

“American” vs “Japanese” Procurement: Evidence from Transaction-Level Trade Data*

Sebastian Heise[†]

Yale Department of Economics

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

July 2015

Abstract

We show theoretically how trade policy conducive to the emergence of long-term relationships between domestic buyers and foreign sellers can lead to the adoption of “Japanese”-style procurement practices, in which buyers ensure the provision of high-quality goods via long-run, just-in-time relationships. Consistent with the predictions of the model, we find that a change in U.S. trade policy eliminating the annual threat of a trade war with China coincides with a shift towards smaller, more frequent shipments between U.S. importers and Chinese exporters *vis a vis* other U.S. trading partners. To the extent that this shift lowers inventory and procurement costs, the results highlight a new source of efficiency gains from trade policies that facilitate the formation of long-term buyer-seller relationships. (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 [your name here] 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 Board of Governors or its research staff. All results have been reviewed to ensure that no confidential information is disclosed.

[†]28 Hillhouse Avenue, New Haven, CT 06511, email: sebastian.heise@yale.edu.

[‡]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 have introduced “Japanese”-style procurement practices in an effort to boost operational efficiency.¹ Under the Japanese system, buyer-seller relationships are characterized by joint learning and information sharing, and buyers motivate sellers to maintain product quality by committing to long-run purchases at a price above sellers’ costs. In the more traditional, “American” system, by contrast, buyers choose the lowest-cost seller for each order via competitive bidding, and low quality is deterred via costly inspection.

Given the increasingly global nature of firms’ supply chains (Baldwin and Lopez-Gonzalez 2013), trade barriers represent a potentially important – yet under-studied – consideration in the formation of buyer-seller relationships. Indeed, if buyer and seller are located in different countries, the possibility of a trade war may prevent foreign sellers from entering into the sort of long-term relationships with domestic buyers that characterize the Japanese system. To the extent that this disincentive to adopting the Japanese system prevents reductions in buyers’ inventory and procurement costs, efficiency may suffer.

In this paper, we examine the role of trade policy in firms’ selection of the optimal procurement system. We propose a theoretical model in which reductions in the probability of a trade war increase firms’ incentives to switch from American- to Japanese-style procurement. Empirically, we show that a shift in U.S. trade policy that permanently normalized trade relations between the United States and China coincides with changes in the pattern of U.S. firms’ imports from Chinese suppliers along the lines suggested by the model.

Our theoretical analysis is built around the framework introduced by Taylor and Wiggins (1997), who demonstrate that because of the fixed cost of inspection in the American system and the need for repeated payment of premia in the Japanese system, shipments between seller and buyer are optimally smaller and more frequent – i.e., more “just-in-time” – under the Japanese model. We extend Taylor and Wiggins (1997) to a setting in which a buyer purchases inputs from a foreign supplier. We assign exogenous beliefs about the probability of a prohibitive increase in import tariffs (i.e., a trade

¹This movement is documented in a series of studies. See, for example, O’Neal (1989), Heide and John (1990), Lyons et al. (1990), Dyer and Ouchi (1993), Han et al. (1993), Helper and Sako (1995) and Liker and Choi (2004). We follow Taylor and Wiggins (1997) in using the term procurement to refer to the structuring of purchases – in our case, imports – over time.

war as in Ossa 2014) to both parties, and demonstrate that the lower the probability assigned by the seller to a trade war, the more likely Japanese- versus American-style procurement is to be taken up. The intuition for this result is straightforward: the lower the probability of a trade war, the greater the seller's confidence that a long-term relationship with a buyer can be sustained. This confidence lengthens the time horizon over which the seller expects to collect price premia in a long-term relationship with the buyer, driving down the premium needed on each shipment to enforce honesty and thereby the relative cost of the Japanese system compared to the American system.

The model's implications guide an empirical analysis of the impact on firms' procurement patterns of the United States' granting of permanent Normal Trade Relations (PNTR) to China in October 2000. Conferral of PNTR was a non-traditional trade liberalization in that it did not change the actual import tariffs the United States applied to Chinese goods but rather ruled out the possibility of tariff increases to potentially prohibitive levels. Specifically, while U.S. imports from China had been subject to the relatively low NTR tariff rates reserved for WTO members since 1980, continued access to these low rates required annual renewals that were uncertain and politically contentious. Absent renewal, U.S. import tariffs on Chinese goods would jump to the non-NTR tariff rates assigned to non-market economies originally established under the Smoot-Hawley Tariff Act of 1930. Importantly, the amount by which tariffs would increase under this scenario, measured as the gap between non-NTR and NTR rates, varied substantially across products, a fact that we employ in our empirical strategy. By permanently setting U.S. duties on Chinese imports to NTR levels, PNTR may have encouraged U.S. importers and Chinese exporters to adopt Japanese-style procurement, particularly for products with high NTR gaps.²

In the empirical analysis, we use transaction-level U.S. import data to estimate the effect of PNTR on several measures of procurement that capture key differences between the American and Japanese systems, including average shipment size, frequency and price. We employ a triple difference-in-differences identification strategy that exploits variation in the gap between non-NTR and NTR rates to assess whether U.S.-China importer-exporter-product procurement patterns change relative to imports from exporters in other countries (first difference) after the change in U.S. policy is implemented (second difference) in products with higher NTR gaps (third difference).³

²Pierce and Schott (2015) show that PNTR coincided with sharp increases in U.S. imports from China, as well as the number of U.S. and Chinese firms engaged in U.S.-China trade.

³In our model, seller and buyer trade a single product, so the probability of a trade war and the

For each procurement measure, we compare outcomes within a series of increasingly broad bins: within importer-exporter-product triplets, within importer-product pairs, and within products.

Consistent with the model's predictions, we find that PNTR is associated with a shift towards Japanese-style procurement for U.S.-China relationships, and that this shift is more pronounced for products where the change in policy is more binding, namely those with larger NTR gaps. In the preferred, within importer-exporter-product specification, our results suggest that a one standard deviation increase in the gap between non-NTR and NTR tariff rates is associated with a relative decline in shipment quantity of 13 percent and an increase in shipment price of 4 percent. Via the lens of the model, this estimated reduction in shipment size implies an increase in procurement efficiency and a commensurate reduction in inventory costs.

To our knowledge this is the first paper in the international trade literature to explicitly model procurement patterns, and the first paper on optimal procurement to consider cross-border relationships.⁴ The model we develop also provides an alternate perspective on the large literature examining contractual frictions in international trade (see the survey by Antras and Helpman 2008). A recent working paper Pflueger and Kukharsky (2010), for example, suggests the problem of hold-up on the decision to outsource may be solved by relationship formation, i.e., the sharing of long-term gains in a repeated game from a sustained relationship. Here, we examine how long-term relationships under the Japanese system can solve frictions related to enforcing the provision of high-quality inputs in buyer-seller relationships. One attractive feature of our approach is that it yields predictions regarding shipment patterns that can be tested using transaction-level trade data.

More broadly, our analysis contributes to several literatures in economics and operations research. Our linking of a change in trade policy to firm import patterns joins a growing number of papers examining the various impacts of specific trade policies on firm and aggregate outcomes (see the survey by Goldberg and Pavcnik 2015). We also contribute to research examining the behavior of importers (e.g., Blaum et al. 2015), the implications of trade wars (e.g. Ossa 2014), information frictions in international trade (e.g., Cristea 2011) and trade policy uncertainty (e.g., Handley 2014).

