwise. 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 the homogeneous good clears as well,

$$\sum_n Z_n = \sum_n a_n L^O_n. \tag{A.26}$$

Finally, labor market clearing in each country requires that

$$L_n = \sum_{n'} \sum_{s} I_{n'n,s}(\omega) \int_0^1 \frac{\bar{\theta}}{\Upsilon_n(\omega)} q_{n'}(\omega) d\omega + f_n \sum_{n'} \sum_{s} I_{n'n,s}(\omega) \int_0^1 \frac{q_{n'}(\omega)}{x^*_{n'n,s}(\omega)} d\omega$$
$$+ \sum_{i} I_{ni,A} \int_0^1 (\omega) m(\omega) \frac{q_n(\omega)}{x^*_{ni,s}(\omega)} d\omega + L_n^O \quad \forall n \in N, \tag{A.27}$$

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 the homogeneous "outside" good 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.

G 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.

$$AC_A(\mathbf{k}) = \min_x \left(\frac{r}{q_k}\right) \frac{k_1 + k_2 x}{\left[1 - e^{-rx/q_k}\right]} \tag{A.28}$$

and

$$AC_J(\mathbf{k}) = \min_x \left(\frac{r}{q_k}\right) \frac{k_1 + e^{(r+k_3)x/q_k} k_2 x}{\left[1 - e^{-rx/q_k}\right]}. \tag{A.29}$$

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 and country n from $G_n(m)$ and by drawing a productivity $\Upsilon_n(\omega)$ from $F_n(\Upsilon)$. 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 for each q_k 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, $x^*_{ni,s}$, 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 (15). 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

$$p_n(\omega; q_k, t) = \left(\frac{\alpha W_n(t)}{q_k} \right)^{\frac{1}{\sigma}} P_n(t)^{\frac{\sigma-1}{\sigma}}. \tag{A.30}$$

   (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 the homogeneous good sector L_n^O from the labor market clearing condition (A.27). 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 (16).

   (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 via interpolation on a grid for efficiency reasons. While it would be possible to solve the buyer's problem for each $ni\omega$ tuple at each iteration, thus combining steps 1-3, in practice the numerical solution of the buyer's 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.

H Parameters and Empirical Moments

H.1 Calibrated Parameters

In this subsection, we provide more details on how we compute the calibrated parameters in Table 12.

1. Elasticity of substitution (σ)

Under CES preferences, firms markups can be used to infer the demand elasticity. Antràs, Fort, and Tintelnot (2017) use data of the Census of Manufactures to compute establishment-level markups, and find a median markup of 35 percent across establishments. This estimate implies an elasticity of substitution of $\sigma = 3.85$.

2. Consumption share of manufactured goods (α)

We construct this parameter based on estimates by Duarte (2020), who uses detailed data on individual consumption expenditure by households from the International Comparisons Programs (ICP) to compute consumption expenditures and relative prices for manufactured goods and for services for a large number of countries. Duarte (2020) computes a real share of manufactured goods consumption in all consumption expenditures of $45\% - 50\%$ for high-income countries (Table 4).

3. Dispersion of productivities (ζ)

We set this parameter based on Eaton and Kortum (2002), who estimate it from a gravity equation that relates bilateral trade flows to the characteristics of the trading partners and the distance between them.

4. Wages (w_n)

We estimate the average wage in China and in the rest of the world from the World Bank World Development Indicators (WDI) in 2016. For each country we obtain GDP in current USD (series NY.GDP.MKTP.CD) and the total size of the labor force (series SL.TLF.TOTL.IN). Assuming that two thirds of a worker's output goes to labor, we construct each country's average wage as two thirds times GDP divided by the size of the labor force (i.e., GDP per worker). For the rest of the world, we take an average across the U.S.' top-ten trading partners (listed in Table 3) using U.S. imports from each country in 2016 as weight, where we obtain U.S. imports from the U.S. Census Bureau. We compute relative wages of China and the Rest of the World by dividing by the U.S. wage.

5. Labor force (L_n)

We obtain the size of the labor force from the World Bank World Development Indicators (WDI) in 2016 (series SL.TLF.TOTL.IN). For the rest of the world, we set the total labor force equal to the labor force of the top ten U.S. trading partners in the period 1992-2016. These countries are shown in Table 3. We compute the relative labor force by dividing by the size of the U.S. labor force.

