Trade Policy and the Structure of Supply Chains*

Sebastian Heise†

Federal Reserve Bank of New York

Justin R. Pierce‡

Board of Governors of the Federal Reserve System

Georg Schaur§

University of Tennessee, Knoxville

Peter K. Schott¶

Yale School of Management & NBER

PRELIMINARY and INCOMPLETE

October 2016

Abstract

We model the impact of changes in trade policy on supply chains and show that a reduction in the probability of a trade war can foster the adoption of "Japanese"-style procurement practices, in which domestic buyers ensure the provision of high-quality inputs from foreign suppliers via long-term, just-in-time relationships. Empirically, we show that the model provides a useful framework for analyzing shipments between U.S. importers and foreign exporters, and find that a change in U.S. trade policy that eliminated the possibility of substantial increases in U.S. tariffs on Chinese goods coincides with a shift towards "Japanese" procurement. (JEL F13, F14, F15, F23) (Keywords: Supply Chain, Uncertainty, Trade War, Procurement)

*Schott and Heise thank the National Science Foundation (SES-1427027) for research support. We thank seminar participants at Indiana, Oregon and Toronto 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.

†33 Liberty Street, New York, NY 10045, email: sebastian.heise@ny.frb.org.

‡20th & C Streets NW, Washington, DC 20551, tel: (202) 452-2980, email: justin.r.pierce@frb.gov.

§519 Stokely Management Center, Knoxville, TN tel: (865) 974-1710, email: gschaur@utk.edu.

¶135 Prospect Street, New Haven, CT 06520, tel: (203) 436-4260, email: peter.schott@yale.edu.

1 Introduction

Motivated by the success of Japanese manufacturers such as Toyota, many firms around the world have introduced "Japanese"-style procurement practices in an effort to enhance operational efficiency.[1] Procurement is often referred to as "Japanese" if it is characterized by close, long-term relationships, which facilitate just-in-time procurement (Liker and Choi (2004)).[2] Given the increasingly global nature of firms' supply chains (Baldwin and Lopez-Gonzalez (2015)), trade policy represents a potentially important – yet under-studied – consideration in the formation of cross-border buyer-seller relationships: if buyer and seller are located in different countries, a high probability of a trade war can inhibit foreign sellers from entering into the sort of long-term relationships with domestic buyers that characterize the Japanese system.[3] This disincentive can adversely affect firm performance and welfare in several ways. For example, by causing firms to choose longer inventory replenishment cycles, buyers' inventory costs may increase.

This paper examines the role of trade policy in firms' selection of procurement systems. In the first part of the paper, we develop a model in which buyers face a trade-off between two procurement systems. Under the "Japanese" system, buyers motivate sellers to maintain product quality by committing to long-run purchases at a price above sellers' costs. The opposing "American" system, by contrast, has buyers choosing the lowest-cost seller for each order via competitive bidding, and using costly inspection to deter cheaters from shipping low quality. We demonstrate that changes in sellers' beliefs about the probability of a trade war can induce firms to switch between the American and Japanese systems. In the second part of the paper we show that our model captures key features of transaction-level U.S. import data, and demonstrate that a change in U.S. trade policy that eliminated the possibility of substantial tariff increases on imports from China coincides with a relative shift towards Japanese-style procurement between Chinese sellers and U.S. buyers. In the final part of the paper we combine the results of the empirical analysis with numerical simulations of the

[1] This movement is documented in a series of studies. See, for example, 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).

[2] More broadely, "Japanese"-style buyer-seller relationships are also characterized by joint learning and information sharing, though we do not examine these elements in this paper.

[3] A new but growing literature uses the detailed importer-exporter information in the U.S. trade data to observe the structure of supply chains and buyer-seller relationships, including Monarch (2015), Boehm, Flaaen, and Pandalai-Nayar (2015), Monarch and Schmidt-Eisenlohr (2015) and Heise (2015).

estimated model to explore potential welfare gains associated with the change in U.S. policy.

Our theoretical analysis and nomenclature are based on the framework introduced by Taylor and Wiggins (1997), where buyers under the American system pay a fixed cost per shipment to inspect goods for quality, while buyers under the Japanese system incentivize sellers to provide high-quality goods through repeated payment of a price premium over the sellers' costs. Taylor and Wiggins (1997) demonstrate that shipments between seller and buyer are optimally smaller and more frequent – i.e., more "just-in-time" – under the Japanese system. Here, we embed the Taylor and Wiggins (1997) framework in an Eaton and Kortum (2002) style general equilibrium model of trade to extend the analysis to international procurement.

We demonstrate that the higher are sellers' (exogenous) beliefs about the probability of trade peace with another country, the more likely they are to enter into Japanese-style procurement relationships with buyers from that country. The intuition for this result is straightforward: the higher the belief about the probability of trade peace, the greater the seller's confidence that a long-term relationship with a particular buyer can be sustained. This increased confidence lengthens the time horizon over which the seller expects to collect a premium over their costs from exporting their intermediate good to the buyer, driving down the premium needed to enforce honesty and thereby the relative cost of the Japanese system compared to the American system.

In our empirical analysis, we examine some of the fundamental features of the Taylor and Wiggins (1997) model using transaction-level U.S. import data. Through the lens of the model, we classify importers as using either Japanese- or American-style procurement based on the number of foreign suppliers from which they purchase goods within a product-country bin. Purchases from many suppliers are interpreted as evidence of American-style procurement while purchases from a small number or even a single supplier are deemed evidence of a Japanese-style relationship. We then show that transactions classified as American do indeed exhibit larger, less frequent shipments at lower prices. To our knowledge, these results provide the first systematic evidence supporting the key insights in Taylor and Wiggins (1997).

We then examine the core implication of our extended model, i.e., whether a change in U.S. trade policy that eliminated the possibility of a sudden spike in U.S. tariffs on Chinese goods resulted in a relative shift towards Japanese procurement. This analysis exploits variation in the exposure of U.S. import products to the U.S. extension of Per-

manent Normal Trade Relations (PNTR) to China in October 2000. Following Pierce and Schott (2016), we measure the exposure of a product to this trade liberalization as the potential jump in the tariff rate that could have occurred before the change in policy. Our triple difference-in-differences specification asks whether U.S.-China transactions within importer-exporter-product bins change after the policy is implemented (first difference) relative to bins for other countries (second difference) in products with greater exposure (third difference).[4] Consistent with the model's predictions, we find that U.S.-Chinese shipments of more-exposed products become relatively smaller, relatively more frequent, and relatively higher priced – that is, more Japanese-style – after the change in policy. Coefficient estimates suggest that a one standard deviation increase in the *ex ante* potential jump in tariff rates is associated with a relative decline in average shipment quantity of 13 percent and an increase in average shipment price of 4 percent.

In the final part of the paper we present quantitative simulations of the model that incorporate changes in shipment patterns highlighted in our empirical analysis. These simulations reveal that the change in procurement patterns induced by the policy change increases U.S. imports from China by approximately 20 percent relative to previous levels, partly at the expense of other trading partners that were not subject to the policy change. The change in procurement patterns also has implications for U.S. welfare, which increases by 0.2 percent via a decline in final goods prices. This analysis suggests that changes in trade policy can have a meaningful impact on trade flows and welfare by inducing firms to re-optimize with respect to procurement, even in the absence of other forces such as tariff changes or wage and productivity differences that are commonly associated with welfare gains from trade..

This paper makes contributions to several fields. The model we develop is to our knowledge the first to link trade policy to choice of procurement patterns, and provides an alternate perspective on the large literature examining contractual frictions in international trade.[5] A working paper by Kukharskyy and Pflüger (2010), for example, suggests that the problem of hold-up in the decision to outsource may be solved by

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

[5]See, for example, the survey by Antràs and Helpman (2008). Procurement within countries is a subject of considerable research in the industrial organization literature. See, for example, Tadelis and Zettelmeyer (2015), Cicala (2015) and Bajari et al. (2014).

relationship formation, i.e., the sharing of long-term gains in a repeated game from a sustained relationship. Here, we examine how long-term, "Japanese" relationships can overcome frictions associated with guaranteeing the provision of high-quality inputs. One attractive feature of our approach is that it yields predictions regarding shipment patterns that can be tested using transaction-level trade data.[6]

More broadly, our paper contributes to research examining the behavior of importers (e.g., Blaum, Lelarge, and Peters (2015)), the implications of trade wars (e.g., Ossa (2014)), information frictions in international trade (e.g., Cristea (2011)), trade policy uncertainty (e.g., Handley and Limão (2013), Handley (2014)), importer-exporter relationships in international trade (Heise (2015), Monarch and Schmidt-Eisenlohr (2015)), and the impact of supply-chain disruptions on output (e.g., Boehm, Flaaen, and Pandalai-Nayar (2015)).

The remainder of this paper proceeds as follows. Section 2 outlines our theoretical model. Section 3 describes the data and presents our empirical analysis. Section 4 contains our quantitative simulations. Section 5 concludes. An online appendix contains additional results.

2 Theoretical Model

Incomplete contracts, information asymmetries and contract enforcement are common problems when domestic buyers procure products from foreign suppliers. Observed organizational forms and contract structures are the result of firms optimally structuring their supply chain and procurement systems. Existing models in the international trade literature focus on the trade-offs associated with forming related party (within-firm) versus arm's-length outsourcing relationships to allocate property rights and solve hold-up problems (Antràs (2003, 2005); Antràs and Helpman (2008); Feenstra and Hanson (2005); Fisman and Wang (2010)). In the case of asymmetric information, where effort to successfully produce components by the foreign supplier is not perfectly observed, upstream integration reduces the costs of monitoring (Grossman and Helpman (2004); Spencer (2005)). Integration is therefore one solution to mitigate contract and information frictions in international trade.

[6]Our model also contrasts with existing models of heterogeneous firms and trade, in which producers balance fixed and variable costs in determining whether to export or engage in foreign direct investment (e.g., Melitz 2003, Bustos 2011). Here, as in Taylor and Wiggins (1997), however, the fixed and variable costs are endogenous to firms' choice of a procurement system.

In this paper we provide a perspective on how firms engaged in international trade can mitigate contractual and information frictions when vertical integration is not an option, perhaps due to legal barriers. Taylor and Wiggins (1997) – hereafter TW – show that in this setting firms can solve a quality control problem using one of two procurement strategies, which they label the "American" and "Japanese" systems. Under the American system, buyers use competitive bidding to select the lowest-cost supplier for each shipment, and use the threat of inspection to deter provision of low quality goods. Under the Japanese system, buyers incentivize honesty by purchasing exclusively from a single seller and indefinitely paying this seller a premium over her fixed and variable costs. TW demonstrate that shipments under the American system are larger and less frequent than under the Japanese system for two reasons. First, the fixed costs associated with inspection under the American system encourage buyers to minimize the number of orders. Second, sellers under the Japanese system have an incentive to order more frequently as a way of minimizing the payoff to a deviating seller. TW show that the optimal procurement choice depends on the ratio of the seller's fixed cost of producing each shipment to the buyer's fixed cost of inspecting each shipment. Intuitively, the lower the ratio of these fixed costs, the cheaper the Japanese system and the more likely it is to be embraced.

We use TW's model as a starting point to study how changes in trade policy affect firms' choice of procurement systems. We extend TW's model by assuming that firms evaluate future rent streams not only based on their rate of time preference, but also the likelihood of a trade war, defined as the imposition in the buyer's country of a prohibitively high tariff on the sellers' output. In this setting, trade policy can influence sellers' decisions to enter into Japanese-style procurement relationships by affecting their beliefs about the probability of a trade war taking place.

We embed this structure into a general equilibrium model of international trade, following Eaton and Kortum (2002). Our model differs from previous work mainly in that cross-country income differences arise from variation in the likelihood of a trade war across countries, which affect the choice of procurement system, rather than from firm-level productivity differences. Countries that have a relatively higher likelihood of a trade war are less likely to form Japanese-style procurement relationships since future rents are discounted more heavily, which shifts firms towards the American system and on average raises procurement costs. A change in trade policy that makes trade peace more likely reduces procurement costs by making the Japanese system

relatively cheaper, thus increasing consumer welfare. We develop this argument in this section, and test the model's empirical implications in Section 3.

2.1 Households

There exists a finite number of countries N indexed by $n = 1, ..., N$. Each country n is populated by a set of households, which purchase a continuum of goods q_n^j from final goods producers in their country. These goods are indexed by $j \in [0, 1]$. Households aggregate these varieties according to a Dixit-Stiglitz aggregator

$$Q_n = \left[\int_0^1 (q_n^j)^{(\sigma-1)/\sigma} dj\right]^{\sigma/(\sigma-1)}, \tag{1}$$

where σ is the elasticity of substitution. Households provide labor in exchange for a wage ω_n. Labor is completely immobile across countries and normalized to $L_n = 1$. We denote households' total income per period by W_n, and discuss its link with the wage below. The households' objective is to maximize their total consumption Q_n subject to the budget constraint

$$\int_0^1 p_n^j q_n^j dj \leq W_n, \tag{2}$$

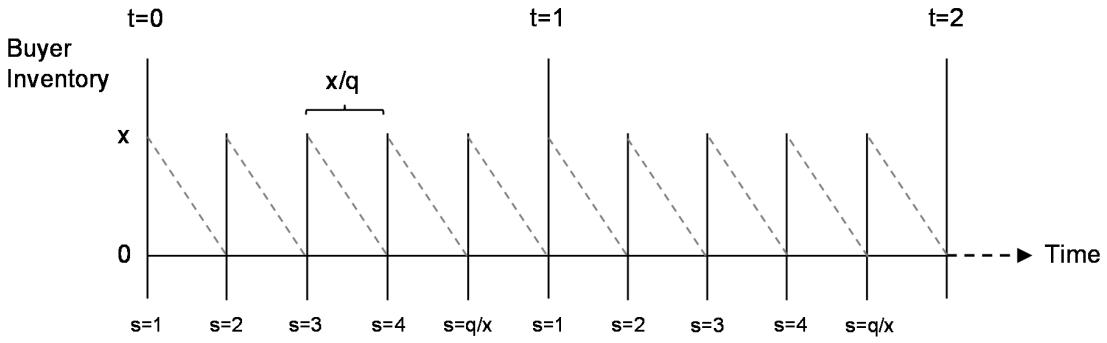
where p_n^j is the final goods price of good j in country n, taken as given by households. A household's optimal consumption choice is thus

$$q_n^j = \left(\frac{p_n^j}{P_n}\right)^{-\sigma} Q_n, \tag{3}$$

where $P_n = \left[\int_0^1 (p_n^j)^{1-\sigma}\right]^{1/(1-\sigma)}$ is country n's price index.

2.2 Firms

For each variety j, in each country n, there exists a perfectly competitive sector of final goods firms. In each period of time t, the representative firm in each country-variety purchases the quantity $q_n^j(t)$ from a seller firm, which may be located abroad, where $q_n^j(t)$ is determined by household demand in country n. Since there will be no time-varying states in our model and therefore the quantities purchased $q_n^j(t)$ are the same in each period, we omit the time index from now on, and solve for an infinitely repeated static equilibrium.



Notes: Buyer uses inputs purchased from a seller at rate *q* over each time interval, e.g., *t*=0 to *t*=1. There are *s*={1,2,..*q/x*} shipments of size *x* per time interval, each arriving *x/q* time intervals apart. Between each shipment, the buyer's inventory falls from *x* to 0.

Figure 1: Model Setup

A final goods (buyer) firm receives a given product q_n^j in each period in a series of symmetric *shipments* of size x_n^j from the seller firm. As a result, there are q_n^j/x_n^j shipments during each time interval. The size and number of these shipments are chosen optimally by the buyer. Denoting the length of each period by Δt, these shipments arrive $\Delta t/(q_n^j/x_n^j)$ time intervals apart. We normalize $\Delta t = 1$, e.g., 1 year, so that the length of time passing between shipments is x_n^j/q_n^j. This shipment pattern is visualized in Figure 1.

We assume the following timing. At the beginning of each period, households announce their consumption plan for each variety q_n^j to the final goods firms. The firms then choose their optimal shipment pattern for the period and receive shipments, incurring costs at each point a shipment arrives. To finance these expenditures, the final goods firms borrow their expected working capital requirements from (exogenous) banks at the beginning of the period at continuously compounded interest rate r. Perfect insurance markets redistribute funds between firms that did not experience a trade war and those that did, so that firms are always able to repay their loans. At the end of the period, the quantities q_n^j are shipped to households, the firms receive payment, and the working capital loan is paid off.

As we show below, buyers in each country choose to source each product j from exactly one origin country. We refer to a buyer purchasing a given product from a given seller as a *sales relationship* for that product. Sellers can choose to produce output of either low or high quality. Let $\theta \in \{\underline{\theta}, \bar{\theta}\}$ index the quality level of the output. Buyers

require high quality, e.g., an acceptably low defect rate among the units shipped.[7] The sellers' problem is to determine whether to provide high- or low-quality goods for each shipment sent to the buyer.