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.

⁴Procurement 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).

The remainder of this paper proceeds as follows. Section 2 outlines our theoretical model. Section 3 describes the data used in our empirical analysis. Sections 4 and 5 contain our empirical analysis. Section 6 concludes. An online appendix contains additional results.

2 Theoretical Model

This section develops a model of optimal procurement when buyer and seller are located in different countries. We start with the framework developed by Taylor and Wiggins (1997) – hereafter TW – in which buyers solve a quality control problem with respect to their inputs using one of two procurement strategies. 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 consider a setup in which a buyer has a fixed procurement need for a single good, and faces the problem of determining the optimal procurement pattern. They 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 cause buyers under the American system 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.

Given these distortions, both the Japanese and American systems are more costly than the first best, where sellers supply high-quality inputs without need for threat or incentive. TW show that while American and Japanese procurement may co-exist as local solutions to the buyer’s quality-control problem, the global optimum 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 generalize TW to allow for exogenous beliefs about changes in the expected level of buyer-country import tariffs as well as explicit inventory costs. We then demonstrate that eliminating the possibility of a trade war provides a seller with a greater incentive to adopt Japanese-style procurement by reducing her effective discount rate. This in turn reduces inventory costs. The model yields empirical predictions that we examine

in Section 5.

2.1 The Setting

Over time interval Δt a “buyer” (B) uses total inputs q purchased from a “seller” (S). Without loss of generality, we normalize $\Delta t = 1$, e.g., 1 year. The buyer receives his inputs in a series of symmetric shipments of size x . As a result, there are q/x shipments during each time interval, each arriving $\Delta t/(q/x) = x/q$ time intervals apart. During the “order cycle” between each shipment, the buyer’s inventory falls from x to 0. This setting is described visually in Figure 1, where $s = 1, 2, \dots, q/x$ indexes the shipments during each time interval Δt and $t = \{1, 2, \dots\}$ indexes time intervals.

The exogenous interest rate over time interval Δt is r , so that the discount rate between orders is rx/q . Let $h = r/q$. With continuous discounting, the discount rate between two shipments is then $\delta(x) = (e^{-hx})$. If $w(x)$ denotes the cost of each shipment x , the present value of an order placed $T = ts + s$ shipments in the future is $\delta(x)^T w(x) = (e^{-hx})^T w(x)$.

Let $\theta \in \{\underline{\theta}, \bar{\theta}\}$ index the low or high quality of the input produced by the seller. The buyer requires high quality, e.g., an acceptably low defect rate among the units shipped.⁵ The seller’s problem is to determine whether to provide high- or low-quality goods for each shipment sent to the buyer.

The buyer can inspect each shipment at cost m per shipment before accepting and paying for it. Let α be the probability that such an inspection occurs. If the buyer chooses to inspect and the quality is low, the relationship 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 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.⁶

⁵In 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 h sufficiently small, the optimal level of quality under both the Japanese and American systems is arbitrarily close to the first best optimal level of quality.

⁶This assumption is a simplification. In practice, practitioners of Japanese procurement tend to

TW do not consider inventory costs explicitly. Here, we assume that the instantaneous inventory cost increases with quality θ and that it is proportional to the inventory utilization rate,

$$v\theta \frac{(x - qz)}{x}, \quad (1)$$

where v is an exogenous constant and $z \in [0, x/q]$ denotes a time within an order cycle. With this formulation, total inventory costs over an order cycle are

$$\int_0^{x/q} v\theta \frac{(x - qz)}{x} dz = \frac{v\theta x}{2q}. \quad (2)$$

These costs increase with order size x because larger, less frequent orders increase the amount of time inventories remain closer to capacity. Smaller, more frequent shipments, by contrast, reduce inventory costs by increasing the speed of inventory throughput relative to capacity. We assume inventory costs for a given order are paid upon acceptance of an order, so the net present cost of storing an order purchased T shipments in the future is

$$\delta(x)^T \frac{v\theta x}{2q}. \quad (3)$$

As inventory costs over an order cycle are a function of the time between shipments (x/q), changes in that interval are directly proportional to changes in inventory costs if inventory holding costs and the cost of producing quality remain constant.

2.2 The Seller's Problem

A seller produces batches of quantity x with variable cost θ and fixed cost f per batch, where f encompasses the fixed cost of both setting up and delivering a production run. The seller receives order value $w_i(x, \theta)$ per shipment, where i indicates whether the payment is under an American or Japanese system. We assume the seller does not have any bargaining power and fills an order only if she at least breaks even,

$$w_i(x, \theta) \geq f + \theta x. \quad (4)$$

We assume free trade between the buyer's and seller's countries, but that a trade

 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).

war 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.⁷ The seller's exogenous belief about the probability of continued *peaceful* trade, and therefore that the relationship will continue, is $0 < k_S < 1$.⁸ The seller's discount factor for an order placed T time intervals in the future is $\delta_S(x)^T$, where the subscript indicates that this is potentially specific to the seller. Given that the stationary environment described above (and summarized in Figure 1) implies a continuous repetition of order cycles over time, the net present value to the seller of supplying batches of x to the buyer as $T \rightarrow \infty$ is

$$\frac{w_i(x, \theta) - f - x\theta}{1 - \delta_S(x)k_S}. \quad (5)$$

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.,

$$\frac{w_i(x, \bar{\theta}) - f - x\bar{\theta}}{1 - \delta_S(x)k_S} \geq (1 - \alpha)w(x, \bar{\theta}) - f - x\underline{\theta}. \quad (6)$$

As this expression makes clear, decreases in shipment size, x , as well as increases in the seller's belief about continued trade peace, k_S , raise the seller's discount factor, $\delta_S(x)k_S$, thereby strengthening the seller's incentive to provide high-quality shipments.⁹

2.3 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 simply 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 $\alpha_J = 0$.¹⁰ In that case, under the American system, the seller just

⁷We provide a micro-foundation for this belief in Section A of the online appendix. 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.

⁸The model considers trade in a single product. An alternate interpretation of k_S 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.

⁹An alternative approach to incorporating trade policy uncertainty would be to include exogenous parameter k_S as part of the discount rate, e.g., $\delta_S(x) = e^{-hx+ks}$.

¹⁰TW show that optimal inspection under the American system is a function of shipment size and quality, $\alpha_A^* = \bar{\alpha}(x, \theta) > 0$, while under the Japanese system inspections do not occur, $\alpha_J^* = 0$.

breaks even on each order,

$$w_A(x, \bar{\theta}) = f + \bar{\theta}x. \quad (7)$$

As there is no expectation of a long-term relationship under the American system, this order 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 order value under the Japanese system is

$$w_J(x, \bar{\theta}) = f + \bar{\theta}x + \left(\frac{1}{\delta_S(x)k_S} - 1 \right) (\bar{\theta} - \underline{\theta})x. \quad (8)$$

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 order value paid under the American system, $w_A(x)$, 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 buyer's belief in trade peace, k_S , falls.

The buyer discounts future payments with discount factor $\delta_B(x)$ and assigns a probability (potentially different from the seller), k_B , to the continuation of the relationship. To simplify the analysis, we set $k_B = 1$ under the assumption that the buyer is able to replace a lost relationship with an approximately equal relationship by switching to a different country. Alternatively, U.S. buyers may be more confident about avoiding trade wars because they have more information about existing lobbying efforts and policy practices compared to foreign partners.