6. Rate of trade wars U.S.-China ($\rho_{US,CN}$)

We calibrate this parameter by taking all buyer-seller ($mxhcz$) quintuples in the data that appear to be associated with "Japanese" procurement, identified as those where the associated $mhcz$ quadruple is in the first quartile of the within-country-product-mode SPS distribution in the entire dataset. For

these relationships, we compute the separation hazard rate $H_{c\tau}$, i.e., the probability that a relationship separates after τ quarters conditional on having lasted that long, separately for c =China and c =Rest of the World:

$$H_{c\tau} = \frac{\sum_{mxhzt} \mathbb{I}^T(\tau_{mxhczt} = \tau)}{\sum_{mxhzt} \mathbb{I}(\tau_{mxhczt} = \tau)}. \tag{A.31}$$

Here, $\mathbb{I}(\tau_{mxhczt} = \tau)$ is an indicator that is equal to one if quintuple $mxhcz$ is at age $\tau_{mxhczt} = \tau$ quarters in quarter t, where the age is defined as the number of quarters passed since the first transaction of $mxhcz$. The indicator $\mathbb{I}^T(\tau_{mxhczt} = \tau)$ is equal to one for all such quintuples that additionally trade for the last time in quarter t. Both the numerator and denominator sum these indicators across all quintuples and quarters in the data, separately for China and for the rest of the world. We exclude 2016 from this computation due to the censoring of our data.

We fit the exponential decay function $e^{-\psi_{ni}t}$ to the empirically estimated hazard curves to minimize the squared deviation the two for each country. We fit this function from quarter two of the quintuple onwards since a large share of quintuples break up after only one transaction for exogenous reasons. Since we assume that all break-ups with the rest of the world are due to normal churn, we take the excess probability of break-ups with Chinese suppliers as reflecting the added probability of a trade war, and hence set $\rho_{US,CN} = 0.0264$.

H.2 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 manufacturing 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[67], and obtain gross output in the manufacturing sector in 2016 from the Bureau of Economic Analysis (BEA).[68] We also obtain total U.S. imports and exports of goods in 2016 from the Census Bureau. We then construct the target

$$IP_{CN} = \frac{\text{Imports}_{CN}}{\text{Domestic output} + \text{Total imports} - \text{Total exports}}, \tag{A.32}$$

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 manufacturing 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

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

from the Census Bureau. We then construct the target

$$IP_{RoW} = \frac{\text{Imports}_{RoW}}{\text{Domestic output} + \text{Total imports} - \text{Total exports}}, \tag{A.33}$$

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.

3. Moments from regression (19)

We start from our baseline sample for the main classification regression (8) and restrict the sample to only those $mhcz$ quadruples with WBS_{mhcz} in the first or in the fourth quartile of the distribution of WBS_{mhcz} within a given country by product by mode bin. We then run two regressions of the form

$$ln(\overline{WBS}_{mhcz}) = \beta_0 + \beta_1 1\{WBS_{mhcz} = Q4\} + \beta_2 \ln(QPW_{mhcz}) + \beta_3 beg_{mhcz} + \beta_4 end_{mhcz} + \lambda_{hcz} + \epsilon_{mhcz}, \tag{A.34}$$

separately for China and for the rest of the world. In this regression, $\bar{WBS}_{mhcz}$ is the average weeks between shipments for the $mhcz$ quadruple, constructed as described in Appendix C, and $1\{WBS_{mhcz} = Q4\}$ is a dummy indicating whether WBS_{mhcz} falls into the fourth quartile of its distribution, as opposed to the first quartile (the omitted category). The variable QPW_{mhcz} is the total quantity traded by the quadruple over the course of its life divided by the number of weeks it exists, and beg_{mhcz} and end_{mhcz} are continuous variables indicating the week numbers that the quadruple starts trading and ceases, respectively. We include product-country-mode fixed effects, λ_{hcz}. We target $\hat{\beta}_0 + \hat{\beta}_1$ to match the fixed costs of shipping and $\hat{\beta}_1$ to match the dispersion of inspection costs.

We additionally target the standard deviation of the residuals from this regression for "American" quadruples. These residuals capture the dispersion in shipping frequencies across quadruples due to different inspection cost draws, after controlling for differences in the quantity, time of trade, or fixed product-country-mode factors. We construct this moment by taking the residuals $\hat{\epsilon}_{mhcz}$ for only those quadruples that are "American", i.e., $1\{WBS_{mhcz} = Q4\} = 1$. We then take the mean of these residuals for each HS10 product to remove within-product variation and compute the standard deviation across products, separately for China and for the rest of the world.

I Additional Details on the Estimation and Robustness

In this section, we provide additional details on the estimation. In Section I.1 we provide additional details on the estimation algorithm and on identification of the parameters. In Section I.2, we re-estimate the model but assume a Fréchet distribution for inspection costs instead of a Pareto distribution.

I.1 Baseline Estimation

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

$$\arg\min_{\phi\in\mathbb{F}} \sum_x T(\mathcal{M}_x(\phi), \hat{\mathcal{M}}_x) \tag{A.35}$$

where $T(\cdot)$ is the percentage difference between the model, $(\mathcal{M}_x(\phi))$, and data, $(\hat{\mathcal{M}}_x)$, 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 $T((\mathcal{M}_x(\phi), (\hat{\mathcal{M}}_x))$ we follow Jarosch (2016) and Lise, Meghir, and Robin (2016) and minimize the sum of the percentage deviations between model-generated and empirical moments.