Buyers can inspect each shipment at a variety-specific real cost m^j per shipment before accepting and paying for it. This cost is independent of the country of origin and captures, e.g., the complexity of the product shipped. Let α be the probability that such an inspection occurs. If a buyer chooses to inspect and the quality is low, the relationship with the current seller is terminated and the seller receives no payment from the buyer. We assume that goods are specific to the buyer, so that the seller cannot sell them to an alternative partner. Furthermore, if the seller ships low-quality goods and is found out her reputation is harmed and she is excluded from the market forever. If the buyer does not inspect, the order is accepted and the seller is paid. If the order subsequently turns out to be of low quality, the relationship is terminated. In that case, the buyer cannot recover payment from the seller but can substitute contemporaneous and future orders from an alternate seller. Here, too, a seller found shipping low quality is excluded from the market forever.[8]

The Seller's Problem

We denote by x_{ni}^j the quantity shipped of product j from country i to country n. The seller produces output according to the production function

$$x_{ni}^j = \Upsilon_i l_i^j,$$

where Υ_i is the productivity level in country i, and l_i^j represents the labor input needed to produce quantity x_{ni}^j. Workers are paid a wage ω_i per unit of labor provided, which is paid after a batch is shipped. We assume that the wage costs are proportional to the quality parameter θ, capturing the fact that producing high quality goods is more costly. Thus, sellers are faced with a total real cost of $\theta\omega_i/(\Upsilon_i P_i)$ per unit of output at each shipment, capturing labor costs and a quality premium.

Finally, sellers incur a fixed cost f_{ni}^j per batch shipped, which encompasses the fixed

[7] In an extension of their basic setup, TW consider the output market into which buyers sell and have buyers choose the optimal level of θ. They show that for a sufficiently small discount rate, the optimal level of quality demanded by buyers is arbitrarily close to the first best optimal level of quality.

[8] This assumption is a simplification. In actuality, practitioners of Japanese procurement tend to reduce orders to suppliers that ship sub-standard goods but do not eliminate them unless violations are egregious or not corrected. See, for example, Liker and Choi (2004).

cost of both setting up and delivering a production run.[9] These costs can be further decomposed as

$$f_{ni}^{j} = f^{j} + d_{ni}^{j},$$

where f^j is a real cost associated with preparing product j for shipment, and d_{ni}^{j} is the average shipping cost of product j from origin country i to destination country n.

The seller receives a real order value of $\tau_{ni,s}^{j}\left(x_{ni,s}^{j}, \theta\right)$ per *shipment*, where s indicates whether the payment is under an American or Japanese system. In addition to the shipping system (s), the expression depends on the origin (i) and the destination (n) country as well as on the variety shipped (j). For simplicity, we omit these indices from now on, and simply write $\tau_s\left(x_{ni,s}^{j}, \theta\right)$. We assume the seller does not have any bargaining power and fills an order only if she at least breaks even,

$$\tau_s(x_{ni,s}^{j}, \theta) \geq f_{ni}^{j} + \theta \frac{\omega_i}{\Upsilon_i P_i} x_{ni,s}^{j}. \tag{4}$$

Firms discount payments at the exogenously given per-period interest rate r. We assume free trade between the buyer's and seller's countries, but that a trade war (i.e., a prohibitive increase in the import tariffs as in Ossa (2014)) is possible. In the event of a trade war, the import tariff on the input rises enough to sever existing buyer-seller relationships between the affected countries.[10] Sellers' exogenous belief about the probability of continued peaceful trade between countries i and n in a given period, and therefore that the relationship will be maintained, is $e^{-\rho_{ni}}$, where $\rho_{ni} > 0$ reflects the rate at which trade war shocks arrive. We assume that $\rho_{ni} = \rho_{in}$.[11] With continuous discounting, the discount factor between shipments is then given by $e^{-(r+\rho_{ni})x_{ni,s}^{j}/q_n^j}$.

Given that the stationary environment described above implies a continuous repetition of order cycles over time, the net present value to the seller of supplying shipments of $x_{ni,s}^{j}$ to the buyer as $T \to \infty$ is

[9] Thus, we ignore any transportation costs which depend on shipment size or value. We note that uncertainty over these costs may also inhibit the formation of long-term relationships.

[10] Recent research (Ossa (2014)) indicates that the optimal tariffs countries might set in the event of a trade war are substantial, averaging 63 percent worldwide.

[11] The model considers trade in a single product. An alternate interpretation of ρ_{ni} that brings the model closer to our data analysis below is that it reflects both the probability of a trade war (which is the same for all products) and the subsequent rise in tariffs (which might vary across products) for the particular good being traded. The probability of breakup is rising in the latter.

$$\frac{\tau_s(x^j_{ni,s},\theta) - f^j_{ni} - \theta\frac{\omega_i}{\Upsilon_i P_i}x^j_{ni,s}}{1-e^{-(r+\rho_{ni})x^j_{ni,s}/q^j_n}}. \tag{5}$$

As a result, the seller ships high quality ($\theta = \bar{\theta}$) if and only if expression (5) is at least as great as the one-time profit from cheating by supplying low quality ($\theta = \underline{\theta}$), i.e.,

$$\frac{\tau_s(x^j_{ni,s},\bar{\theta}) - f^j_{ni} - \bar{\theta}\frac{\omega_i}{\Upsilon_i P_i}x^j_{ni,s}}{1-e^{-(r+\rho_{ni})x^j_{ni,s}/q^j_n}} \geq (1-\alpha)\tau_s(x^j_{ni,s},\bar{\theta}) - f^j_{ni} - \underline{\theta}\frac{\omega_i}{\Upsilon_i P_i}x^j_{ni,s}. \tag{6}$$

As this expression makes clear, decreases in shipment size $x^j_{ni,s}$, as well as decreases in the arrival rate of trade war shocks, ρ_{ni}, raise the seller's discount factor, thereby strengthening the seller's incentive to provide high-quality shipments.[12]

The Buyer's Problem

The buyer chooses to conduct procurement either under the American (A) or the Japanese (J) system. Under the American system, buyers select the lowest cost supplier and use inspections to deter cheating. To simplify the problem we assume buyers under the American system always inspect while buyers in the Japanese system never inspect, so that $\alpha_A = 1$ and $\alpha_J = 0$.[13] To focus on how cross-country differences in trade peace, rather than productivity, affect shipment patterns, we also assume that $\Upsilon_i = 1$ for all i. The extension to include Υ_i is straightforward and will be analyzed in more detail in Section 4. Given our assumptions, under the American system, the seller just breaks even on each shipment,

$$\tau_A(x^j_{ni,A},\bar{\theta}) = f^j_{ni} + \bar{\theta}\frac{\omega_i}{P_i}x^j_{ni,A}. \tag{7}$$

As there is no expectation of a long-term relationship under the American system, this shipment value satisfies the seller's incentive compatibility constraint (equation (6)).

Under the Japanese system, buyers obtain seller honesty through repeat purchases and by paying sellers a premium over their costs. The shipment value under the Japanese system is

[12] An alternative approach to incorporating trade policy uncertainty would be to multiply the discount factor by an exogenous probability of trade peace $(1-\rho_{ni})$. However, a drawback of this approach is that then the probability of relationship separation over a given time period is not independent of the number of shipments made.

[13] TW show that optimal inspection under the American system is a function of shipment size and quality, $\alpha^*_A = \overline{\alpha}(x,\theta) > 0$, while under the Japanese system inspections do not occur, $\alpha^*_J = 0$.

$$\tau_J(x^j_{ni,J},\bar{\theta}) = f^j_{ni} + \bar{\theta}\frac{\omega_i}{P_i}x^j_{ni,J} + \left(\frac{1}{e^{-(r+\rho_{ni})x^j_{ni,J}/q^j_n}-1}\right)(\bar{\theta}-\underline{\theta})\frac{\omega_i}{P_i}x^j_{ni,J}. \tag{8}$$

This equation holds with strict equality given the assumption that the buyer holds all the bargaining power, but is still incentive compatible for the seller. The third term on the right hand side reflects the premium over the shipment value paid under the American system, $\tau_A(x^j_{ni,s},\bar{\theta})$, that a buyer under the Japanese system pays to incentivize the seller to sustain high quality over a long-term relationship. Intuitively, this premium rises as the rate at which trade wars arrive, ρ_{ij}, rises.

Let $\Psi(x^j_{ni,s}) = e^{-(\rho_{ni}/q^j_n)x^j_{ni,s}}$ be the per shipment probability of continued peaceful trade. We assume that in the event of a trade war, the relationship is terminated and buyers switch to an alternate supplier from a different country. We define the net present discounted cost of procurement from that alternate supplier by $\hat{C}$. The net present cost to the $s = \{A, J\}$ buyer of procuring a total quantity q^j_n from country i using batch size $x^j_{ni,s}$ for each shipment is then (excluding the interest costs on the working capital loan):

$$C_{i,s}(x^j_{ni,s}(q^j_n)) = \frac{\Psi(x^j_{ni,s})\left[\tau_s(x^j_{ni,s},\bar{\theta}) + \alpha_s m^j\right] + \left(1-\Psi(x^j_{ni,s})\right)\hat{C}}{1-e^{-(r+\rho_{ni})x^j_{ni,s}/q^j_n}}, \tag{9}$$

where we write $x^j_{ni,s}(q^j_n)$ to make explicit the dependence of the batch size of the total quantity demanded per period. The equation states that at each shipment, with probability $\Psi(x^j_{ni,s})$ no trade war has occurred and the buyer makes a payment of $\tau_s(x^j_{ni,s},\bar{\theta})$ plus, under the American system, the inspection cost. With complementary probability, the relationship ends forever and the buyer receives her outside option (cost) $\hat{C}$. The discount factor in the denominator takes into account the probability of a trade war since the payments are only made up to that point.

The outside option value $\hat{C}$ may be chosen to reflect a number of scenarios. If $\hat{C} = 0$, then a trade war forces the buyer to stop purchasing the product forever. Going forward, we assume that $\hat{C} = C_{i,s}(x^j_{ni,s}(q^j_n))$, so that the buyer is able to replace the lost relationship with a substitute that allows her to continue purchases under the same shipment pattern. However, buyers have to pay a shipping system-specific fixed cost κ_s associated with finding an alternate supplier in the event of a trade war.[14] Since

[14]The assumption that the same shipping pattern can be continued holds for example if there is a continuum of countries so that in the event of a trade war another country with an arbitrarily close f^j_{ni} and ρ_{ni} is available, and if the wage impact of a trade war in the supplier country is negligible.

the expressions are in real terms, κ_s represents the fraction of output lost in the event of a break-up. We let $\kappa_A = 0$ since under the American system buyers regularly switch suppliers to find the lowest cost supplier for each order, and denote by $\kappa \equiv \kappa_J > 0$ the cost under the Japanese system.[15] Given these assumptions, we have that

$$C_{i,s}(x^j_{ni,s}(q^j_n)) = \frac{\tau_s(x^j_{ni,s}, \bar{\theta}) + \alpha_s m^j + \left(1 - \Psi(x^j_{ni,s})\right)(1-\alpha_s)\,\kappa}{1 - e^{-r x^j_{ni,s}/q^j_n}}, \tag{10}$$

where now, since purchases are always possible, the payments in each period do not depend on the probability of a trade war and discounting takes place at rate r. We define $\delta(x^j_{ni,s}) = e^{-r x^j_{ni,s}/q^j_n}$.

The buyer's problem is to choose the optimal order size $(x^j_{ni,s})^*$, the optimal procurement system $s = \{A, J\}$, and the supplier country i for each variety j to maximize profits, taking wages ω_i, the interest rate r, and price indices P_i as given. Since buyer firms are perfect competitors in their output market, they take the market price for the final good p^j_n as given. Given prices, they also take household demand q^j_n as given, provided that marginal profits are non-negative. The buyer's discounted expected real profits of procuring quantity q^j_n of variety j in batches of $x^j_{ni,s}$ using system s from country i are thus

$$\pi^j_{ni,s}(x^j_{ni,s}(q^j_n)) = \frac{\left(\frac{p^j_n}{P_n} q^j_n\right) e^{-r}}{1 - e^{-r}} - C_{i,s}(x^j_{ni,s}) - C_{i,s}(x^j_{ni,s})(1 - e^{-r}), \tag{11}$$

where revenues $(p^j_n/P_n)q^j_n$ accrue at the end of each period when households make purchases, and are hence discounted by e^{-r}. The last term in equation (11) reflects the borrowing costs for the working capital. At the beginning of each period, buyers take up a loan of $C_{i,s}(x^j_{ni,s})(1-e^{-r})$ to cover their expected working capital requirements for the period.[16] Given interest costs of $e^r - 1$, buyers expect net present interest payments of $C_{i,s}(x^j_{ni,s})(1-e^{-r})$ at the beginning of each period when choosing their procurement strategy.

An alternative assumption is that the trade war is short-lived and hence another supplier from the same country is available after payment of κ_s. In the absence of this assumption, the wage changes arising from a trade war between any two countries will, in general equilibrium, affect all countries' sourcing. In that case, the probabilities in equation (9) must represent the probabilities of a trade war between any two countries.

[15] In Section 4 below we outline a procedure for estimating κ using transaction-level trade data.

[16] Complete insurance markets imply that borrowings are redistributed between firms that experienced a trade war and those that did not.

Since final goods prices and quantities are taken as given by buyers and the interest rate is exogenous, maximizing profits is equivalent to minimizing the cost function. In Appendix A.1, we show that the cost function is strictly convex in $x^j_{ni,s}$ if r and κ/q are small. Then, the tradeoff associated with choosing lower- versus higher-frequency procurement can be seen by setting the first order condition of the cost minimization problem to zero, yielding

$$\frac{\tau'_s(x^j_{ni,s},\bar{\theta}) - (1-\alpha_s)\,\kappa\Psi'(x^j_{ni,s})}{1-\delta(x^j_{ni,s})} = \frac{-\delta'(x^j_{ni,s})}{\left(1-\delta(x^j_{ni,s})\right)^2}\left(\tau_s(x^j_{ni,s},\bar{\theta}) + \alpha_s m^j + (1-\alpha_s)\,\kappa\left(1-\Psi(x^j_{ni,s})\right)\right), \tag{12}$$

where

$$\tau'_s(x^j_{ni,s},\bar{\theta}) = \begin{cases} \bar{\theta}\dfrac{\omega_i}{P_i} & \text{if } s = A \\ \bar{\theta}\dfrac{\omega_i}{P_i} + \left(\bar{\theta}-\underline{\theta}\right)\dfrac{\omega_i}{P_i}\left\{\left[1+\left(\dfrac{r+\rho_{ni}}{q^j_n}\right)x^j_{ni,s}\right]\dfrac{1}{e^{-(r+\rho_{ni})\frac{x^j_{ni,s}}{q^j_n}}} - 1\right\} & \text{if } s = J \end{cases} \tag{13}$$

The left hand side of equation (12) represents the discounted value of higher costs associated with a small increase in order size (i.e., a small decrease in order frequency). The right hand side measures the savings from an increased discount factor due to spacing orders further apart in time. Note that fixed order costs, f^j_{ni} – a parameter of $\tau_s(x^j_{ni,s},\bar{\theta})$ – and m^j, appear only on the right hand side of the expression: the higher these costs, the greater the benefit of raising order size (i.e., a small decrease in order frequency).

Given the real order costs $C_{i,s}(x^j_{ni,s}(q^j_n))$ at the optimal batch size under each system, buyers choose the system $s \in \{A, J\}$ for each country which minimizes costs. Define $C_i(q^j_n) = \min\left\{C_{i,A}(x^j_{ni,s}(q^j_n)^*), C_{i,J}(x^j_{ni,s}(q^j_n)^*)\right\}$ as country n's minimum cost of procuring quantity q^j_n from country i. This expression incorporates that buyers first find the cost-minimizing batch size within each system, and then find the minimum across systems. As in Eaton and Kortum (2002), buyers in each country n order each good j from the seller in the lowest-cost country. Thus, we define country n's actual procurement cost as

$$C(q_n^j) = \min\left\{C_i(q_n^j);\ i = 1, ..., N\right\}. \tag{14}$$

The actual batch size supplied will be

$$x_{ni,s}^j = \begin{cases} (x_{ni,s}^j)^* & \text{if } C(q_n^j) = C_i(q_n^j) \\ 0 & \text{if } C(q_n^j) \neq C_i(q_n^j) \end{cases}.$$

These order sizes imply per-period shipping quantity q_n^j from country i to country n if country i is the low-cost producer, and a shipping quantity of zero otherwise.

The model we propose resembles Eaton and Kortum (2002) in that each country chooses to procure each good from exactly one source country. However, while in Eaton and Kortum (2002) the choice of supplier is driven by cross-country productivity differences, we build a general equilibrium model in which not productivity differences but the probability of a trade war and the choice of optimal procurement system drive shipping patterns. As we show in the quantitative exercise in Section 4, changes in the probability of a trade war can give rise to switches in procurement strategies and optimal supplier choice which may entail significant welfare effects.

2.3 Equilibrium

Since buyers are perfect competitors in their output markets, each period they set the (real) price

$$\frac{p_n^j}{P_n} = \frac{C(q_n^j)\left(1 - e^{-r}\right)\left(2 - e^{-r}\right)}{q_n^j e^{-r}}, \tag{15}$$

which leaves them with zero expected profits in each period. We normalize prices by setting $p_n^1 = 1$ as the numeraire, for each country. Labor market clearing implies that households' labor supply equals labor demand in each country in each period,

$$1 = \sum_{n \in \tilde{N}} \sum_{j \in \tilde{J}_n} l_i^j(x_{ni,s}^j)\left(q_{ni}^j / x_{ni,s}^j\right) = \sum_{n \in \tilde{N}} \sum_{j \in \tilde{J}_n} q_{ni}^j, \tag{16}$$

where $\tilde{N}$ is the set of countries that country i exports to, $\tilde{J}_n$ is the set of products exported to country n, and q_{ni}^j is the quantity of good j purchased by country n from country i in each period. We write $l_i^j(x_{ni,s}^j)$ to emphasize that the labor supply is for the production of a batch of size x_{ni}^j. In equilibrium, each country purchases each good

j from exactly one exporter, and thus $q_{ni}^j = q_n^j$ for $j \in \tilde{J}_n$. The second equality uses the fact that $l_i^j(x_{ni,s}^j) = x_{ni,s}^j$ when $\Upsilon_i = 1$.