Including inventory costs, the net present cost to the $i = A, J$ buyer of continual ordering under the two systems is then

$$C_i = \frac{w_i(x, \bar{\theta}) + \alpha_i m}{1 - \delta_B(x)k_B} + \frac{v\bar{\theta}x}{2q(1 - \delta_B(x)k_B)}, \quad (9)$$

where m is the fixed cost of inspecting a shipment and α_i is the probability of inspection.

The buyer under each procurement system chooses the optimal order size x_i^* to minimize equation (9),

$$x_i^* = \operatorname{argmin} \left(\frac{w_i(x, \bar{\theta}) + \alpha_i m}{1 - \delta_B(x)k_B} + \frac{v\bar{\theta}x}{2q(1 - \delta_B(x)k_B)} \right). \quad (10)$$

The tradeoff associated with choosing lower- versus higher-frequency procurement can

be seen by setting the first order condition for this problem to zero, yielding

$$\frac{1}{1 - \delta_B(x)k_B} \left(w'_i(x, \bar{\theta}) + \frac{v\bar{\theta}}{2q} \right) = \frac{-\delta'_B(x)k_B}{(1 - \delta_B(x)k_B)^2} \left(w_i(x, \bar{\theta}) + \alpha_i m + \frac{v\bar{\theta}x}{2q} \right). \quad (11)$$

where,

$$w'_i(x, \bar{\theta}) = \begin{cases} \bar{\theta} & \text{if } i = A \\ \bar{\theta} + \left(\frac{1}{\delta_S(x)k_S} - 1 \right) (\bar{\theta} - \underline{\theta}) - \frac{\delta'_S(x)}{\delta_S(x)^2 k_S} (\bar{\theta} - \underline{\theta})x & \text{if } i = J \end{cases} \quad (12)$$

The left hand side of equation (11) 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 – a parameter of $w_i(x, \bar{\theta})$ – and m , 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).

2.4 Numerical Solutions

In this section we use numerical solutions of the model to provide intuition for its key relationships and to motivate the empirical analysis in Section 5. These solutions assume the baseline parameters listed in Table 1 unless otherwise noted. We provide analytic solutions for some of these relationships in the next section.

Figure 2 shows that the overall cost of the Japanese system falls with the seller's probability of trade peace (k_S). The intuition for this relationship is straightforward: as k_S rises, Japanese sellers require less compensation to refrain from providing low quality goods, driving overall costs lower. Two other features of Figure 2 are worth noting. First, it shows that even if $k_S = 1$ the cost of the Japanese system does not drop to that of the first-best (FB) scenario, where neither inspection nor payment premia are required to deter provision of low-quality goods.¹¹ 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 k_S (arbitrarily equal to 0.91 in the figure), which we denote k_S^{Switch} , the cost of the

¹¹Optimal order size is independent of k under both the American system and the first-best scenario. We solve for the x^{FB} as x_A^* conditional on the fixed inspection cost (m) being zero.

Japanese system drops below that of the American system. At that point, buyer and seller switch from the American to the Japanese system.

The left and right panels of Figure 3 demonstrate that the optimal order size under the Japanese system rises with the seller's probability of trade peace (k_S), while the optimal order price, $w_J(x, \bar{\theta})/x_J^*$, falls. The decline in order price as k_S rises reflects the just-noted drop in the seller's rent. Given that decline, the buyer shifts towards larger and less frequent shipments to reduce payment of fixed costs. However, joint consideration of Figures 2 and 3 reveals that if an increase in the probability of trade peace causes k_S to jump from below k_S^{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 5 are consistent with PNTR leading to a switch to the Japanese system in U.S.-China procurement.

Figure 4 reveals that the optimal order size under both systems increases with the seller's fixed cost f , while the optimal order size in the American system increases with per-shipment inspection cost m . In both cases, buyers seek to minimize incurring larger fixed costs by reducing shipments, thereby increasing order size.

Finally, Figure 5 shows that optimal order size under both the American and Japanese systems declines with the marginal costs of high quality ($\bar{\theta}$) and inventory (v). As the cost to produce high quality rises, buyers have an incentive to push purchases further into the future via more frequent, smaller orders. When inventory costs are high, buyers economize on inventory costs by ordering more frequently.

2.5 Analytic Solutions

This section makes use of additional simplifying assumptions to derive a series of analytical implications from the model. In particular, we seek to relate reductions in the probability of a trade war – i.e., U.S. conferral of PNTR on China – to a switch from American- to Japanese-style procurement.

We shut down the channel by which changes in shipment size (and therefore frequency) affect sellers' incentives to cheat by assuming *for the seller* that $hx \rightarrow 0$, so that $\delta_S(x) \rightarrow 1$.¹² Our interpretation of this assumption is that sellers care more about

¹²TW assume $\delta(x) \rightarrow 1$ for both buyer and seller for much of their analysis.

the likelihood that their relationship with the buyer might be broken by trade policy than they do about the interest rate.¹³ As in our numerical solutions, we set $\underline{\theta} = 0$.

With these assumptions, we solve for x_i^* using the same Padé approximation of $\delta_B(x)$ employed by TW, i.e., $\delta_B(x) = \frac{2-hx}{2+hx}$.¹⁴ Substituting equations (7) and (8) into equation (11), we find optimal order sizes

$$x_i^* = \begin{cases} \sqrt{\frac{2q(f+m)}{\gamma r \theta}} & \text{if } i = A \\ \sqrt{\frac{2qf}{\lambda r \theta}} & \text{if } i = J \end{cases} \quad (13)$$

where $\gamma \equiv 1 + \frac{v}{2q} < \frac{1}{k_S} + \frac{v}{2q} \equiv \lambda$.

This result yields two propositions which highlight the key differences between the two systems.

Proposition 1. *In the Japanese system order sizes are smaller (and therefore more frequent) than in the American system.*

Proof. This result follows directly from equation (13). \square

With the optimal order sizes in hand, it is easy to order compare order unit values, $w_i(x)/x$, under the two systems.

Proposition 2. *All else equal, order unit values are greater under the Japanese system than under the American system, $w_J^*(x, \bar{\theta})/x_J^* > w_A^*(x, \bar{\theta})/x_A^*$.*

Proof. From Proposition 1 we have $x_J^* < x_A^*$. Therefore, because $0 < k < 1$, $w_J^*(x, \bar{\theta})/x_J^* = \frac{1}{k}\bar{\theta} + \frac{f}{x_J^*} > \bar{\theta} + \frac{f}{x_A^*} = w_A^*(x, \bar{\theta})/x_A^*$. \square

Buyers under the Japanese system pay a premium for smaller, more frequent orders compared to the American system in order to incentivize the provision of high quality inputs.

In our empirical analysis, we are able to compare the procurement patterns of buyers and sellers trading at arm's length versus those of related parties.¹⁵ We conjecture that the latter require neither inspections nor order value premia to solve the quality-control

¹³The interest rate becomes more important as $k_S \rightarrow 1$, i.e., as our model converges to that of TW.

¹⁴As noted in TW footnote 17 (and Judd 1997), the Padé expansion is more convenient and often more accurate for obtaining closed-form solutions than a Taylor-series expansion.

¹⁵By law, U.S. 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 (see Section 402(e) of the Tariff Act of 1930).

problem. As a result, related parties' procurement patterns may correspond to those of the first-best.