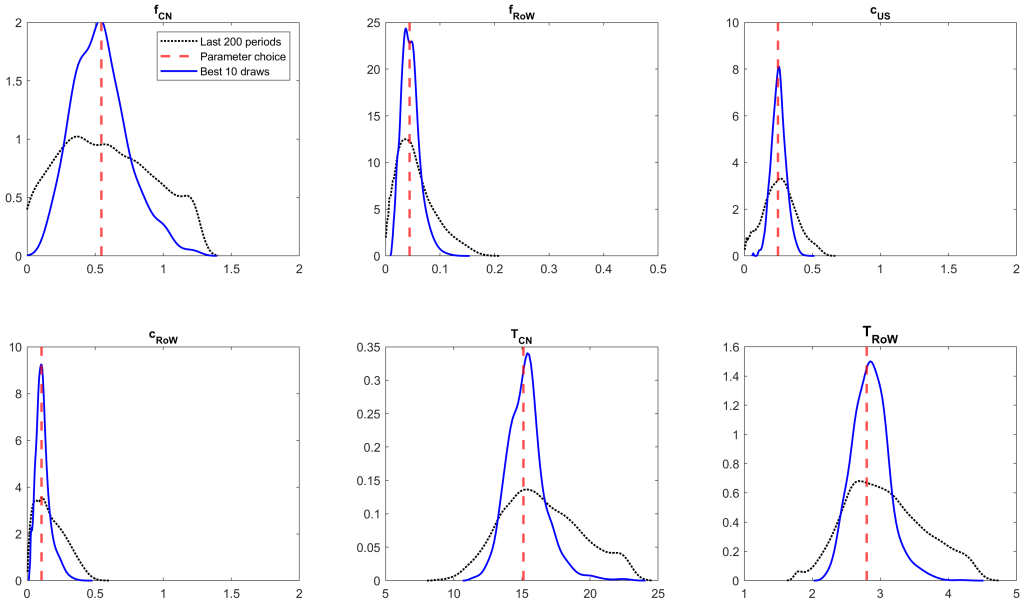
The minimization algorithm that we use to solve the problem combines the approaches of Jarosch (2016) and Lise, Meghir, and Robin (2016), and Engbom and Moser (2021), adapted to our needs.

We simulate, using Markov Chain Monte Carlo for classical estimators as introduced in Chernozhukov and Hong (2003), 100 strings of length 1,000 (+ 200 initial scratch periods used only to calculate posterior variances) starting from 100 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 Gelman, Carlin, Stern, and Rubin (2013). The average parameter values across the 20 strings with the lowest values of the objective function provide a first estimate of the vector of parameters. We then repeat the same MCMC procedure, but we start each of our 100 strings from these parameter estimates.

Figure A.1 illustrates our approach. The black dotted line shows the density function of the parameter values associated with the last 200 iterations of our 100 strings. We pick the optimal parameters following Engbom and Moser (2021) and select the vector of parameters that minimizes the objective function among all our draws, by taking the average across the 100 best outcomes across all the draws. Our estimates are shown with red dashed lines in the figure. These correspond to the estimates reported in Table 13 in the main text. Finally, for comparison, the blue density functions shows the density of the 10 best outcomes of each string, computed across all strings. This density provides an alternative way to select the best parameter values. We find that all the densities are single-peaked, which suggests that the model is, at least locally, identified. Moreover, our chosen parameter values are generally very close to the peak of the densities.

Figure A.2 provides more detail on how each parameter is identified based on our moments. We start from the optimal parameter values (red dashed lines in the previous figure) and vary each of the six parameters one-by-one on a grid of 200 values. For each parameter combination we solve the model 100 times with different productivity and inspection cost draws and compute the average value of each moment. The panels in Figure A.2 plot the different values of each parameter (rows) against the values of the eight moments (columns). The moments identifying the parameters are mostly along the diagonal. For example, the top left panel plots different values of the productivity scale parameter for China, T_{CN}, against the share of manufacturing consumption from China. There is a strong negative relationship between the two. The red horizontal line represents the value of the