Finally, households' income expectations at the beginning of the period are consistent with their actual wage payments throughout the period,

$$W_i = \sum_{n \in \tilde{N}} \sum_{j \in \tilde{J}_n} \omega_i l_i^j(x_{ni,s}^j) \sum_{\tau=1}^{q_{ni}^j / x_{ni,s}^j} e^{r(1-\tau(x_{ni,s}^j / q_{ni}^j))}. \tag{17}$$

At each shipment to country n of good j, households earn income of $\omega_i l_i^j(x_{ni,s}^j)$, which is invested at interest rate r until the end of the period when purchases occur. The summation term embeds the interest income earned at the end of the period from the wages paid at each shipment.

Definition 1. A competitive equilibrium consists of a vector

$$\left\{ \left\{ x_{ni,s}^j \right\}_{n,i,j,s}, \left\{ q_{ni}^j \right\}_{n,i,j}, \left\{ p_n^j, q_n^j \right\}_{n,j}, \{\omega_n, W_n\}_n \right\}$$

such that

1. Given prices p_n^j and income W_n, households choose quantities q_n^j to maximize (1) subject to (2)
2. Given household demand q_n^j and seller country wages ω_i, buyer firms choose order size $x_{ni,s}^j$ and a procurement system to minimize costs, according to (14); this choice implies a quantity imported per period from each country q_{ni}^j
3. Prices p_n^j are such that firms make zero profits, according to equation (15), where $p_n^1 = 1$ is the numeraire
4. Wages ω_i are such that there is labor market clearing according to (16).
5. Households' income expectations W_n are consistent with their wage payments ω_n according to (17)

2.4 Optimal Choice of Procurement System

We now provide analytical results for the optimal shipment pattern and comparative statics of the model, and illustrate our findings numerically. To demonstrate the model

properties more clearly, we begin by focusing on a relationship between a single buyer and a seller, who take the quantity demanded by households q_n^j, wages ω_n, and price indices P_n as given. Given this setup with only a single relationship, we omit the country and product indices in this section. The fully calibrated quantitative general equilibrium model is presented in Section 4. Throughout the rest of this paper, we normalize $\underline{\theta} = 0$.

To derive our analytical results, in this section we assume that κ is small so that $\kappa/q \approx 0$. The numerical simulations presented in this section are based on the value of κ actually used in the quantitative analysis, and indicate that our results are not very sensitive to the choice of κ.[17]

We begin our analysis by proving that the optimal order size under the American system, x_A^*, is larger than the optimal order size in the first-best (FB) scenario, x_{FB}^*, where neither inspection nor payment premia are required to deter provision of low-quality goods.[18]. On the other hand, the optimal order size under the Japanese system, x_J^* is smaller than the optimal order size in the first-best case.

Proposition 1. *The optimal shipment size satisfies* $x_A^* > x_{FB}^* > x_J^*$.

Proof. See Appendix A.2. □

The intuition for this result is straightforward. Under the American system, the buyer encourages the seller to supply high quality by paying a fixed inspection cost at the arrival of each shipment. In order to economize on these costs, the buyer optimally orders less frequently, which leads to larger batch sizes. Under the Japanese system, on the other hand, the seller is incentivized to provide high quality via premia, which she collects at each shipment. Smaller, more frequent shipments lead to lower discounting of future rent streams, which improves the seller's incentives.

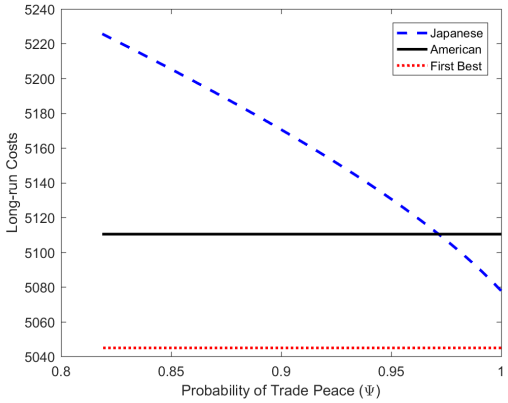
Using equation (10) at the optimal order quantity, we find after application of the envelope theorem that

$$\frac{\partial C_A(q)}{\partial \rho} = 0 \tag{18}$$

[17] In general, the switching cost acts like a (probability-weighted) fixed cost that has to be paid on each order, similar to the inspection cost under the Japanese system. When κ becomes unboundedly large, firms' desire to save on switching costs outweighs their desire to order frequently to provide incentives, and the Japanese system becomes like the American system, with κ playing the role of m in pushing buyers towards large, infrequent orders. We assume that this threshold is not met.

[18] The quantity x_{FB}^* is the solution to the problem under the American system conditional on the fixed inspection cost (m) being zero.

Figure 2: Overall Cost vs Continuation Probability (Ψ)



under the American system and

$$\frac{\partial C_J(q)}{\partial \rho} = \frac{\frac{x_J^*}{q}\frac{1}{\delta(x_J^*)}\bar{\theta}\frac{\omega}{P}x_J^*}{1-\delta(x_J^*)} + \frac{\frac{x_J^*}{q}\kappa\Psi(x_J^*)}{1-\delta(x_J^*)} > 0 \tag{19}$$

under the Japanese system. These equations highlight two key properties of our model. First, an increase in the likelihood of a trade war ρ (a decrease in the probability of trade peace) does not affect costs under American procurement, since under that system incentives are provided via inspections and there are no switching costs. Second, a higher likelihood of a trade war raises total procurement costs under the Japanese system via two channels. First, when relationships are more likely to break up, it also becomes harder for the buyer to provide incentives to the seller, forcing her to pay a higher premium over marginal costs (the first term in (19)). Second, as separations occur more frequently, buyers incur the switching cost more regularly, which drives up total procurement costs (the second term in (19)).

We illustrate this property of our model graphically in Figure 2. We impose the baseline parameters listed in Table A.1 in Appendix B.1 in this simulation. Define $\Psi \equiv e^{-\rho}$ as the per-period probability of trade peace. Figure 2 confirms that costs are unaffected by an increase in the probability of trade peace under the American system and in the first-best scenario, while they decline under the Japanese system.

Two other features of Figure 2 are worth noting. First, it shows that even if $\Psi = 1$ the cost of the Japanese system does not drop to that of the first-best scenario, The reason for this outcome is that even when trade peace is assured, the seller must be compensated for discounting if, as is the case here, $r > 0$. Second, Figure 2 reveals

that beyond some threshold level for Ψ, which we denote Ψ^{Switch} (arbitrarily equal to 0.97 in the figure), the cost of the Japanese system drops below that of the American system. At that point, buyer and seller switch from the American to the Japanese system.

Proposition 2. *For a given set of parameters and $m \geq \underline{m} > 0$, where $\underline{m}$ is a threshold value, there exists $0 < \Psi^{Switch} \leq 1$ such that the American system is chosen for $\Psi < \Psi^{Switch}$ and the Japanese system is chosen for $\Psi > \Psi^{Switch}$. If $m < \underline{m}$ then the American system is always chosen.*

Proof. See Appendix A.3. □

To see the intuition for this proposition, note that for $m = 0$ the American system corresponds to the first-best solution and is therefore preferred to Japanese-style procurement since the payment of incentive premia can be avoided. However, since the procurement cost under the Japanese system is finite when $\Psi = 1$, and since procurement costs under the American system are independent of Ψ, there must be a threshold level $\underline{m}$ such that the Japanese system is preferred once the inspection cost exceeds this threshold level. As m is increased beyond $\underline{m}$, the Japanese system becomes preferred for additional values of $\Psi < 1$ and the switching point shifts to the left.

We now turn to how the optimal shipment size under the Japanese system changes with the probability of trade peace.

Proposition 3. *The optimal order size under the Japanese system, x_J^*, increases with the probability of trade peace Ψ.*

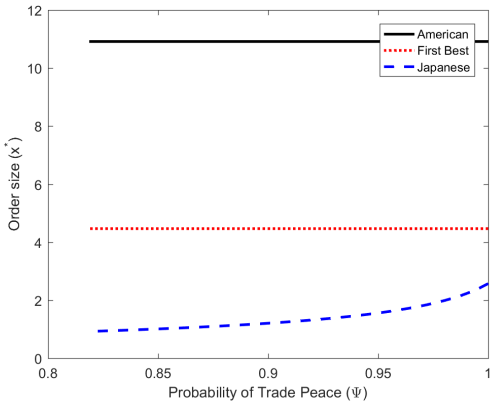
Proof. See Appendix A.4. □

We illustrate this property of the model in the left panel of Figure 3. When the probability of trade peace increases, it becomes easier for the buyer to provide incentives under the Japanese system, which enables her to economize on the shipping costs f by ordering less frequently. Under the American system and under the first-best scenario, the optimal shipment size does not depend on the probability of trade peace.

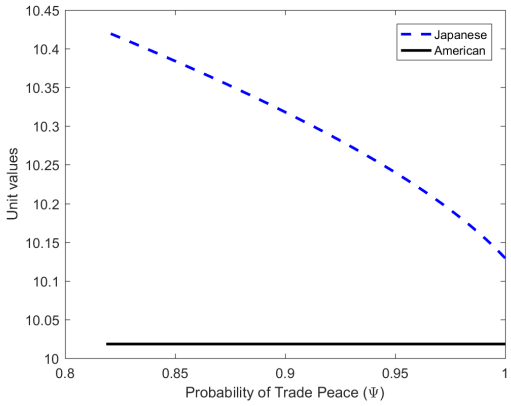
The price under the Japanese system, $\tau_J(x, \bar{\theta})/x_J^*$, falls when the probability of trade peace rises. The proof of this claim is straightforward. On the one hand, x_J^* is declining in the probability of trade peace, as shown in Proposition 3. On the other hand, $\partial C_J(q)/\partial\rho > 0$, as shown in equation (19). Since both terms in that equation are positive, it follows that $\partial\tau_J(x_J^*, \bar{\theta})/\partial\rho > 0$. This effect is shown in the right panel

Figure 3: Order Size and Price vs Continuation Probability (Ψ)

(a) Order Size



(b) Price



of Figure 3. The rise in order price as Ψ falls reflects the increase in the seller's rent necessary to incentivize her to produce high quality.

The key relationships for our empirical analysis come from joint consideration of Figures 2 and 3. Together, they reveal that if an increase in the probability of trade peace causes Ψ to jump from below Ψ^{Switch} to above this level, observed order size *falls* and observed order price *rises* as buyer and seller switch from the American to the Japanese system, i.e., from the solid black lines in the figure to the dashed blue lines. This implication of the model allows us to distinguish empirically between a change within a given procurement system and a switch of systems. The empirical results reported in Section 3 are consistent with PNTR leading to a switch to the Japanese system in U.S.-China procurement in 2001.

The comparative statics with respect to the other model parameters are summarized in the following proposition.

Proposition 4. *The optimal order size satisfies the following properties :*

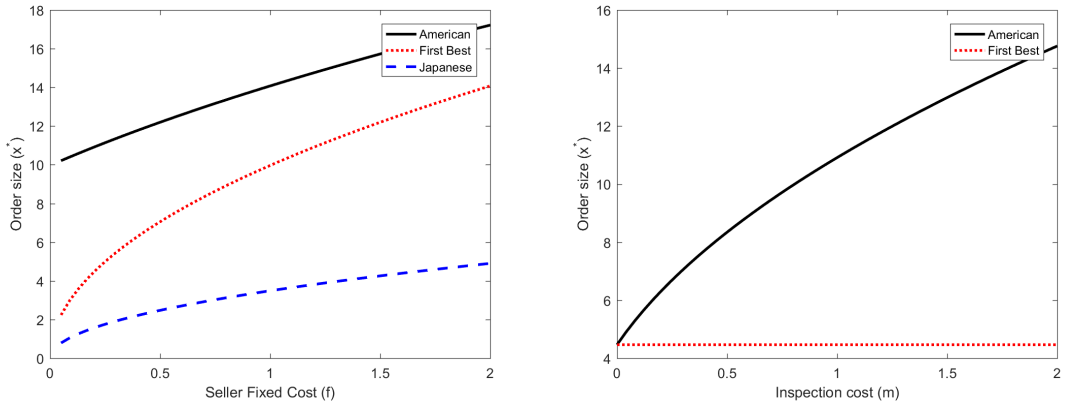
1. *Under both systems, the optimal order size x^* is increasing with the seller's fixed cost f*

2. *Under the American system, the optimal order size is increasing with the per-shipment inspection cost m*

3. *Under both systems, the optimal order size is decreasing in the marginal costs of high quality $\overline{\theta}$*

Figure 4: Order Size vs Fixed Costs (f, m)

(a) Shipment Cost f (b) Inspection Cost m



Proof. See Appendix A.5. □

The left panel of Figure 4 presents the relationship between order size and the seller's fixed cost f under both systems, while the right panel shows the relationship between order size and per-shipment inspection cost m under the American system. In both cases, buyers seek to minimize incurring larger fixed costs by reducing shipments, thereby increasing order size.

Finally, the left panel of Figure 5 shows that optimal order size under both the American and Japanese systems declines with the marginal costs of high quality $(\overline{\theta})$. As the cost to produce high quality rises, buyers have an incentive to push purchases further into the future via more frequent, smaller orders. The right panel of Figure 5 shows the effect of switching costs (κ) on order size under the Japanese system.[19] When switching costs are high, buyers seek to reduce the likelihood of being affected by a trade war by ordering larger lot sizes less frequently. The figure shows that the size of the switching cost has almost no impact on order frequency.

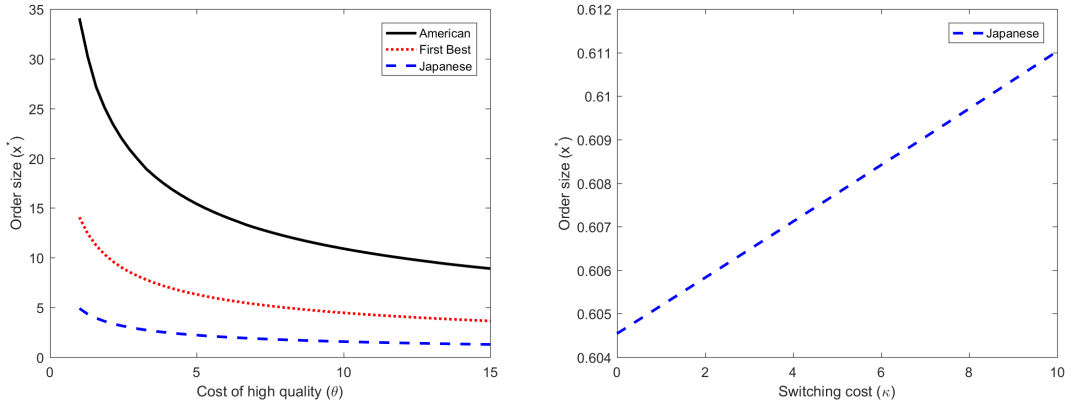
3 Empirical Analysis

We use transaction-level U.S. import data to examine the implications of our theoretical model. First, we develop a procedure for classifying U.S. importers as users of either

[19]For this figure, we have increased the likelihood of trade wars ρ to 0.5 from its baseline value of 0.05, since otherwise the effect of changing switching costs is almost zero.

Figure 5: Order Size vs Variable Costs and Switching Costs $(\bar{\theta}, \kappa)$

(a) Cost of High Quality $\bar{\theta}$ (b) Switching Cost κ



the Japanese or American procurement systems and examine whether purchases by these firms differ along the dimensions suggested by the model. We then focus on a trade policy change that increased the likelihood of continued peaceful trade between the U.S. and China and demonstrate that this change led to a shift of procurement patterns toward the Japanese system.

3.1 Description of the Data

The U.S. Census Bureau's Longitudinal Foreign Trade Transaction Database (LFTTD) tracks every U.S. import transaction from 1992 to 2011. Data available include the dates the shipment left the exporting country and arrived in the United States, identifiers for the U.S. and foreign firm conducting the trade and whether they are related or at arm's length, the transaction value and quantity, a ten-digit Harmonized System (HS) code classifying the product traded, and the country of origin of the exporter.[20]

We refine the data as follows. First, we drop all transactions that are warehouse entries, so that our dataset represents imports for consumption. Second, we remove all transactions that do not include an importer identifier, an exporter identifier, an HS code, a value, a quantity or a valid transaction date. Third, we use the procedure

[20] As noted above, import transactions are defined to be between related parties if either party owns, controls or holds voting power equivalent to 6 percent of the outstanding voting stock or shares of the other organization. We classify observations with a missing related party identifier as related. For further information on the LFTTD, see Bernard, Jensen, and Schott (2009) and Kamal, Krizan, and Monarch (2015).

suggested by Pierce and Schott (2012b) to create time-consistent HS codes, and correct an inconsistency in U.S. importing firms' identification codes over time by mapping firms in the LFTTD into the Longitudinal Business Database (LBD) and using the identifiers in the latter.[21] Fourth, we deflate transaction values using the quarterly GDP deflator from the FRED database maintained by the Federal Reserve Bank of Saint Louis. Finally, we collapse the refined version of the data by U.S. importer (m), foreign exporter (x), origin country (c), week the export left the foreign country (w) and ten-digit Harmonized System product category (h).

We summarize the importer-exporter-product relationships observed in the data along several dimensions relevant to the model presented in the previous section. After excluding triplets with just a single shipment, we compute the total shipment value across the relationship ($Value_{mxh}$), the total length of the relationship in terms of the number of weeks between the first and last observed shipment ($Length_{mxh}$) and the total number of weeks in which a shipment occurs ($Shipments_{mxh}$) during the length of the relationship. We note that $Length_{mxh}$ is potentially subject to both left and right censoring.