Proposition 3. *The first-best order size is $x^{FB} = \sqrt{\frac{2qf}{\gamma r\theta}}$.*

Proof. Intuitively, as is evident from examination of equation (13), this is tantamount to $k_s \rightarrow 1$ under the Japanese system and $m = 0$ under the American system.¹⁶ For $k_s < 1$, $x_J^* < x^{FB}$ because smaller orders reduce the gain to the seller from cheating on a particular order.¹⁷ . Here, as in TW, $x_A^* > x^{FB} > x_J^*$ if $k_s < 1$. \square

We now turn to a key result of the model, the extent to which the probability of trade peace influences adoption of the American versus Japanese systems. The buyer adopts the system that minimizes the costs of procurement. For the Japanese system, we substitute the optimal order size x_J^* into the cost function from equation (9) to obtain the net present value, C_J ,

$$C_J = \frac{\left(\sqrt{2r\lambda\bar{\theta}f/q} + rf/q\right) \left(\sqrt{2r\lambda\bar{\theta}f/q} + 2\lambda\bar{\theta}\right) \sqrt{2}}{4(r/q)\sqrt{r\lambda\theta f/q}}. \quad (14)$$

Lemma 1. *The cost of the Japanese system strictly decreases as k_S rises and approaches infinity as $k \rightarrow 0$.*

Proof. Substitute $\lambda = \frac{1}{k_S} + \frac{1}{2} \frac{v}{q}$ and take the derivative with respect to k_S to obtain

$$\frac{\partial C_J}{\partial k_S} = - \left(\frac{\left(fhk_Sq + \sqrt{k_Sv + 2q} \sqrt{\bar{\theta}} fk_Sq \right) \sqrt{\bar{\theta}}}{h^{3/2} \sqrt{k_Sv + 2q} \sqrt{f} k_S^{5/2} \sqrt{q}} \right) < 0 \quad (15)$$

$$\lim_{k_S \rightarrow 0} \text{Costs Japanese System} = \infty \quad (16)$$

\square

For the American system, substitute the optimal order size x_A^* into the cost function to obtain the long run expected costs

¹⁶As $k_S \rightarrow 1$, the discount rate becomes more important than the probability of trade peace, which is more consistent with the closed-economy version of the model developed by TW. In our setup, it is plausible to think that k_S can never actually equal 1, i.e., there is always some possibility of a trade war occurring, however small it is.

¹⁷To the extent the seller cares about discounting, x_J^* also is below x^{FB} because smaller orders raise the net present value of future orders via the discount rate.

$$C_A = \frac{\left(\sqrt{2}\sqrt{h\gamma\bar{\theta}(f+m)} + hf + mh\right)\left(\sqrt{2}\sqrt{h\bar{\theta}\gamma(f+m)} + 2\gamma\bar{\theta}\right)\sqrt{2}}{4h\sqrt{\gamma h\bar{\theta}(f+m)}}. \quad (17)$$

Lemma 2. *The cost of the American system increases in the fixed inspection cost m .*

Proof. Take the derivative of the cost with respect to m to obtain

$$\frac{\partial C_A}{\partial m} = \frac{1}{2} \left(\frac{\gamma\sqrt{2\bar{\theta}} + \sqrt{\gamma h\bar{\theta}(f+m)}}{\sqrt{h\bar{\theta}\gamma(f+m)}} \right) > 0. \quad (18)$$

□

Proposition 4. *For a finite $m > 0$, there exists a unique k_S^{Switch} such that if $k_S < k_S^{Switch}$ the buyer adopts the American system and if $k_S > k_S^{Switch}$ the buyer adopts the Japanese system. Furthermore, as suggested by Proposition 3, if $k_S \rightarrow 1$ and $m = 0$, then the firm is indifferent between the two systems.*

Proof. Let $\Delta C(k_S, m) = C_A - C_J$. Substitute $m = 0$ and $k_S = 1$ to show that $\Delta C(1, 0) = 0$. Given Lemma 2 this means that $\Delta C(1, m > 0) > 0$. By Lemma 1 we have that $\frac{\partial \Delta C(1, m > 0)}{\partial k_S} < 0$ and $\lim_{k \rightarrow 0} \Delta C(1, m > 0) = -\infty$. Therefore, for any finite $m > 0$, there must be a unique k_S^{Switch} such that $\Delta C(k_S < k_S^{Switch}, m) < 0$ and $\Delta C(k_S > k_S^{Switch}, m) > 0$. □

With an increase in k_S , the seller is more confident about the relationship continuing, so adoption of the Japanese system becomes more likely. In the extreme, if $k_S = 1$ and $m > 0$, all firms adopt the Japanese system. This result mirrors the numerical solution displayed in Figure 3.

As noted in the previous section, we use this implication of the model to motivate an empirical analysis of whether a substantial, exogenous shock to the continuation probability k_S (i.e., PNTR) can cause a shift to Japanese procurement (i.e., increased order frequency and price, smaller order size). We expect these changes to be larger for relationships encompassing goods where the change in the continuation probability is the most pronounced.

3 Transaction-Level U.S. Import Data

We use transaction-level U.S. import data from the U.S. Census Bureau to identify the international procurement patterns of U.S.-based importing firms. The 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.¹⁸

We refine the data as follows. First, we drop all transactions that are warehouse entries, so that our dataset represents all imports used 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 suggested by Pierce and Schott (2012) 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.¹⁹ 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.

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

¹⁸As 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).

¹⁹The 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.

censoring.

The averages and standard deviations of these attributes are reported in Table 2, where the left panel contains results for arm's-length (AL) relationships and the right panel shows results for related-party (RP) relationships.²⁰ These unconditional comparisons reveal three trends. First, AL relationships are smaller in terms of overall value traded and number of shipments received, and more short-lived in terms of their overall length compared to related-party relationships. Second, individual shipments within AL relationships also appear smaller and more frequent than RP shipments in terms of average value per shipment and length per shipment. Finally, the large standard deviations reported in the table indicate that attributes of both arm's-length and related-party relationships exhibit substantial variation across all dimensions of activity.

4 Arm's-Length versus Related-Party Shipments

In this section we provide a formal comparison of the procurement attributes of arm's-length (AL) versus related-party (RP) relationships that controls for potential variation in the firms and products they encompass. In contrast with the unconditional summary statistics in Table 2, we find here that RP shipments are smaller and more frequent than AL shipments. To the extent that RP relationships approximate the first best scenario discussed in Section 2, these results may reflect the fact that AL relationships are predominantly American, an interpretation that is consistent with the fact they are relatively short-lived and small compared to RP relationships.

We estimate a series of descriptive OLS regressions that examine procurement attributes within importer-product-country bins,

$$\begin{aligned} \ln(\bar{Y}_{mxhcp}) = & \alpha_0 + \alpha_1 \{RP_{mxhc} = 1\} + \alpha_2 \ln(\text{Total Value}_{mxcht}) \\ & + \delta_{mch} + \delta_t + \delta_l + \epsilon_{mxhct} \end{aligned} \quad (19)$$

where m , x , c , h and t index U.S. importers, foreign exporters, origin country, ten-digit HS products and years. The regression sample is restricted to importer-exporter-

²⁰Results 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 et al. 2006).

product triplets that have at least two transactions in each $mxhct$ bin and that engage solely in arm's length or solely in related-party (RP) transactions over the full 1992 to 2011 sample period. \bar{Y}_{mxhct} represents one of three attributes: VPS_{mxcht} , the average value per shipment in the bin; QPS_{mxcht} , the average quantity per shipment in the bin; and WBS_{mxcht} , the average weeks between shipments in the bin. The indicator variable $1\{RP_{mxhc} = 1\}$ distinguishes arm's-length from related parties, and $TotalValue_{mxcht}$ accounts for the total value of shipments in the bin. By including this variable we compare AL relationships with RP relationships importing the same total value in the same year. Importer-product-country, year and relationship-length fixed effects are represented by δ_{mhc} , δ_t and δ_l , where δ_l categorizes relationships with lengths from 1 to 6 years plus a final category of 7-or-more years. The coefficient of interest is α_1 , which estimates the relative difference between related and arm's-length procurement in log points.