Figure A.1: Estimation Outcomes

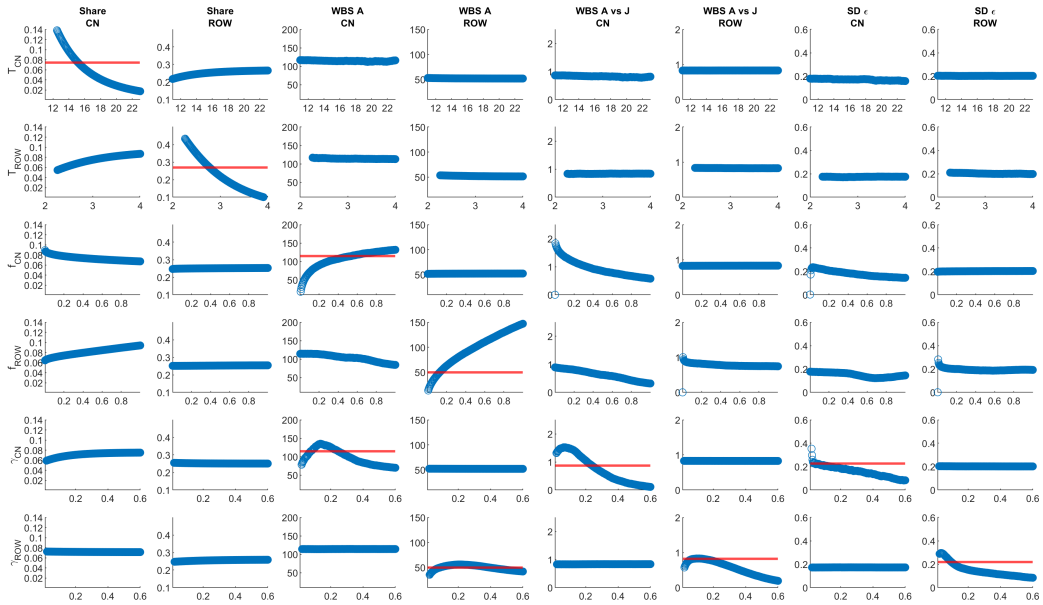


Source: Author's calculations, based on the estimation procedure described. Each panel shows the estimated parameter values for the parameter indicated in the title, under the assumption of a Pareto distribution for inspection costs. The black dotted line shows the density function of the parameter values associated with the last 200 iterations of our 100 strings. The red dashed line shows the average parameter values across the 100 best outcomes from all the draws. The blue density functions shows the density of the 10 best outcomes of each string, computed across all strings.

moment in the data. Thus, the intersection point between the red line and the blue dotted curve gives the parameter value that would lead the model to perfectly match this moment. The relationships between the first four parameters and their identifying moments are monotonic, and hence there is a unique parameter value that leads the model to perfectly match the targeted moment. For the last two parameters (the dispersion of inspection costs, γ_n) the relationships with some of the targeted moments are hump-shaped, and hence in principle there could be multiple values for each of these parameters that match a given moment equally well. We therefore target two sets of moments: the difference between the shipping frequency for shipments in the first and in the fourth quartile of the WBS distribution, obtained as $\hat{\beta}_1$ from specification 19 (columns 5 and 6 of the figure), and the standard deviation of the shipping time residuals (columns 7 and 8). In addition, the frequency of shipments for shipments that are most likely "American", obtained as $\exp(\hat{\beta}_0 + \hat{\beta}_1)$ from specification 19 (columns 3 and 4 of the figure), are also informative. As shown in Figure A.1 above, this strategy yields a unique value for these parameters that minimizes the objective function.

As a final exercise, we vary all parameters jointly and plot the relationships between parameters and moments. In Figure A.3, we draw 100,000 different combinations of parameters and provide a bin scatter of the parameter values against the value of each moment. The combinations of the six parameter values are obtained as quasi random numbers drawn from a Sobol sequence. In the figure, lighter colors indicate combinations of moment and parameter values that arise more frequently across the 100,000 estimations. The figure shows similar relationships as Figure A.2 where we varied

Figure A.2: Identification of Parameters



Source: Author's calculations, based on the estimation procedure described. Each panel plots different values of the parameter indicated on the row against the moment indicated on the column, keeping all other parameters fixed at their optimal value. The blue dots show the averaged moment value across 100 runs with the given parameter choice, where the averaging is needed since the inspection cost and productivity draws differ across runs. The red horizontal lines represent the value of the moment in the data. We add these only for the main panels used to identify a given parameter in the data.

only one parameter at a time, although the associations are noisier. In particular, there are strong and monotone relationships between the first four parameters and their targeted moments, and more hump-shaped relationships for the final two parameters.

Overall, these exercises highlight that our parameters of interest are well-identified from the moments we target.

I.2 Fréchet Distribution of Inspection Costs

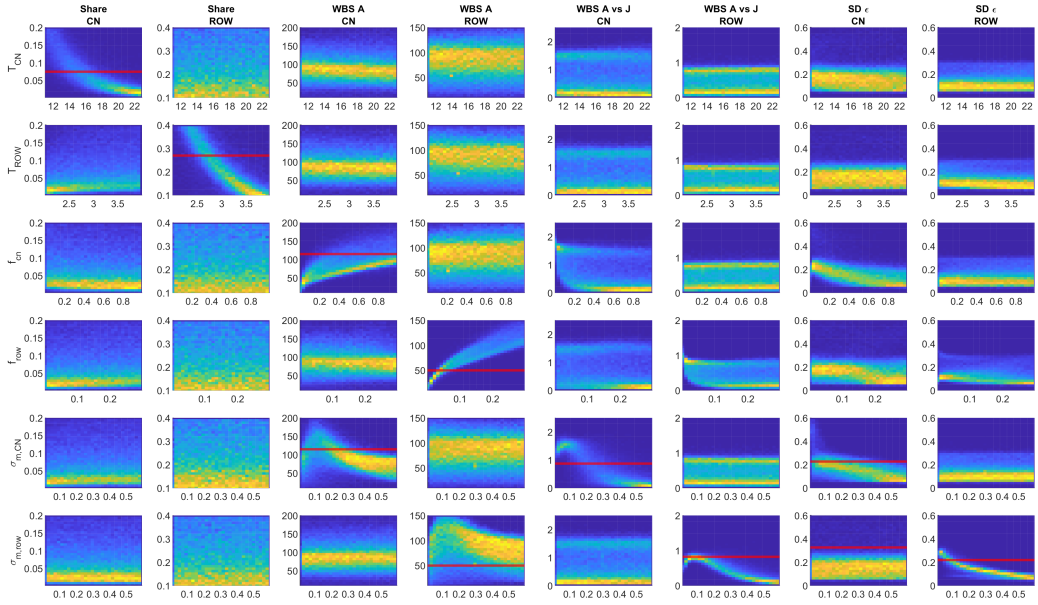
In this section, we re-estimate the model using a Fréchet distribution instead of a Pareto distribution for the inspection costs. Specifically, we assume that inspection costs are distributed according to