The averages and standard deviations of these attributes are reported in Table 1, where the left panel contains results for arm's-length (AL) relationships and the right panel shows results for related-party (RP) relationships.[22] The table highlights that the average AL relationship lasts for more than two years, with shipments on average every six weeks. The large standard deviations illustrate that there is considerable variation in the length and depth of relationships.

3.2 Classifying Japanese and American Relationships

We use the information on AL relationships from the data together with our model to classify arms'-length relationships into Japanese and American. Based on the model, importers purchasing the same HS10 product from many suppliers are more likely to be in American relationships, since such behavior is indicative of switching to the

[21] The inconsistency arises due to a change in single-unit firms' identification codes in 2002. We drop observations for invalid exporter identifiers, e.g., those that do not begin with a letter (it should start with the country name) or that have fewer than the requisite number of characters.

[22] Results for AL relationships are restricted to relationships that never report an RP shipment. Results for RP relationships encompass all other relationships. We do not summarize the prices of AL vs RP relationships due to the potential influence of transfer pricing (see Bernard, Jensen, and Schott (2006)).

Table 1: Relationship summary statistics

	Relationship Type	
	Arm's-Length	Related-Party
Total Value Traded	228,874	1,757,764
	(11,720,829)	(79,918,870)
Overall Length (Months)	32	66
	(77)	(130)
Total Number of Shipments	4	10
	(11)	(34)
Value/Shipment (VPS)	43,257	65,379
	(601,379)	(1,091,935)
Length/Shipment (LPS)	6	10
	(15)	(22)
Number of Relationships	24,138,500	7,523,500

Notes: Table reports the mean and standard deviation of each attribute across relationships, which are defined as importer by exporter ten-digit Harmonized System category triplets observed across the 1992 to 2011 sample period. First column summarizes arm's-length relationships and second column summarizes related-party relationships (see text). Observations are restricted to relationships with more than one transaction. Value, Length, and Shipments refer to the total real value of imports observed during the relationship, the duration of the relationship in weeks, and the total number of shipments observed during the relationship. Number of observations has been rounded to the nearest 100 as per U.S. Census Bureau Disclosure Guidelines.

lowest-price supplier for each transaction. On the other hand, monogamous importer-exporter relationships are more likely to be Japanese. We show that the relationships we identify as Japanese and American based on this insight indeed differ along the dimensions suggested by the model.

We begin our classification by assigning each transaction to a bin defined by the importer, the HS10 product, the country of the supplier, and the mode of transporation (z).[23] For each of these bins, we count the total number of transactions conducted across the entire dataset and the number of distinct foreign suppliers used for these. We refer to the ratio between these two variables as Suppliers per Shipment (SPS), which reflects whether an importer uses few or many suppliers to obtain a given product. The variable attains its maximum of one if an importer uses a different foreign supplier for every shipment he obtains of a given good from a given country using a given

[23]The four main modes of transporation are 10=vessel, 20=rail, 30=road, and 40=air. A small fraction of transporations transported by other means, e.g., hand-carried by passengers, is dropped.

mode of transporation. Since bins with few transactions are difficult to classify and are suggestive of importers trying out a new product or country or ordering only irregularly, we keep only importer-HS10-country-mode bins with at least $t \in \{5, 15\}$ transactions.[24] We choose to stratify the data by product, country, and mode of transporation because each of these dimensions may affect the number of suppliers used by an average U.S. importer without necessarily being related to a specific procurement system.

To determine the cutoffs for classifying a transaction as American or Japanese, we compute for each of our two samples the deciles of the SPS-distribution within HS10-mode pairs across all countries. We then classify an importer-HS10-country-mode as transacting under the American system if SPS is above the 90th percentile, and hence the importer uses a large number of suppliers relative to the average. We classify the importer as procuring under the Japanese system if SPS is below the 10th percentile. Cells that are neither American nor Japanese are dropped. As an alternative, we also perform our analyses below using the raw SPS figure. Our choice of computing the deciles across countries, rather than within countries separately, is motivated by the idea that importers may choose to form Japanese relationships with suppliers from some countries but not with suppliers from others. By defining deciles across countries, we obtain cross-country variation in the share of Japanese-style relationships.

With this classification in hand, we run a regression of the average value per shipment within each cell (VPS_{mhcz}), the average number of weeks between successive shipments (WBS_{mhcz}), and the average unit value per shipment ($Price_{mhcz}$) on an indicator of whether the importer uses the American system (d^A_{mhcz}), or alternatively on SPS_{mhcz}. The idea behind these regressions is that, based on our model, importers classified as American based on the number of suppliers should on average purchase more per shipment, order less frequently, and pay lower prices. Our regression also includes several controls. First, we include HS10-country fixed effects as well as mode of transportation fixed effects. Second, we include the total quantity transacted within the importer-HS10-country-mode cell, $Quantity_{mhcz}$. This variable captures that cells transacting more in total may transact more on each shipment, or might have systematically different prices (for example because of scale effects). Including these variables allows us to compare Japanese versus American-style importers obtaining the same total quantity of the same product, from the same country and mode of transportation. We also include the week of the first trade of the cell (beg_{mhcz}) and the week of the last

[24] We have also run regressions using cutoffs of $t = 10$ and $t = 20$. The results are very similar.

Table 2: Classification regressions at the importer level, for $t = 5$

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable	ln(VPS)	ln(WBS)	ln(Price)	ln(VPS)	ln(WBS)	ln(Price)
d^A_{mhcz}	1.221***	1.301***	−0.4803***			
	(0.003)	(0.003)	(0.005)			
$\ln(SPS_{mhcz})$				0.473***	0.499***	−0.185***
				(0.001)	(0.000)	(0.001)
$\ln(Quantity_{mhcz})$	0.756***	−0.242***	−0.355***	0.783***	−0.219***	−0.385***
	(0.001)	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)
beg_{mhcz}	0.002***	−0.002***	−0.001***	0.002***	−0.002***	−0.002***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
end_{mhcz}	−0.002***	0.002***	0.001***	−0.002***	0.002***	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	388,000	388,000	388,000	2,239,000	2,239,000	2,239,000
R-Squared	0.957	0.739	0.844	0.952	0.579	0.816
Fixed Effects	hc, z	hc, z	hc, z	hc, z	hc, z	hc, z

Notes: Superscripts *, **, and *** indicate statistical significance at the 10, 5, and 1 percent levels, respectively. Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines.

trade of the cell (end_{mhcz}) to capture time effects and effects related to the duration of importing. Thus, we estimate

$$Y_{mhcz} = \beta_0 + \beta_1 d^A_{mhcz} + \beta_2 \ln(Quantity_{mhcz}) + \beta_3 beg_{mhcz} + \beta_4 end_{mhcz} + \lambda_{hc} + \lambda_z + \epsilon_{mhcz}, \tag{20}$$

where Y_{mhcz} is the dependent variable of interest, λ_{hc} are the product-country fixed effects, and λ_z are mode of transportation fixed effects. Standard errors are clustered at the HS10-country level.

The results of the regression for the cut-off $t = 5$ are presented in Table 2. The first three columns show the coefficients of the regressions using the dummy d^A_{mhcz}. We find that, as expected, both the value per shipment and the number of weeks passed between shipments are significantly higher for importers that we classify as procuring under the American system than for importers that procure under the Japanese system. Going from an importer we classify as purchasing under the Japanese system to one that purchases under the American system raises VPS_{mhcz} and WBS_{mhcz} by more than one log point. The third column highlights that the transaction price is significantly lower under American-style procurement.

The last three columns of Table 2 show the coefficients of the regression using SPS_{mhcz} directly. In line with our previous results, we find that increasing the number of suppliers per shipment by 1% raises the value traded per shipment by 0.47% and the number of weeks between shipments by 0.50%. On the other hand, the average transaction price falls by 0.19%. Our results indicate that classifying importers based on the number of foreign suppliers per transaction – one dimension by which American and Japanese-style procurement can be distinguished – yields results for the average order size, frequency, and order price for the two groups that are consistent with our theory. Importers purchasing the same product from many suppliers order larger lot sizes less frequently and at lower prices, while importers purchasing from few suppliers obtain inputs in smaller lot sizes, more frequently, and at higher prices. The results for $t = 15$ are presented in Table A.2 in Appendix B.1 and are similar.

One shortcoming of the previous analysis is that we do not control for the identity of the exporter. Our theory predicts that suppliers under the Japanese system obtain incentive rents, which are reflected in a positive mark-up over marginal costs. Suppliers under the Japanese system should therefore charge higher prices, holding costs fixed. However, if different exporters within the same country have different costs, some suppliers might be in the Japanese system yet charge overall lower prices than suppliers under the American system due to the different cost structure. To examine this effect, we estimate equation (20) at the importer-exporter-country-product-mode level, and include exporter-product-country fixed effects. As before, the classification into American and Japanese is done for each importer-product-country-mode cell. However, the regression now investigates whether two importers purchasing the same product from the same supplier using the same mode of transportation, but under the two different systems, differ systematically. We also introduce the additional variable $Rellength_{mxhcz}$, which captures the average length of an importer-exporter-product-mode of transportation relationship in weeks, where the average is taken across the length of the relationship at each transaction. This variable captures that relationships might display different trading patterns when they are young than when they are old, regardless of the shipping system chosen (Heise (2015)). The specification is:

$$\begin{aligned} Y_{mxhcz} = \beta_0 + \beta_1 d^A_{mhcz} + \beta_2 \ln(Quantity_{mxhcz}) + \beta_3 beg_{mxhcz} + \beta_4 end_{mxhcz} \\ + \beta_5 \ln(Rellength_{mxhcz}) + \lambda_{xhc} + \lambda_z + \epsilon_{mxhcz}. \end{aligned} \tag{21}$$

We now cluster standard errors at the exporter-HS10-country level.

Table 3: Classification regressions at the importer-exporter level, for $t = 5$

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable	ln(VPS)	ln(WBS)	ln(Price)	ln(VPS)	ln(WBS)	ln(Price)
d^A_{mhcz}	0.417***	1.296***	−0.088***			
	(0.021)	(0.040)	(0.024)			
SPS_{mhcz}				0.288***	0.551***	−0.077***
				(0.001)	(0.002)	(0.001)
$\ln(Quantity_{mhcz})$	0.549***	−0.233***	−0.105***	0.659***	−0.176***	−0.151***
	(0.001)	(0.002)	(0.001)	(0.000)	(0.001)	(0.001)
beg_{mhcz}	0.000***	−0.000***	−0.000***	0.001***	0.000***	−0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
end_{mhcz}	−0.001***	0.000***	0.000***	−0.001***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\ln(Rellength_{xmhcz})$	−0.275***	0.208***	0.077***	−0.195***	0.271***	0.060***
	(0.002)	(0.005)	(0.003)	(0.001)	(0.002)	(0.001)
Observations	183,000	183,000	183,000	1,686,000	1,686,000	1,686,000
R-Squared	0.981	0.790	0.969	0.980	0.688	0.957
Fixed Effects	xhc, z	xhc, z	xhc, z	xhc, z	xhc, z	xhc, z

Notes: Superscripts *, **, and *** indicate statistical significance at the 10, 5, and 1 percent levels, respectively. Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines.

The results for $t = 5$ are presented in Table 3. They support our theory. An importer whose number of suppliers is in the 90th pecentile of the SPS distribution purchases on average 41% more per shipment and receives shipments spaced on average more than twice as far apart than does another importer who purchases from the same supplier but is in the 10th pecentile of the SPS distribution. Moreover, the supplier charges a price that is on average 8.8% lower to an importer who we classify as procuring under the American system than to an importer using the Japanese system. Our results are similar using SPS_{mhcz} directly. The results for $t = 15$ are presented in Table A.3 in Appendix B.1, and are also consistent with our predictions.

3.3 The Effect of PNTR on the Choice of Procurement System

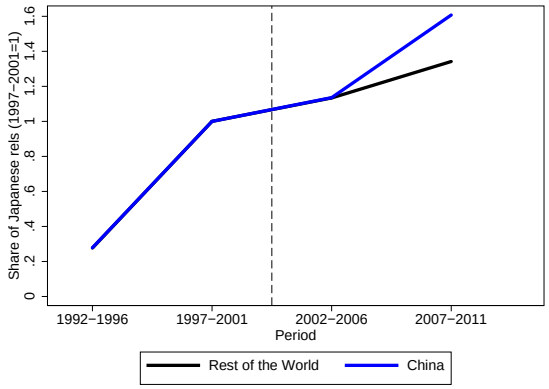
The model presented in Section 2 suggests that the share of American and Japanese procurement relationships in the economy is not constant, and can be affected by trade

policy. In particular, an increase in the probability of peaceful trade can induce buyer and seller to switch from the American to the Japanese system. Such a switch may lower procurement costs and increase consumer welfare.

We study the model's predictions for the effects of trade policy on procurement using a natural policy experiment. The U.S. granting of PNTR to China in October 2000 substantially reduced the possibility of a trade-war-like hike in U.S. import tariffs on Chinese goods. U.S. imports from non-market economies such as China are generally subject to non-NTR tariff rates originally set under the Smoot-Hawley Tariff Act of 1930. These rates are typically substantially larger than the NTR rates the U.S. offers fellow members of the World Trade Organization (WTO) – 32 percentage points higher on average, as discussed below. The U.S. Trade Act of 1974 allows the President to grant NTR tariff rates to non-market economies on an annually renewable basis subject to Congressional approval, and U.S. Presidents began granting such a waiver to China in 1980. While these waivers kept the actual tariff rates applied to Chinese goods low, the need for annual approval by Congress created uncertainty about whether the low tariffs would continue, particularly during the 1990s. In our model, this implies that the probability that relationships with U.S. firms can continue depends on the uncertainty that NTR tariffs are renewed and the amount that tariffs would increase if China's NTR status was withdrawn. A low probability of renewal associated with a high pre-NTR tariff results in a lower probability that U.S.-China relationships are long lasting.

To assess the impact of PNTR on the structure of supply chains, we begin by plotting the share of Japanese-style relationships in U.S.-China trade over time. For this purpose, we divide our dataset into four time windows: 1992-1996, 1997-2001, 2002-2006, and 2007-2011. For each of these four windows, we then compute the number of suppliers per shipment for each importer-HS10-country-mode cell, SPS_{mhcz}, keeping only cells with at least 5 transactions. For the time period 1997-2001, which runs just up to PNTR, we then compute the 10th percentile of the SPS-distribution across importers and countries for each HS10-mode bucket, and classify cells as Japanese if SPS_{mhcz} falls within the 10th percentile. By construction, across all transactions in that window, 10% of cells are Japanese. We take the values of the SPS cut-off values we found for the 1997-2001 window and apply them to the other three windows to obtain shares of Japanese relationships for these periods, which depend on how procurement patterns evolve over time. We take a value-weighted average across HS10

Figure 6: Share of Japanese-Style Relationships in U.S. Trade



codes and modes of transportation for each window, both for U.S. trade with China and with the Rest of the World.

Taking a simple average across windows, we find based on our classification methodology that Japanese-style relationships account for approximately 4.4% of all U.S.-China imports in the pre-PNTR period. This share is significantly smaller than for trade with the Rest of the World, where Japanese-style relationships account for 8.6% of trade. Reassuringly, the share of Japanese-style relationships for imports from Japan is even higher, at 11.1%. Figure 6 shows the evolution of the share of Japanese-style relationships in total trade relative to the 1997-2001 period, which we normalize to one. The figure shows a strong increase in the share of Japanese relationships over time. In the 2002-2006 period, which immediately follows PNTR, the share of Japanese relationships grew by about 20% compared to the 1997-2001 period, both for U.S.-China trade and for trade with the Rest of the World. While our theory predicts a relatively faster increase in the share of Japanese relationships with China due to PNTR, the lack of an immediate effect could be due to a large number of U.S. importers exploring importing from China and forming new relationships after 2001.[25] We find that the share of Japanese relationships grew significantly more rapidly for trade with China than for trade with the Rest of the World during the 2007-2011 period, to about 61% and 34% above the baseline level in 1997-2001, respectively.

To assess the effects of PNTR on procurement patterns for imports from China more carefully, we perform a differences-in-differences regression. We define the NTR

[25]In fact, we show below that after PNTR a large number of new relationships were formed with Chinese suppliers and these were often short-lived.

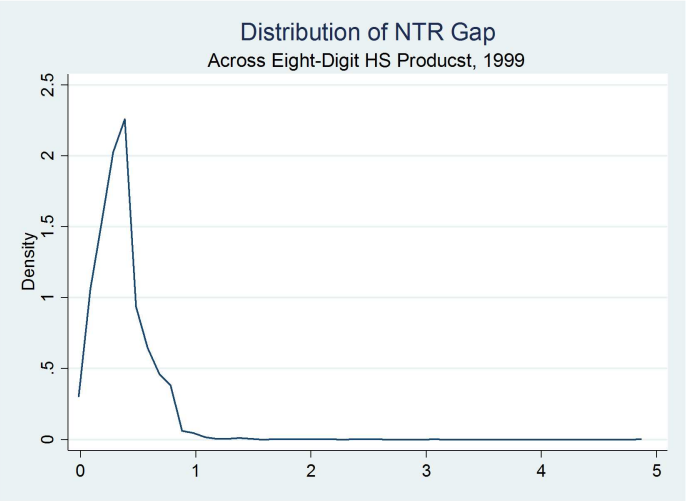


Figure 7: Distribution of the NTR Gap Across Import Products

gap for eight-digit HS import product h as the difference between non-NTR and NTR tariff rates,

$$NTR\,Gap_h = Non\,NTR\,Rate_h - NTR\,Rate_h, \tag{22}$$

using *ad valorem* equivalent tariff rates provided by Feenstra, Romalis and Schott (2002) for 1999, the year before passage of PNTR in the United States.[26] As indicated in Figure 7, these gaps vary widely across products, and have a mean and standard deviation of 0.32 and 0.23. Our identification strategy exploits this variation in the NTR gap to determine whether U.S.-China procurement patterns change relative to procurement patterns with exporters from other source countries (first difference) after the change in U.S. policy is implemented (second difference) in industries with higher NTR gaps (third difference). The last difference captures the fact that industries with larger NTR gaps experience a larger increase in the relationship continuation probability than industries with smaller gaps. We expect the largest shifts toward Japanese-style procurement after PNTR to occur in U.S. imports of high-gap products from China.