Results for α_1 and α_2 are reported in Table 3, where each column contains the regression for a different relationship attribute. As indicated by the negative and statistically significant point estimates in the first two columns of the table, conditional on procuring the same total value, related parties use smaller shipments and order more frequently than arm's-length parties within the noted fixed-effect dimensions. The -0.04 and -0.18 point estimates for α_1 reported in the first two columns of the table indicate that average shipment size in terms of value and quantity for related parties are approximately 4 and 18 percent lower compared to an arm's-length relationship of the same size. Results in the final column of the table indicate that related-party shipments arrive approximately 4 percent more frequently than arm's length shipments. Via the lens of the model developed in Section 2 (equation 1), the more just-in-time nature of related-party shipments implies relatively lower inventory costs.

5 Procurement and PNTR

The numerical and analytic results in Section 2 indicate that an increase in the seller's belief in peaceful trade can induce buyer and seller to switch from the American to the Japanese system. In this section, we examine the relationship between the U.S. granting of PNTR to China in October 2000 – which substantially reduced the possibility of a trade-war-like hike in U.S. import tariffs on Chinese goods – and several outcome variables that capture differences between American and Japanese procurement. We

provide a brief summary of the change in U.S. policy before presenting our empirical strategy and describing the results in detail.

5.1 Empirical Strategy

U.S. imports from non-market economies such as China are subject to non-NTR tariff rates originally set under the Smoot-Hawley Tariff Act of 1930. These rates are often substantially larger than the NTR rates the U.S. offers fellow members of the World Trade Organization (WTO). 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.

The U.S. Congress passed a bill granting China permanent NTR (PNTR) status in October 2000, which was implemented on January 1, 2002 as part of China's entry into the WTO in December 2001. By eliminating the threat of sudden spikes in U.S. import tariffs on Chinese goods, this change in U.S. policy likely encouraged greater adoption of Japanese-style procurement between U.S. importers and Chinese exporters. Via the lens of the model developed in Section 2, this encouragement was stronger for firms trading products with relatively large NTR gaps, as the probability that a U.S. buyer would abandon a Chinese seller in the event of a failed annual renewal would be higher in these products.

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

$$NTR\ Gap_h = Non\ NTR\ Rate_h - NTR\ Rate_h, \quad (20)$$

using *ad valorem* tariff rates provided by Feenstra, Romalis and Schott (2002) for 1999, the year before passage of PNTR in the United States.²¹ As indicated in Figure 6, 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

²¹While 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.

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.

5.2 Main Empirical Results

We analyze the relationship between PNTR and procurement patterns within increasingly broad bins across three specifications.

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

$$\begin{aligned} \ln(\bar{Y}_{mxhcp}) = & \alpha_0 + \alpha_1 1\{p = Post\} * 1\{c = China\} * NTR Gap_p + \gamma \chi_{mxhcp} \quad (21) \\ & + \alpha_2 \ln(\text{Total Value}_{mxhcp}) + \delta_{mxh} + \delta_c + \delta_p + \epsilon_{mxhcp} \end{aligned}$$

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 five-year windows around 2001, either 1995 to 2000 (pre period) or 2002 to 2007 (post period).

\bar{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.²² 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 α_1 . $TotalValue_{mxhcp}$ is

²²The 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.

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, α_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 $\alpha_1 < 0$ for VPS_{mxhcp} , QPS_{mxhcp} and WBS_{mxhcp} , and $\alpha_1 > 0$ for $Price_{mxhcp}$ and $Length_{mxhcp}$ if PNTR induced a switch from the American to the Japanese system.

The second specification ignores exporter identity and analyzes shipments within importer-products across periods,

$$\begin{aligned} \ln(\bar{Y}_{mhcp}) = & \alpha_0 + \alpha_1 1\{p = Post\} * 1\{c = China\} * NTR Gap_p + \gamma \chi_{mhcp} \\ & + \alpha_2 \ln(\text{Total Value}_{mhcp}) + \delta_{mh} + \delta_c + \delta_p + \epsilon_{mhcp} \end{aligned} \quad (22)$$

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,

$$\begin{aligned} \ln(\bar{Y}_{hcp}) = & \alpha_0 + \alpha_1 1\{p = Post\} * 1\{c = China\} * NTR Gap_h + \gamma \chi_{hcp} \\ & + \alpha_2 \ln(\text{Total Value}_{hcp}) + \delta_h + \delta_c + \delta_p + \epsilon_{hcp} \end{aligned} \quad (23)$$

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 α_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. Through the lens of our model, the decline in shipment quantity implies a commensurate drop in inventory costs.

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

6 Conclusion

This paper analyzes the impact of trade policy on firms' procurement patterns. We develop a theoretical model in which firms' choice of how to structure shipments along a supply chain responds to their beliefs regarding the probability of continued peaceful trade. This model reveals that reductions in the likelihood of a trade war can allow domestic buyers to reduce inventory costs by forming long-term, just-in-time relationships with foreign sellers.

We examine the model's implications empirically by estimating 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 transaction-level U.S. import data and a triple difference-in-differences specification, we show that PNTR is associated with a movement toward more Japanese-style procurement along the lines suggested by the model.

The results suggest that an important but under-examined aspect of trade agreements in a world with already low tariffs may be their effect on relationship formation. That is, trade agreements promoting institutions which 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.²³ The extent to which such

²³Indeed, improving the efficiency of trade relationships is a goal of the recent WTO agreement on

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.

References

- [1] Antràs, Pol, and Elhanan Helpman. 2008. "Contractual Frictions and Global Sourcing." In *The Organization of Firms in a Global Economy*, edited by Elhanan Helpman, T Verdier, and D Marin, 9-54. Cambridge, MA: Harvard University Press.
- [2] Autor, David H., David Dorn and Gordon H. Hanson. 2014. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review*.
- [3] Autor, David H., David Dorn, Gordon H. Hanson and Jae Song. 2013. "Trade Adjustment: Worker Level Evidence." NBER Working Paper 19226.
- [4] Bajari, Patrick, Stephanie Houghton and Steven Tadelis. 2014. Bidding for Incomplete Contracts: An Empirical Analysis of Adaptation Costs. *American Economic Review* 104(4): 1288-1319.
- [5] Baldwin, Richard and Javier Lopez-Gonzalez. 2013. Supply-Chain Trade: A Portrait of Global Patterns and Several Testable Hypotheses. NBER Working Paper 18957.
- [6] Bernard, Andrew, J. Bradford Jensen and Peter K. Schott. 2006. Transfer Pricing by U.S.-Based Multinationals. NBER Working Paper 12493.
- [7] Bernard, Andrew, J. Bradford Jensen and Peter K. Schott. 2009. Importers, Exporters and Multinationals: A Portrait of Firms in the U.S. that Trade Goods in T. Dunne, J.B. Jensen and M.J. Roberts (eds.), *Producer Dynamics: New Evidence from Micro Data* (University of Chicago Press, 2009). Joint w/ Bernard, Jensen. This chapter is a revised version of NBER Working Paper 11404.)
- [8] Blaum, Joaquin, Claire Lelarge and Michael Peters. 2015. "The Gains From Input Trade in Firm-Based Models of Importing." Mimeo.

trade facilitation. See https://www.wto.org/english/thewto_e/minist_e/mc9_e/desci36_e.htm.