$$G_n(m) = e^{-\Phi m^{-\gamma_n}}, \tag{A.36}$$

which Φ and γ_n parameters to be estimated. The other model parameters are set as before. In practice, we find that the scale parameter of the Fréchet distribution, Φ, is not well identified. We therefore set it to $\Phi = 1$, and re-estimate the model with different values for Φ below. We find that all the estimated parameter values and all results are nearly unchanged for different values of Φ.

Figure A.4 presents our estimated parameter values analogously to Figure A.1. The black dotted line shows the density function of the parameter values associated with the last 200 iterations of our 100 strings. The red dashed lines show the estimated parameter values. The blue density functions

Figure A.3: Joint Identification of Parameters

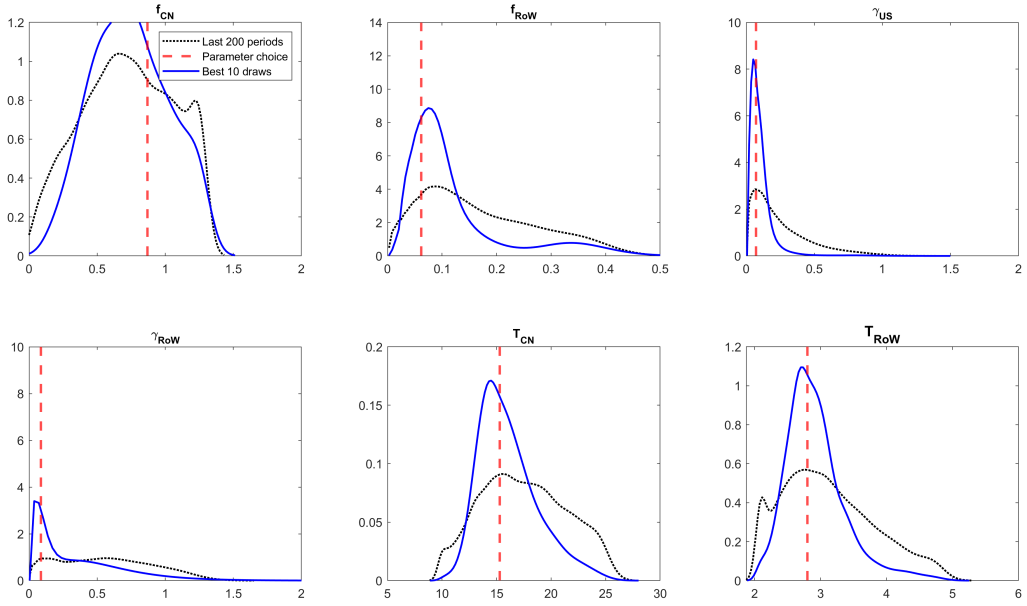


Source: Author's calculations, based on the estimation procedure described. Each panel plots different values of the parameter indicated on the row against the moment indicated on the column, where all parameters vary jointly based on 100,000 random parameter draws from a Sobol sequence. Lighter colors indicate more frequently observed combinations of parameter values and moment values. The red horizontal lines represent the value of the moment in the data. We add these only for the main panels used to identify a given parameter in the data.

shows the density of the 10 best outcomes of each string, computed across all strings, for comparison. We find that all the densities are single-peaked and quite tightly estimated. Our chosen parameter values are very close to the peak of the densities.

Table A.21 presents the estimated parameter values and the values of the targeted moments in the simulations and in the data. The moments are generally well-matched. The model generates shares of Chinese and RoW imports in U.S. manufacturing consumption that are close to the data, and generates shipping frequencies in line with their empirical analogues. The only moment the model does not match well is the difference in shipping frequencies between the first and the fourth quartile for shipments from China. In the data, the difference in shipping times between the first and the fourth quartile of the WBS_{mhcz} distribution is relatively small, while the dispersion of shipping times within the first quartile is relatively large. To match the latter the model estimates a high volatility of inspection costs (low γ_{CN}), which causes the model to overshoot the former moment for China. For the rest of the world, the two moments are relatively well matched. Due to this deviation from the targeted moments, we prefer the Pareto distribution as our baseline, which matches all moments better due to its different shape.