Our first, preferred specification compares shipments within importer-exporter-product triplets across two symmetric time intervals around the change in U.S. trade policy, $p \in \{Pre, Post\}$,

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

$$ln(\overline{Y}_{mxhcp}) = \beta_0 + \beta_1 1\{p = Post\} * 1\{c = China\} * NTR\,Gap_p + \gamma\chi_{mxchp} + \beta_2 ln(Total\,Value_{mxhcp}) + \lambda_{mxh} + \lambda_c + \lambda_p + \epsilon_{mxhcp} \quad (23)$$

where subscripts m, x, h and p index U.S. importers, exporters from country c, ten-digit HS products and time period. The regression sample consists of all shipments by "always-arm's-length" parties, i.e., parties that engage solely in arm's length transactions over the entire 1992 to 2011 sample period, so long as there is at least one shipment in each period. Periods are one of two distinct five-year windows around 2001, either 1995 to 2000 (pre period) or 2002 to 2007 (post period).

$\overline{Y}_{mxhcp}$ represents one of several attributes of shipment patterns within an $mxhcp$ bin deemed relevant by the model developed in Section 2: WBS_{mxhcp} is the average number of weeks between shipments, VPS_{mxhcp} is the average value per shipment, QPS_{mxhcp} is the average quantity per shipment, $Price_{mxhcp}$ is the average unit value per shipment, and $Length_{mxhcp}$ is the average length in weeks of the importer-exporter-product relationships appearing within the $mxhcp$ bin.[27] The matrix χ_{mxhcp} represents the full set of interactions of the NTR gap, the post dummy variable ($1\{p = Post\}$) and the China dummy variable ($1\{c = China\}$) required to identify β_1. $TotalValue_{mxhcp}$ is the total value of all shipments occurring within the $mxhcp$ bin; its inclusion accounts for the varying scale of imports across bins. Relationship (mxh), country and period fixed effects are represented by δ_h, δ_c and δ_p. The difference-in-differences coefficient of interest, β_1, measures the log difference in activity for shipments from China versus other countries after the change in U.S. policy versus before for products with higher versus lower NTR gaps. From the model presented in Section 2, we expect $\beta_1 < 0$ for VPS_{mxhcp}, QPS_{mxhcp} and WBS_{mxhcp}, and $\beta_1 > 0$ for $Price_{mxhcp}$ and $Length_{mxhcp}$ if PNTR induced a switch from the American to the Japanese system.[28]

The second specification ignores exporter identity and analyzes shipments within

[27] The length of each relationship is defined as the number of weeks between the first observed transaction during the period and the last observed transaction during the period.

[28] There are several motivations for why PNTR might *not* have induced procurement patterns to become more Japanese. For example, if PNTR caused only a small increase in Chinese exporters' assessment trade peace, switching from American to Japanese procurement would be minimal. As a result, import patterns for American-style importers would remain unchanged while those for Japanese-style importers would become slightly larger at a slightly higher price. Or, the most salient impact of PNTR and China joining the WTO might have been to give U.S. importers greater confidence in the enforceability of contracts. In that case, a shift from Japanese to American might be expected.

importer-products across periods,

$$\begin{aligned} ln(\overline{Y}_{mhcp}) =& \beta_0 + \beta_1 1\{p = Post\} * 1\{c = China\} * NTR\,Gap_p + \gamma\chi_{mhcp} \\ &+ \beta_2 ln(Total\,Value_{mhcp}) + \delta_{mh} + \delta_c + \delta_p + \epsilon_{mhcp} \end{aligned} \tag{24}$$

Here, too, the regression sample includes all shipments by "always-arm's-length" parties so long as there is at least one shipment for each $mhcp$ bin. After the procurement attributes are computed, the $mxhcp$ data are collapsed to the $mhcp$ level so that there is one observation – the average – in the regression for each $mhcp$ bin.

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

$$\begin{aligned} ln(\overline{Y}_{hcp}) =& \beta_0 + \beta_1 1\{p = Post\} * 1\{c = China\} * NTR\,Gap_h + \gamma\chi_{hcp} \\ &+ \beta_2 ln(Total\,Value_{hcp}) + \delta_h + \delta_c + \delta_p + \epsilon_{hcp} \end{aligned} \tag{25}$$

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

Results for the first, second and third specifications are reported in the corresponding three columns of Table 4, where each row reports the estimated DID term coefficient and standard error for a different relationship attribute. Starting with the preferred, within-mxh results reported in column 1, we find that all estimates of β_1 are consistent with a switch towards Japanese procurement: point estimates for value per shipment, quantity per shipment and weeks between shipments are all negative, though statistically significant only for the first two, while they are positive and statistically significant for shipment price and overall length. In terms of economic significance, these results imply that a one standard deviation increase in the NTR gap (0.23) is associated with relative declines in shipment value and shipment quantity of 1.6 and 3.0 percent after the change in U.S. policy. Shipment price and relationship length, by contrast, rise by 0.9 and 2.3 percent, respectively.

Comparison of the within-relationship results in column 1 with the within-product results in column 3 provides further intuition for our theoretical framework, and illuminates our findings in the initial analysis of the share of Japanese-style relationships. For example, the relatively large (in absolute terms) DID point estimates for VPS_{hcp},

Table 4: PNTR and procurement

Dependent Variable	Within Importer-Exporter-Product	Within Importer-Product	Within Product
ln(Value per Shipment)	-0.07^{***}	-0.05^{***}	-0.17^{***}
	0.01	0.01	0.05
ln(Quantity per Shipment)	-0.13^{***}	-0.04^{**}	-0.04
	0.02	0.02	0.10
ln(Price per Shipment)	0.04^{**}	-0.04^{**}	-0.11
	0.02	0.02	0.09
ln(Weeks between Shipments)	-0.04	-0.06^{***}	-0.36^{***}
	0.03	0.02	0.07
ln(Overall Relationship Length)	0.10^{***}	0.00	-0.34^{***}
	0.04	0.03	0.08
Observations	752,600	1,011,700	324,300
Sample	$mxhcp$	$mhcp$	mcp
Fixed Effects	mxh, c, p	mh, c, p	h, c, p

Notes: Table summarizes the results of generalized differences-in-differences regressions of relationship attributes on a DID coefficient representing the interaction of the NTR gap and dummy variables representing the post-PNTR period and trade with China (see text). Each cell in the table represents the result of a different regression. Data are collapsed to the importer-exporter-product-country-period ($mxhcp$) level in column 1, the importer-product-country-period ($mhcp$) level in column 2, and the product-country-period (hcp) level in column 3. Sample is restricted to bins with at least two observations for the pre- and post period. Dependent variables are computed with respect to the noted sample bins. Results for fixed effects and other covariates needed to identify the DID coefficient of interest are suppressed. Superscripts *, **, and *** indicate statistical significance at the 10, 5, and 1 percent levels, respectively. Number of observations has been rounded to the nearest 100 as per U.S. Census Bureau Disclosure Guidelines.

WBS_{hcp} and $Length_{hcp}$ reflects the fact that the change in U.S. policy gave rise to many new relationships. Since many of these relationships involved firms that had not imported from China before (see Pierce and Schott 2015), it is unsurprising that they were short-lived and perhaps encompass smaller, trial-size shipments.

4 Quantitative Analysis

In this section we quantify the welfare gains due to changes in procurement driven by an exogenous increase in the probability of trade peace. These welfare effects have two sources. First, the lower costs of Japanese-style procurement from China lead to a switch towards this procurement system for products already sourced from China under the American system. Second, new trade relationships are formed for products where China is now the lowest cost origin country.

We first discuss how we estimate the model parameters. We then simulate the model to assess the implications of PNTR. Since the model can only be solved using numerical methods, we consider a setup with $N = 3$ countries and $J = 100$ products, where $n = 1$ represents the U.S., $n = 2$ is China, and $n = 3$ is the Rest of the World.

4.1 Estimation and Identification

Estimation Strategy

We set a number of the model's parameters based on existing literature. First, we consider an annual time horizon, and set $r = 0.02$. We choose an elasticity of substitution $\sigma = 4$ as in Nakamura and Steinsson (2008), which delivers a mark-up of buyers for their final goods close to estimates by Berry, Levinsohn, and Pakes (1995). For the variable costs, we set $\bar{\theta} = 10$.

We set the switching cost κ based on the estimates in Heise (2015), who studies the average number of months it takes a U.S. importer to find a new supplier after an exogenous break-up of a relationship that has lasted 24 months or more. Heise (2015) shows that it takes the average importer about 10.7 months more than her regular order time to obtain a product after an exogenous break-up of such a long-term relationship.[29]. Since relationships of length 24 months or more account for on average 36% of imports, assuming that the replacement of a younger relationship is costless we obtain across all relationships an average time lag of about 4.0 months to find a new partner. We therefore set $k = \frac{1}{3}\bar{q}$, where $\bar{q}$ is an unweighted average over all q_{ni}^j. This parameter choice implies that in the event of a trade war, the average importer loses about one third of her annual production.

[29]Thus, if an importer purchases a product on average every 3 months, after a break-up it takes her on average 13.7 months until she purchases the product again.

Table 5: Simulation Parameters

	Baseline
Interest Rate (r)	0.02
Elasticity of Substitution (σ)	4
High Quality ($\bar{\theta}$)	10
Switching cost (κ)	$\frac{1}{3}\bar{q}$
Trade war with self ($\Psi_{i,i}$)	1
Trade war U.S.-RoW ($\Psi_{US,RoW}$)	0.98

We assume that the per-period probabilities of trade peace, $\Psi_{ni} = e^{-\rho_{ni}}$ are symmetric, and set own-country probabilities to one. Furthermore, we assume that a trade war between the U.S. and the Rest of the World is unlikely, and set the annual probability of trade peace to $\Psi_{US,RoW} = 0.98$. These parameters values are presented in Table 5.

Three sets of parameters remain to be estimated: the shipment cost parameters f^j and d^j_{ni}, the inspection costs m^j, and the remaining per-period probabilities of trade peace, Ψ_{ni}. We estimate these parameters via a simulated method of moments procedure using moments observed in the LFTTD data and in external data. While the parameters are jointly estimated, we proceed to describe the empirical moments targeted and the underlying identification assumptions for each parameter in turn.

The first set of parameters to be estimated are the remaining probabilities of trade peace. To set the pre-PNTR probability of trade peace between the U.S. and China, $\Psi^{pre}_{US,China}$, we target the observed share of Japanese-style relationships in U.S.-China trade, which was estimated in Section 3 to be 4.4%. Since the share of Japanese relationships in US-Chinese trade is strictly increasing in $\Psi_{US,China}$, this parameter is well identified for a given level of inspection costs, which we set below. We set the post-PNTR probability of trade peace, $\Psi^{post}_{US,China}$, to generate the increase in the share of Japanese-style trade between the U.S. and China that we observe in the data post PNTR. To account for the fact that the share of Japanese relationships is generally increasing across all countries, we target only the differential increase in Japanese-style relationships with China relative to the Rest of the World. Based on our previous results, we obtain an increase in the share of Japanese-style relationships due to PNTR of about 27%. We set the probability of trade peace between China and the Rest of the World to the same level as $\Psi^{pre}_{US,China}$.

We estimate parameter values for m^j, f^j, and d^j_{ni} for 18 broad product categories,

which are based on the first two digits of the HS10 code.[30] Thus, we estimate moments for 18 tuples of the form $(m^j, f^j, \left\{d_{ni}^j\right\}_{n,i})$ based on our transaction-level data. In the simulation, we will then take J draws from these product bins with probability weights in proportion to their value share in U.S. imports..

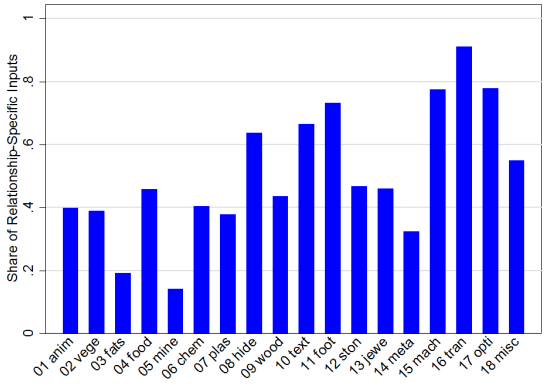
We estimate the 18 inspection cost parameters m^j using a two-step procedure. First, we estimate the average inspection cost $\bar{m}$ using the average share of Japanese relationships observed in the LFTTD across all countries in the pre-PNTR period. This average share was estimated to be 8.6% in Section 2. The parameter is well-identified since the total share of Japanese relationships was not used in the estimation of the Ψ, which used the share of Japanese relationships in U.S.-China trade only. In the second step, we obtain a distribution for m^j using the measure of contract intensity by Nunn (2007). His paper estimates for each 6-digit BEA industry code the share of intermediate inputs that is relationship-specific, where specificity is defined as the share of inputs that is neither sold on an organized exchange nor reference priced, based on the classification by Rauch (1999). Our assumption is that products that exhibit a higher degree of relationship specificity are more likely to be complex, which makes them more costly to inspect. We use the liberal classification measure provided, and map BEA codes to NAICS codes and from there to HS10 codes using the concordance by Pierce and Schott (2012a). We then aggregate the shares of relationship-specific inputs to the 18 HS2 product categories, taking a value-weighted average using the import value of each industry in 2002 from the U.S. Census.

Figure 8 shows the estimated share of relationship-specific inputs by product category. The value-weighted average share across all categories of 0.616. To obtain the distribution of m^j, we calculate the ratio between a category's relationship specificity and the mean of 0.616, and apply these ratios to the mean inspection cost $\bar{m}$.

We pin down the shipment cost parameters f^j and d_{ni}^j for the 18 product categories using the shipping frequencies observed in the data. Previous work seeking to estimate how distance affects shipment costs such as Limao and Venables (2001) and Hummels (2007) has used cost data such as shipping company quotes or air fares to estimate the elasticity of transport costs with respect to distance and other covariates. The disadvantage of this approach is that transport costs are usually based on industry

[30]These categories are: 1. Animal products (HS2=01-05), 2. Vegetables (06-14), 3. Fats (15), 4. Food (16-24), 5. Minerals (25-27), 6. Chemicals (28-38), 7. Plastics (39-40), 8. Hide (41-43), 9. Wood (44-49), 10. Textiles (50-63), 11. Footwear (64-67), 12. Stones (68-70), 13. Jewelry (71), 14. Metals (72-83), 15. Machinery (84-85), 16. Transportation (86-89), 17. Optics (90-92), 18. Miscellaneous (94-96).

Figure 8: Relationship Specificity Distribution



surveys focusing on a standardized good or broad sectors, for shipments between a few destinations. Here, we take a different approach and use the frequency of shipments to estimate an elasticity of shipment costs with respect to distance. We exploit the fact that in our model, higher values of f^j and d_{ni}^j are associated with less frequent, larger shipments (Proposition 4). Thus, we can use variation in the number of shipments used to transact the same total value of the same HS10 good via the same mode of transportation from different countries to the U.S. to determine a good-specific cost component that is independent of distance, f^j, and a distance-dependent part d_{ni}^j.

The computation of the empirical moments proceeds along several steps. First, for each of the 18 product categories, we run a regression of the form

$$\ln(WBS_{mhcz}) = \beta_0 + \beta_1 \cdot \ln(Dist_c) + \beta_2 \cdot \ln(Value_{mhcz}) + \beta_3 \cdot \ln(SPS_{mhcz}) + \lambda_h + \lambda_z + \epsilon_{mhct}, \tag{26}$$

where WBS_{mhcz} is the average number of weeks between shipments for a given importer m purchasing HS10 product h from country c using mode of transportation z, $Dist_c$ is the great circle distance between the most important city / agglomeration in the exporting country and in the U.S. provided by the CEPII GeoDist database, $Value_{mhcz}$ is the total value purchased by the importer-product-country-mode cell, λ_h are HS10 fixed effects, and λ_z are mode of transportation fixed effects. Running the regression for each product category separately seeks to account for variation in shipment patterns that is due to product-specific factors.[31] Based on our result that shipping patterns

[31] Limao and Venables (2001) and Hummels (2007) point out other variables affecting shipment costs, such as the weight of the item shipped. The split into different product categories seeks to capture variation across products in this dimension.

Table 6: Frequency moments

HS2 code	Product category	Dist (β_1)	SPS (β_3)	$\hat{WBS}^0_{10}$	$\hat{WBS}^0_{90}$
	All				
01-05	Animal products				
06-14	Vegetables				
15	Fats				
16-24	Food				
25-27	Minerals				
28-38	Chemicals				
39-40	Plastic				
41-43	Hide				
44-49	Wood				
50-63	Textiles				
64-67	Footwear				
68-70	Stones				
71	Jewelry				
72-83	Metals				
84-85	Machinery				
86-89	Transportation				
90-92	Optics				
94-96	Miscellaneous				

are sensitive to the procurement system used, we also include controls for suppliers per shipment (SPS) to compare shipments within a given procurement system. The regressions yield for each product category an estimate of the elasticity of shipment frequency with respect to distance. These regression coefficients for β_1 and β_3 and their standard errors are presented in the first two columns of Table 6.