- [9] Bustos, Paula. 2011. Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms. *American Economic Review* 101(1): 304-40.
- [10] Cicala, Steve. 2015. When Does Regulation Distort Costs? Lessons from Fuel Procurement in US Electricity Generation. *American Economic Review* 105(1):411-44.
- [11] Cristea, Anca. 2011. Buyer-Seller Relationships in International Trade: Evidence from U.S. State Exports and Business-Class Travel. *Journal of International Economics*, 2011, 84:2, 207-220.
- [12] Dyer, Jeffrey H. and William G. Ouchi. 1993. Japanese-Style Partnerships: Giving Companies a Competitive Edge. *MIT Sloan Management Review* 35(1): 51-63.
- [13] Feenstra, Robert, John Romalis and Peter K. Schott. 2002. US Imports, Exports and Tariff Data, 1989 to 2001. NBER Working Paper 9387.
- [14] Boehm, Christoph, Aaron Flaaen and Nitya Pandali-Nayar. 2015. Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tohoku Earthquake. Mimeo.
- [15] Goldberg, Pinelopi K. and Nina Pavcnik. 2015. The Consequences of Trade Policy. Mimeo.
- [16] Han, Sang-Lin, David T. Wilson, and Shirish P. Dant. 1993. Buyer-Seller Relationships Today. *Industrial Marketing Management* 22(4): 331-338.
- [17] Handley, Kyle. 2014. Exporting Under Trade Policy Uncertainty: Theory and Evidence. *Journal of International Economics* 94(1):50-66.
- [18] Heide, Jan B. and George John. 1990. Alliances in Industrial Purchasing: The Determinants of Joint Action in Buyer-Seller Relationships. *Journal of Marketing Research* 27(1): 24-36.
- [19] Helper, Susan R. and Mari Sako. 1995. Supplier Relationships in Japan and the United States: Are They Converging? *MIT Sloan Management Review* 36(3): 77-84.)
- [20] Judd, Kenneth L. 1997. *Numerical Methods in Economics*. Cambridge: MIT Press.

- [21] Kamal, Fariha, C.J. Krizan and Ryan Monarch. 2015. Identifying Foreign Suppliers in U.S. Merchandise Import Transactions. Mimeo.
- [22] Kranton, Rachel and Anand V. Swamy. 2008. Contracts, Hold-Up, and Exports: Textiles and Opium in Colonial Bengal. *American Economic Review* 98(3): 967-89.
- [23] Kukharskyy, Bohdan and Michael Pflueger. 2010. Relational Contracts and the Economic Well-Being of Nations. *IZA Discussion Paper* 5394.
- [24] Liker, Jeffrey and Thomas Y. Choi. 2004. Building Deep Supplier Relationships. *Harvard Business Review* (December).
- [25] T.F. Lyons, 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 (Spring 1990).
- [26] O'Neal, Charles R. 1989. JIT Procurement and Relationship Marketing. *Industrial Marketing Management*. 18(1): 55-63.
- [27] Ossa, Ralph. 2014. Trade Wars and Trade Talks with Data. *American Economic Review* 104(12): 4104-4146.
- [28] Pierce, Justin R. and Peter K. Schott. 2012. Concording U.S. Harmonized System Categories Over Time. *Journal of Official Statistics* 28(1):53-68 (March 2012).
- [29] Pierce, Justin R. and Peter K. Schott. 2015. The Surprisingly Swift Decline of U.S. Manufacturing Employment. Mimeo.
- [30] Tadelis, Steven and Zettelmeyer, Florian. 2015. Information Disclosure as a Matching Mechanism: Theory and Evidence from a Field Experiment. *American Economic Review* 105(2): 886-905
- [31] Taylor, Curtis, and Steven Wiggins, "Competition or Compensation: Supplier Incentives Under the American and Japanese Subcontracting Systems," *American Economic Review*, 1997, 87:4, 598-618.

Table 1: Default Parameters for Numerical Solutions

Baseline	
Buyer, Seller Continuation Probability (k_B, k_S)	(1, 0.99)
Inventory Cost (v)	0.02
Order Quantity (q)	0.02
Interest Rate (r)	0.02
Low, High Quality ($\underline{\theta}, \bar{\theta}$)	(0, 0.02)
Seller Fixed Cost (f)	0.01
Buyer Inspection Cost (m)	0.02

	Arm's-Length	Related-Party
Total Value Traded	228,874	1,757,764
	11,720,829	79,918,870
Overall Length	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
Observations	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 by 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 over the relationship, the duration of the relationship in weeks, and the 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.

Table 2: Relationship Summary Statistics

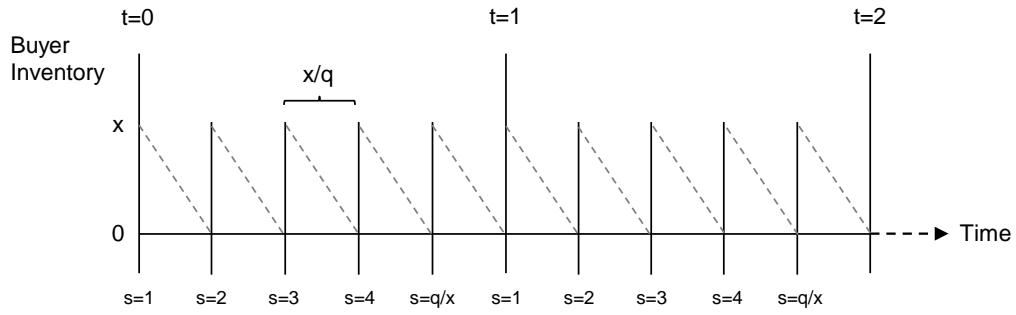
	ln(VPS _{mxhct})	ln(QPS _{mxhct})	ln(WBS _{mxhct})
1{RP _{mxhct} =1}	-0.04 *** 0.001	-0.18 *** 0.002	-0.04 *** 0.001
ln(Total Value _{mxch})	0.580 *** 0.0001	0.550 *** 0.0002	-0.2 *** 0.0001
Observations	16,491,300	16,491,300	16,491,300
R-Squared	0.93	0.93	0.50
Fixed Effects	mh,c,t,l	mh,c,t,l	mh,c,t,l
Notes: Table presents results of importer-exporter-product-country-year (<i>mxhct</i>) level OLS regressions of relationship attributes on a dummy variable indicating related-party status (RP) and a covariate (Total Value) indicating the size of the relationship in terms of value shipped. Regression sample restricted to relationships with at least two observations per <i>mxhct</i> bin that engage solely in arm's length or solely in related-party (RP) transactions over the sample period. WBS is the average number of weeks between shipments. VPS and QPS are the average value and quantity per shipment. Results for fixed effects 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.			