Figure A.4: Estimation Outcomes



Source: Author's calculations, based on the estimation procedure described. Each panel shows the estimated parameter values for the parameter indicated in the title. The black dotted line shows the density function of the parameter values associated with the last 200 iterations of our 100 strings. The red dashed line shows the average parameter values across the 100 best outcomes from all the draws. The blue density functions shows the density of the 10 best outcomes of each string, computed across all strings.

Table A.21: Estimated Parameters and Targeted Moments

	(1)	(2)	(3)	(4)	(5)
	Parameter	Estimated Value	Moment that Primarily Identifies the Parameter	Moment in Data	Moment in Model
(1)	Productivity China (T_{CN})	15.288	Share of Chinese imports in domestic sales	0.074	0.070
(2)	Productivity RoW (T_{RoW})	2.800	Share of RoW imports in domestic sales	0.270	0.240
(3)	Fixed costs, China (f_{CN})	0.870	$\exp(\hat{\beta}_0 + \hat{\beta}_1)$ from (19) for CN	114.71	116.28
(4)	Fixed costs, RoW (f_{RoW})	0.062	$\exp(\hat{\beta}_0 + \hat{\beta}_1)$ from (19) for RoW	49.83	53.68
(5)	Dispersion of inspection	0.072	$\hat{\beta}_1$ from (19) for China	0.871	1.399
(6)	costs, China (γ_{CN})		Sd of $\hat{\epsilon}$ from (19) for China	0.227	0.209
(7)	Dispersion of inspection	0.086	$\hat{\beta}_1$ from (19) for RoW	0.822	0.761
(8)	costs, RoW (γ_{RoW})		Sd of $\hat{\epsilon}$ from (19) for RoW	0.219	0.261

Source: LFTTD and authors' calculations. Column (1) lists the parameters estimated for the model. Column (2) contains the estimated parameter values. Column (3) reports the moment targeted to identify the parameter. Column (4) presents the value of the moment in the data, and Column (5) presents the value of the moment computed in our simulated model.

Table A.22 shows selected moments from our baseline equilibrium and the counterfactual without "Japanese" relationships. Compared to the equilibrium with a Pareto distribution, the estimated share of "Japanese" relationships is significantly higher for both China and for the rest of the world. This higher share of "Japanese" relationships than in the baseline results from the higher dispersion

of inspection costs in this estimation. More dispersed inspection costs generate more high-cost inspection cost draws, leading "Japanese" sourcing to be cheaper than "American" sourcing for more products. 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 3. As a result of the higher share of "Japanese" relationships, the welfare losses from removing the possibility to form such relationships rise by almost two percentage points compared to the baseline, yielding a 3.5 percent welfare loss. These estimates thus provide a sensitivity check of our baseline results. They illustrate that if the share of "Japanese" relationships were higher than what we estimate in the baseline, welfare losses would rise.

Table A.22: Comparison of Equilibria with Fréchet Distribution

		(1)	(2)	(3)	(4)
		Baseline Equilibrium	**Equilibrium Without Japanese Sourcing**	**Autarky**	**PNTR**
(1)	Value imported from China (%)	7.0%	4.7%	.	7.8%
(2)	- of which, "Japanese"	57.2%	.	.	64.8%
(3)	Value imported from ROW (%)	23.9%	10.0%	.	23.7%
(4)	- of which, "Japanese"	73.6%	.	.	73.6%
(5)	Value imported from U.S. (%)	69.1%	85.3%	100.0%	68.5%
(6)	Avg. inspection costs	0.2%	1.6%	.	0.2%
(7)	Avg. fixed costs (imports)	5.2%	3.3%	.	5.3%
(8)	Manufacturing price index	1.000	1.058	1.105	0.998
(9)	Utility	1.000	0.965	0.944	1.0002

Table shows various statistics of the equilibrium under the assumption of a Fréchet distribution for inspection costs. The first column presents the statistics for the baseline equilibrium, using the parameters that minimize the objective function. The second column shows the same statistics for a counterfactual economy in which the formation of "Japanese" relationships is not possible due to $\rho \to \infty$. The third column shows an autarky economy in which trade is not possible. The fourth column shows a counterfactual economy in which we reduce the arrival rate of trade wars from China to zero. Rows 1-5 show the share of U.S. manufacturing sales, $P_{US}Q_{US}$, that is from China, from the rest of the world, and from the U.S., respectively, and the share of these manufacturing sales that is sourced under the "Japanese" system. Row 6 presents the average inspection costs as a share of the import value, computed over all imports, including under the "Japanese" system. Row 7 shows the average fixed costs as a share of the import value. Row 8 shows the manufacturing price index, P_{US}, normalized to one in the baseline. Row 9 shows total utility, $W_{US} = Q_{US}^{\alpha} Z_{US}^{1-\alpha}$, normalized to one in the baseline.