As the second step, we construct the predicted value of $\ln(WBS_{mhcz})$ for each transaction under the assumption that $\ln(Dist_c) = 0$ and SPS is at the 10th percentile for the product category. This estimate, which we call $\ln(\hat{WBS}^0_{mhcz,10})$, is the counterfactual shipment frequency under the Japanese system if the distance for a given transaction had been zero. We similarly construct $\ln(\hat{WBS}^0_{mhcz,90})$ for the implied shipment frequency at the 90th percentile (corresponding to the American system). By taking a value-weighted average of these two variables across all transactions within each product category, we obtain product category-specific counterfactual shipment frequencies at distance zero under the Japanese system and under the American system. We set the f^j to match these 36 moments as closely as possible.

The estimation of the parameters d_{ni}^{j} proceeds similarly. We compute the distance between the U.S. and the Rest of the World as a value-weighted average distance between the largest agglomeration in the U.S. and in each exporting country, using the value of imports in 2002 from the U.S. Census as weights. We similarly estimate the distance between China and the Rest of the World using trade flows from the UN COMTRADE database as weights. For the distance of the Rest of the World with itself, we calculate the value-weighted distance between all country pairs. These distances, together with the results from regression (26), imply estimated shipment frequencies for each of the 18 product categories. We generate the frequencies again for both the American and the Japanese system, which are set as the shipment frequencies at the 10th and the 90th percentile of the SPS distribution. We choose the d_{ni}^{j} to match these implied shipment frequencies under each system.

Our procedure targets in total 2 moments for probabilities of trade peace, one moment for inspection costs, 36 moments for counterfactual shipping probabilities at distance zero, and $2 \times 4 \times 18 = 144$ moments for distances, which sums to a total of 183 moments. Let the true values of the parameters in the data be Θ, and denote the estimated parameters by $\hat{\Theta}$. We denote the vector of data moments and model moments by $G(\Theta)$ and $G(\hat{\Theta})$, respectively. We estimate a total of 93 parameters by choosing the parameter values that minimize

$$J = \min_{\hat{\Theta}} E\left[(G(\Theta) - G(\hat{\Theta}))'(G(\Theta) - G(\hat{\Theta}))\right]. \tag{27}$$

Estimation Results

Work in progress. We can show our distance cost estimates and how they compare to earlier work, etc.

4.2 Effects of a Change in the Probability of Trade Peace

[The results here are illustrative only] Given the estimated model, we are now in a position to simulate the effect of PNTR on U.S. trade flows by increasing the probability of trade peace from $\Psi_{US,China}^{pre}$ to $\Psi_{US,China}^{post}$. Table 7 presents the results of these simulations. The first three columns present statistics before PNTR, for the U.S., China, and the Rest of the World. The last three columns of the table refer to the same countries post-PNTR.

The first two rows of Table 7 illustrate the trade diversion effects of PNTR. We compute the fraction of the value imported into each region from China,

$$V_{nc} = \frac{\int p_{nc}^j q_{nc}^j}{\int p_{ni}^j q_{ni}^j},$$

where $i = c$ if the origin country is China. Prior to PNTR, the U.S. imports about half of its value from the rest of the world, and nothing from China. PNTR lowers the costs of imports from China under the Japanese system, leading some products to be switched towards Chinese suppliers. This switch raises the share of value imported from China to nearly one fifth, at the expense of imports from the rest of the world, which fall by 15 percentage points, and at the expense of domestic U.S. suppliers. The consequences of the switch of domestic production towards imports from China have been documented in Pierce and Schott (2015).

The shift towards the Japanese system is illustrated in the third row of Table 7. The value share of products sourced by U.S. importers under the Japanese system is about 46% pre-PNTR, but rises to 51% afterwards, across all countries. In line with this shift towards Japanese-style procurement, the average number of shipments, for the average product, rises in the U.S. from 0.443 shipments per period to 0.489 shipments per period (row 4). Rows 5-6 show that this increase in shipment frequency is almost solely due to products switching from American- to Japanese-style procurement. The small increase in for non-switchers is due to the increase in total quantity traded, q_{ni}^j, which makes more frequent shipments less costly. Since the probability of trade peace with China does not change for the Rest of the World, these countries do not experience a shift in their procurement patterns.

The effects of less costly procurement from China on the U.S. consumer price index is presented in row 7 of Table 7. All prices are expressed relative to the U.S. pre-PNTR price index. We find that consumer prices in the U.S. fall by 0.2% as a result of PNTR. There is also a small decrease in the Chinese price index, as Chinese importers are now also able to procure more easily under the Japanese system from the U.S. The flipside of this movement is an increase in aggregate consumption (row 8). Aggregate U.S. consumption, and hence consumer welfare, rises by 0.2% as a direct consequence of PNTR. This result highlights that non-traditional trade policies on their own can generate welfare increases, in the case of PNTR as a result of differences in the likelihood of maintaining relationships, even without the presence of fundamental differences in

Table 7: Simulated Effects of PNTR

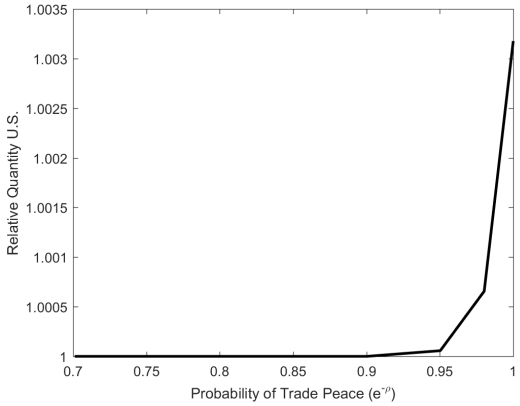
	Before PNTR			**After PNTR**		
	United States	China	Rest of World	United States	China	Rest of World
Value from U.S. (%)						
- of which, "Japanese"						
Value imported from China (%)	0.000	1.000	0.000	0.199	0.802	0.000
- of which, "Japanese"						
Value imported from ROW (%)	0.495	0.000	0.702	0.346	0.000	0.702
- of which, "Japanese"						
Fraction sourced under J	0.456	0.505	0.500	0.505	0.500	0.500
Average number of shipments	0.443	0.430	0.428	0.489	0.434	0.428
- Products switching to J system	0.182	–	–	1.078	–	–
- Products that do not switch	0.456	0.430	0.428	0.459	0.434	0.428
Aggregate price index (U.S.=1)	1.000	1.004	1.007	0.998	1.003	1.007
Aggregate quantity (U.S.=1)	1.000	0.996	0.993	1.002	0.997	0.993
Wage level (U.S.=1)	1.000	0.996	0.993	1.002	0.997	0.993

productivity.

The last row of Table 7 shows the effect of PNTR on real wages. Since labor supply is one in each country, the real wage is exactly equal to the total quantity consumed. Note that the wage level in China is slightly below the wage level in the U.S. The relatively high probability of a trade war with China deters other countries from importing Chinese goods, which necessitates a drop in Chinese real wages to stimulate exports in order to clear the labor market. However, while Chinese exporters are hurt by the higher likelihood of trade wars with the U.S., they also benefit from the fact that by symmetry it is also more difficult for U.S. exporters to export to China. Since there are no cross-country productivity differences in the model, wages in the more closed Chinese economy remain relatively similar to the wages in the U.S. If productivity differences were introduced into the model by choosing a lower value of Υ_i for China, it would become easier for foreign firms to break into the Chinese market, causing the wage level there to drop even further to equalize labor supply and demand.

Figure 9 plots aggregate consumption in the U.S. against different probabilities of trade peace, where we normalize the quantity imported when $\Psi = 0.7$ to one. The figure shows that the benefits of a higher likelihood of trade peace accrue in a non-

Figure 9: Aggregate Consumption vs. Continuation Probability (Ψ)



linear fashion. Changing the probability of trade peace from 0.7 to 0.9 has no impact on aggregate consumption, since no imports are made from China for such levels of Ψ. As the probability of trade peace increases further, however, the costs of importing from China fall more and more sharply, leading to higher and higher benefits from switching to China. Aggregate U.S. consumption when the probability of trade peace is 0.95 is only 0.01% higher than when Ψ is 0.7, while at $\Psi = 1$ the benefits exceed 0.3%. This exercise highlights that reducing the likelihood of a trade war only has significant effects if it goes all the way towards ruling out trade wars.

5 Conclusion

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

We examine the model's key implications using transaction-level U.S. import data. We begin by classifying importer-exporter relationships as American- or Japanese-

style and show that these relationships differ along the dimensions – such as shipment size, shipment frequency and shipment size – emphasized in the model. Next we the effect of the U.S. granting of Permanent Normal Trade Relations – which substantially reduced the possibility of a U.S.-China trade war – on the procurement patterns of U.S.-based firms. Using triple difference-in-differences specification, we show that PNTR is associated with a movement toward more Japanese-style procurement among U.S. importers and Chinese exporters along the dimensions highlighted by the model.

Our findings suggest that an important but under-examined aspect of trade agreements in a world with already low tariffs may be their affect 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.[32] 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.

[32]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

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 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.

Baldwin, R. E., and J. Lopez-Gonzalez (2015): "Supply-Chain Trade: A Portrait of Global Patterns and Several Testable Hypotheses," *The World Economy*, 38(11), 1682–1721.

Bernard, A. B., J. . B. Jensen, and P. K. Schott (2006): "Transfer Pricing by U.S.-Based Multinationals," NBER Working Paper No. 12493.

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.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63(4), 841–890.

Blaum, J., C. Lelarge, and M. Peters (2015): "The Gains from Input Trade in Firm-Based Models of Importing," mimeo.

Boehm, C. E., A. Flaaen, and N. Pandalai-Nayar (2015): "Multinationals, Offshoring, and the Decline of U.S. Manufacturing," mimeo, University of Michigan.

Cristea, A. D. (2011): "Buyer-Seller Relationships in International Trade: Evidence from U.S. States' Exports and Business-Class Travel," *Journal of International Economics*, 84(2), 207–220.

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.

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.

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

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

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

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

Handley, K., and N. Limão (2013): "Policy Uncertainty, Trade and Welfare: Theory and Evidence for China and the U.S.," NBER Working Paper No. 19376.

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. (2015): "Firm-to-Firm Relationships and the Cost of Recessions," mimeo, Yale.

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.

Hummels, D. (2007): "Transportation Costs and International Trade in the Second Era of Globalization," *Journal of Economic Perspectives*, 21(3), 131–154.

Kamal, F., C. J. Krizan, and R. Monarch (2015): "Identifying Foreign Suppliers in U.S. Merchandise Import Transactions," US Census Bureau Center for Economic Studies Paper No. CES-WP-15-11.

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.

Limao, N., and A. J. Venables (2001): "Infrastructure, Geographical Disadvantage, Transport Costs and Trade," *World Bank Economic Review*, 15(3), 451–479.

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.

Monarch, R. (2015): ""It's Not You, It's Me": Breakups in U.S.-China Trade Relationships," Federal Reserve Board, mimeograph.

Monarch, R., and T. Schmidt-Eisenlohr (2015): "Learning and the Value of Relationships in International Trade," SSRN 2594035.

Nakamura, E., and J. Steinsson (2008): "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, 123(4), 1415–1464.

Nunn, N. (2007): "Relationship-Specificity, Incomplete Contracts, and the Pattern of Trade," *Quarterly Journal of Economics*, 122(2), 569–600.

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 (2012a): "A Concordance between Ten-Digit U.S. Harmonized System Codes and SIC/NAICS Product Classes and Industries," *Journal of Economic and Social Measurement*, 37(1,2), 61–96.

——— (2015): "The Surprisingly Swift Decline of U.S. Manufacturing Employment," mimeo.

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

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

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

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.

Appendix

A Analytical Results

A.1 Proof that the Cost Function is Convex

American system

We omit country and product indices for simplicity. Under the American system, the cost function is given by

$$C_A(x) = \frac{f + \bar{\theta}\frac{\omega}{P}x + m}{1 - \delta(x)},$$

The second derivative of the cost function with respect to x is

$$C_A''(x) = \frac{-2\left(\frac{r}{q}\right)\delta(x)\bar{\theta}\frac{\omega}{P}}{\left[1-\delta(x)\right]^2} + \frac{\left(\frac{r}{q}\right)^2\delta(x)\left[1+\delta(x)\right]\left[f + \bar{\theta}\frac{\omega}{P}x + m\right]}{\left[1-\delta(x)\right]^3}.$$

Re-writing this, we obtain

$$C_A''(x) = \frac{\left(\frac{r}{q}\right)\delta(x)\bar{\theta}\frac{\omega}{P}\left[-2\left(1-\delta(x)\right) + \left(\frac{r}{q}\right)\left[1+\delta(x)\right]\left[x + \frac{f+m}{\bar{\theta}(\omega/P)}\right]\right]}{\left[1-\delta(x)\right]^3}.$$

Thus, the function is convex if and only if

$$\left[1+\delta(x)\right]\left[x + \tfrac{f+m}{\bar{\theta}(\omega/P)}\right] > 2\frac{1-\delta(x)}{(r/q)}.$$

Consider the case of $r \to 0$. If x/q is finite, then this condition converges to

$$2\left[x + \tfrac{f+m}{\bar{\theta}(\omega/P)}\right] > 2x,$$

where we have used the expression for $\delta(x) = e^{-(r/q)x}$ and L'Hopital's rule. This expression holds since $(f+m) > 0$.

Japanese system

Define $\tilde{\delta}(x) = e^{-(r+\rho)x/q}$. Under the Japanese system, the cost function is then given by

$$C_J(x) = \frac{f + \underline{\theta}\frac{\omega}{P}x + \frac{1}{\tilde{\delta}(x)}(\bar{\theta} - \underline{\theta})\frac{\omega}{P}x + (1 - \Psi(x))\kappa}{1 - \delta(x)}.$$

The second derivative of this cost function with respect to x is

$$\begin{aligned} C_J''(x) = & \frac{2\left(\frac{r}{q}\right)^2 \delta(x)^2 \left[f + \underline{\theta}\frac{\omega}{P}x + \frac{1}{\tilde{\delta}(x)}(\bar{\theta} - \underline{\theta})\frac{\omega}{P}x + (1 - \Psi(x))\kappa\right]}{\left[1 - \delta(x)\right]^3} \\ & + \frac{\left(\frac{r}{q}\right)^2 \delta(x) \left[f + \underline{\theta}\frac{\omega}{P}x + \frac{1}{\tilde{\delta}(x)}(\bar{\theta} - \underline{\theta})\frac{\omega}{P}x + (1 - \Psi(x))\kappa\right]}{\left[1 - \delta(x)\right]^2} \\ & - \frac{2\left(\frac{r}{q}\right)\delta(x)\left[\underline{\theta}\frac{\omega}{P} + \frac{1}{\tilde{\delta}(x)}(\bar{\theta} - \underline{\theta})\frac{\omega}{P}\left(1 + \left(\frac{r+\rho}{q}\right)x\right) + \frac{\rho}{q}\Psi(x)\kappa\right]}{\left[1 - \delta(x)\right]^2} \\ & + \frac{\left(\frac{r+\rho}{q}\right)^2 \frac{1}{\tilde{\delta}(x)}(\bar{\theta} - \underline{\theta})\frac{\omega}{P}x + 2\left(\frac{r+\rho}{q}\right)\frac{1}{\tilde{\delta}(x)}(\bar{\theta} - \underline{\theta})\frac{\omega}{P} - \left(\frac{\rho}{q}\right)^2 \Psi(x)\kappa}{1 - \delta(x)}. \end{aligned}$$

Combining terms and using the assumption that $\kappa/q \approx 0$, we obtain

$$\begin{aligned} C_J''(x) = & \frac{\left(\frac{r}{q}\right)^2 \delta(x) \left[f + \left(\underline{\theta} + \frac{1}{\tilde{\delta}(x)}(\bar{\theta} - \underline{\theta})\right)\frac{\omega}{P}x + (1 - \Psi(x))\kappa\right]\left[1 + \delta(x)\right]}{\left[1 - \delta(x)\right]^3} \\ & - \frac{2\left(\frac{r}{q}\right)\delta(x)\left[\underline{\theta}\frac{\omega}{P} + \frac{1}{\tilde{\delta}(x)}(\bar{\theta} - \underline{\theta})\frac{\omega}{P}\left(1 + \left(\frac{r+\rho}{q}\right)x\right)\right]\left[1 - \delta(x)\right]}{\left[1 - \delta(x)\right]^3} \\ & + \frac{\left(\frac{r+\rho}{q}\right)\frac{1}{\tilde{\delta}(x)}(\bar{\theta} - \underline{\theta})\frac{\omega}{P}\left[2 + \left(\frac{r+\rho}{q}\right)x\right]\left[1 - \delta(x)\right]^2}{\left[1 - \delta(x)\right]^3}. \end{aligned}$$

Hence, for the cost function to be convex, the numerator of the expression must be greater than zero. This implies the condition

$$\begin{aligned} & \delta(x)\left[f + \left(\underline{\theta} + \frac{1}{\tilde{\delta}(x)}(\bar{\theta} - \underline{\theta})\right)\frac{\omega}{P}x + (1 - \Psi(x))\kappa\right]\left[1 + \delta(x)\right] \\ & + \left(\frac{r+\rho}{q}\right)\frac{1}{\tilde{\delta}(x)}(\bar{\theta} - \underline{\theta})\frac{\omega}{P}\left[2 + \left(\frac{r+\rho}{q}\right)x\right]\frac{\left[1 - \delta(x)\right]^2}{\left(\frac{r}{q}\right)^2} \\ & > 2\delta(x)\left[\underline{\theta}\frac{\omega}{P} + \frac{1}{\tilde{\delta}(x)}(\bar{\theta} - \underline{\theta})\frac{\omega}{P}\left(1 + \left(\frac{r+\rho}{q}\right)x\right)\right]\frac{1 - \delta(x)}{\left(\frac{r}{q}\right)}. \end{aligned}$$

Taking $r \to 0$ and applying L'Hopital's rule, we obtain

$$\begin{aligned}
&2\left[f + \left(\underline{\theta} + e^{(\rho/q)x}(\bar{\theta} - \underline{\theta})\right)\frac{\omega}{P}x + (1 - \Psi(x))\kappa\right] \\
&+ \left(\tfrac{\rho}{q}\right)e^{(\rho/q)x}(\bar{\theta} - \underline{\theta})\tfrac{\omega}{P}\left[2 + \left(\tfrac{\rho}{q}\right)x\right]x^2 \\
&> 2\left[\underline{\theta}\tfrac{\omega}{P}x + e^{(\rho/q)x}(\bar{\theta} - \underline{\theta})\tfrac{\omega}{P}\left(x + \left(\tfrac{\rho}{q}\right)x^2\right)\right],
\end{aligned}$$

which simplifies to

$$2f + (1 - \Psi(x))\kappa + \left(\tfrac{\rho}{q}\right)^2 e^{(\rho/q)x}(\bar{\theta} - \underline{\theta})\tfrac{\omega}{P}x^3 > 0.$$

Since $f > 0$, this condition holds. Therefore, the cost function is convex.