Table 3: Related-Party versus Arm's Length Transactions

Dependent Variable	Within Importer- Exporter- Product	Within Importer- Product	Within Product
$\ln(VPS_{hcp})$	-0.07 *** 0.01	-0.05 *** 0.01	-0.18 *** 0.05
$\ln(QPS_{hcp})$	-0.13 *** 0.02	-0.04 ** 0.02	-0.04 0.10
$\ln(Price_{hcp})$	0.04 ** 0.02	-0.04 ** 0.02	-0.11 0.09
$\ln(WBS_{hcp})$	-0.04 0.03	-0.06 *** 0.02	-0.36 *** 0.07
$\ln(Length_{hcp})$	0.10 *** 0.04	0.00 0.03	-0.34 *** 0.08
Observations	752,600	1,011,700	324,300
Sample	mxhcp	mhcp	hcp
Fixed Effects	mxh,c,p	mh,c,p	h,c,p

Notes: Table summarizes the results of generalized difference-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*) 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 noted sample bins. WBS, VPS and QPS are weeks between shipments, value per shipment and quantity per shipment. Price is the average price per shipment. Length is the length of the importer-exporter-product relationship. 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.

Table 4: PNTR and Procurement



Notes: Buyer uses inputs purchased from a seller at rate q over each time interval, e.g., $t=0$ to $t=1$. There are $s=\{1,2,\dots,q/x\}$ 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

Figure 2: Overall Cost vs Continuation Probability (k)

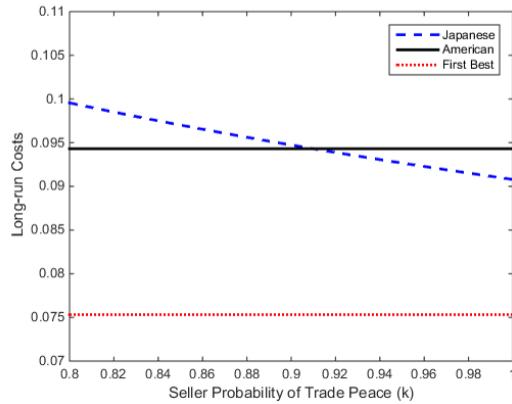


Figure 3: Order Size and Price vs Continuation Probability (k)

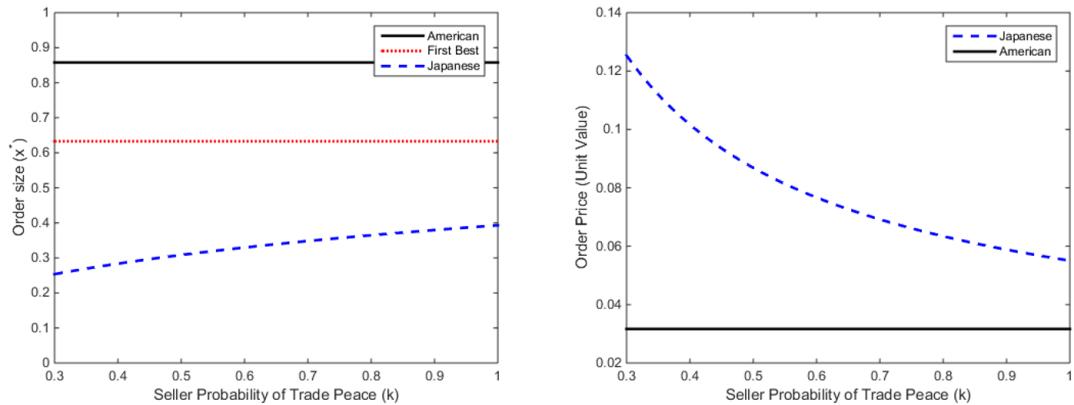


Figure 4: Order Size vs Fixed Costs (f, m)

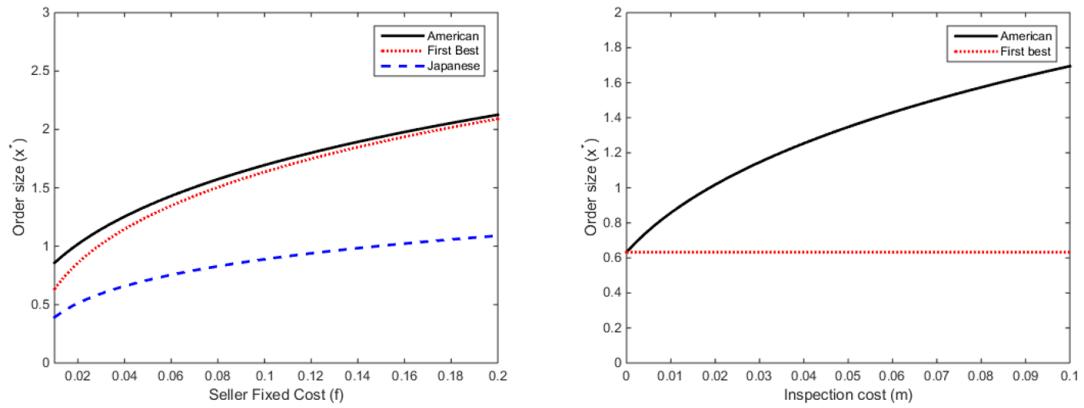
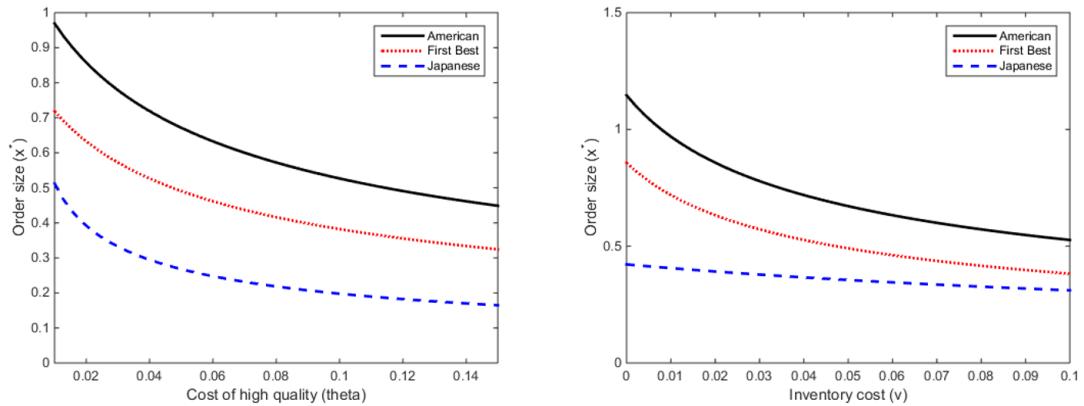


Figure 5: Order Size vs Variable Costs $(\bar{\theta}, v)$



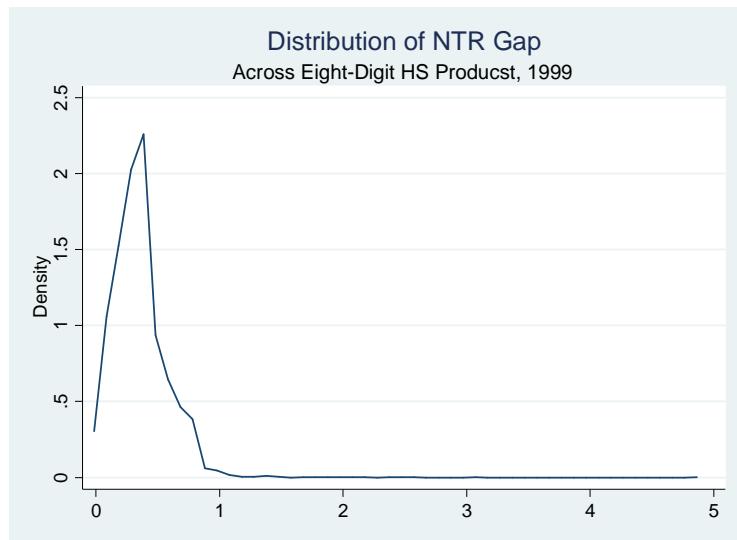


Figure 6: Distribution of the NTR Gap Across Products

Online Appendix (For Online Publication)

A Micro-foundation for k_S

In Section 2.2 of the main text, we assume that in the event of a trade war, the import tariff on the product the buyer imports from the seller rises to a prohibitive level, with the result that the buyer-seller relationship is severed. In this section, to provide a closer link between the model and our empirical analysis, we offer a short micro-foundation for the seller's belief k_S about the probability that the relationship will continue that is a function of the change in tariffs and therefore product-specific.

Assume that the initial *ad valorem* import tariff on the traded input, τ_0 , is equal to zero, but rises to $\tau_{TW} > 0$ in the event of a trade war. In that case, the buyer may seek an alternate seller from another country. Let the difference in the net present value for the buyer thereafter sourcing q from this alternate seller be

$$NPV(q)_{Seller} - NPV(q)_{Alternative}. \quad (A.1)$$

We assume the seller does not know the value of the second term and therefore treats it as a random variable ϵ , where ϵ is independent of the true probability of a trade war, $(1-\pi)$. The seller's probability that the relationship with the buyer continues in any period, k_S , is then

$$k_S = (1 - \pi)P(\epsilon > (1 + \tau_{TW})NPV(q)_{Alternative}) + \pi P(\epsilon > NPV(q)_{Alternative}). \quad (A.2)$$

Given that $P(\epsilon > (1 + \tau_{TW})NPV(q)_{Alternative}) < P(\epsilon > NPV(q)_{Alternative})$, k_S decreases with the true probability of a trade war, $1 - \pi$, as well as the trade-war tariff, τ_{TW} . In our empirical analysis, we interpret τ_{TW} as the NTR gap. In that case, in the event of a trade war, buyer-seller relationships oriented around a particular product are less likely to survive the higher that product's spike in tariffs. We note that in this setup, the continuation probability, k_S , is independent of the buyer's choice of optimal order size, consistent with the idea that the buyer is true to his commitment.

B Additional Analytic Propositions

This section highlights additional implications of the simplified model discussed in Section 2.5.

Proposition 5. *Under both the American and Japanese procurement systems, optimal order size rises (optimal frequency falls) with an increase in fixed production cost f and decreases in variable cost $\bar{\theta}$ or inventory cost v . Under the Japanese system, optimal order size rises (optimal frequency falls) with an increase in the continuation probability k_S . Under the American system, optimal order size rises (optimal frequency falls) with an increase in order inspection costs m .*

Proof. It is straightforward to show that $\frac{\partial x_J^*}{\partial \lambda} < 0$, $\frac{\partial \lambda}{\partial k} < 0$ and $\frac{x_J^*}{x_A^*} = \sqrt{\frac{\gamma f}{\lambda(f+m)}} < 1$ \square

In order to see the intuition behind this, hold the discount rate fixed, and assume that the variable cost to produce high quality $\bar{\theta}$ or inventory variable cost v falls. This lowers the present value of variable order costs relative to fixed order costs. In that case, buyers reduce order frequency to raise variable order costs at the expense of costs of lower discounted future order costs. Likewise, holding the discount rate fixed, an increase in the fixed order cost f raises discounted future fixed order costs. The buyer therefore balances the increase in fixed order costs by increasing lowering order frequency. The continuation probability k_S affects order size only under the Japanese system. An increase in k_S lowers the incentive premium the buyer pays the seller in each order cycle. At a given discount rate, this reduces the variable discounted order costs and buyers increase shipment sizes and reduce order frequencies to re-optimize on order costs.

Proposition 6. (i) *An increase in the probability of peaceful trade k_S raises order size (lowers order frequency) in the Japanese system relative to the American System; (ii) the greater the inventory cost v the less elastic are relative order size (and relative order frequency) with respect to a change in k_S ; and (iii) variable and fixed production and delivery costs do not affect the elasticity of relative order size (or shipping frequency) with respect to k_S .*

Proof. The elasticity of relative order size with respect to k_S , $\frac{d(x_J^*/x_A^*)/k_S}{dk_S/(x_J^*/x_A^*)}$ is $\epsilon \equiv \frac{q}{(k_S v + 2q)k_S} > 0$. \square

This proposition summarizes the impact of a switch from the American to the Japanese system due to an exogenous decline in the seller's perceived probability of a trade war. As a result of such a switch, order size falls and shipping frequency rises, with the magnitude of these changes falling in the inventory cost v . Combined with proposition 1 this implies that those firms that had the lowest order frequencies before the increase in the continuation will switch to the Japanese system and see an increase in order frequencies, while those firms with the highest order frequencies and are already in a Japanese style contract will see a decrease in the order frequency.

Proposition 7. *Under both the American system and the Japanese system, order unit values increase in fixed cost f , variable production cost $\bar{\theta}$, and variable inventory cost v . Under the American system, they also increase in the fixed inspection cost m . In the Japanese system, they decrease as the probability of trade peace k_S increases.*

Proof. Substitute the optimal shipping quantities x_J^* and x_A^* into the order unit values and take the derivatives with respect to the appropriate parameters. \square

Within relationships already organized according to the Japanese system, a rise in the probability of peaceful trade induces a decline in order unit value. Within relationships previously organized according to the American system, however, a rise in the probability of peaceful trade induces a rise in the order unit value if the relationship switches to the Japanese system. Again, this is similar as above. It suggests a non-linearity in the dependent variable. The higher order values see a decrease (those are Japanese before the liberalization) and the lowest unit values see an increase (those are american but may switch to Japanese.)

Proposition 8. *An increase in the inspection cost m will lower the cutoff k_S^* at which point the buyer switches from the American to the Japanese system.*

Proof. For any given value of m , k_S^* is the cutoff such that $\Delta C(k_S^*, m) = 0$, where the buyer is just indifferent between the two systems. Apply the implicit function theorem to consider only arm's length transactions. show that

$$\frac{\partial k(m)^*}{\partial m} = -\frac{\frac{\partial \Delta(k_S, m)}{\partial m}}{\frac{\partial \Delta(k_S, m)}{\partial k_S}} = -\frac{1}{2} \left(\frac{\left(\sqrt{(2q + v)(f + m)r\bar{\theta}} + 2q\bar{\theta} + \bar{\theta}v \right) rk_S^{5/2} \sqrt{(k_S v + 2q)f}}{q\sqrt{(f + m)(2q + v)} \left(frk_S + \sqrt{(k_S v + 2q)f}k_S r\bar{\theta} \right) \bar{\theta}} \right) < 0 \quad (\text{A.3})$$

\square

For a given increase in the probability of trade peace k_S , a buyer under the American procurement system with a high m is more likely to switch to the Japanese system.