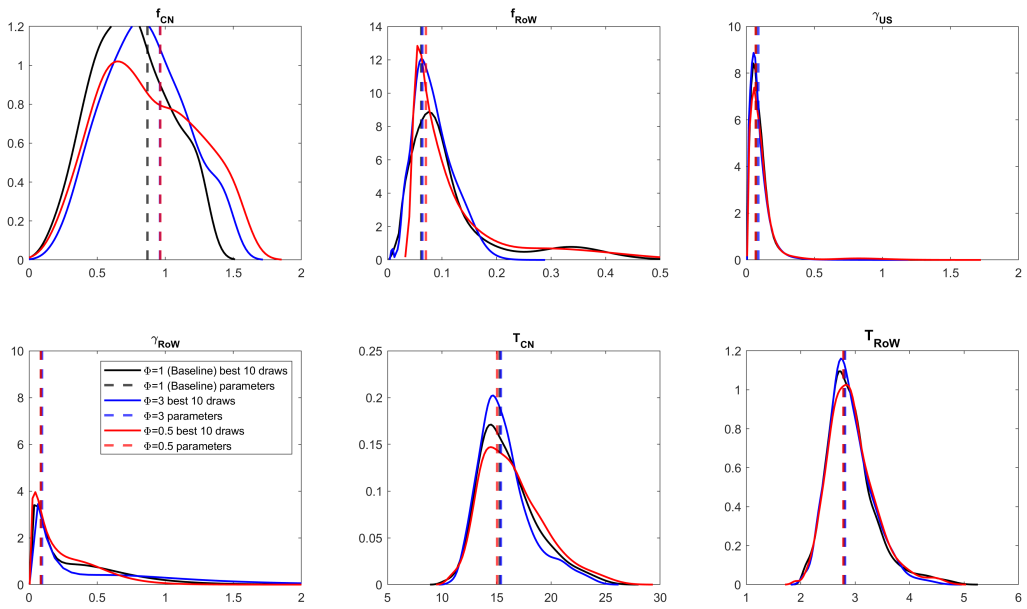
Sensitivity of the Results to the Choice of Φ

We next examine how the choice of Φ, the parameter governing the scale of the inspection cost distribution, affects our results. So far, we have fixed this parameter to $\Phi = 1$ and have argued that it does not significantly affect the estimation outcomes. We now verify this assertion.

We re-run the full estimation but set $\Phi = 3$ or $\Phi = 0.5$ instead of its baseline value of one. Figure A.5 compares the estimated parameter values under these alternatives to the baseline estimation. The black solid line shows the density of the 10 best outcomes of each string, computed across all

strings, from the baseline estimation. The black dashed line shows the optimal parameters for the baseline, computed as the average over the 100 best outcomes across all draws. The blue lines show the corresponding density and optimal parameter values for the case of $\Phi = 3$, and the red lines show the corresponding results for $\Phi = 0.5$. We find that the optimal parameter values across all three estimations are very similar.

Figure A.5: Estimation Outcomes for Different Values of Φ



Source: Author's calculations, based on the estimation procedure described. Each panel shows the estimated parameter values for the parameter indicated in the title. The dashed vertical lines show the average parameter values across the 100 best outcomes from all the draws for the estimations with different values of Φ. The density functions show the density of the 10 best outcomes of each string, computed across all strings, for the estimations with different values of Φ.

Table A.23 presents the estimated parameter values across the three estimations, and re-confirms that they are very close. Finally, in table A.24, we show some moments of interest from the three estimations, replicating Table A.22, and compute the counterfactual welfare losses without "Japanese" relationships. The table illustrates that the share of "Japanese" relationships from each country is very similar across estimations, about $54\% - 58\%$ for China and about $71\% - 76\%$ from the rest of the world. The welfare loss from removing "Japanese" relationships is between 3.4% and 3.7%, with a manufacturing price index that increases around $5\% - 6\%$. Thus, the choice of Φ does not significantly affect the estimation results.

Table A.23: Estimated Parameters

		(1)	(2)	(3)
	Parameter	$\Phi = 1$	$\Phi = 3$	$\Phi = 0.5$
(1)	Productivity China (T_{CN})	15.288	15.395	15.059
(2)	Productivity RoW (T_{RoW})	2.800	2.813	2.781
(3)	Fixed costs, China (f_{CN})	0.870	0.962	0.964
(4)	Fixed costs, RoW (f_{RoW})	0.062	0.064	0.070
(5) (6)	Dispersion of inspection costs, China (γ_{CN})	0.072	0.088	0.070
(7) (8)	Dispersion of inspection costs, RoW (γ_{RoW})	0.086	0.096	0.087

Source: LFTTD and authors' calculations. Column 1 lists the parameters estimated in the baseline estimation with $\Phi = 1$. These are also indicated by the dashed vertical line in Figure A.5. Column 2 shows the estimates with $\Phi = 3$, and column 3 shows the estimates with $\Phi = 0.5$.

Table A.24: Comparison of Equilibria

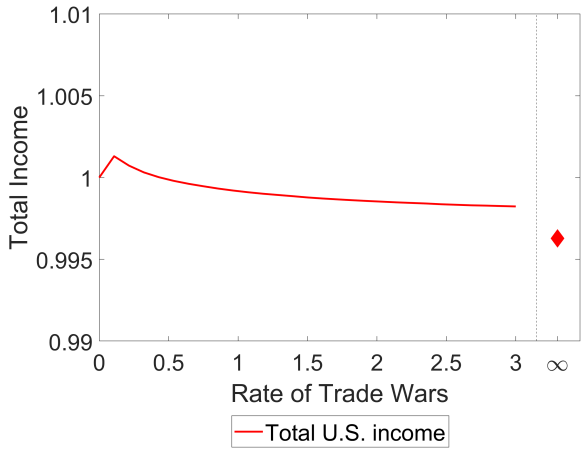
		(1)	(2)	(3)	(4)	(5)	(6)
		$\Phi = 1$		$\Phi = 3$		$\Phi = 0.5$	
		Baseline	**No J**	**Baseline**	**No J**	**Baseline**	**No J**
(1)	Value imported from China (%)	7.0%	4.7%	6.8%	4.9%	7.3%	4.5%
(2)	- of which, "Japanese"	57.2%	.	54.5%	.	58.0%	.
(3)	Value imported from ROW (%)	23.9%	10.0%	23.7%	10.7%	24.3%	9.8%
(4)	- of which, "Japanese"	73.6%	.	71.3%	.	76.1%	.
(5)	Avg. inspection costs	0.2%	1.6%	0.3%	1.6%	0.2%	1.8%
(6)	Avg. fixed costs (imports)	5.2%	3.3%	5.2%	3.3%	5.7%	3.5%
(7)	Manufacturing price index	1.000	1.058	1.000	1.053	1.000	1.063
(8)	Utility	1.000	0.965	1.000	0.966	1.000	0.963

Table shows various statistics of the equilibrium for the three alternative estimations with $\Phi = 1$ (rows 1-2), $\Phi = 3$ (rows 3-4), and $\Phi = 0.5$ (rows 5-6). Columns 1, 3, and 5 present the statistics for the baseline equilibrium, using the parameters that minimize the objective function for the given value of Φ shown in Table A.23. We normalize the price index and the utility level for each alternative to one. Columns 2, 4, and 6 present the same statistics for a counterfactual economy in which the formation of "Japanese" relationships is not possible due to $\rho \to \infty$. Rows 1-4 show the share of U.S. manufacturing sales, $P_{US}Q_{US}$, that is from China and from the rest of the world, respectively, and the share of this consumption that is sourced under the "Japanese" system. Row 5 shows the average inspection costs as a share of the import value, computed over all imports, including under the "Japanese" system. Row 6 shows the average fixed costs as a share of the import value. Row 7 shows the manufacturing price index, P_{US}, normalized to one in the baseline. Row 8 shows total utility, $W_{US} = Q_{US}^{\alpha} Z_{US}^{1-\alpha}$, normalized to one in the baseline.

J Effects on Income

A source of the welfare effects arising from changes in the arrival rate of trade wars is that "Japanese" exports generate additional income due to the incentive premium (the second term on the right-

Figure A.6: Effect of Trade War Arrival Rate on Income



Notes: The figure shows U.S. total income, i.e., wage income plus incentive premia, normalized to one for the baseline case, as a function of the trade war arrival rate ρ.

hand side in equation (16)). For exports sold under the "Japanese" system, the exporting country appropriates the incentive premium instead of having the foreign buyer country inspect the goods. As the arrival rate of trade wars rises, the number of products sourced under the "Japanese" system falls. At the same time, a higher arrival rate of trade wars increases the incentive premium for each good that is still shipped under the "Japanese" system.[69] The overall effect of these two opposing forces on U.S. income, W_n, is highlighted by Figure A.6. There is an interior point which maximizes total U.S. income, highlighting a potentially interesting avenue for trade policy. It is beneficial for a country to be associated with a lower arrival rate of trade wars, as this will allow its exporters to ship more under the "Japanese" system and to collect the incentive premium. However, as the arrival rate of trade wars becomes too low, in our estimated equilibrium the reduction of the incentive premium dominates the extensive margin effect of additional products shipped under the "Japanese" system. Thus, some trade policy uncertainty can be good to allow exporters to collect incentive premia. In our model the trade war arrival rate is symmetric for any country pair, and since importers strictly benefit from a lower arrival rate of trade wars overall welfare strictly falls with ρ, as shown earlier in Figure 6b in the main text. However, in a more general model in which $\rho_{n,i} \neq \rho_{i,n}$, a country would want to be perceived as slightly uncertain to maximize exporters' incentive income, while it would simultaneously want to commit its trade partners to never start a trade war to reduce import costs.

[69] Note from equation (3) that the incentive premium is positive even for $\rho = 0$.