A.2 Proof of Proposition 1

Proof of $x_A^* > x_{FB}^*$

For simplicity, we omit country and product indices. From equation (12) and using the expression for $\delta(x_A^*)$, we have that optimality under the American system requires

$$\frac{\bar{\theta}\frac{\omega}{P}}{1 - \delta(x_A^*)} = \frac{\frac{r}{q}\delta(x_A^*)}{(1 - \delta(x_A^*))^2}\left[f + m + \bar{\theta}\frac{\omega}{P}x_A^*\right]. \tag{A.1}$$

Re-arranging this expression yields

$$\bar{\theta}\tfrac{\omega}{P} = \frac{\frac{r}{q}\delta(x_A^*)}{1 - \delta(x_A^*)}\left[f + m + \bar{\theta}\frac{\omega}{P}x_A^*\right]. \tag{A.2}$$

Similarly, optimality under the first-best scenario, where $m = 0$, requires

$$\bar{\theta}\tfrac{\omega}{P} = \frac{\frac{r}{q}\delta(x_{FB}^*)}{1 - \delta(x_{FB}^*)}\left[f + \bar{\theta}\frac{\omega}{P}x_{FB}^*\right]. \tag{A.3}$$

Since the left-hand side of equations (A.2) and (A.3) is the same, we can set them equal and obtain

$$\delta(x_A^*)\frac{f + m + \bar{\theta}\frac{\omega}{P}x_A^*}{1 - \delta(x_A^*)} = \delta(x_{FB}^*)\frac{f + \bar{\theta}\frac{\omega}{P}x_{FB}^*}{1 - \delta(x_{FB}^*)}.$$

From this expression, we obtain the sequence of inequalities:

$$\begin{aligned}\delta(x_A^*)\frac{f+m+\bar{\theta}\frac{\omega}{P}x_A^*}{1-\delta(x_A^*)} &= \delta(x_{FB}^*)\frac{f+\bar{\theta}\frac{\omega}{P}x_{FB}^*}{1-\delta(x_{FB}^*)} \\ &< \delta(x_{FB}^*)\frac{f+\bar{\theta}\frac{\omega}{P}x_A^*}{1-\delta(x_A^*)} \\ &< \delta(x_{FB}^*)\frac{f+m+\bar{\theta}\frac{\omega}{P}x_A^*}{1-\delta(x_A^*)},\end{aligned}$$

where the first inequality follows because the fraction is exactly the cost function $C(q)$ in the first-best case from (10), which is minimized at x_{FB}^*. The second inequality follows since m is a positive constant. Comparing the left-hand side and the right-hand side of the expression yields $\delta(x_A^*) < \delta(x_{FB}^*)$, and therefore $x_A^* > x_{FB}^*$, as claimed.

Proof of $x_J^* < x_{FB}^*$

The proof proceeds along the same lines as in the case of the shipment quantities under the American system. Define $\tilde{\delta}(x_J^*) \equiv e^{-(r+\rho)x_J^*/q}$. From equation (12) and using the expression for $\tilde{\delta}(x_J^*)$ and $\underline{\theta} = 0$, we have that optimality under the Japanese system requires

$$\begin{aligned}&\frac{\frac{\omega}{P}\bar{\theta}\frac{1}{\tilde{\delta}(x_J^*)} + \frac{(r+\rho)/q}{\tilde{\delta}(x_J^*)}\bar{\theta}\frac{\omega}{P}x_J^* - \kappa\Psi'(x_J^*)}{1-\delta(x_J^*)} \\ =&\frac{\frac{r}{q}\delta(x_J^*)}{\left(1-\delta(x_J^*)\right)^2}\left[f+\bar{\theta}\frac{\omega}{P}\frac{x_J^*}{\tilde{\delta}(x_J^*)} + \kappa(1-\Psi(x_J^*))\right].\end{aligned}$$

This expression can be simplified to

$$\frac{\frac{\frac{\omega}{P}\bar{\theta}}{\tilde{\delta}(x_J^*)}\left[1+\frac{r+\rho}{q}x_J^*\right] - \kappa\Psi'(x_J^*)}{1-\delta(x_J^*)} = \frac{\frac{r}{q}\delta(x_J^*)}{\left(1-\delta(x_J^*)\right)^2}\left[f+\bar{\theta}\frac{\omega}{P}\frac{x_J^*}{\tilde{\delta}(x_J^*)} + \kappa(1-\Psi(x_J^*))\right]. \quad \text{(A.4)}$$

Re-arranging and using the expression for $\Psi'(x)$ yields

$$\bar{\theta}\frac{\omega}{P} = \frac{\frac{r}{q}\delta(x_J^*)}{(1-\delta(x_J^*))}\frac{f\tilde{\delta}(x_J^*) + \frac{\omega}{P}\bar{\theta}x_J^* + \kappa\tilde{\delta}(x_J^*)(1-\Psi(x_J^*))}{1+\frac{r+\rho}{q}x_J^*} - \frac{\kappa\frac{\rho}{q}\delta(x_J^*)}{1+\frac{r+\rho}{q}x_J^*}\Psi(x_J^*). \quad \text{(A.5)}$$

As in the case of the American system, we set equation (A.5) equal to the expression for the first-best solution (A.3) and obtain

$$\delta(x^*_{FB})\frac{f+\frac{\omega}{P}\bar{\theta}x^*_{FB}}{1-\delta(x^*_{FB})} = \delta(x^*_J)\frac{f\tilde{\delta}(x^*_J)+\frac{\omega}{P}\bar{\theta}x^*_J+\kappa\tilde{\delta}(x^*_J)(1-\Psi(x^*_J))}{(1-\delta(x^*_J))\left(1+\frac{r+\rho}{q}x^*_J\right)} - \kappa\frac{\rho}{q}\left[\frac{\tilde{\delta}(x^*_J)}{1+\frac{r+\rho}{q}x^*_J}\Psi(x^*_J)\right]..$$

Using the assumption $\kappa/q \approx 0$, the last term in the previous expression disappears. Dividing both sides by $\tilde{\delta}(x^*_J)$ then yields

$$\begin{aligned}
\frac{\delta(x^*_{FB})}{\tilde{\delta}(x^*_J)}\frac{f+\frac{\omega}{P}\bar{\theta}x^*_{FB}}{1-\delta(x^*_{FB})} &= \delta(x^*_J)\frac{f+\frac{\omega}{P}\bar{\theta}\frac{x^*_J}{\tilde{\delta}(x^*_J)}+\kappa(1-\Psi(x^*_J))}{(1-\delta(x^*_J))\left(1+\frac{r+\rho}{q}x^*_J\right)} \\
&< \delta(x^*_J)\frac{f+\frac{\omega}{P}\bar{\theta}\frac{x^*_J}{\tilde{\delta}(x^*_J)}+\kappa(1-\Psi(x^*_J))}{1-\delta(x^*_J)} \\
&< \delta(x^*_J)\frac{f+\frac{\omega}{P}\bar{\theta}\frac{x^*_{FB}}{\tilde{\delta}(x^*_{FB})}+\kappa(1-\Psi(x^*_{FB}))}{1-\delta(x^*_{FB})} \\
&= \frac{\delta(x^*_J)}{\tilde{\delta}(x^*_{FB})}\frac{f\tilde{\delta}(x^*_{FB})+\frac{\omega}{P}\bar{\theta}x^*_{FB}+\tilde{\delta}(x^*_{FB})\kappa(1-\Psi(x^*_{FB}))}{1-\delta(x^*_{FB})} \\
&< \frac{\delta(x^*_J)}{\tilde{\delta}(x^*_{FB})}\frac{f+\frac{\omega}{P}\bar{\theta}x^*_{FB}+\kappa(1-\Psi(x^*_{FB}))}{1-\delta(x^*_{FB})},
\end{aligned}$$

where the first inequality holds because $1+\frac{r+\rho}{q}x^*_J > 1$, and the second inequality follows from the fact that the fraction is exactly the cost function $C(q)$ in the Japanese system from (10), which is minimized at x^*_J. The final inequality holds because $\tilde{\delta}(x^*_{FB}) < 1$. If κ is sufficiently small, then the term involving κ is negligible and can be disregarded (note that the probability of a trade war multiplying κ is likely also small). In that case, comparing the left-hand side and the right-hand side yields the condition

$$\tilde{\delta}(x^*_{FB})\delta(x^*_{FB}) < \tilde{\delta}(x^*_J)\delta(x^*_J),$$

from which it follows immediately that $x^*_J < x^*_{FB}$. Thus, if there are no switching costs the quantity ordered under the Japanese system is always smaller than the quantity ordered under the American system. For small switching costs, the relationship still holds, but if κ becomes too large then the desire to save on switching costs outweighs

the advantage from ordering more frequently to provide incentives, and firms order less frequently than under first-best. In that case the Japanese system becomes likely the American system.

A.3 Proof of Proposition 2

Consider the case of the American system with $m = 0$. Given per period demand q, the net present value of costs under the optimal order quantity x_A^* is

$$C_A(x_A^*) = \frac{f + \frac{\omega}{P}\bar{\theta}x_A^*}{1 - \delta(x_A^*)}.$$

Since this cost function is the same as in the first-best case, the American solution corresponds to the first-best solution and therefore $x_A^* = x_{FB}^*$ and $C_A(x_A^*) = C_{FB}(x_{FB}^*)$.

Consider now the Japanese system with $\rho = 0$, and hence $\Psi = 1$. As demonstrated in the main text, costs are declining in Ψ and thus are lowest under the Japanese system when $\Psi = 1$. Furthermore, assume no switching costs, and hence $\kappa = 0$. The costs under the Japanese system must satisfy

$$C_J(x_J^*) = \frac{f + \frac{\omega}{P}\bar{\theta}\frac{x_J^*}{\delta(x_J^*)}}{1 - \delta(x_J^*)} > \frac{f + \frac{\omega}{P}\bar{\theta}x_J^*}{1 - \delta(x_J^*)} > \frac{f + \frac{\omega}{P}\bar{\theta}x_{FB}^*}{1 - \delta(x_{FB}^*)} = C_{FB}(x_{FB}^*) = C_A(x_A^*), \quad \text{(A.6)}$$

where the first inequality holds because $\delta(x_J^*) < 1$ since $r > 0$, and the second inequality holds because x_J^* is not the cost-minimizing batch size in the first-best cost function. Hence, costs under the Japanese system are strictly greater than under the American system. For $\kappa > 0$, costs under the Japanese system are even greater and therefore must also be higher than under the American system.

Since the cost function under the American system is monotonely increasing in m, there must exist $\underline{m}$ such that for $\rho = 0$ the inequalities in (A.6) become an equality:

$$\frac{f + \frac{\omega}{P}\bar{\theta}\frac{x_J^*}{\delta(x_J^*)}}{1 - \delta(x_J^*)} = \frac{f + \frac{\omega}{P}\bar{\theta}x_A^* + \underline{m}}{1 - \delta(x_A^*)}. \quad \text{(A.7)}$$

This equation implicitly defines $\underline{m}$. If $\kappa > 0$, the threshold level must increase. Finally, if $m > \underline{m}$, then the left-hand side of equation (A.7) must be strictly smaller than the right-hand side. Since the costs under the Japanese system are increasing in ρ (decreasing in Ψ), and since costs under the Japanese system diverge to infinity as

$\rho \to \infty$, for any finite $m > \underline{m}$ there must exist a Ψ^{Switch} such that the costs under both system are equal.

A.4 Proof of Proposition 3

Recall that the probability of trade peace Ψ is inversely related to the arrival rate of trade wars ρ. Define $\tilde{\delta}(x_J^*) \equiv e^{-(r+\rho)x_J^*/q}$. Applying the implicit function theorem to the optimality condition (12) under the Japanese system, we obtain that $dx_J^*/d\rho = A/B$, where the numerator A is given by

$$A = \bar{\theta}\frac{\omega}{P}\frac{x_J^*}{q} + r\frac{x_J^*}{q}\delta(x_J^*)\tilde{\delta}(x_J^*)\frac{\frac{f}{q} + \frac{\kappa}{q} - 2\frac{\kappa}{q}\Psi(x_J^*)}{1-\delta(x_J^*)} + \left[\frac{q - 2\rho x_J^*}{q}\right]\frac{\kappa}{q}\tilde{\delta}(x_J^*)\Psi(x_J^*), \quad \text{(A.8)}$$

and the denominator B equals

$$B = -C + D - E - F + G + H - I, \quad \text{(A.9)}$$

where

$$C = \bar{\theta}\frac{\omega}{P}\frac{r+\rho}{q}, \quad D = \frac{\rho(2\rho + r)}{q}\frac{\kappa}{q}\tilde{\delta}(x_J^*)\Psi(x_J^*),$$

$$E = \frac{\left(\frac{r}{q}\right)^2 \delta(x_J^*)\left[f\tilde{\delta}(x_J^*) + \bar{\theta}\frac{\omega}{P}x_J^* + \kappa(1-\tilde{\delta}(x_J^*))\Psi(x_J^*)\right]}{1-\delta(x_J^*)},$$

$$F = \frac{\frac{r}{q}\tilde{\delta}(x_J^*)\delta(x_J^*)(r+\rho)\left(\frac{f+\kappa}{q}\right)}{1-\delta(x_J^*)}, \quad G = \frac{\frac{r}{q}\delta(x_J^*)\bar{\theta}\frac{\omega}{P}}{1-\delta(x_J^*)},$$

$$H = \frac{\frac{r}{q}\delta(x_J^*)\frac{\kappa}{q}(2\rho + r)\tilde{\delta}(x_J^*)\Psi(x_J^*)}{1-\delta(x_J^*)},$$

$$I = \frac{\left(\frac{r}{q}\right)^2 [\delta(x_J^*)]^2\left[f\tilde{\delta}(x_J^*) + \bar{\theta}\frac{\omega}{P}x_J^* + \kappa(1-\tilde{\delta}(x_J^*))\Psi(x_J^*)\right]}{[1-\delta(x_J^*)]^2}.$$

Under the assumption that $\kappa/q \approx 0$, term A simplifies to

$$A = \bar{\theta}\frac{\omega}{P}\frac{x_J^*}{q} + r\frac{x_J^*}{q}\delta(x_J^*)\tilde{\delta}(x_J^*)\frac{\frac{f}{q}}{1-\delta(x_J^*)} > 0.$$

To show that $dx_J^*/d\rho < 0$, it therefore remains to prove that B is negative. With $\kappa/q \approx 0$, terms D and H are approximately zero, and term F becomes

$$F = \frac{\frac{r}{q}\tilde{\delta}(x_J^*)\delta(x_J^*)(r+\rho)\frac{f}{q}}{1-\delta(x_J^*)}.$$

Thus, we need to show that

$$C + E + F + I > G. \tag{A.10}$$

Since r is small and $x < q$, we can use the Taylor approximation $\delta(x_J^*) \approx 1 - \frac{r}{q}x_J^*$. Applying this approximation, we obtain

$$E \approx \frac{\frac{r}{q}\delta(x_J^*)\left[f\tilde{\delta}(x_J^*) + \bar{\theta}\frac{\omega}{P}x_J^*\right]}{x_J^*},$$

where the κ-term disappears because $\kappa/q \approx 0$,

$$F \approx \frac{\tilde{\delta}(x_J^*)\delta(x_J^*)(r+\rho)\frac{f}{q}}{x_J^*},$$

$$G \approx \frac{\bar{\theta}\frac{\omega}{P}}{x_J^*} - \bar{\theta}\frac{\omega}{P}\frac{r}{q},$$

and

$$\begin{aligned} I \approx & \frac{f\tilde{\delta}(x_J^*) + \bar{\theta}\frac{\omega}{P}x_J^* + \kappa(1-\tilde{\delta}(x_J^*))\Psi(x_J^*)}{(x_J^*)^2} \\ & + \left(\frac{r}{q}\right)^2\left[f\tilde{\delta}(x_J^*) + \bar{\theta}\frac{\omega}{P}x_J^* + \kappa(1-\tilde{\delta}(x_J^*))\Psi(x_J^*)\right] \\ & - 2\frac{r}{q}\left[\frac{f\tilde{\delta}(x_J^*)}{x_J^*} + \bar{\theta}\frac{\omega}{P}\right]. \end{aligned}$$

Adding together the approximated terms E and F yields

$$E + F = 2\frac{\left(\frac{r}{q}\right)\tilde{\delta}(x_J^*)\delta(x_J^*)f}{x_J^*} + \frac{\left(\frac{\rho}{q}\right)\tilde{\delta}(x_J^*)\delta(x_J^*)f}{x_J^*} + \frac{r}{q}\delta(x_J^*)\bar{\theta}\frac{\omega}{P}.$$

We can apply another approximation to the first term of this expression to obtain

$$2\frac{\left(\frac{r}{q}\right)\tilde{\delta}(x_J^*)\delta(x_J^*)f}{x_J^*} \approx 2\frac{\left(\frac{r}{q}\right)\tilde{\delta}(x_J^*)f}{x_J^*} - 2\left(\frac{r}{q}\right)^2\tilde{\delta}(x_J^*)f.$$

Plugging all these expressions into the condition (A.10) gives

$$\begin{aligned}\frac{\bar{\theta}\frac{\omega}{P}}{x_J^*} - \bar{\theta}\frac{\omega}{P}\frac{r}{q} <& \bar{\theta}\frac{\omega}{P}\frac{r+\rho}{q} + 2\left(\frac{r}{q}\right)\frac{\tilde{\delta}(x_J^*)f}{x_J^*} - 2\left(\tfrac{r}{q}\right)^2\tilde{\delta}(x_J^*)f \\ &+\frac{\left(\frac{\rho}{q}\right)\tilde{\delta}(x_J^*)\delta(x_J^*)f}{x_J^*} + \tfrac{r}{q}\delta(x_J^*)\bar{\theta}\tfrac{\omega}{P} \\ &+\frac{f\delta(x_J^*) + \kappa(1-\tilde{\delta}(x_J^*))\Psi(x_J^*)}{(x_J^*)^2} + \frac{\bar{\theta}\frac{\omega}{P}}{x_J^*} \\ &+\left(\frac{r}{q}\right)^2\left[\tilde{\delta}(x_J^*)f + \bar{\theta}\tfrac{\omega}{P}x_J^* + \kappa(1-\tilde{\delta}(x_J^*))\Psi(x_J^*)\right] \\ &-2\left(\frac{r}{q}\right)\left[\frac{\tilde{\delta}(x_J^*)f}{x_J^*} + \bar{\theta}\tfrac{\omega}{P}\right].\end{aligned}$$

Cancelling terms, the expression simplifies to

$$\begin{aligned}0 <& \bar{\theta}\frac{\omega}{P}\frac{\rho}{q} - \left(\frac{r}{q}\right)^2\tilde{\delta}(x_J^*)f + \frac{\left(\frac{\rho}{q}\right)\tilde{\delta}(x_J^*)\delta(x_J^*)f}{x_J^*} + \tfrac{r}{q}\delta(x_J^*)\bar{\theta}\tfrac{\omega}{P} \\ &+\frac{1}{(x_J^*)^2}\tilde{\delta}(x_J^*)f + \frac{\kappa(1-\tilde{\delta}(x_J^*))\Psi(x_J^*)}{(x_J^*)^2} + \left(\frac{r}{q}\right)^2\left[\bar{\theta}\tfrac{\omega}{P}x_J^* + \kappa(1-\tilde{\delta}(x_J^*))\Psi(x_J^*)\right].\end{aligned}$$

Note that only the second term in this expression is negative. A sufficient condition for it to hold is

$$\frac{1}{(x_J^*)^2} > \frac{r^2}{q^2}.$$

Since $r < 1$and $q > x_J^*$, this condition is satisfied and hence $dx_J^*/d\rho < 0$, as claimed.

A.5 Proof of Proposition 4

1a. Proof of $dx_A^*/df > 0$

From the implicit function theorem, we have that

$$\frac{dx_A^*}{df} = \frac{\frac{r}{q}\delta(x_A^*)}{B},$$

where

$$B = \left(\frac{r}{q}\right)^2 \delta(x_A^*)\left[f + m + \tfrac{\omega}{P}\bar{\theta}x_A^*\right] - \left(\frac{r}{q}\right)\delta(x_A^*)\tfrac{\omega}{P}\bar{\theta} + \frac{1}{1-\delta(x_A^*)}\left(\frac{r}{q}\right)^2 [\delta(x_A^*)]^2\left[f + m + \tfrac{\omega}{P}\bar{\theta}x_A^*\right].$$

Dividing B by the numerator $\frac{r}{q}\delta(x_A^*)$ and combining the first and the last term in B yields $dx_A^*/df = 1/C$, where

$$C = \frac{1}{1-\delta(x_A^*)}\left(\frac{r}{q}\right)\left[f + m + \tfrac{\omega}{P}\bar{\theta}x_A^* - (1-\delta(x_A^*))\tfrac{q}{r}\tfrac{\omega}{P}\bar{\theta}\right].$$

Using the approximation $\delta(x_A^*) \approx 1 - \frac{r}{q}x_A^*$, the term in parentheses can be simplified to yield

$$C = \frac{1}{1-\delta(x_A^*)}\left(\frac{r}{q}\right)[f+m].$$

Therefore, we obtain

$$\frac{dx_A^*}{df} = \frac{1}{\frac{1}{1-\delta(x_A^*)}\left(\frac{r}{q}\right)[f+m]} > 0. \tag{A.11}$$

1b. Proof of $dx_J^*/df > 0$

Using the implicit function theorem and equation (12), we obtain that

$$\frac{dx_J^*}{df} = \frac{\tilde{\delta}(x_J^*)}{B},$$

where

$$B = \left(\frac{r+\rho}{q}\right)\tfrac{\omega}{P}\bar{\theta}\left(\frac{q}{r}\right)\frac{1-\tilde{\delta}(x_J^*)}{\tilde{\delta}(x_J^*)} - \rho\frac{\kappa}{q}\Psi(x)\delta(x)\left(\frac{2\rho+r}{r}\right)\frac{1-\tilde{\delta}(x_J^*)}{\tilde{\delta}(x_J^*)} + \frac{r}{q}\frac{1}{1-\tilde{\delta}(x_J^*)}\left[f\delta(x_J^*) + \tfrac{\omega}{P}\bar{\theta}x_J^* + \kappa\left(1-\Psi(x_J^*)\right)\delta(x_J^*)\right] + f\left(\frac{r+\rho}{q}\right)\delta(x_J^*) - \tfrac{\omega}{P}\bar{\theta} + \frac{\kappa}{q}\delta(x_J^*)\left[\rho + r - (2\rho+r)\Psi(x_J^*)\right].$$

Using the assumption that $\kappa/q \approx 0$, the expression simplifies to

$$B = \left(\frac{r+\rho}{q}\right)\frac{\omega}{P}\bar{\theta}\left(\frac{q}{r}\right)\frac{1-\tilde{\delta}(x_J^*)}{\tilde{\delta}(x_J^*)} + f\left(\frac{r+\rho}{q}\right)\delta(x_J^*) - \frac{\omega}{P}\bar{\theta} + \frac{r}{q}\frac{1}{1-\tilde{\delta}(x_J^*)}\left[f\delta(x_J^*) + \frac{\omega}{P}\bar{\theta}x_J^* + \kappa\left(1-\Psi(x_J^*)\right)\delta(x_J^*)\right].$$

Using the approximation $\tilde{\delta}(x_J^*) \approx 1 - \frac{r}{q}x_J^*$, the expression becomes

$$B = \left(\frac{r+\rho}{q}\right)\frac{\omega}{P}\bar{\theta}\left(\frac{q}{r}\right)\frac{1-\tilde{\delta}(x_J^*)}{\tilde{\delta}(x_J^*)} + f\left(\frac{r+\rho}{q}\right)\delta(x_J^*) + f\frac{\delta(x_J^*)}{x_J^*} + \kappa\left(1-\Psi(x_J^*)\right)\frac{\delta(x_J^*)}{x_J^*} > 0.$$

Thus, both the numerator and the denominator are positive, and therefore $dx_J^*/df > 0$, as claimed.

2. Proof of $dx_A^*/dm > 0$

Since f and m appear in equation (12) in the same way, the proof follows exactly the same steps as the proof that $dx_A^*/df > 0$. Thus, $dx_A^*/dm > 0$.

3a. Proof of $dx_A^*/d\bar{\theta} < 0$

Applying the implicit function theorem to equation (12), we obtain

$$\frac{dx_A^*}{d\bar{\theta}} = \frac{1 - \frac{\frac{r}{q}\delta(x_A^*)}{1-\delta(x_A^*)}x_A^*}{C},$$

where

$$C = -\frac{\left(\frac{r}{q}\right)^2\delta(x_A^*)}{\left[1-\delta(x_A^*)\right]^2}\left[f + \frac{\omega}{P}\bar{\theta}x_A^* + m\right] + \frac{\left(\frac{r}{q}\right)\delta(x_A^*)}{\left[1-\delta(x_A^*)\right]}\frac{\omega}{P}\bar{\theta}.$$

Dividing both the numerator and the denominator by $\left(\frac{r}{q}\right)\delta(x_A^*)/\left[1-\delta(x_A^*)\right]$ yields

$$\frac{dx_A^*}{d\bar{\theta}} = \frac{\frac{1-\delta(x_A^*)}{\frac{r}{q}\delta(x_A^*)} - x_A^*}{-\frac{(r/q)}{1-\delta(x_A^*)}\left[f + \frac{\omega}{P}\bar{\theta}x_A^* + m\right] + \frac{\omega}{P}\bar{\theta}}.$$

Using the approximation $\delta(x_A^*) \approx 1 - \frac{r}{q}x_A^*$ to the $1 - \delta(x_A^*)$ terms, we obtain

$$\frac{dx_A^*}{d\bar{\theta}} \approx \frac{\frac{x_A^*}{\delta(x_A^*)} - x_A^*}{-\frac{f+m}{x_A^*}} = -\frac{(1-\delta(x_A^*))\,(x_A^*)^2}{(f+m)\,\delta(x_A^*)} < 0,$$

as required.

3b. Proof of $dx_J^*/d\bar{\theta} < 0$

Applying the implicit function theorem to equation (12) yields

$$\frac{dx_J^*}{d\bar{\theta}} = \frac{1 + \left(\frac{r+\rho}{q}\right)x_J^* - \frac{\frac{r}{q}\delta(x_J^*)}{1-\delta(x_J^*)}x_J^*}{D},$$

where

$$\begin{aligned} D = &-\left(\frac{r+\rho}{q}\right)\frac{\omega}{P}\bar{\theta} + \rho\frac{\kappa}{q}\Psi(x)\tilde{\delta}(x)\left(\frac{2\rho+r}{q}\right) \\ &- \frac{\left(\frac{r}{q}\right)^2\delta(x_J^*)}{[1-\delta(x_J^*)]^2}\left[f\tilde{\delta}(x_J^*) + \tfrac{\omega}{P}\bar{\theta}x_J^* + \kappa(1-\Psi(x_J^*))\tilde{\delta}(x_J^*)\right] \\ &- \frac{\left(\frac{r}{q}\right)\delta(x_J^*)}{[1-\delta(x_J^*)]}\left[f\left(\frac{r+\rho}{q}\right)\tilde{\delta}(x_J^*) - \tfrac{\omega}{P}\bar{\theta}x_J^*\right] \\ &- \frac{\left(\frac{r}{q}\right)\delta(x_J^*)}{[1-\delta(x_J^*)]}\frac{\kappa}{q}\tilde{\delta}(x_J^*)\left[\rho + r - (2\rho+r)\,\Psi(x_J^*)\right]. \end{aligned}$$

Using the fact that $\kappa/q \approx 0$ and multiplying both the numerator and the denominator by $(r/q)\delta(x_J^*)/1-\delta(x_J^*)$, we obtain

$$\frac{dx_J^*}{d\bar{\theta}} = \frac{\left[1 + \left(\frac{r+\rho}{q}\right)x_J^*\right]\frac{1-\delta(x_J^*)}{\frac{r}{q}\delta(x_J^*)} - x_J^*}{E},$$

where

$$\begin{aligned} E = &-\left(\frac{r+\rho}{q}\right)\frac{\omega}{P}\bar{\theta}\frac{1-\delta(x_J^*)}{\frac{r}{q}\delta(x_J^*)} - f\left(\frac{r+\rho}{q}\right)\tilde{\delta}(x_J^*) + \tfrac{\omega}{P}\bar{\theta}x_J^* \\ &- \frac{\left(\frac{r}{q}\right)}{[1-\delta(x_J^*)]}\left[f\tilde{\delta}(x_J^*) + \tfrac{\omega}{P}\bar{\theta}x_J^* + \kappa(1-\Psi(x_J^*))\tilde{\delta}(x_J^*)\right]. \end{aligned}$$

We apply the approximation $\delta(x_A^*) \approx 1 - \frac{r}{q}x_A^*$ to the $1-\delta(x_A^*)$ terms to obtain, for the numerator,

$$\left[1+\left(\tfrac{r+\rho}{q}\right)x_J^*\right]\tfrac{1-\delta(x_J^*)}{\frac{r}{q}\delta(x_J^*)} - x_J^* \approx \left[1+\left(\tfrac{r+\rho}{q}\right)x_J^*\right]\frac{x_J^*}{\delta(x_J^*)} - x_J^* > 0,$$

where the inequality holds because $\tilde{\delta}(x_J^*) < 1$. For the denominator we obtain

$$\begin{aligned} E \approx &-\left(\frac{r+\rho}{q}\right)\frac{\omega}{P}\bar{\theta}\tfrac{x_J^*}{\delta(x_J^*)} - f\left(\frac{r+\rho}{q}\right)\tilde{\delta}(x_J^*) \\ &- f\frac{\tilde{\delta}(x_J^*)}{x_J^*} - \kappa(1-\Psi(x_J^*))\frac{\tilde{\delta}(x_J^*)}{x_J^*} < 0. \end{aligned}$$

Therefore, $dx_J^*/d\bar{\theta} < 0$, as claimed.

B Additional Tables and Figures

B.1 Tables

Table A.1: Simulation parameters

	Baseline
Probability of trade peace ($p = e^{-\rho}$)	0.95
Order quantity (q)	10
Interest rate (r)	0.05
Low, high quality ($\underline{\theta}, \bar{\theta}$)	(0, 10)
Seller fixed cost (f)	0.2
Buyer inspection cost (m)	1
Switching cost (κ)	1

Table A.2: Classification regressions at the importer level, for $t = 15$

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable	ln(VPS)	ln(WBS)	ln(Price)	ln(VPS)	ln(WBS)	ln(Price)
d^A_{mhcz}	0.761^{***}	0.847^{***}	-0.114^{***}			
	(0.004)	(0.004)	(0.007)			
$\ln(SPS_{mhcz})$				0.293^{***}	0.315^{***}	-0.031^{***}
				(0.001)	(0.001)	(0.001)
$\ln(Quantity_{mhcz})$	0.732^{***}	-0.254^{***}	-0.314^{***}	0.744^{***}	-0.253^{***}	-0.334^{***}
	(0.001)	(0.001)	(0.002)	(0.000)	(0.001)	(0.000)
beg_{mhcz}	0.001^{***}	-0.002^{***}	-0.001^{***}	0.001^{***}	-0.002^{***}	-0.001^{***}
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
end_{mhcz}	-0.001^{***}	0.002^{***}	0.001^{***}	-0.001^{***}	0.002^{***}	0.001^{***}
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	158,000	158,000	158,000	927,000	927,000	927,000
R-Squared	0.960	0.728	0.869	0.952	0.593	0.844
Fixed Effects	hc, z	hc, z	hc, z	hc, z	hc, z	hc, z

Notes: Superscripts *, **, and *** indicate statistical significance at the 10, 5, and 1 percent levels, respectively. Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines.

Table A.3: Classification regressions at the importer-exporter level, for $t = 15$

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable	ln(VPS)	ln(WBS)	ln(Price)	ln(VPS)	ln(WBS)	ln(Price)
d^A_{mhcz}	0.427***	1.307***	−0.102***			
	(0.010)	(0.018)	(0.010)			
SPS_{mhcz}				0.246***	0.531***	−0.056***
				(0.001)	(0.002)	(0.001)
$\ln(Quantity_{mhcz})$	0.555***	−0.208***	−0.108***	0.637***	−0.142***	−0.140***
	(0.002)	(0.002)	(0.002)	(0.000)	(0.001)	(0.001)
beg_{mhcz}	0.000***	−0.000***	−0.000***	0.001***	0.000***	−0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
end_{mhcz}	−0.001***	0.000***	0.000***	−0.001***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\ln(Rellength_{xmhcz})$	−0.261***	0.221***	0.078***	−0.211***	0.2002***	0.063***
	(0.005)	(0.009)	(0.003)	(0.001)	(0.002)	(0.001)
Observations	72,000	72,000	72,000	1,325,000	1,325,000	1,325,000
R-Squared	0.982	0.789	0.970	0.979	0.652	0.959
Fixed Effects	xhc, z	xhc, z	xhc, z	xhc, z	xhc, z	xhc, z

Notes: Superscripts *, **, and *** indicate statistical significance at the 10, 5, and 1 percent levels, respectively. Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines.