K.7

“It’s Not You, It’s Me”: Breakups in U.S.-China Trade Relationships

Monarch, Ryan

Please cite paper as:
Monarch, Ryan (2016). “It’s Not You, It’s Me”: Breakups in U.S.-China Trade Relationships. International Finance Discussion Papers 1165.

http://dx.doi.org/10.17016/IFDP.2016.1165

International Finance Discussion Papers

Board of Governors of the Federal Reserve System

Number 1165
May 2016

Board of Governors of the Federal Reserve System

International Finance Discussion Papers

Number 1165

May 2016

"It's Not You, It's Me": Breakups in U.S.-China Trade Relationships

Ryan Monarch

"It's Not You, It's Me": Breakups in U.S.-China Trade Relationships

Ryan Monarch*

Costs to switching suppliers can affect prices by discouraging buyer movements from high to low cost sellers. This paper uses confidential U.S. Customs data on U.S. importers and their Chinese exporters to investigate these costs. I find considerable barriers to supply chain adjustments: 45% of arm's-length importers keep their partner, and one-third of switching importers remain in the same city. Guided by these regularities, I propose and structurally estimate a dynamic discrete exporter choice model. Cost estimates are large and heterogeneous across products. These costs matter for trade prices: halving switching costs reduces the U.S.-China Import Price Index by 14.7%.

Keywords: International Trade, Import Price, Transactional Relationships

JEL classifications: F14, F23, L14, D21;

The author is a staff economist in the Division of International Finance, Board of Governors of the Federal Reserve System, Washington, D.C. 20551 U.S.A. The views in this paper are solely the responsibility of the author(s) and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System. Contact: ryan.p.monarch@frb.gov

*I thank my advisors Andrei Levchenko, Alan Deardorff, Jeremy Fox, Jagadeesh Sivadasan, and Jing Zhang for their advice, patience and energy. I am grateful to Dan Ackerberg, Vanessa Alviarez, Andrew Bernard, Leland Crane, Javier Cravino, Rafael Dix-Carneiro, Kathryn Dominguez, Aaron Flaaen, Fariha Kamal, Ryan Kellogg, Illenin Kondo, C.J. Krizan, Eric Lewis, Logan Lewis, Bill Lincoln, Andrew McCallum, Edie Ostapik, Gisela Rua, Tim Schmidt-Eisenlohr, Evan Starr, and Olga Timoshenko, and participants at the 2015 NBER ITI Spring Meeting, EIIT 2014, the Michigan informal international seminar, the Michigan informal IO seminar, the Penn State Applied Economics conference, the Michigan international seminar, the 2013 RDC Conference, the 2013 Midwest International Meetings, and the U.S. Census Bureau Center for Economic Studies seminar series for helpful comments. I thank Jiandong Ju, Hong Ma, and the Tsinghua University Center for International Economic Research, as well as Loren Brandt, Johannes von Biesebrouck, and Xue Bai for their sharing of concordances. I acknowledge financial support from the Michigan Institute for Teaching and Research in Economics (MITRE), Rackham Graduate School, the Tokyo Foundation, and the Michigan Center for Chinese Studies. The views expressed herein are mine, and should not be interpreted as reflecting the views of the U.S. Census Bureau, the Federal Reserve Board of Governors, or any other person associated with the Federal Reserve System. This work was undertaken while I was under Special Sworn Status at the U.S. Census Bureau. All results have been reviewed to ensure no confidential information is disclosed. Clint Carter, Jim Davis, Maggie Levenstein, Arnold Reznek and William Wisniewski were extremely helpful with data requests and disclosure processes. All errors are mine.

"[T]he initiatives we announced today are aimed at ... accelerating user growth and customer acquisition. Our Quality Supplier Program will allow buyers to trade with greater confidence." – David Wei, CEO, Alibaba.com (November 3, 2008)

"Alibaba ... reported that 2,326 high volume sellers - 'gold suppliers', as they're called- defrauded customers over the course of two years, with the assistance of nearly 100 Alibaba.com employees... As a result of the scandal, Alibaba.com CEO David Wei ... resigned yesterday." – Fortune Magazine (February 22, 2011)

1 Introduction

Choosing the right supplier is one of the most difficult decisions a firm can make. There are many relevant considerations: price, location, "quality" of the supplier, and so on. One of the most important considerations, however, is rarely observable directly: how to know which suppliers are trustworthy? As the quotes above demonstrate, Alibaba.com tried to smooth some of the confidence problems for importers sourcing from China by identifying "gold suppliers", but the challenge proved enormous.

The most basic way to assess reliability is to transact with a seller, and then determine whether to pursue a subsequent transaction. Importantly, the relationship inertia over time implied by such a practice can lead to high prices: if a firm is unaware of acceptable, lower-priced alternatives to its current supplier, or is unwilling to bear the costs and uncertainty involved in such a change, then that firm would be paying higher than the efficient price. In this paper, I use confidential U.S. Customs and Border Protection data to demonstrate that U.S. importer relationships with Chinese suppliers are persistent, and study the implications of this persistence using a model where switching suppliers is costly.[1]

The U.S. customs data reveals several regularities that large costs from changing one's supply network. Supplier-level switching costs are substantial: even with over 30 Chinese suppliers to choose from in the average HS10 product, 45% of importers kept the exact same supplier from one year to the next- accounting for nearly half the total value of U.S. imports from China. Switching costs apply to changing cities as well: one-third of switching importers remained in the same city as their original partner. Beyond these implicit costs, price is an additional factor influencing the decision of whether to switch: those importers who paid the highest prices are much more likely to change their partner in the following year. Thus there is the potential for easier switching to lead to lower import prices.

[1] I use the term importer to refer to a *firm-HS10 product pair*. HS10 is a ten-digit *harmonized system* product code, the most disaggregated product code used by U.S. firms in international trade.

Guided by the above empirical findings, I develop a dynamic discrete *exporter* choice model. The importing firm decides which exporter to use by comparing partner-specific profits across all possible choices, including its current match. Which exporter to use depends not only on the price and quality of each exporting firm, but also a cost from switching, both at the partner and the city level. The key tradeoff for the importer is that switching to either a cheaper or a higher quality partner raises profits, but changing one's partner or location is costly. The model produces closed-form expressions of choice probabilities for each potential outcome, allowing computation of the switching cost parameters via maximum likelihood. I use the Mathematical Programming with Equilibrium Constraints (MPEC) techniques developed by Dubé et al. (2012) and Su and Judd (2012) to estimate the model. I compute the parameters of the model product-by-product (at the HS10 level) and industry-by-industry (at the HS6 level).

The main quantitative results can be summarized as follows. First, model estimates of switching costs are large, heterogeneous across products, and match the characteristics of underlying data well. How large are switching costs? Just to be indifferent between its current partner and an otherwise identical partner, an importer would require an unexplained, partner-specific boost to profits that is two standard deviations higher than average from that new partner.

Second, the impact of switching frictions on importer prices is sizable. To examine this, I alter switching cost estimates and allow importers to re-optimize their supplier choice. The counterfactual U.S.-China Import Price Index that arises from cutting frictions in half is 14.7% lower, a decline achieved by shrinking the percent of staying importers from 51% to 13%.[2] Thus lowering the cost of switching export partners yields substantial efficiency gains. Finally, I estimate the trade flow to a theoretical supplier not subject to any geographic switching friction (such as a new supplier located in the U.S.). If each product category adds a supplier charging the median price and producing a high-quality variety, approximately 4% of total imports from China would flow back to a theoretical U.S. supplier that charged the median price of Chinese exporters in a product. That said, such a price is approximately 57% lower than the price charged by U.S. exporters for the same product mix, showing that there are substantial barriers to "re-shoring" Chinese imports back to U.S. suppliers beyond low Chinese prices.

What drives these large effects of switching costs on exporter choice and import prices? There are three interpretations, each related to a different interpretation of what the model estimates of switching costs actually represent. One is a lack of knowledge about potential

[2]To calculate this, I use 50 of the highest-value HS products between the U.S. and China, and generate a price index from those products.

suppliers- the high cost of switching partners and additional cost of geographic switching implies that importers are simply not aware of other low-price options that are available.[3] In this case, a policy that reduces switching costs would reduce information asymmetry, such as a "gold standard" exporter certification or a directory of exporter and price information for importers to use when selecting partners. A second interpretation is that the existence of long-term trading contracts prevents more efficient matching.[4] Here, an overall improvement in contracting institutions leading to more short-term contracts would be consistent with reductions in switching costs discussed above. A third interpretation is that there is some logistical difficulty in adjusting suppliers in response to price changes, even with knowledge of other alternatives and a freedom to use them, as in Drozd and Nosal (2012). Here, the experiments described above are best thought of as more widespread use of intermediaries- companies specializing in connecting importers with exporters.[5] In summary, the above quantitative exercises can be thought of as policies that reduce matching frictions, and they can lead to significant reduction in input prices.

Although the field of international trade has focused on numerous aspects of firm-level participation in international activity, the study of individual exporter-importer relationships remains relatively nascent. Empirical work on this question began with the study of networks in international trade: Rauch (2001) surveys the potential for transnational cultural networks to help smooth international trade and reduce barriers to entry, while Rauch and Watson (2004) present a general equilibrium model through which economic agents can use their supply of networks to either produce/export more efficiently or to become an intermediary. Recent work has made use of the U.S. Customs data, which provides information about U.S. importers and their foreign exporting partners. Eaton et al. (2014) study the relationship between Colombian exporters and the number of U.S. importers they partner with over time and calibrate a search and matching model to match exporter decisions, including sales, number of clients, and transition probabilities. Kamal and Sundaram (2013) use the same U.S. import data to determine how likely textile producers in Bangladeshi cities are to follow other exporters in their same city to export to a particular partner. Many of these papers put the onus on the exporter to undertake searching behavior buyer by buyer; in this work, I model matching as an importer's choice given information about each exporter. Other work takes advantage of two-sided trade data to study the effects of heterogeneity in trade: Bernard et al. (2014) develop a model of relationship-specific fixed costs to exporting

[3]Allen (2014) shows the importance of information frictions, especially geographically, for Filipino farmers searching for buyers of their product.

[4]These contracts are studied in Kleshchelski and Vincent (2009).

[5]Indeed, intermediaries play an important in the Chinese export market, as described in Ahn et al. (2011) and Tang and Zhang (2012).

using Norwegian buyer-supplier trade data, while Blum et al. (2013) use exporter-importer pair data on Chile to study the effects of intermediaries. Alessandria (2009) is a model of search frictions in international trade that, like my model, generates deviations from the law of one price, without distinguishing importers within a country. Kleshchelski and Vincent (2009) also construct a model of switching frictions, where firms and customers form long-term relationships, showing that prices stabilize as the number of repeat buyers increases. Antràs and Costinot (2011) and Petropoulou (2011) model the effects of trade with costly search frictions, and Allen (2014) estimates a buyer search model on agricultural trade inside the Philippines. I combine the theory of partner choice with data on importer-exporter relationships and geographic location, through which I am able to determine the effects of switching frictions on import prices.

The way I measure these frictions is with a structural demand model that incorporates the factors underpinning importer-exporter switching behavior, including geographic components. To do this, I use a model of dynamic discrete choice, pioneered by Rust (1987) in his study of bus engine replacement. I implement the problem in a similar way, using the MPEC methodology for solving discrete choice problems found in Su and Judd (2012) and Dubé et al. (2012). As in those studies, and as in Lincoln and McCallum (2011), my model includes fixed costs entering into a firm's profit/utility function- namely, supplier switching costs at both the partner and city level. Estimates are retrieved through maximizing a likelihood function based on observed outcomes for importer-exporter switching. The model I estimate is most similar to the model of employer choice utilized by Fox (2010) in his study of Swedish engineers. Similar to the use of wages as a driving force behind employee switching behavior, in my context, one of the main components of the "stay or switch" decisions is the price offered to U.S. importers by a Chinese supplier. Roberts et al. (2012) also study exports from China in a structural context. The dispersion of prices among Chinese exporters found in Byrne et al. (2013), and the implications for exporter heterogeneity they find, match my findings of widespread differences among Chinese suppliers even within disaggregated products. Similarly, the role for Chinese cities in exporting echoes the findings in Head et al. (2014) about the heterogeneity of importing patterns among Chinese cities. The sparsity and discrete choice nature of the supplier-matching problem I set up also bears similarity to the "balls-and-bins" models of trade of Armenter and Koren (2014), as well as the discrete choice framework found in Artuç and McLaren (2012). More generally, there have been a number of studies that estimate the effects of relationship networks in other contexts: Joskow (1985) studies contract length among coal suppliers and power plants, while Atalay et al. (2014) measure the extent to which firms rely on subsidiaries versus outside firms for intermediate input purchase. Bernard et al. (2015) study network formation for domestic firm relationships in Japan, while Egan and Mody (1992) and Kranton and Mineheart (2001) present models on the formation of buyer-seller networks and how the properties of these networks affect economic outcomes.

The rest of the paper is organized as follows. Section 2 describes the data sources used in this paper and summarizes the empirical results. Section 3 presents the dynamic discrete choice model with supply chain adjustment costs. Section 4 describes the implementation of the model and summarizes the baseline results. Section 5 describes the quantitative experiments used to determine the importance of the supplier-switching channel. Section 6 concludes.

2 Data and Stylized Facts

2.1 Importer-Exporter Data

I work with 2002-2008 relationship data from the Longitudinal Foreign Trade Transaction Database (LFTTD), which contains confidential information on all international trade transactions by U.S. firms, and is maintained jointly by the U.S. Census Bureau and U.S. Customs.[6] Every transaction of a U.S. company importing or exporting a product requires filing a form with U.S. Customs and Border Protection, and the LFTTD contains the universe of these transactions. In particular, the import data consists of all the information included in customs documents provided by U.S. firms purchasing goods from abroad, including quantity and value exchanged for each transaction, HS 10 product classification, date of import and export, port information, country of origin, and a code identifying the foreign supplier firm. Known as the *manufacturing ID*, or *MID*, the foreign partner identifier contains limited information on the name, address, and city of the foreign supplier. Kamal et al. (2015)) find substantial support for the validity of the MID and its applicability towards studying buyer-seller relationships.[7] I use this variable to provide stylized facts for the amount of churning in U.S.-China trade relationships and the geographic elements of switching behavior. My unit of observation throughout will be a U.S.-China trade relationship, comprising a U.S. *importer* (the combination of a U.S. importing firm and an HS10 product) and a Chinese

[6]I clean the LFTTD using methods outlined in Bernard et al. (2009)) and Pierce and Schott (2012). As in Bernard et al. (2009), I drop all transactions with imputed quantities or values (which are typically very low-value transactions) or converted quantities or values. I concord HS10 codes over time according to the methodology in Pierce and Schott (2012).

[7]They describe the institutional background for the creation of the variable, propose various cleaning methods, and compare the number and composition of exporters from U.S. data with international data. I include a brief summary of the variable, including its construction and the relevant laws surrounding information provided in trade transactions in Appendix A.

supplier.[8]

I focus on U.S.-China trade relationships for three reasons. First, individual exporter varieties within HS10 codes are more likely to be comparable within a country, and China became the largest exporter to the U.S. in my timeframe (2002-2008). Combined with its small share of related party trade, the bloc of U.S.-China trade relationships forms the largest laboratory to explore arm's-length supplier relationships in trade.[9] Second, I combine my quantitative findings with domestic Chinese firm-level data to illustrate how relationship patterns are affected by Chinese production patterns (Section 4.3). Finally, by manually translating firm names and addresses from Chinese production data, I show that the MID is likely to uniquely identify individual foreign suppliers (Appendix A).

Given this focus, I take special care to remove unreasonable values for the MID related to U.S.-China trade.[10] I restrict the sample to importers with an MID country identifier of China (meaning the producing firm is located in China), rather than a "country of origin" of China (which is often based on where the imported good last departed from). Due to the entrepôt nature of Hong Kong's international trade flows and other thorny data issues, I concentrate solely on Mainland China - deleting any observation that appears to come from Hong Kong, Macau, or Taiwan. For example, a three-letter city code of "HON", even with a country code for mainland China, is likely referring to Hong Kong and thus dropped.[11] I also drop any exporter whose city code does not correspond to any of the top 300 cities of China by population.

2.2 Static Relationships: Import Concentration at One Supplier

I begin by characterizing U.S.-China trade relationships in a single year. The main finding is straightforward: even though one U.S. importer has an average of two Chinese suppliers, trade tends to be overwhelmingly clustered within an importer's *primary supplier*- meaning that supplier from which it has the most imports. First of all, 65% of U.S. importers have only a single Chinese supplier. Furthermore, even when restricting to importers with multiple suppliers, the average trade share at an importer's primary supplier is a sizable 64%, compared to an average trade share of 24% an importer's second-largest relationship.

[8]Thus, a single U.S. firm could be counted as many importers. The median number of HS10 products a U.S. firm imports from China is 3.

[9]Only 16% of importers whose main supplier is Chinese will change to a main supplier from another country (including Hong Kong) in the following year.

[10]I eliminate all related-party transactions, as importers receiving goods from another branch of their firm will likely have very different relationship dynamics than arm's-length exporters. 10.2% of all U.S.-China transactions in 2006 are between related parties, accounting for 18.6% of U.S. imports from China.

[11]Other dropped city codes: "KOW" for Kowloon, Hong Kong; "MAC" for Macau; "AOM" for the *Hanyu Pinyin* spelling of Macau, "Aomen"; "KAO" for Kaohsiung, Taiwan.

Primary relationships are also the most important in the aggregate: they account for 62% of all U.S. imports from China. For all of these reasons, my unit of analysis for studying relationship dynamics is the combination of an importer and its primary supplier, what I refer to as a “primary relationship”.

2.3 Dynamic Relationships: Inertia in Suppliers and Locations

Given that primary relationships dominate in the data, I track how often importers maintain their primary supplier and source over time. I define an importer as “staying” with its partner in a year if its primary supplier remained the same from one year to the next.I track the universe of U.S. importers from China in the years 2003-2008, and for each year, determine whether they (a) also imported one year previously, and if so, (b) whether they continued to import from the same primary supplier and/or geographic location as in the previous year. Figure 1 illustrates the fraction of importers staying with their supplier, staying in the same city, and staying in the same province.

Two facts are clear from Figure 1. First, there is a significant share of U.S. importers who maintain the same supplier over time. Even though the number of potential exporting choices is increasing over this time period, the share of importers using the same supplier year-to-year is 45.9%.[12] As a benchmark, given that there are an average of 30 Chinese exporters to the U.S. per HS10 product in the data, if importers were choosing their partners randomly each year, the probability of staying would be 1/30, or 3%. Thus path dependence is far higher than would be expected if importers were choosing their supplier randomly. Secondly, among those firms who do choose to switch, approximately one-third of all importers remain in the same city as their original supplier. Using a similar benchmark as above, random exporter selection would imply a 12-13% chance of staying in the same city.[13] Thus there is strong inertia keeping firms in their original city, even if they choose not to use the same supplier as before.In sum, supplier choices are highly correlated with previous supplier usage, and decisions of whom to switch to are dependent on geographic considerations. It is these stylized facts related to switching, both geographically and over time, that govern the dynamic discrete choice model I lay out in Section 3.

[12]The share of trade accounted for by these relationships over the years 2003-2008 is approximately 43%.

[13]There is an average of 9 cities for each HS10 product, but the number of exporters are not distributed equally across cities. I first calculate the probability an importer originally in city c picks a random supplier also in c as $\frac{X_{ic}}{\sum_c X_{ic}}$, where X_{ic} is the number of exporters of product i in city c. Therefore, the probability *any* importer picks a random supplier also in their original city is the weighted average of these probabilities, where the weights are the number of importers M_{ic} buying from each city. Calculating this object using 2006 data gives a 12% probability of staying in the same city under random supplier choice. Replacing X_{ic} with the value of trade in a city (meaning the probability of staying depends on the distribution of exports, not exporters) gives a 13% probability of staying.

The stylized facts about importer-exporter relationships described above are robust to a number of alternative checks and specifications. One may be concerned that switching is driven by exit on the exporter side: given the structural changes in the Chinese economy over this period, it is likely that many exporters are entering or exiting. To account for this, in Figure 2 Panel A, I recreate the results using only those matches where the same MID is found in both years. The two main facts carry over: a significant share of importers stay with their suppliers, with previous geographic location an important factor in the decision of where switching importers move to. I also check that the results are not driven by defining "staying" to mean buying from the same main supplier. Figure 2 Panel B shows that when considering "staying" to mean buying from keeping at least one supplier within a product category the same, the results are again very similar. These facts are robust to a number of other specifications, including using only importers classified as manufacturing firms, calling an importer a firm-HS6 product combination or a firm, using other definitions of switching, and following the share of importers not adjusting their supply network over a longer time frame than one year.

It is certainly the case that the above results (and all those that follow) are an artifact of the decision to only follow primary relationships over time. That said, there are a number of justifications to do so. First of all, threshold switching - meaning minimal adjustments in import value causing "switches"- is uncommon. Among all recorded switches from 2005-2006, only 6% had a reduction of 25 or less percentage points in the primary supplier's share of the importer's imports.[14] Secondly, the primary supplier measure catches the majority of "big changes" in a supply network. If a new supplier was added that occupied 25% or more of an importer's total trade, the switch measure picked this up 71% of the time; if a supplier that occupied 25% or more of an importer's total trade was dropped, the switch measure captured this 73% of the time. Thirdly, importers seldom add a supplier without concurrently dropping one (and vice versa), meaning that the single-supplier assumption implicit in following only primary supplier relationships is reflective of how most importers source.[15] Finally, only a minority of switching (21.5%) is to a previous year's minor supplier. Thus, in this context, primary relationships are the most relevant unit of analysis for tracking relationship dynamics.

[14] Only 1% of switches involved a change of 10 percentage points or less.

[15] 14% of importers added a supplier in 2006 without dropping a 2005 supplier (accounting for 14% of trade), and 11% dropped a supplier in 2006 without adding a new supplier (6% of trade). 53% (accounting for 73% of trade) both added and dropped.

2.4 Reduced-Form Regression Results

I next analyze the factors that govern the stay-or-switch supplier decisions. There are a number of potential explanations for switching behavior that I can measure using the LFTTD data. Using U.S.-China trade data from 2002-2008, I use the following linear probability models to estimate the relationship between the decision of a U.S. importer to switch Chinese exporting partners and a variety of potential explanatory variables, including price, size and age of the Chinese partner, U.S. importer size, and the date of entry into importing.

$$Stay^j_{i,t,t+1} = \alpha_0 + \alpha_1 StandardPrice^j_{i,t} + \gamma_1 ExpChar^j_{i,t} + \gamma_2 ImpChar^j_{i,t} + f_j + f_t + v_{i,t} \quad (1)$$

$$\begin{aligned} Stay^j_{i,t,t+1} &= \alpha_0 + \sum_{k=1}^{10} \alpha_k \mathbb{1}\left[StandardPriceCat^j_{i,t} = k\right] \\ &+ \gamma_1 ExpChar^j_{i,t} + \gamma_2 ImpChar^j_{i,t} + f_j + f_t + v_{i,t} \end{aligned} \quad (2)$$

As above, I define firm i importing product j as staying ($Stay^j_{i,t,t+1} = 1$) with its supplier if it maintains the largest percentage of imports from the same supplier in time t and $t+1$. I define the "standardized price" ($StandardPrice$) to be the relationship-specific "unit value" from the LFTTD, minus its HS10 product mean and divided by the standard deviation. Thus prices are comparable across industries. In Equation (2), I allow for non-monotonic effects of price on staying by including it as a categorical variable $StandardPriceCat$ representing each of ten deciles. I omit the 5th decile in the regression, in order to check the effects of having both very low and very high prices. I use both exporter characteristics ($ExpChar$) and importer covariates ($ImpChar$). For exporters, I calculate the "Supplier Size" of a Chinese exporter by summing together its total exports to the U.S., and similarly calculate the "Supplier Age" of a Chinese exporter by calculating the first year a MID appears in the trade data. On the importer side, I construct importer size by summing together total imports from China, and the first year of its entry into the Chinese import market by calculating its first appearance in the Chinese import data. All of the above covariates are assigned using their values in the prior year, using later year data only to determine whether or not an importer switched exporting partners. Finally, I include HS10 level fixed effects f_j and year fixed effects f_t. The results of the Linear Probability Models (1) and (2) are reported in Table 1. Standard errors are clustered at the HS10 level.

It is clear that the higher the price a firm paid in a previous year, the lower the probability that it would stay with its main supplier. Though the effect is not significant for prices near the middle of the distribution, the results in Table 1 make it clear that the importers who received the highest prices were 2-3% more likely to switch their supplier than the omitted group (5th decile). Those importers paying the lowest prices were also more likely to stay with their supplier when accounting for importer and exporter characteristics, as can be seen in Columns (2')-(4'). Figure 3 tells the same story as Table 1, and is generated using the results in Table 1 Column (4'). The shaded regions are 99% confidence intervals for the category-specific coefficients. Those importers paying the highest prices are much more likely to switch their supplier than those in the middle and low price regions, while those paying the lowest prices are more likely to stay with their suppliers.

The effects of the various importer and exporter characteristics are also of interest. Table 1 shows that the older and/or larger a Chinese exporting firm was, the lower the probability that a U.S. importer would switch. In addition, larger U.S. firms were most likely to stay with their supplier. Thus there is substantial room for exporter and importer heterogeneity in explaining the staying decision for U.S. importers with their Chinese exporters.

To summarize the regression results, price is an important factor in the decision of an importer to switch partners, especially the magnitude of the price paid in the previous year. Exporter and importer characteristics more generally are also important factors in the decision of whether or not to stay with one's partner. I use these results to guide the modeling of the exporter choice problem below.

3 Model

This section lays out a dynamic discrete choice framework used to model U.S. importer decisions of exporter choice. Different exporters set different prices for the same product j and have heterogeneous quality. Importers of products in that industry make a decision each period about which firm to import from, a decision that is based both on their current exporter and information about other price/quality menus that are available. Switching exporters involves payment of a set of costs, including both an overall switching cost and an additional cost to be paid if an importer finds a new partner in a previously unused city. Each individual exporter of product j at time t is denoted $x_{j,t}$, and exporters are distinguished both by the price they charge $p_{x,j,t}$ and by the quality of their individual variety $\lambda_{x,j,t}$. If importer m chooses the exporter indexed $x_{j,t}$, I denote this match as $x^m_{j,t}$ and the price paid in that match as $p^m_{x,j,t}$.

3.1 Importers

Importers are final good producers, and demand for their variety m has a constant elasticity of substitution demand curve.

$$Q_m = Bp_m^{-\sigma}$$

In the above equation, B is a demand shifter, p_m is the final good price for variety m and σ is the elasticity of substitution.

Final good producer m requires J inputs, indexed $j = 1, ...J$, in order to produce its final good, and production of final good is Cobb-Douglas in labor and quantity $\{I_j\}_{j=1}^J$ of those intermediates:

$$Q_m = L^\alpha \left(\prod_{j=1}^J I_j^{\gamma_j} \right)^{1-\alpha}$$

Although the production function and final demand for its variety are fixed, importer m can choose which exporter for each input. By considering all possible exporters in the market, importers are able to make a profit-maximizing decision between exporters. There are a number of components that affect the decision of which exporter to use.

Firstly, importers make a decision based in part on the expected price they will pay from any exporter, $\mathbb{E}\left[p^m_{x,j,t}\right]$. In particular, importers form expectations about the price from their original partner, and from every other exporter.[16] Since the expectation differs depending on what partner was used, this expectation is both *importer-specific* (m) and *exporter-specific* (x), which allows the same exporter to charge different prices to different importers. I describe the calibration of this expectation in the next subsection.

Secondly, there are frictions involved in finding a different supplier in the following period, modeled as an additional component of the price paid. There is a cost that is paid from switching exporters $\zeta_{X,j}$. Reflecting the geographic nature of switching discussed above, I also include an additional geographic cost $\zeta_{C,j}$ that is paid if an importer uses a separate partner in a separate city.

I define importer m's expected per-unit cost of purchasing intermediate j from supplier $x_{j,t}$ at time t, incorporating the frictions involved in searching for a supplier, in the following

[16]This assumption allows each firm to observe the entire spectrum of prices, even though the observed data on prices is a selected sample, namely, only successful importer-exporter matches.

manner:

$$\bar{p}^m_{x,j,t} = \mathbb{E}\left[p^m_{x,j,t}\right] \exp\left\{\zeta_{X,j}\ \mathbb{1}\{x^m_{j,t} \neq x^m_{j,t-1}\} + \zeta_{C,j}\ \mathbb{1}\{c^m_{j,t} \neq c^m_{j,t-1}\}\right\} \tag{3}$$

where $\bar{p}^m_{x,j,t}$ is the expected cost from purchasing one unit of the intermediate from seller $x^m_{j,t}$, and the indicator functions are equal to one if an importer picks a different partner $x_{j,t}$ from its current match $x_{j,t-1}$, or another different city $c_{j,t}$ from its current city $c_{j,t-1}$. If final good producer m chooses a new partner in the same city (c_{t-1}) as its old partner, then only $\zeta_{X,j}$ is paid, while if an exporter in a separate city is chosen, $\zeta_{X,j} + \zeta_{C,j}$ is paid. This means that the cost of an input bundle will differ depending on what supplier is chosen, not just because of a higher or lower offered price, but also because of costs of switching one's current partner.

Let $X^m_t = \left\{x^m_{j,t}\right\}^J_{j=1}$ be the vector of supplier choices made by importer m across inputs at time t. Then with wage w, the expected cost of an input bundle for the final good is:

$$c_m\left(X^m_t\right) = w^\alpha \left(\prod_{j=1}^J \left[\bar{p}^m_{x,j,t}\right]^{\gamma_j}\right)^{1-\alpha}$$

Producing one unit of the final good for a final good producer with productivity ϕ requires $\frac{1}{\phi}$ input bundles, each with cost depending on the vector of suppliers X^m_t. I assume that the productivity of a final good producer depends on factors unobserved by the econometrician (such as the quality of the supplier's product) that are particular to its individual supplier match. In particular, productivity for producer m is multiplicative in a common element for that producer ϕ_m and the "quality" of the variety from exporter x, $\lambda_{x,j,t}$.[17]

$$\phi_m\left(X^m_t\right) = \psi_m \prod_{j=1}^J \lambda^\nu_{x,j,t}$$

The marginal cost of an importer m with productivity ϕ_m is:

$$MC\left(X^m_t\right) = \frac{1}{\phi_m\left(X^m_t\right)} c_m\left(X^m_t\right) \tag{4}$$

Maximizing expected profits at time t means that importer m must set the price of their

[17]Given the richness of data available, I implement a model that takes explicit account of "quality" considerations, in particular, those characteristics of an exporting firm that are observed by the potential importer, but unobserved by the econometrician and tend to be correlated with the price. I use the control function approach of Kim and Petrin (2010). I specify the estimating procedure in Section 4.2 below.

final good optimally and choose the optimal vector of exporter choices X_t^m:

$$\pi_t^m = \max_{p_m, X_t^m} p_m Q_m - MC\left(X_t^m\right) Q_m$$

Using the assumption of CES demand, the optimum price of the final good for producer m is a markup over the marginal cost, $p_m = \frac{\sigma}{\sigma-1} MC\left(X_t^m\right)$. Plugging this and our expression for marginal costs (4) into the above profits equation gives the following equation:

$$\pi_t^m = \max_{X_t^m} \frac{1}{\sigma} B \left(\frac{\sigma}{\sigma-1}\right)^{1-\sigma} \left[\phi_m\left(X_t^m\right)\right]^{\sigma-1} c_m\left(X_t^m\right)^{1-\sigma} \tag{5}$$

The multiplicative nature of Equation (5) and its subcomponents means that supplier choices are independent across inputs. To see this, taking logs of (5), plugging in, and defining log expected profits attributable to input j as $\ln \pi_{j,t}^m$, we have:

$$\ln \pi_t^m = A + \ln \pi_{j,t}^m + \sum_{k \neq j} \ln \pi_{k,t}^m$$

where

$$\ln \pi_{j,t}^m = \max_{x_{j,t}^m} \nu\left(\sigma-1\right) \ln \lambda_{x,j,t}$$
$$+ \left(1-\alpha\right)\left(1-\sigma\right)\gamma_j \left[\mathbb{E}\left[\ln p_{x,j,t}^m\right] + \zeta_{X,j} \mathbb{1}\{x_{j,t}^m \neq x_{j,t-1}^m\} + \zeta_{C,j} \mathbb{1}\{c_{j,t}^m \neq c_{j,t-1}^m\}\right] \tag{6}$$

and $A = \ln\left\{\frac{1}{\sigma} B \left(\frac{\sigma}{\sigma-1}\right)^{1-\sigma} w^{\alpha(1-\sigma)} \psi_m^{\sigma-1}\right\}$ captures all the terms not associated with the cost of an input bundle for the final good.[18][19] Since the decision of input j is wholly separate from the decision of other inputs, I now focus attention only on the market for one input and drop the j subscript.

How does an importer decide which exporter maximizes profits? It is a maximization problem of discrete choice, so expected profits are calculated for each choice, and the partner with the highest expected profits will be chosen. I define the *exporter-specific* expected log profit term for any choice x_t^m as $\overline{\pi}_t^m\left(x_t^m, \boldsymbol{\beta}\right)$:

$$\overline{\pi}_t^m\left(x_t^m, \boldsymbol{\beta}\right) = \xi \ln \lambda_{x,t} + \beta_P \mathbb{E}\left[\ln p_{x,t}^m\right] - \beta_X \mathbb{1}\{x_t^m \neq x_{t-1}^m\} - \beta_C \mathbb{1}\{c_t^m \neq c_{t-1}^m\} \tag{7}$$

where $\xi = \nu\left(\sigma-1\right)$, $\beta_P = -\left(1-\alpha\right)\left(\sigma-1\right)\gamma$, $\beta_X = \left(1-\alpha\right)\left(\sigma-1\right)\gamma\zeta_X$, and $\beta_C =$

[18] In Equation (6), I use Jensen's Inequality and the fact that the expected price is almost-surely constant to assert that the log of the expected price is equal to the expected log price.

[19] Intuitively, for elasticities of substitution $\sigma > 1$, higher quality $\lambda_{x,j,t}$ leads to higher profits, while higher prices lead to lower profits.

$(1-\alpha)(\sigma-1)\gamma\zeta_C$. I summarize the vector of unknown parameters as $\boldsymbol{\beta} = \{\beta_P, \beta_X, \beta_C, \xi\}$.

The final element of the importer's decision is a stochastic profit shock. For the dynamic profit problem, as in Rust (1987), I allow for a shock from choosing x_t that is observable to the importer, $\epsilon^m_{x,t}$. This profit shock stands in for all of the other unobservable traits not included in (7) that may result in a match.

Equation (7) is the cornerstone of my estimation strategy.[20] Starting from a general model of importing behavior, I have derived a choice-specific profit equation that can be estimated using MPEC.[21] To summarize this equation in words, if importer m chooses a different exporter than they used in the previous period, then this firm must they must pay a fixed cost β_X, while if they use a different exporter in a different city, they pay $\beta_X + \beta_C$. The parameter β_p is a measure of how sensitive switching is to changes in price. Estimating $\{\beta_P, \beta_X, \beta_C\}$ product-by-product will provide a measure of the frictions firms face in switching partners and locations, and enable the posing of counterfactual experiments.

Estimating the unknowns in Equation (7) can be achieved by calculating expected log profits from each potential importer-exporter match, using observed outcome data to find the most likely values. Put differently, it is possible to rank an importer's expected profits from choosing any exporter x_t, meaning that the observed exporter choice in the data should be the one such that expected profits from that exporter are higher than all other potential choices. I use data on observed outcomes, prices, and estimated exporter quality as well as assumptions about the evolution of the stochastic profit shock to solve for the parameters via maximum likelihood estimation of the firm's dynamic profit maximization problem.

In order to illustrate which state variables enter into the dynamic estimation of Equation (7), I first lay out how price expectations for any potential importer-exporter relationship are formed.

3.2 Prices

3.2.1 Exporters

Within any product category j there are numerous exporters x producing individual varieties. They set the price for their variety at time t based upon their firm-specific marginal cost, which in turn depends on their quality choice $\lambda_{x,t}$. I follow the same functional form for exporter quality as in Hallak and Sivadasan (2013), and continue to drop the j subscript.

[20] In Section 3.5, I describe a variant of Equation (7) with an additional term capturing whether the importer and exporter are in an ongoing relationship, allowing me to assess any potential bias in my estimates from serial correlation in exporter choice.

[21] This equation resembles the worker utility function from choosing different employers discussed in Fox (2010).

The optimal price is:

$$p_{x,t} = \mu_x MC_{x,t} = \mu_x \frac{w_{x,t}}{z_x} (\lambda_x) \tag{8}$$

where z is the idiosyncratic productivity of exporter x, w is the wage, and, as above, λ is quality. I assume that exporters simply set a constant markup μ_x over time. Finally, I allow for an additive error term to affect the exporter's observed price to an actual importer, $u^m_{x,t}$. This means that the price an individual importer will pay to an exporter over time can be written:

$$\ln p^m_{x,t} = \ln p^m_{x,t-1} + (\ln w_{x,t} - \ln w_{x,t-1}) + u^m_{x,t} \tag{9}$$

3.2.2 Price Evolution

Equation (9) means that the importer can partially predict every exporter's future price that they will face. The first component of the importer's expected price is $\ln p^m_{x,t-1}$. I generate this object for each possible match by assuming that importers expect this component to be the price they paid last period if they keep their partner, and the average price all other importers paid from a given supplier if they switch to that supplier. The second component of the importer's expected price is the change in wages, $\ln w_{x,t} - \ln w_{x,t-1}$. I cannot measure this directly, but proxy for this by including a term to capture city-specific trends in price changes. Finally, there is the relationship-specific component $u^m_{x,t}$. I generate the second and third components by tracking deviations from price stability over time, and attribute them to city and relationship-specific terms.

If importer m stays with its current partner x in city c, the price process is:

$$\ln p^m_{x,t} = \ln p^m_{x,t-1} + \eta_{c,t} + u^m_{x,t}, \quad u^m_{x,t} \sim \mathcal{N}\left(0, \sigma^2_{x,t}\right) \tag{10}$$

Importing firms know there is a city-specific change in prices $\eta_{c,t}$, as well as the distribution of an exporter-specific realization of a shock. The city-specific price shocks are correlated for firms in the same city, so there is a city component and an exporter-specific component.

If importer m decides to use a different partner $\widetilde{x}$, then the price process is:

$$\ln p^m_{\widetilde{x},t} = \frac{1}{N_{\widetilde{x}}} \sum_{n=1}^{N_{\widetilde{x}}} \ln p^n_{\widetilde{x},t-1} + \eta_{\widetilde{c},t} + u^m_{\widetilde{x},t} \quad u^m_{\widetilde{x},t} \sim \mathcal{N}\left(0, \sigma^2_{\widetilde{x},t}\right) \tag{11}$$

$N_{\widetilde{x}}$ is the number of firms who imported from firm $\widetilde{x}$ from the previous period, and they are indexed $n = 1, ...N_{\tilde{x}}$. Each price paid by importer n is $p^n_{\tilde{x},t-1}$. As above, there is both a city shock to prices and an importer-specific realization of the exporter price shock.

Given the specification of these different prices based on shocks calibrated to data in period $t-1$ and t, I can write down a density function for prices $f\left(p^m_{x,t} \middle| \mathbf{p_{t-1}}, x^m_{t-1}, x^m_t\right)$, where $\mathbf{p_{t-1}}$ is the vector comprising each exporter's price at time $t-1$. In other words, given state variables $\mathbf{p}$ and x, and a potential future choice x', the price p' from that x' can be predicted by any importer. Notably, as a result of different price expectations for an exporter depending on whether an importer used them previously, the expected price has an m superscript, $p^m_{x,t}$ (the same exporter could have different predicted prices depending on whether the importer used them in the prior period).

The parameters $\{\eta_{c,t}\}^C_{c=1}$ (where C is the total number of cities) are the mean changes in price for any city, and are known by the importing firm. I assume the stochastic parameters $u^m_{x,t}$ are normally distributed with mean zero and standard deviations $\sigma_{x,t}$ (each exporter has a particular distribution of price shocks known to importers, but the specific value of which is observed only after the match occurs). I use the LFTTD data to calibrate the parameters $\{\{\eta_{c,t}\}^C_{c=1}, \{\sigma_{x,t}\}^X_{x=1}\}$ by using observed prices in both pre- and post-periods and estimating equations (10)-(11) for each product.[22]

3.3 Value Function

Individual importers make their choice of exporter based on price concerns, quality concerns and any added costs involved from changing their current exporter. Entering period t, importer m has a collection of state variables that affect its optimal choice x^m_t, following the exporter-specific profit term given by $\overline{\pi}^m_t\left(x^m_t, \boldsymbol{\beta}\right)$ in Equation (7). Those state variables are the identity of the exporter used last period, x^m_{t-1} with location c^m_{t-1}, and (in order to form price expectations) the set of prices charged by all exporters last period $\mathbf{p}_{t-1}$. Based on these state variables, the profit shock $\epsilon^m_{x,t}$, the costs of switching one's current exporter, and the other components of $\boldsymbol{\beta}$, the importing firm will choose which exporter to use in the current period, x^m_t. After all choices are made, the true vector of prices $\mathbf{p_t}$ is realized, as is the matrix of next period's profit shocks $\boldsymbol{\epsilon_{t+1}}$. These states evolve according to the joint density $h\left(\mathbf{p_t}, \boldsymbol{\epsilon_{t+1}}\right)$.

Infinitely-lived importer m chooses an exporter x in each period in order to maximize the present discounted stream of expected profits. With single-period expected profits described

[22] If no prior year information is available for a potential supplier- i.e. an importer chooses a supplier that did not exist in the previous year- I allow the expected price to be the average price among all exporters in that city in the previous period. If there is no city information in the previous period, I drop that exporter. If an exporter is only found in the pre-period, then I calibrate $\{\eta_{c,t}, \sigma_{x,t}\}$ using all other firms and use them to form the expected price from using that exporter.

by Equation (7), the infinite-time problem for any importer (dropping the m superscript) is summarized by the following value function:

$$V\left(x_{t-1}, \mathbf{p_{t-1}}, \boldsymbol{\epsilon_t}\right) = \max_{\{x_t, x_{t+1}, \ldots\}} \mathbb{E}\left[\sum_{\tau=t}^{\infty} \delta^{\tau-t}\left(\overline{\pi}_\tau\left(x_\tau, \mathbf{p_{\tau-1}}, x_{\tau-1}, \boldsymbol{\beta}\right) + \epsilon_{x,\tau}\right)\right] \tag{12}$$

where the expectation operator is taken over the importer's knowledge about the possible evolution of $(p_{x,t}, \epsilon_{x,t+1})$, governed by the density $h\left(p_{x,t}, \epsilon_{x,t+1}\right)$ at every period t. Recall that the price from choosing exporter x_t is not known before making the choice, but is predicted based on $\mathbf{p_{t-1}}, x_{t-1}$ and x_t, according to the density function $f\left(p_{x,t}|\mathbf{p_{t-1}}, x_{t-1}, x_t\right)$.

Writing the one-step ahead value of any variable a as a', the value function in (12) can be rewritten as a Bellman Equation:

$$V\left(x, \mathbf{p}, \boldsymbol{\epsilon'}\right) = \max_{x'} \quad \overline{\pi}\left(x', \mathbf{p}, x, \boldsymbol{\beta}\right) + \epsilon'\left(x'\right) + \delta EV\left(x', \mathbf{p}, x, \boldsymbol{\epsilon'}\right) \tag{13}$$

for

$$EV\left(x', \mathbf{p}, x, \boldsymbol{\epsilon'}\right) = \int_{\mathbf{p'}} \int_{\boldsymbol{\epsilon''}} V\left(x', \mathbf{p'}, \boldsymbol{\epsilon''}\right) h\left(\mathbf{p'}, \boldsymbol{\epsilon''}|\mathbf{p}, x, x', \boldsymbol{\epsilon'}\right) d\mathbf{p'} d\boldsymbol{\epsilon''}.$$

At this point, I make a key assumption about the joint density of the state variables and the profit shock: that they evolve separately from each other.

Assumption 1 (Conditional Independence) *The joint transition density of $p_{x,t}$ and $\epsilon_{x,t+1}$ can be decomposed as:*

$$h\left(p_{x,t}, \epsilon_{x,t+1}\right) = g\left(\epsilon_{x,t+1}\right) f\left(p_{x,t+1}|\mathbf{p_t}, x_t, x_{t+1}\right)$$

I also assume that the profit shock ϵ is distributed according to a multivariate extreme value distribution, with known parameters:

Assumption 2 *The profit shock is distributed Type I Extreme Value (Gumbel). The cumulative distribution function G is*

$$Pr\left(\epsilon_t < y\right) = G\left(y\right) = \exp\left\{-\exp\left\{-y - \gamma\right\}\right\}$$

for $\gamma = 0.577...$ (Euler's constant).

These two assumptions permit the computation of conditional choice probabilities observing any particular outcome[23]:

Proposition 1 *Let the value of a present time variable a one period ago be written as a_{-1}, and one period in the future be written as a'. Given Assumptions 1 and 2, and grouping together the state variables as $s = \{\mathbf{p_{-1}}, x_{-1}\}$, the probability of observing a particular exporter choice x^C conditional on state s and parameters $\boldsymbol{\beta}$, $P\left(x^C|s,\boldsymbol{\beta}\right)$, is:*

$$P\left(x^C|s,\boldsymbol{\beta}\right) = \frac{\exp\left[\overline{\pi}\left(x^C,s,\boldsymbol{\beta}\right) + \delta EV\left(x^C,s\right)\right]}{\sum_{\widehat{x}\in X}\exp\left[\overline{\pi}\left(\widehat{x},s,\boldsymbol{\beta}\right) + \delta EV\left(\widehat{x},s\right)\right]} \tag{14}$$

where the function $EV(x,s)$ is the solution to the fixed point problem:

$$EV(x,s) = \int_{s'} \log\left\{\sum_{x'\in X}\exp\left[\overline{\pi}\left(x',s',\boldsymbol{\beta}\right) + \delta EV\left(x',s'\right)\right]\right\} f\left(s'|s,x\right) \tag{15}$$

Proof *See Appendix B.*

3.4 Maximum Likelihood Estimation

The parameters $\boldsymbol{\beta}$ can be computed via maximum likelihood estimation. Let x_t^m be the actual choice of exporter at time t for importer m from data. Then the likelihood of observing importer m choosing exporter x_t^m is:

$$L\left(x_t^m|\mathbf{p_{t-1}},x_{t-1}^m,\boldsymbol{\beta}\right) = P\left(x_t^m|x_{t-1}^m,\mathbf{p_{t-1}},\boldsymbol{\beta}\right)\cdot f\left(p_{x,t}^m|\mathbf{p_{t-1}},x_{t-1}^m,x_t^m\right)$$

And thus the total likelihood function for the set of importer choices $(m = 1,...M)$ at time t is:

$$\mathcal{L}\left(\boldsymbol{\beta}\right) = \prod_{m=1}^{M} P\left(x_t^m|x_{t-1}^m,\mathbf{p_{t-1}},\boldsymbol{\beta}\right)\cdot f\left(p_{x,t}^m|\mathbf{p_{t-1}},x_{t-1}^m,x_t^m\right)$$

The constraints for the maximization problem are the system of fixed point equations defined by Equation (15). To solve this problem, I follow the MPEC approach as described in Su and Judd (2012) and Dubé et al. (2012), namely an inner loop for solving the fixed point problem in (15) for the constraint vector **EV** and $\boldsymbol{\beta}$, and testing each candidate $\boldsymbol{\beta}$ within the (log) likelihood function to see where the function is maximized. Thus the problem to

[23] Artuç and McLaren (2012) derive a very similar conditional choice probability function for their model of mobility costs for workers affected by offshoring.

solve is:

$$\max_{\boldsymbol{\beta}} \log \mathcal{L}(\boldsymbol{\beta}) =$$

$$\max_{\boldsymbol{\beta}} \sum_{m=1}^{M} \log \frac{\exp\left[\overline{\pi}_t^m\left(x_t^m, s_t^m, \boldsymbol{\beta}\right) + \delta EV\left(x_t^m, s_t^m\right)\right]}{\sum_{\widehat{x}_t^m \in X} \exp\left[\overline{\pi}_t^m\left(\widehat{x}_t^m, s_t^m, \boldsymbol{\beta}\right) + \delta EV\left(\widehat{x}_t^m, s_t^m\right)\right]} + \sum_{m=1}^{M} \log f\left(p_{x,t}^m | s_t^m, x_t^m\right) \tag{16}$$

$$s.t.$$

$$EV\left(x_t, s_t\right) = \int_{s_{t+1}} \log \left\{ \sum_{x_{t+1} \in X} \exp\left[\overline{\pi}_{t+1}\left(x_{t+1}, s_{t+1}, \boldsymbol{\beta}\right) + \delta EV\left(x_{t+1}, s_{t+1}\right)\right] \right\} f\left(s_{t+1} | s_t, x_t\right) \tag{17}$$

Solving this problem produces maximum likelihood estimates for the vector of parameters $\boldsymbol{\beta}$ for each product.

3.5 Extension: Long-Term Relationships

One concern from using Equation (7) in estimating the dynamic discrete choice model of the above section is that the independent and identically distributed profit shock $\epsilon_{x,t}^m$ that captures unobserved factors governing relationship choice could be serially correlated. For example, if a multi-year importer-exporter relationship encourages an importer to continue buying from that exporter, then switching cost estimates would be biased upwards from not considering this dimension.[24] Indeed, for U.S.-China imports in 2005, 23% of relationships are at least one year old.

In order to address this issue, I estimate a variant of Equation (7) that allows for long-term relationships to matter:

$$\overline{\pi}_t^m\left(x_t^m, \boldsymbol{\beta}\right) = \xi \ln \lambda_{x,t} + \beta_P \mathbb{E}\left[\ln p_{x,t}^m\right] - \beta_X \mathbb{1}\{x_t^m \neq x_{t-1}^m\} - \beta_C \mathbb{1}\{c_t^m \neq c_{t-1}^m\} + \omega \mathbb{1}\{Rel_{x,t}^m = 1\} \tag{18}$$

where $Rel_{x,t}^m$ is an indicator for whether importer m and exporter x had a pre-existing relationship- meaning that the use of x^m at time $t-1$ was not the first period the importer and exporter were matched. The equation is identical to Equation (7), but includes an extra term allowing for a long-term relationship between importer m and exporter x to affect the importer's supplier decision.

[24] I thank Andrew Bernard for this suggestion.

4 Estimation

This section has three objectives: first, I describe the specific assumptions involved in discretizing the state space for the constrained maximization problem in Equations (16)-(17). Second, I specify the process used to calculate quality. Third, I present analysis of the raw structural parameters obtained from the solution of the constrained maximization problem using U.S.-China trade data.

4.1 Implementation

In order to estimate the above model, I need to solve the system of equations defined by (17) for the unknown elements EV and $\boldsymbol{\beta}$. To do this, I discretize the price state space into N intervals, allowing me to rewrite the fixed point equation (17) as:

$$EV\left(x_t, \widehat{s}_t\right) = \sum_{\widehat{s}_{t+1}=1}^{N} \log\left\{\sum_{x_{t+1}\in X} \exp\left[\overline{\pi}_{t+1}\left(x_{t+1}, \widehat{s}_{t+1}, \boldsymbol{\beta}\right) + \delta EV\left(x_{t+1}, \widehat{s}_{t+1}\right)\right]\right\} Pr\left(\widehat{s}_{t+1}|\widehat{s}_t, x_t\right) \tag{19}$$

where $\widehat{s}_t = \{\widehat{p}_{t-1}, x_{t-1}\}$, and $\widehat{p}$ is the midpoint of each price interval, chosen such that $\frac{1}{N}$ of all firms are in each interval. My use of MPEC in solving the maximum likelihood model follows the description from Su and Judd (2012) and Dubé et al. (2012). The MPEC maximization protocol uses values of the vector $\boldsymbol{\beta}$ that satisfy the fixed point Equation (17), given expected prices and price transition probabilities for each potential choice, and selects the vector that delivers the highest likelihood. As a simple example to fix ideas, suppose there are 30 exporters in an industry and N discrete price states. Then there are $30N$ possible state values and 30 possible choices, meaning that the vector **EV** contains $900N$ elements, one for each value of $EV\left(x_t, \widehat{s}_t\right)$. Thus the constraint set in (19) is a fixed point problem of $900N$ equations and $900N + 4$ unknowns, where the additional 4 unknowns are $\boldsymbol{\beta} = \{\beta_P, \beta_X, \beta_C, \xi\}$. Each of the possible values of $\boldsymbol{\beta}$ and **EV** that satisfy these constraints are tested in the objective function (16) to see which give the closest match between the estimated probabilities and the true data. Asymptotic normal standard errors are calculated following Rust (1994).

4.2 Data Preparation

In this section, I describe how I bring the LFTTD data to the model, in order to run the MLE problem of Equation (16) with the constraints in Equation (19). To do so, I must calculate the one-period log expected profits from an importer choosing each potential exporter x_t as given in Equation (7). There are four elements to this profit equation:

- Whether x_t is different from an importer's previous partner.
- Whether x_t is located in a different city from an importer's previous partner.
- The expected price of x_t.
- The quality of x_t.

The first two are easily identified using the MID variable discussed in Section 2. I described the process for generating expected prices in Section 3.2.1. It remains to specify the process through which I estimate the quality of an exporter, λ. I consider exporter quality as an estimate of exporter heterogeneity, with intuition related to Kim and Petrin (2010) and Khandelwal (2010): if exporters are very similar in terms of observables but one charges a higher price, then that one has a higher quality.

Specifically, I use the control function methodology of Kim and Petrin (2010) to account for unobserved supplier heterogeneity that is likely to be positively correlated with the price, via the following regression:

$$\ln p_{x,t} = Z_{x,t} + \ln \lambda_{x,t} \tag{20}$$

This equation follows from (8), the specification of exporter prices described in Section 3.2. $Z_{x,t}$ is a vector of firm-specific regressors. From the trade data, I cannot observe the productivity or wage of individual Chinese exporters. However, there are a number of exporter-specific variables in the data that I use to capture supplier heterogeneity, including total U.S. exports, exporter market share within an HS code, number of HS products exported, number of years exporting to the U.S., number of import partners, and number of transactions. For each industry and together with time fixed effects, I regress the exporter's average offered price (firm-level unit value) on these variables. I then take the residual from this regression and use it as a measure of quality. The resulting approach is similar to the "quality ladder" estimation for varieties used by Khandelwal (2010).[25]

Finally, I describe the final pieces of the puzzle necessary to run the MLE problem above. Firstly, some U.S. importers use multiple exporters each year. As before, rather than counting every possible permutation of exporters as a discrete choice, I restrict attention to that exporter from which a U.S. importer obtained the plurality (highest percentage) of its

[25] As a check, I estimate the model on goods that are considered highly homogeneous across sellers. The results from including quality and leaving out quality in these industries are qualitatively similar.

imports from each year. Thus the "choice" in the discrete choice model is which exporter the firm imports the most from, rather than which exporter the firm uses.

Additionally, given the fact that not every exporter is found in both periods, I have to take a stand on the set of potential exporters X. I define the set of possible exporter choices at $t + 1$ broadly, consisting of a) any exporter found in time t and b) any new exporter in time $t + 1$, as long as I know what price they charged in time t. As described above, I am making the exporter choice one of "where do I get my majority of imports from", meaning it is possible that we have some "new" exporters found in time $t+1$ that have price information from time t, even though they did not actually appear as any importer's main supplier in time t.

The last step is to clean the LFTTD by eliminating unreasonable prices. Before averaging prices across transactions, I eliminate any transactions with prices both a) outside the 95th percentile range and b) ten times the median price for each product/industry.[26] I follow the above procedure to estimate the model for a large number of products and industries, using data on U.S.-China trade from 2005-2006, and the software TOMLAB / KNITRO to implement MPEC.

4.3 Estimation Results

A key empirical finding confirmed by the estimation procedure is the large size of switching costs. The left panel of Table 2 illustrates the summary results from running the model on 50 of the most-imported HS10 products from China in 2005.[27] The value-weighted average of switching costs across the sample of products is $\beta_X = 2.75$ and $\beta_C = 2.80$.[28] The numeric results are interpreted in units of the Type I Extreme Value shock, which has mean 0 and standard deviation $\sqrt{\frac{\pi^2}{6}} \approx 1.29$, giving the following implication: for an importer to be indifferent between its current partner and some other potential partner in the same city charging the same price, the new partner must provide a shock to profits that is approximately $2.75/1.29 = 2.1$ standard deviations above the mean. If that new partner is located in a separate city, then that new partner must provide a positive shock to profits

[26] Unit values in the LFTTD are particularly prone to wildly unreasonable outliers, sometimes caused by firms writing down a quantity of 1 instead of the standard quantity that should be used for a product, for example. This procedure allows me to reduce egregious outliers while still maintaining prices that are not too far away from the median values.

[27] Taken together, account for 15% of U.S. imports from China. Appendix C describes how products are selected. The list of products and their trade shares are listed in Table C1.

[28] The weighted average over products is my preferred summary measure, as it matches the construction method of the U.S.-China Import Price Index I use for the counterfactual experiments in Section 5.

that is $(2.75 + 2.80)/1.29 = 4.3$ standard deviations from the mean.[29] I do the identical exercise for 50 randomly selected HS6 industries (the right panel of Table 2) and find results of very similar magnitude. Interestingly, the weighted-average of the city cost is smaller for more-aggregated industries than for products, meaning HS6 industry exports are more disaggregated spatially. This intuition dovetails with the well-known increase in industrial agglomeration within Chinese cities over this time period, as shown in Lu and Tao (2009).[30]

Table 3 shows the results of the extension to account for relationships, following Equation (18). Again using the weighted average measure, the benefit from being in a continuing relationship ω does appear to be marginally positive, but does not substantially alter the switching cost estimates. Although 23% of U.S.-China trade relationships in 2005 were at least one year old, the model is capturing the result that such relationships, especially from China, become increasingly unlikely to survive as time goes on (Monarch and Schmidt-Eisenlohr (2015)).

Costs are highly heterogeneous across products and industries. Rather than presenting the full battery of results (available upon request), in this subsection I present estimates of the parameters in illustrative industries.[31] These estimates are presented in Table 4.

I begin by presenting results for industries that have noteworthy spatial characteristics related to the location of exporters.[32] The HS6 industries "Hand Pumps for Liquids" (HS6 841320) and "Files, Rasps, and Similar Tools" (HS 820310) are both characterized by fairly low degrees of out-city switching (30% and 15% from 2005 to 2006, respectively). Thus the estimates should reflect the fact that switching cities is more costly through higher city-switching costs relative to partner-switching costs. On the other hand, the HS6 industries "Portable Digital ADP Machines" (HS6 847130) and "Motorcycles, Side-cars, Reciprocating Engine of a cylinder capacity greater than 50 cc but not exceeding 250 cc " (HS 871120) are characterized by very high levels of inter-city switching among switching firms: 72% find a partner in another city in HS 847130, while 86% switch cities in HS 871120. Thus the size of β_C relative to β_X should be much smaller than in the previous two industries, given that calculations of these parameters take into account how much switching is actually occurring.

[29] Given that the price term enters into the profit equation (7) non-linearly and is discretized, it is not possible to assign a "dollar value" for these costs. Since there are very different prices for the products I use, however, the measure of deviations from the mean shock is appealing in that costs across products (and differences in costs between products and industries) can be analyzed meaningfully.

[30] This result also represents the flip-side of the results from Head et al. (2014) that within narrow product categories, sourcing behavior of Chinese cities is very heterogeneous, and is often driven by buyer-supplier shocks. The same appears true for export patterns of Chinese cities.

[31] Not all industries discussed below are included in the sample of 50 industries above.

[32] Asymptotically normal standard errors are calculated following Rust (1994), and are available upon request. All estimates in this subsection and that are used in the counterfactuals are significantly different from zero according to the 99% threshold.

The first panel of Table 4 demonstrates this to be true: those industries with relative low levels of city switching have higher relative levels of city-switching costs, and the opposite for industries with greater city switching.

Another illustrative comparison is to examine the "slackness" of the market for imports-how many available exporters there are compared to the number of importers. Though most industries tend to have similar numbers of importers and exporters, the industry "Floor Coverings, Wall or Ceiling Coverings, of Polymers of Vinyl Chloride" (HS6 391810) has many more importers than available exporters: 58 to 42. Thus this is a market where importers are truly competing for exporters, a trend that should be reflected in low switching and high partner-switching costs. The middle panel of Table 4 demonstrates this to be true: the cost of switching exporters is high relative to other industries and products, and represents a larger fraction of the total costs borne by switching both partner and city.

Next, I present selected results for textile imports from China, for the purpose of illustrating differences in elasticities of substitution. According to estimates for HS10 categories in U.S. imports provided in Broda and Weinstein (2006), the industries "Gloves, Impregnable Plastic 4Chtt less than 50% Cotton, Man-Made Fibers, kt" (HS10 6116106500) and "Footwear, soles of rubber/plastic/leather, upper leather other protective toe-cap" (HS10 6403406000) have particularly high elasticities of substitution (between 9 and 15, where 20 is the generally accepted upper bound on feasibility of the estimate). High elasticities of substitution mean importers buying these products are very sensitive to changes in price, and are more likely to change their behavior in response to price changes. Although different elasticities of substitution are not included explicitly in the model, the characteristics of highly elastic industries should imply a high sensitivity to price changes, and indeed a strong response of switching, measured through β_P. On the other hand, industries such as "Men's Underpants and Briefs of Manmade Fibers, Knit" (HS10 6107110010) and "Ski/Snowmobile Gloves of Synthetic Fibers" (HS10 6116930800) are industries with particularly low elasticities of substitution (between 1 and 2, where the non-inclusive lower bound from the estimation procedure is 1). In contrast to the industries described above, we would expect very weak responses of switching to price changes: these firms will not adjust their partners in response to price changes. The lower panel of Table 4 confirms these conclusions about switching and the elasticity of substitution: we see negative values for β_P in the industries earmarked as having very high substitution elasticities, whereby an increase in price of one log point implies lower profits and thus, according to the model, a switch more likely. On the other hand, the high values of β_P imply that importers in those industries are not sensitive to price changes. This unresponsiveness shows that these firms are simply not likely to switch, as higher prices do not alter their choices.

I further make use of concordances developed by Brandt et al. (2013) between HS codes and the China Industry Code (CIC) system to analyze different types of industries based on their China-based production characteristics. Using China National Bureau of Statistics firm-level data for 2005, I isolate CIC industries where exporters have particular characteristics potentially related to the importer-exporter partnership decision. I then use the HS6-CIC concordances to estimate the switching behavior parameters among all firms importing in that CIC code. I summarize the estimates according to their underlying traits in Table 5.

First, the labor productivity of workers in different exporting industries influences switching. Chinese exporters in the industry "Manufacturing of Teaching Aids and Materials" (CIC 2413) are in the lower tail of value added per worker relative to other industries, while exporters in the industry "Rolling and Processing of Rare Earths" have very high levels of value added per worker. As can be seen from the top panel of Table 5, it is more common for U.S. importers to switch partners in the industry that uses with low-skilled workers, where the relationship-specific benefit is likely to be smaller. We see a similar result for the size of the Chinese exporting firm: Exporters in the industry "Other Ward Care and Medical Equipment" (CIC 3689) have low levels of employment compared to exporters in the industry "Arms and Ammunition" (CIC 3663), and subsequently higher switching costs. This evidence suggests that smaller firms generally seem to be better tailored to specific needs of importing firms, which is in line with earlier findings in the literature, such as Blum et al. (2013).

In summary, the results are broadly what we would expect of the estimates *ex ante*: higher exporter switching costs appear in products and industries with low levels of inter-city switching, many importers, and highly skilled workers. Lower exporter switching costs are found in industries with high levels of inter-city switching, and a high proportion of large firms. The next section uses the set of quantitative estimates to perform counterfactual experiments about the role of these frictions in import prices and trade flows.

5 Counterfactual Experiments

5.1 Changes in switching costs

Switching costs in this model can be interpreted as import market frictions, by which firms would like to import from particular other firms, but for some reason (lack of information, poor logistics, etc.) do not actually import from these partners. There are potential efficiency gains to be realized if these costs were reduced, and importers could enjoy lower-priced alternatives rather than remaining "stuck" with their previous exporting partner. The structural model I estimated above allows me to assess how U.S. import prices from China would change in response to reduced switching costs.

I follow the procedure outlined by the BLS Handbook of Methods to calculate the Import Price Index for my sample (Bureau of Labor Statistics (2013)).[33] Using the same samples of products and industries as above, I then generate data according to a new set of parameters for each industry that reflect differences in switching costs, keeping the state variables the same for each firm (supplier and price in the previous period).

To make the price index, I first generate the product/industry price index P_j, which sums together firm-level prices, weighted by the share of one firm's imports in total industry imports:

$$P_j = \sum_{i=1}^{M_j} \omega_i p_i \tag{21}$$

In the above, p_i is the median of firm i's received price across 1000 replications of the counterfactual, and M_j is the number of importers of product j.[34] I weight each firm i by the value of its imports relative to total imports in that industry, ω_i. Given these industry level price ratios, I aggregate up using the share of industry imports in total trade ω_j across the industries in my sample:

$$P = \sum_{j \in J} \omega_j P_j \tag{22}$$

The result is an price index that accounts for firm size and industry size. I create the same index for each different simulation and compare it to the distribution of prices generated by the original parameters.[35]The results are in Table 6.

The main thought experiment is to reduce both β_X (the partner-switching cost) and β_C (the city-switching cost) by half for all products, and determine the size of the efficiency gain when more importers can separate and/or find better matches. The results are found

[33]I make one deviation from the BLS methodology, as I compute the index for the counterfactual and then compare, while the BLS measure compares individual prices first before aggregating to a comparative index. This is because I am comparing model simulations to other simulations. Results are qualitatively similar, but more subject to simulation outliers, if I compare each price first and make one index, rather than making two indices and then comparing.

[34]Above, I used log prices to estimate the model. Since log price is potentially negative in certain industries, I exponentiate the price in each run of the generated data.

[35]In Appendix D, I assess Model Fit by the same procedure, but comparing generated data from the originally estimated parameters to the true data.

in Table 6 Panel A. I find that the U.S.-China Import Price Index decreases by 14.7% in response to such a change. Since the fixed switching cost is measured in units of the Type I Extreme Value ϵ shock, and results from Section 4 implied that a two standard-deviation shock would be necessary to countenance switching to an identical partner, this can be thought of as a reduction in the size of the shock necessary to switch partners by about one standard deviation. Another way to think about this reduction is at the original parameter values (and in the data) approximately 51% of firms stayed with their partner. The reduction results in only about 13% of all total importers now staying with their partners. It is also apparent that the reductions in the exporter cost have a greater effect on exporter switching than reductions in the city cost on switching, as can be seen from the smaller drop in city-staying that accompanies the experiment.[36] I show the price index for each of 1000 Monte Carlo simulations under both the original and the adjusted parameter set that leads to these results, and present the kernel densities in Figure 4 Panel A. The corresponding results for HS6 industries are qualitatively identical, as can be seen from Table 6 Panel B (a reduction in prices of 12.5%) and Figure 4 Panel B.

One can see the same pattern in individual industries as well. Figure 5 shows the distribution of industry price indices (Equation (21)) in separate HS codes with higher and lower switching costs. In many cases, the distribution is more skewed to the left, meaning prices are typically lower after allowing for more switching. However, there are also more cases of higher prices, such as in industry HS 610432, as a reduction in switching costs can also lead to worse matches. Importers in this industry tend to be insensitive to price in their final exporter decision, and thus an increase in switching often leads to higher prices than in the case with higher switching costs.

How to interpret such a decrease in switching costs? The Chinese government is well known for its investment in capital projects, especially its national development strategy focusing on infrastructure and development of inland provinces. One plausible scenario is that distribution networks to inland cities will improve greatly as China's economy further develops, exactly the type of advance that would lower the cost of adjusting import supply chains. A second is to think of these costs as information frictions, where importers are simply not aware of the alternative exporting options available to them. In this case, a reduction in switching costs would be interpreted as the establishment of a registry where all exporters of a particular HS product would list their prices jointly, thus eliminating information frictions. A "gold standard" system where national governments ensure that producers are known and marketed together is another way to reduce switching costs. A third example would be better contracting institutions in China, allowing importers to adapt

[36]Further experiments (available on request) that reduce one cost versus the other confirm this observation.

short-term contracts while still remaining confident in the ability to find quality inputs at acceptable prices over the long-term.

The results of these counterfactual experiments point more broadly to the importance of importer-exporter dynamics in considering the gains from trade over time. If, as is typically assumed in trade models, importers equally pay the lowest price available in a market, this presents the best scenario for welfare. Any buyer's divergence from the lowest price will necessarily lower estimates of the gains from trade. On the other hand, there are clear policy implications for importer efficiency from improving the general knowledge of the exporter base and helping importers have a better understanding of all possible options. Reducing information and contract frictions in practice can have a major impact on prices of goods, as outsize efficiency gains through lower prices are a robust prediction of the model I estimate.

5.2 Potential for Re-Shoring

Many companies such as Apple have recently announced policies to move production of intermediate inputs back to the U.S. I can use the model to estimate how low prices would have to go in order for importers from China to switch to another potential supplier to which different frictions apply. Specifically, I increase the size of the exporter choice set X by a single firm, and assign it a price. I eliminate the geographic switching cost that must be paid to switching this new firm (though it remains costly to switch from one's previous firm). I also assume this firm has the median estimate for quality in that category. I then re-solve the fixed point equation in (19) for each scenario, and see how many importers would choose to switch to this new firm over 1000 simulations. By using each firm's total share of imports in that industry, I can then determine what fraction of trade accrues to this new firm, i.e. how much trade would be "re-shored" given the existence of a firm with those prices and favorable switching costs. The experiment is one way of thinking about the existence of a highly favorable supplier located in the U.S., about which U.S. firms would tend to have much better information. The results are found in Table 7.

The results demonstrate the clear inertia involved in rousting importers from their Chinese partners, even with the elimination of geographic switching costs. If the hypothetical firm in each industry offered a price in the 75th percentile of the price distribution, only about 2% of trade value would come to this new U.S. firm. However, this price is already far lower than the average U.S. exporter price for firms in the same HS6 product. Furthermore, while it is possible to retrieve 3-4% of Chinese imports back to the U.S. by a hypothetical firm offering the mean or median price in each industry, this price is quite far away from the prevailing prices charged by U.S. exporters: a decrease of approximately 57% compared to

U.S. exporters producing the same product. Thus efforts to return imports from the U.S. are significantly more difficult than simply offering a competitive price- the considerable benefits involved in maintaining existing relationships means that only a small share of imports would be able to move back to this hypothetical supplier.

6 Conclusion

In this paper, I have documented empirically and analytically that frictions from switching suppliers are large, and have important effects for import prices. U.S.-China importer-exporter relationships are characterized by a lack of turnover: 45% of importers remain with their supplier from one year to the next, and one-third of switchers switch within the same city. I estimate a model of dynamic discrete exporter choice, which uses partner switching costs and geographic switching costs in the context of U.S. decisions to import from Chinese exporters. I derive an exporter-specific profit function for importers from a heterogeneous firm model of international trade, and compute the parameters of interest structurally. Switching costs are large, and heterogeneous across products and industries. I then present a number of counterfactuals, including the effects on import prices from improved distribution channels and better information. Specifically, reducing switching costs such that U.S. importers can have better matches leads to 14.7% lower prices. Such a finding can be used to assess the effects of more complete partner information and lower distribution costs in an exporting country on welfare and aggregate productivity for the importing country. Indeed, this paper has shown that better partner options are often available for U.S. importers, but they are not always used. Increasing efficiency of matches will lead to higher gains from trade than are generally considered in models where price decisions occur at the country level, and presents a clear avenue for improving the productivity of U.S. firms through importing. The regional dimension of exporter choice decisions is also much stronger than has generally been known.

This project is merely the first step in a robust area for growth in the study of international trade transactions. Further research can augment this study that uses U.S.-China data and understand when importers change their country of importing, and where they go when they change. Switching costs across countries could play a role in explaining the slow response of exports to exchange rate shocks importers may be unable to quickly switch to more favorable import sources. In addition, future work can assess the impact of specific regional policies on importer behavior, such as the formation of special export zones in cities such as Shenzhen. Finally, the increasing availability of firm-level datasets puts the possibility of firm-to-firm linkages through trade transactions between the production data of separate countries closer to being realized, providing the most complete analysis of the micro-underpinnings of international trade.

References

Ahn, JaeBin, Amit K Khandelwal, and Shang-Jin Wei, "The role of intermediaries in facilitating trade," *Journal of International Economics*, 2011, *84* (1), 73–85.

Alessandria, George, "Consumer Search, Price Dispersion, And International Relative Price Fluctuations," *International Economic Review*, 2009, *50* (3), 803–829.

Alibaba.com, "Alibaba.com Announces Major Initiatives for Buyers and Suppliers," `http://www.alibaba.com/aboutalibaba/releases_081103.html` November 3, 2008.

Allen, Treb, "Information Frictions in Trade," *Econometrica*, 2014, *82* (6), 2041–2083.

Antràs, Pol and Arnaud Costinot, "Intermediated trade," *The Quarterly Journal of Economics*, 2011, *126* (3), 1319–1374.

Armenter, Roc and Miklós Koren, "A Balls-and-Bins Model of Trade," *American Economic Review*, 2014, *104* (7), 2127–2151.

Artuç, Erhan and John McLaren, "Trade Policy and Wage Inequality: A Structural Analysis with Occupational and Sectoral Mobility," Policy Research Working Paper WPS6194, World Bank 2012.

Atalay, Enghin, Ali Hortaçsu, and Chad Syverson, "Vertical Integration and Input Flows," *American Economic Review*, 2014, *104* (4), 1120–1148.

Bernard, Andrew B., Andreas Moxnes, and Karen Helene Ulltveit-Moe, "Two-sided Heterogeneity and Trade," NBER Working Paper 20136 2014.

_ , _ , and Yukiko Saito, "Production Networks, Geography and Firm Performance," NBER Working Paper 21082 2015.

_ , J. Bradford Jensen, and Peter K. Schott, "Importers, Exporters and Multinationals: A Portrait of Firms in the US that Trade Goods," in "Producer dynamics: New evidence from micro data," University of Chicago Press, 2009, pp. 513–552.

Blum, Bernardo S., Sebastian Claro, and Ignatius J. Horstmann, "Occasional and Perennial Exporters," *Journal of International Economics*, 2013, *90* (1), 65–74.

Brandt, Loren, Johannes Van Biesebroeck, Luhang Wang, and Yifan Zhang, "WTO Accession and Performance of Chinese Manufacturing Firms," Discussion Paper DP9681, CEPR 2013.

Broda, Christian and David Weinstein, "Globalization and the Gains from Variety," *The Quarterly Journal of Economics*, 2006, *121* (2), 541–585.

Bureau of Labor Statistics, "Chapter 15," in "BLS Handbook of Methods" 2013.

Byrne, David, Brian K. Kovak, and Ryan Michaels, "Price and Quality Dispersion in an Offshoring Market: Evidence from Semiconductor Production Services," NBER Working Paper 19637 2013.

Drozd, Lukasz A. and Jaromir B. Nosal, "Understanding International Prices: Customers as Capital," *American Economic Review*, 2012, *102* (1), 364–395.

Dubé, Jean-Perre, Jeremy T. Fox, and Che-Lin Su, "Improving the Numerical Performance of Static and Dynamic Aggregate Discrete Choice Random Coefficients Demand Estimation," *Econometrica*, 2012, *80* (5), 2231–2267.

Eaton, Jonathan, Marcela Eslava, David Jinkins, C.J. Krizan, and James Tybout, "A Search and Learning Model of Export Dynamics," 2014.

Egan, Mary Lou and Ashoka Mody, "Buyer-Seller Links in Export Development," *World Development*, 1992, *20* (3), 321–334.

Feenstra, Robert, Zhiyuan Li, and Miaojie Yu, "The Role of Intermediates in Facilitating Trade," *The Review of Economics and Statistics*, 2014, *96* (3), 729–744.

Fortune Magazine, "Why Alibaba's CEO had to go," `http://tech.fortune.cnn.com/2011/02/22/why-alibabas-ceo-had-to-go/` February 22, 2011.

Fox, Jeremy, "Estimating the Employer Switching Costs and Wage Responses of Forward-Looking Engineers," *Journal of Labor Economics*, 2010, *28* (2), 357–412.

Hallak, Juan Carlos and Jagadeesh Sivadasan, "Product and Process Productivity: Implications for Quality Choice and Conditional Exporter Premia," *Journal of International Economics*, 2013, *91* (1), 53–67.

Head, Keith, Ran Jing, and John Ries, "Import Sourcing of Chinese Cities: Order versus Randomness," Mimeo 2014.

Joskow, Paul, "Vertical Integration and Long-Term Contracts: The Case of Coal-Burning Electric Generating Plants," *Journal of Law, Economics and Organization*, 1985, *1* (1), 33–80.

Kamal, Fariha and Asha Sundaram, "Buyer-Seller Relationships in International Trade: Do Your Neighbors Matter?," 2013. Mimeo.

Kamal, Fariha, C.J. Krizan, and Ryan Monarch, "Identifying Foreign Suppliers in U.S. Merchandise Import Transactions," CES Working Paper 15-11, U.S. Census Bureau 2015.

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

Kim, Kyoo and Amil Petrin, "Control Function Corrections for Omitted Attributes in Differentiated Product Models," 2010.

Kleshchelski, Isaac and Nicolas Vincent, "Market Share and Price Rigidity," *Journal of Monetary Economics*, 2009, *56* (3), 344–352.

Kranton, Rachel E. and Deborah F. Mineheart, "A Theory of Buyer-Seller Networks," *American Economic Review*, 2001, *91* (3), 485–508.

Krizan, C.J., "The Nuts and Bolts of International Trade: A Data User's Primer," Technical Report, U.S. Census Center for Economic Studies 2012.

Lincoln, William and Andrew McCallum, "Entry Costs and Increasing Trade," CES Discussion Paper 11-38, U.S. Census Bureau 2011.

Lu, Jiangyong and Zhigang Tao, "Trends and Determinants of China's Industrial Agglomeration," *Journal of Urban Economics*, 2009, *65* (2), 167–180.

McFadden, Daniel, "Econometric Models of Probabilistic Choice," in "Structural Analysis of Discrete Data" 1981.

Monarch, Ryan and Tim Schmidt-Eisenlohr, "Learning and the Value of Relationships in International Trade," Working Paper, SSRN 2015.

Petropoulou, Dimitra, "Information Costs, Networks and Intermediation in International Trade," Working Paper 76, Globalization and Monetary Policy Institute 2011.

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

Rauch, John, "Business and Social Networks in International Trade," *Journal of Economic Literature*, 2001, *39*, 1177–1203.

Rauch, John and Joel Watson, "Network Intermediates in International Trade," *Journal of Economics and Management Strategy*, 2004, *13* (1), 69–93.

Roberts, Mark J., Daniel Yi Xu, Xiaoyan Fan, and Shengxing Zhang, "A Structural Model of Demand, Cost, and Export Market Selection for Chinese Footwear Producers," Working Paper 17725, NBER 2012.

Rust, John, "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 1987, *55* (5), 999–1033.

_ , "Structural Estimation of Markov Decision Processes," in "Handbook of Econometrics, Volme IV," Elsevier Science B.V., 1994, pp. 3081–3143.

Su, Che-Lin and Kenneth L. Judd, "Constrained Optimization Approaches to Estimation of Structural Models," *Econometrica*, 2012, *80* (5), 2213–2230.

Tang, Heiwai and Yifan Zhang, "Quality Differentiation and Trade Intermediation," Technical Report 2012.

Figure 1: Year-to-Year "Staying" Percentages of U.S. Importers from China, 2003-2008

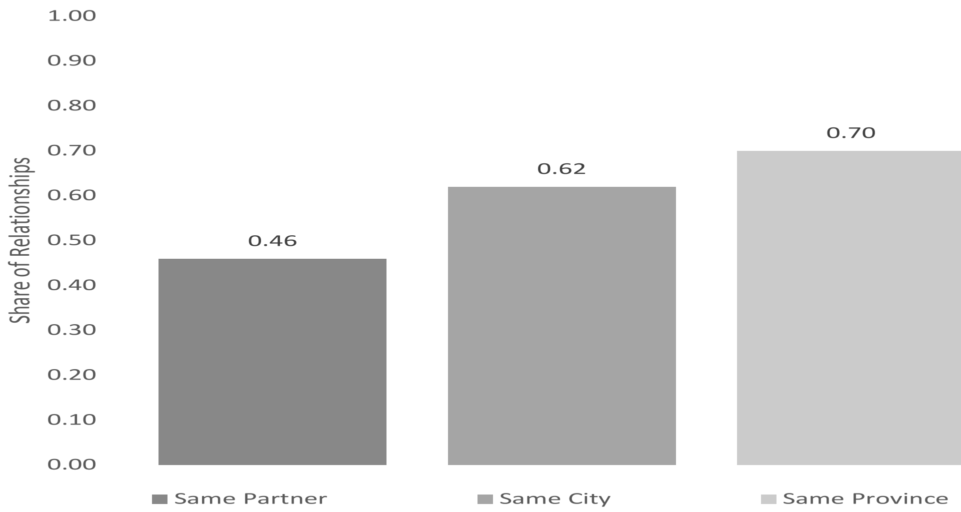

Notes: To determine if a U.S. importer (firm + HS10 product) kept the same exporting partner from one year to the next, I calculate the main partner for each importer in each year, using the value imported from each manufacturing ID in the Longitudinal Foreign Trade Transaction Database (LFTTD). If this main partner remained the same from year-to-year, then the importer "stayed" with its partner. If the city of the main partner remained the same, then the importer stayed in its city, and if the main partner province remained the same, then the importer stayed in the same province. I apply the panel concordance for HS10 products developed by Pierce and Schott (2012).

Figure 2: Year-to-Year "Staying Percentages 2003-2008, Robustness

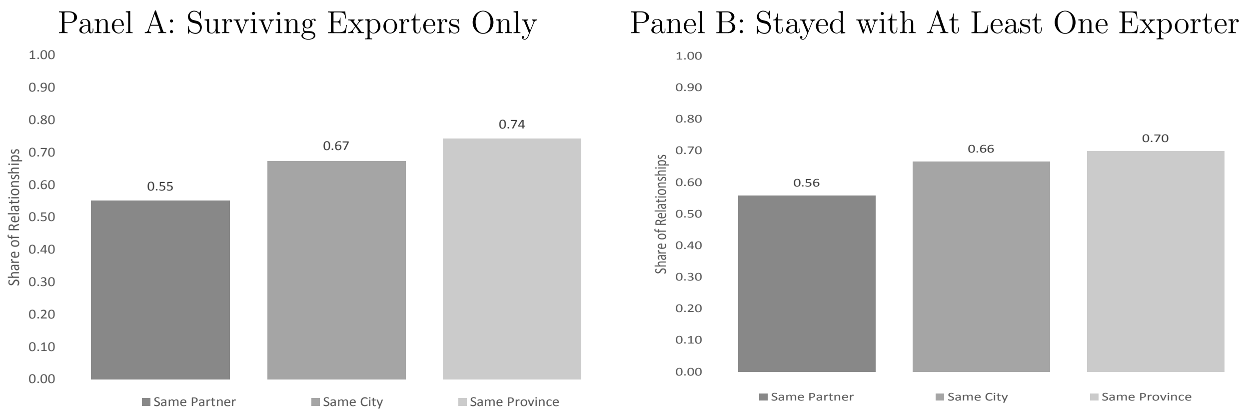

Notes: In Panel A, an importer is considered to have switched (not stayed with) its exporter only if two conditions are met: (a) the main partner changed from one year to the next, and (b) the main partner in the original year is still found exporting to someone else. In Panel B, an importer is considered to have switched from its exporter if it kept any one of its partners in any one of its HS10 imported products.

Figure 3: Role of Price in Supplier Stay/Switch Decision

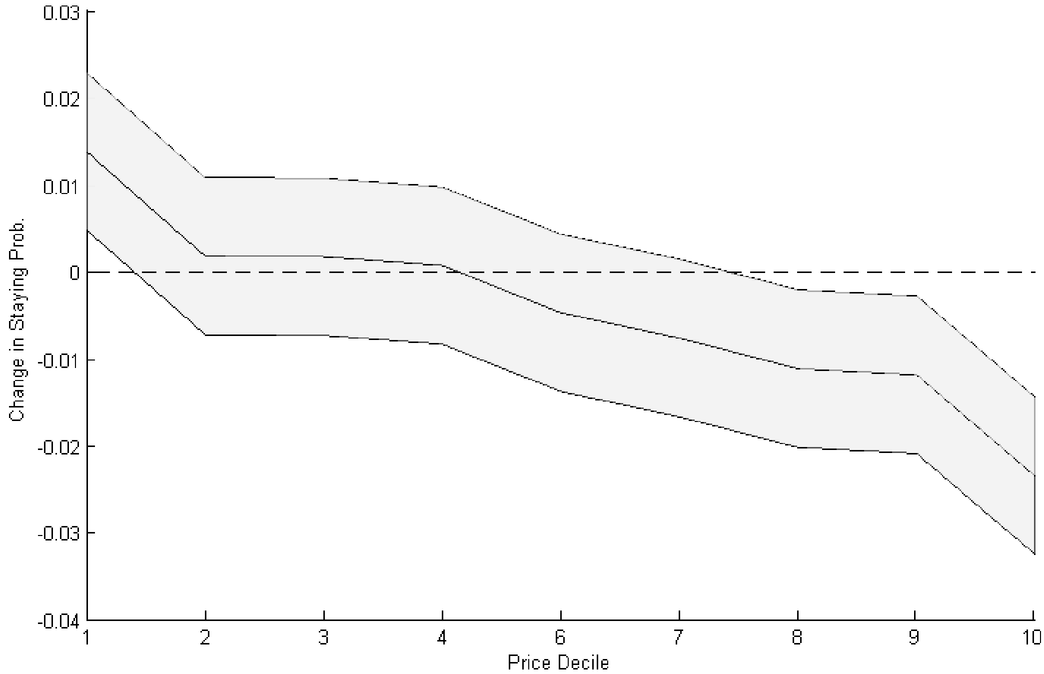

Notes: Log price is the log average unit value across transactions with its main partner in the previous year, standardized across products by subtracting the HS10 mean and dividing by the standard deviation. This variable is split into deciles, and used as an independent variable in a linear probability model of importer staying status. The outer lines are a 99% confidence interval, calculated with robust standard errors clustered at the HS10 level. The sample is the universe of U.S. importers from China who are found two years in a row over the years 2002-2008. Any importer that has the same share of imports from two separate Chinese suppliers is dropped. HS10 and year fixed effects are included. The fifth category is excluded.

Figure 4: Kernel Density Plots, Original $\boldsymbol{\beta}$ vs. Costs Divided by Half

Panel A: HS10 Sample

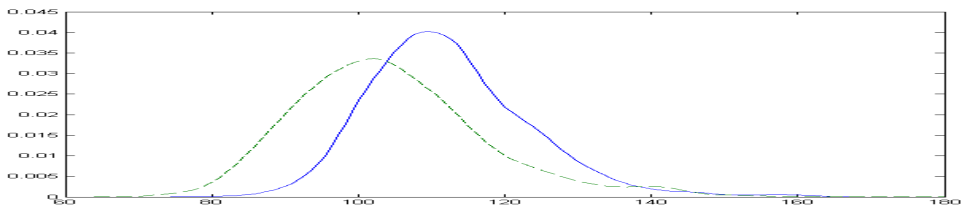

Panel B: HS6 Sample

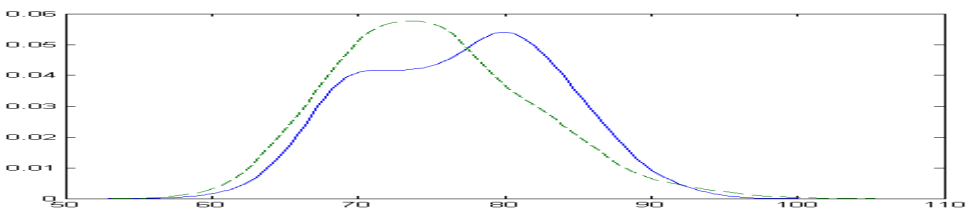

Notes: This figure presents kernel densities for the weighted average Import Price Index (calculated following Bureau of Labor Statistics (2013) and Equations (21)-(22)) for importer-exporter matches predicted under the original parameters (solid line) and reducing switching costs by half (dashed line). I run 1000 replications of the model for each product and parameter set, and calculate the price index for the matches predicted in each replication.

Figure 5: Kernel Density Plots, Original $\boldsymbol{\beta}$ vs. Costs Divided by Half, Selected Industries

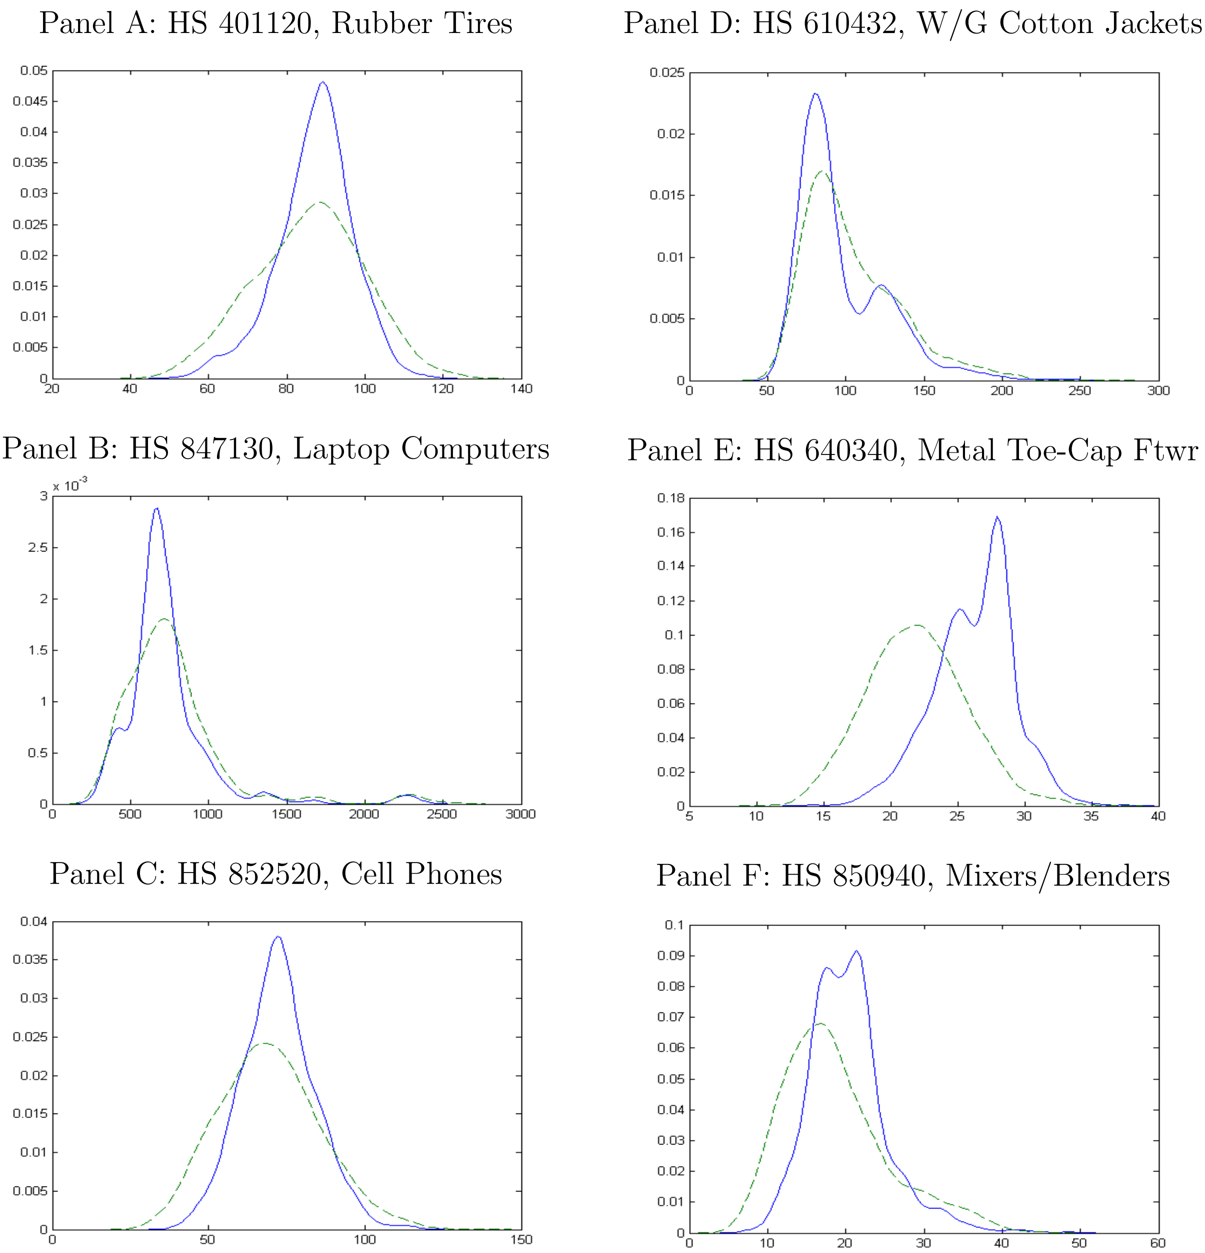

Notes: This figure contains kernel densities for the Industry Import Price Index for importer-exporter matches predicted under the original parameters (solid line) and reducing switching costs by half (dashed line), for individual industries.

Table 1: Determinants of Supplier Stay/Switch Decision

Dependent Variable: Stayed with Chinese Exporter Year-to-Year, 2002-2008								
	(1)	(1')	(2)	(2')	(3)	(3')	(4)	(4')
Log Price	-0.008**		-0.010**		-0.010**		-0.011**	
	(0.001)		(0.001)		(0.001)		(0.001)	
1st Decile		-0.009*		0.017**		0.014**		0.014**
		(0.004)		(0.003)		(0.003)		(0.003)
2nd Decile		-0.005		0.003		0.002		0.002
		(0.003)		(0.003)		(0.003)		(0.003)
3rd Decile		-0.002		0.002		0.002		0.002
		(0.003)		(0.003)		(0.003)		(0.003)
4th Decile		0.001		0.001		0.001		0.001
		(0.003)		(0.003)		(0.003)		(0.003)
6th Decile		-0.004		-0.004		-0.005		-0.005
		(0.003)		(0.003)		(0.003)		(0.003)
7th Decile		-0.006*		-0.007*		-0.007*		-0.008**
		(0.003)		(0.003)		(0.003)		(0.003)
8th Decile		-0.011**		-0.001**		-0.011**		-0.011**
		(0.003)		(0.003)		(0.003)		(0.003)
9th Decile		-0.014**		-0.009**		-0.011**		-0.012**
		(0.003)		(0.003)		(0.003)		(0.003)
10th Decile		-0.036**		-0.019**		-0.023**		-0.023**
		(0.003)		(0.003)		(0.003)		(0.003)
Supplier Size			0.040**	0.040**	0.064**	0.064**	0.064**	0.064**
			(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Supplier Age			-0.002**	-0.002**	-0.003**	-0.003**	-0.003**	-0.003**
			(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Importer Size					-0.032**	-0.032**	-0.031**	-0.031**
					(0.001)	(0.001)	(0.001)	(0.001)
Constant	0.435**	0.444**	0.019**	0.021**	0.101**	0.105**	0.082**	0.086**
	(0.003)	(0.003)	(0.007)	(0.007)	(0.007)	(0.008)	(0.007)	(0.008)
Entry FE	No	No	No	No	No	No	Yes	Yes
R^2	0.07	0.07	0.09	0.09	0.10	0.09	0.10	0.10

Notes: Robust standard errors clustered at the HS10 level in brackets. ** significant at the 1% level, * significant at the 5% level. HS10 and year fixed effects are included. The sample is the universe of U.S. importers (HS10 product code and fim combination) from China who are found two years in a row. The dependent variable is equal to 1 if the U.S. importer had the largest (plurality) share of its total import value from the same Chinese supplier in both years, and 0 if not. Log price is the log average unit value across transactions with its majority partner in the previous year, standardized across products by subtracting the HS10 mean and dividing by the standard deviation. Supplier size (in logs) is the total estimated exports of a Chinese supplier in the HS10 product code in the prior year, based on cross-section summation of total exports to the U.S. Supplier Age is calculated using the first year the Chinese supplier appears in the U.S. customs data, and subtracting it from the prior year. Importer size is the total size of imports in that HS10 product code in the prior year for any U.S. firm. Entry FE is a fixed effect for whether the observation year is the first year a U.S. importer is found importing from China. Any importer that has the same share of imports from two separate Chinese suppliers is dropped.

Table 2: Model Estimates: Summary

		HS10 Results				HS6 Results	
		Weighted Average	Mean			Weighted Average	Mean
β_X	Partner Cost	2.75	2.67	β_X	Partner Cost	2.99	2.77
β_C	City Cost	2.80	3.28	β_C	City Cost	1.61	2.25

Notes: These estimates are the weighted average (by product/industry value) and mean across the 50 HS10 products and 50 HS6 industries used in the analysis. All estimates are different from zero at the 1% level when calculating asymptotic standard errors as in Rust (1994). The results are in units of the Type I Extreme Value Distribution, with mean zero and standard deviation ≈ 1.29, meaning that (using the weighted average for HS10 products) an importer requires a profit shock of approximately $\frac{2.75}{1.29} = 2.1$ standard deviations above the mean to be indifferent between their current supplier and an otherwise identical supplier in the same city.

Table 3: Model Estimates: Relationships (HS10)

		Weighted Average	Weighted Average	Mean	Mean
β_X	Partner Cost	2.75	2.74	2.67	2.70
β_C	City Cost	2.80	2.79	3.28	3.32
ω	Relationship Benefit		0.01		-0.02

Notes: These estimates are the weighted average (by product/industry value) and mean across the same 50 HS10 products as above, with an additional "long-term relationship" indicator in the profit calculation (See Equation (18).

Table 4: Selected Quantitative Estimates, HS Industrial Classification

HS Product/Industry	β_P	β_X	β_C	β_C/β_X
Low City Switching				
Hand Pumps for Liquids	0.08	1.67	3.95	2.36
Files, Rasps, and Similar Tools	−0.06	2.74	2.95	1.08
High City Switching				
Portable Digital ADP Machines (Laptops)	−0.22	3.00	0.43	0.14
Motorcycles, Side-Cars, Engine $\geq$ 50 cc, $<$ 250 cc	−0.13	3.91	0.19	0.05
Many more Importers than Exporters				
Floor Coverings, Wall or Ceiling Coverings, of Polymers of Vinyl Chloride	0.08	3.69	1.38	0.37
Pencils and Crayons	−0.04	3.47	1.21	0.35
High Elasticity of Substitution				
Men's Underpants and Briefs of Manmnade Fibers, Knit	−0.06	3.56	0.97	0.27
Ski/Snowmobile Gloves of Synthetic Fibers, Knit	−0.05	2.82	0.69	0.24
Low Elasticity of Substitution				
Gloves, Impregnable Plastic, 4 chtt, less than 50% cotton, manmade fiber, kt	0.51	1.45	1.64	1.13
Footwear, Sole Rubber/Plastic/Leather, Upper Leather Protective Toe-Cap	0.25	2.76	1.75	0.63

This table contains estimates of the vector of parameters $\boldsymbol{\beta}$ computed through the maximum likelihood problem described in Equations (16) and (19) for various industries and products. Information on industry and product characteristics comes from the LFTTD and Broda and Weinstein (2006).

Table 5: Selected Quantitative Estimates, China Industry Code (CIC) Classification

CIC Industry	β_P	β_X	β_C	β_C/β_X
Low Skilled Workers				
Manufacturing of Teaching Aids and Materials	0.05	1.86	3.47	1.87
High Skilled Workers				
Rolling and Processing of Rare Earths	−0.01	2.63	1.86	0.70
Large Firms				
Arms and Ammunition	−0.03	1.67	2.66	1.58
Small Firms				
Other Medical and Ward Care Equipment	−0.00	2.98	1.30	0.44

This table contains estimates of the vector of parameters $\boldsymbol{\beta}$ computed through the maximum likelihood problem described in Equations (16) and (19) for various China Industrial Categories. Information on industry characteristics comes from data from China's National Bureau of Statistics. CIC industries categories are concorded with HS6 industries following Brandt et al. (2013).

Table 6: Counterfactual Results (I)

HS10 Results		
	Original Sample	$\beta_X \downarrow 50\%$ $\beta_C \downarrow 50\%$
Price Index	-	-14.7%
Staying	*51%*	13%
City Staying	*65%*	31%

HS6 Results		
	Original Sample	$\beta_X \downarrow 50\%$ $\beta_C \downarrow 50\%$
Price Index	-	-12.5%
Staying	*57%*	18%
City Staying	*75%*	43%

Notes: Objects computed by the model simulated with the originally estimated parameters are compared to the same objects after reducing the partner cost and city cost each by half. To compute the Price Index, I take the median received price across 1000 simulations for each importer, then weight each importer by its size within the industry. I then apply industry weights based on total trade among along simulated industries. The staying and city staying percentages are also estimated under the new parameter estimates.

Table 7: Counterfactual Results (II)

Price	Mean Sim. Trade Share	Median Sim. Trade Share	Reduction from average U.S. Exporter Price
Median	3.9%	1.8%	56.3%
Mean	3.8%	1.5%	47.5%
75th Pct.	3.2%	0.8%	34.0%

Notes: This table describes the flow of trade that would go to a hypothetical supplier not subject to geographic switching costs. The first column describes the size of the price charged by this hypothetical supplier compared to alternative Chinese suppliers. The second and third columns describe the percent of imports that would flow to this supplier, while the fourth column compares the price charged to the prevailing price charged by U.S. exporters for the same product.

Appendix

Appendix A Validity of the MID

In this appendix, I describe the foreign exporter identifier in more detail. As shown in Figure A1, two characters on the country of the manufacturer, six characters related to the name of the manufacturer, four characters (in certain circumstances) related to the address of the manufacturer, and three characters related to the city of the manufacturer make up the exporter identification variable. The MID is assembled by the U.S. importing firm (or more likely, by a specialty customs broker utilized by the importing firm) according to an exhaustive list of regulations found in the instructions to the baseline U.S. Customs Document CBP Form 7501, along with the other particulars of the import transaction.[37] I use this identifier to study the behavior of U.S. importers over time, namely what exporter they choose, where the exporting firm is located, and what guides the decision for what partner U.S. importers will choose in the future.

Clearly, the reliability of this variable is important for the stylized facts laid out above. Kamal et al. (2015) go into great detail in their analysis of this variable, but in order to address specific potential concerns related to this work, I still present some brief background on how the U.S. government encourages honest construction of this variable. According to the U.S. Customs and Border Protection, over 99% of entry summary transactions are filed electronically, reducing the risk of misread or misspelled codes. As mentioned above, these forms are also overwhelmingly filed by professional customs brokers well aware of the rules for constructing these codes. Another concern is that the code does not capture the actual producer of a good, but rather some "middle-man", the use of which are very common among firms importing from China (Tang and Zhang (2012)). Importantly, even if a U.S. importer makes use of an intermediary to help them find an exporting firm, information about the actual source of the product is carried through on the final invoice through the entire process.[38] It should also be noted that importers are explicitly warned by the U.S. CBP to make sure that the MID they assemble is reflective of the true producer of the good, not any type of intermediary or processing firm:

"Trading companies, sellers other than manufacturers, etc. cannot be used to create MIDs. Entries and entry summaries in which the first two characters of the MID do not meet the country of origin ISO code, or are created from a company that is known to be a trading house or agent and not a manufacturer, will be rejected for failure to properly construct a MID...Repetitive errors in the construction of MIDs for entries of textile or apparel products will result in the assessment of broker and importer penalties for failure to exercise reasonable care." — U.S. Customs and Border Protection

I augment these facts with a number of checks on this variable by utilizing a rich panel dataset on Chinese firms. This comprehensive dataset from China's National Bureau of Statistics covers all state-owned enterprises (SOEs) and non-SOEs whose annual sales are more than five million *renminbi*, and includes more than 100 financial variables listed in the main accounting sheets of firms.[39] Industries are classified according to the China Industry Code (CIC). Sadly, due to confidentiality and security concerns, the datasets cannot be merged at the firm-to-firm level at this time, despite the availability of plausibly consistent identifiers in both datasets related to name and

[37]See CBP Form 7501 Instructions, p. 30-32 for the exact details.

[38]Krizan (2012), p.10-11, makes clear that this information is available at all stages of the trade transaction.

[39]For more information on this database, see Feenstra et al. (2014).

address. However, this dataset has many other uses in the context of studying importer-exporter behavior.

One application where the NBS industrial database is useful is I can follow the rules laid out for how to construct Manufacturer IDs and assess how commonly multiple firms in an industry possess the same MID- a type of outside check on the uniqueness of the foreign exporter identifier. I do this for five industries in 2005, with uniqueness statistics illustrated in Table A1 Panel A. Although this analysis is subject to some qualification- namely, the NBS data is not the entire universe of Chinese firms, nor is there any guarantee that the name of the firm in Chinese characters (as in the NBS data) is the same as the romanized version of the name of the firm- it appears that the MID does a good job of uniquely identifying foreign firms at the industry level.

An additional complication for studying geographic switching behavior is that only three letters of the city are given in the MID. For example, a city code of "SHE" would be assigned to both Shenyang and Shenzhen, both major cities of more than 8 million people. Again, I use the China Industrial Database in 2005 to check how widespread the problem would be in particular industries. Table A1 Panel B shows that such cases do indeed occur, but not with fatal frequency. It should be noted too that the figures on city-switching will only be misspecified if a U.S. importer switches from a city to another city that happens to start with the same first three letters.

A final concern raised by the construction of the MID is that an importing firm may in fact choose to stay with a supplier, but if the supplier changes its name or address, a new MID means that I will classify that importer as a switching firm. The China Industrial Database tracks firms over time with a unique firm identifier, so I can collapse the data into a panel and see how many firms would fall into this hypothetical scenario by having a change in name or address from 2005 to 2006 that changes their MID. The results of this test are in Table A1 Panel C. Again, though such situations do happen, the vast majority of Chinese exporting firms in the NBS data do not undergo such a change.

Appendix B Proof of Proposition 1

Assumption 1 (Conditional Independence) *The joint transition density of p_t and ϵ_t can be decomposed as:*

$$h\left(p_{t+1}, \epsilon_{t+1} | p_t, x_t, x_{t+1}, \epsilon_t\right) = g\left(\epsilon_{t+1}\right) f\left(p_{t+1} | p_t, x_t, x_{t+1}\right)$$

I also assume that the profit shock ϵ is distributed according to a multivariate extreme value distribution, with known parameters:

Assumption 2 *The distribution of the profit shock is*

$$Pr\left(\epsilon_t < y\right) = G\left(y\right) = \exp\left\{-\exp\{-y-\gamma\}\right\}$$

for $\gamma = 0.577...$ (Euler's constant).

These two assumptions permit the computation of choice probabilities for any particular outcome :

Proposition 1 *Let any present time variable a at one period prior be written as a_{-1}, and one period in the future be written as a'. Given Assumptions 1 and 2, and grouping together the state variables as $s = \{p_{-1}, x_{-1}\}$, the probability of observing a particular exporter choice x^C conditional on state variables s and cost parameters β, $P\left(x^C | s, \beta\right)$, is:*

$$P\left(x^C | s, \beta\right) = \frac{\exp\left[\overline{\pi}\left(x^C, s, \beta\right) + \delta EV\left(x^C, s\right)\right]}{\sum_{\widehat{x} \in X} \exp\left[\overline{\pi}\left(\widehat{x}, s, \beta\right) + \delta EV\left(\widehat{x}, s\right)\right]} \tag{23}$$

where the function $EV(x,s)$ is the solution to the fixed point problem:

$$EV\left(x, s\right) = \int_{s'=0}^{\infty} \log\left\{\sum_{x' \in X} \exp\left[\overline{\pi}\left(x', s', \beta\right) + \delta EV\left(x', s'\right)\right]\right\} f\left(s' | s, x\right) \tag{24}$$

Proof: Let any present time variable a at one period in the past be represented as a_{-1}, and one period in the future be written as a'. Group the state variables together as $s = \{p_{-1}, x_{-1}\}$.

Theorem 1 in Rust (1987) states that, using Assumption 1, for the social surplus function defined as

$$\begin{aligned} & S\left(\left[\overline{\pi}\left(s, \beta\right) + \delta EV\left(s\right)\right]\right) \\ \equiv & \int_{\epsilon} \max_x \left[\overline{\pi}\left(x, s, \beta\right) + \delta EV\left(x, s\right)\right] g\left(\epsilon\right) \end{aligned} \tag{25}$$

the choice probability of any particular exporter choice x occurring can be written

$$P\left(x | s, \beta\right) = S_x\left(\left[\overline{\pi}\left(s, \beta\right) + \delta EV\left(s\right)\right]\right)$$

where G_x is the derivative of S with respect to $\overline{\pi}(x, s, \beta)$. Furthermore, the function $EV(x,s)$ can be written as the contraction mapping:

$$EV\left(x, s\right) = \int_{s'} S\left(\left[\overline{\pi}\left(s', \beta\right) + \delta EV\left(s'\right)\right]\right) f\left(s' | s, x\right)$$

Therefore, we need to compute the social surplus function S given the specific functional form of the density of ϵ.
The location parameter μ for a random variable ϵ with multivariate extreme value distribution is defined such that μ satisfies:

$$Pr\left(\epsilon < y\right) = \exp\{-\exp\{-\left(y-\mu\right)\}\}$$

Additionally, the expectation of ϵ is $\mu+\gamma$, where γ is Euler's Constant. Following a procedure similar to the one in McFadden (1981), Assumption 2 means that the location parameter for the multivariate extreme value distribution of the profit shock ϵ is equal to $-\gamma$. This means that the expectation of ϵ is equal to 0, and we can rewrite the integral in (25) as:

$$\int_{\epsilon} \max_{x} \left[\overline{\pi}\left(x,s,\beta\right) + \delta EV\left(x,s\right) + \epsilon\left(x\right)\right] g\left(\epsilon\right) = \mathbb{E}_{\epsilon}\left\{\max_{x} \mu_{x} + \epsilon\left(x\right)\right\} \tag{26}$$

So the social surplus function will be the expectation of the expression inside the brackets.

For any n independent random variables, $\{\epsilon_1,...\epsilon_n\}$:

$$\begin{aligned} Pr\left(\max\left\{\epsilon_1,...,\epsilon_n\right\} < y\right) &= Pr\left(\epsilon_1 < y,\ldots,\epsilon_n < y\right) \\ &= Pr\left(\epsilon_1 < y\right)\cdots Pr\left(\epsilon_n < y\right). \end{aligned}$$

Thus for any n independent random variables distributed according to the multivariate extreme value distubtion with location parameters $\mu_1,\ldots,\mu_n$, with cumulative distribution function in Assumption 2:

$$\begin{aligned} Pr\left(\max\left\{\epsilon_1,...,\epsilon_n\right\} < y\right) &= Pr\left(\epsilon_1 < y\right)\cdots Pr\left(\epsilon_n < y\right) = \prod_{i=1}^{n} \exp\left\{-\exp\left\{-\left(y-\mu_i\right)\right\}\right\} \\ &= \exp\left\{-\sum_{i=1}^{n} \exp\left\{-y\right\}\exp\left\{\mu_i\right\}\right\} \\ &= \exp\left\{-\left(\exp\left\{-y\right\}\exp\left[\log\sum_{i=1}^{n}\exp\left\{\mu_i\right\}\right]\right)\right\} \\ &= \exp\left\{-\exp\left\{-\left(y-\log\sum_{i=1}^{n}\exp\left\{\mu_i\right\}\right\}\right)\right\} \end{aligned}$$

Thus the maximum of n random variables $\{\epsilon_i\}_{i=1}^{n}$ distributed mulitvariate extreme value with location parameters $\{\mu_i\}_{i=1}^{n}$ is distributed multivariate extreme value with location parameter $\log\sum_{i=1}^{n}\exp\{\mu_i\}$. The expression inside the brackets in Equation (26) is therefore distributed multivariate extreme value with location parameter $-\gamma+\log\sum_{x\in X}\exp\left(\mu_x\right)$. Since the expectation of any random variable distributed multivariate extreme value with location parameter μ is $\mu+\gamma$, the social surplus function from (25) can be written as:

$$\mathbb{E}\left\{\max_{x} \mu_x + \epsilon\left(x\right)\right\} = \log\sum_{x\in X}\exp\left(\mu_x\right) = \log\sum_{x\in X}\exp\left[\overline{\pi}\left(x,s,\beta\right) + \delta EV\left(x,s\right)\right]$$

Following Theorem 1 in Rust (1987), the derivative of the social surplus function is the choice probability:

$$\begin{aligned}
P\left(x^{C}|s,\beta\right) &= S_{x^{C}}\left(\left[\overline{\pi}\left(s,\beta\right)+\delta EV\left(s\right)\right]\right) \\
&= \frac{1}{\sum_{x\in X}\exp\left[\overline{\pi}\left(x,s,\beta\right)+\delta EV\left(x,s\right)\right]}\cdot\exp\left[\overline{\pi}\left(x^{C},s,\beta\right)+\delta EV\left(x^{C},s\right)\right]
\end{aligned}$$

, and the function EV satisfies the fixed point equation:

$$\begin{aligned}
EV\left(x,s\right) &= \int_{s'} S\left(\left[\overline{\pi}\left(s',\beta\right)+\delta EV\left(s'\right)\right]\right) f\left(s'|s,x\right) \\
&= \int_{s'=0}^{\infty}\log\left\{\sum_{x'\in X}\exp\left[\overline{\pi}\left(x',s',\beta\right)+\delta EV\left(x',s'\right)\right]\right\} f\left(s'|s,x\right)
\end{aligned}$$

as desired. ■

Appendix C Selection of Products and Industries

For the model I implement, some products and industries cannot be estimated. For example, products with very few importers or exporters will give wildly high estimates of switching costs. Those with many importers or exporters (≈ 150 or more) also require too huge of a computational burden using current capabilities for the model to be estimated. Thus I pick subsets of HS10 products and HS6 industries, and estimate parameters of the model for each product or industry.

For HS10 products, I pick 50 of the highest-value categories that the U.S. imported from China in 2005. I rank products by their total import value from China, and estimate the model product-by-product. If it is not sensible (fewer than 10 firms) or possible (more than 150 firms) to estimate the model, then I move to the next product. I also do not allow any one product to represent more than 10% of the total value of the sample. The list of products used, and their within-sample trade shares, is in Table C1. Together, these products represent 15% of U.S. imports from China.

For HS10 products, I randomly pick 50 categories among the categories imported from China in 2005 using a random number generator, and estimate the model industry-by-industry. I follow the same rules as above for dismissing potential industries from the sample. The list of industries used, and their within-sample trade shares, is in Table C2. One of the industries is too small to pass disclosure review, and thus its within-sample trade share (along with another's) is not written explicitly in the table.

Appendix D Model Fit

In this appendix, I check how well the estimated parameters do at matching the underlying data used to generate those parameters. Compared to the size of the discrete choice problem, the simple model I estimate is unlikely to match specific importer-exporter outcomes exactly. Thus I check model fit in three areas: how well prices match, how well the percent of switching importers match, and how well the percent of city-switching importers match. I begin by comparing prices.

As can be seen in Table D1, the model with the estimated parameters underpredicts the true price index in the data. In most cases, the pattern is repeated at the industry level- in other words, each industry price index predicted by the model tends to be lower than its real-world counterpart. This is occurring for three reasons: first, the discrete choice model places no distinction on different sizes of the importers- as a precondition of solving the model, the fixed point problem (17) is solved assuming that any two importers with the same state will make the same decision. However, empirical results above show a statistically significant difference in the likelihood of switching based on importer firm size. Thus the model may predict a particular large firm to switch to a lower priced exporter, while in the data, this same firm is in fact less likely to do so. Secondly, the decision of which exporter to use is based on *expected prices* that are predicted with some error, rather than the true actual prices, again giving the potential for prices to be misaligned. Thus the true received price is not an object that I am trying to match through estimating parameters, and is rather an outcome based on a probability distribution. Finally, I discretize the price space into $N = 5$ intervals to estimate the model, applying the midpoint price for each interval, rather than actual price data. This introduces another dimension for the model to fall short.

The results for switching and city switching are more straightforward. For each case, I simply calculate the overall number of firms in an industry predicted to switch for each Monte Carlo run, and take either the median or the mean of that industry percentage for each of 1000 runs. I then translate that into how many total firms are predicted to switch in each industry, and sum together across industries to create an overall measure of switching and city switching behavior. It is clear to see that I match the percentage of firms switching extremely well. I match less well the number of firms switching city, underpredicting the true number by approximately 10%. This is likely because predicting the city puts more pressure on the model of exporter choice to pick the exporter more correctly, while the overall switching percentage does not have to match the chosen exporter in the data as well.

Appendix Tables and Figures

Table A1: Analysis of MIDs as Constructed from China Industrial Production Data

Panel A: Uniqueness of the "MID", 2005

Industry (CIC)	# of Exporters	# of "MID"s	%
CIC 3663	39	38	97.4
CIC 3689	27	26	97.3
CIC 3353	37	37	100
CIC 3331	35	35	100
CIC 4154	74	73	98.6

This panel uses name, address, and city information from China NBS firm data to construct a "MID" for each firm, according to the rules laid out in U.S. CBP Form 7501. In constructing the name of the firm in English, I use the Hanyu Pinyin romanization of Chinese characters, with two to three characters per word of the English name. The second column states the number of firms with positive export values in the given industry in 2005. The third column states the number of unique constructed "MID"s.

Panel B: Uniqueness of the City Code

Industry (CIC)	# of Cities	# of City Codes,2005	%
CIC 3663	22	21	95.5
CIC 3689	15	14	93.3
CIC 3353	28	24	85.7
CIC 3331	15	13	86.7
CIC 4154	19	18	94.7

Panel B uses city information from China NBS firm data to construct city information as found in the MID, where only the first three letters of city are given. The second column states the true number of cities with at least one exporting firm in the data from 2005, while the third column states the number of unique city codes.

Panel C: Changes in the "MID" over Time, 2005-2006

Industry (CIC)	# of Exporters	# of with Identical "MID"	%
CIC 3663	33	33	100
CIC 3689	26	26	100
CIC 3353	31	28	90.3
CIC 3331	20	17	85.0
CIC 4154	63	62	98.4

Panel C uses name, address, and city information from China NBS firm data to track whether constructed "MID"s change over time for the same firm, identified here using the *"faren daima"* firm identifier from the NBS data. The second column states the number of exporting firms found in both 2005 and 2006, while the third column states the number of firms that have identical "MID"s in both 2005 and 2006.

Source: China National Bureau of Statistics.

Figure A1: Sample Invoice

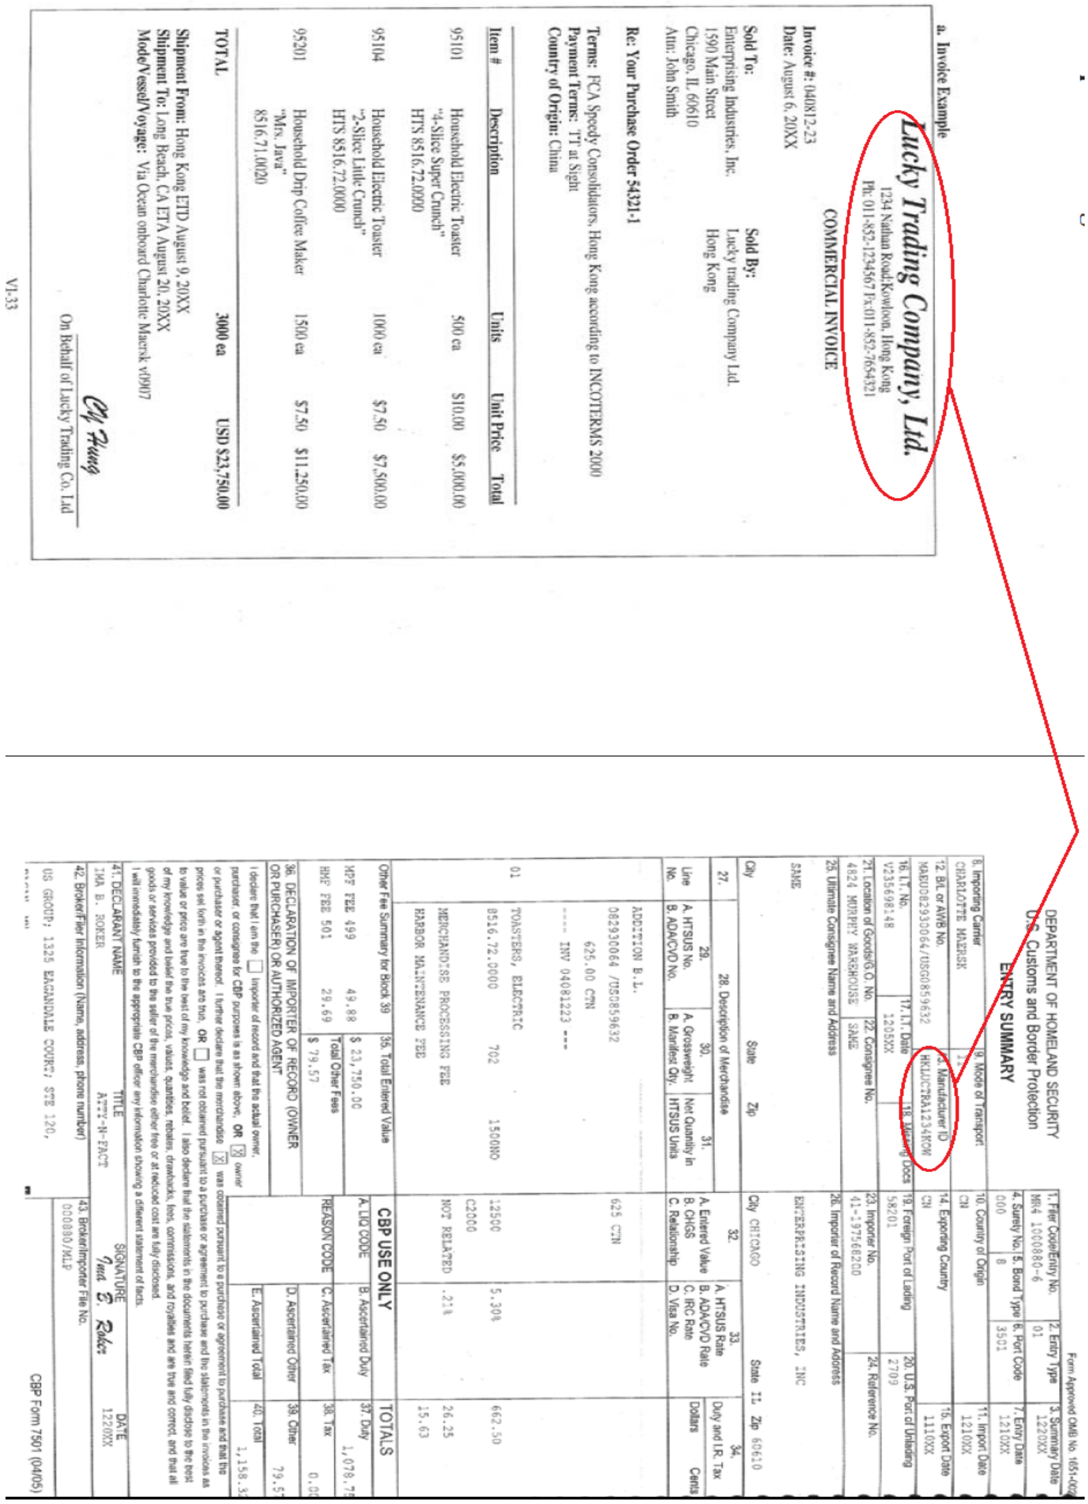

a. Invoice Example

Lucky Trading Company, Ltd.
1234 Nathan Road;Kowloon, Hong Kong
Ph: 011-852-1234567 Fx:011-852-7654321

COMMERCIAL INVOICE

Invoice #: 040812-23
Date: August 6, 20XX

Sold To:
Enterprising Industries, Inc.
1590 Main Street
Chicago, IL 60610
Attn: John Smith

Sold By:
Lucky trading Company Ltd.
Hong Kong

Re: Your Purchase Order 54321-1

Terms: FCA Speedy Consolidators, Hong Kong according to INCOTERMS 2000
Payment Terms: TT at Sight
Country of Origin: China

Item #	Description	Units	Unit Price	Total
95101	Household Electric Toaster "4-Slice Super Crunch" HTS 8516.72.0000	500 ea	$10.00	$5,000.00
95104	Household Electric Toaster "2-Slice Little Crunch" HTS 8516.72.0000	1000 ea	$7.50	$7,500.00
95201	Household Drip Coffee Maker "Mrs. Java" 8516.71.0020	1500 ea	$7.50	$11,250.00
TOTAL		3000 ea		USD $23,750.00

Shipment From: Hong Kong ETD August 9, 20XX
Shipment To: Long Beach, CA ETA August 20, 20XX
Mode/Vessel/Voyage: Via Ocean onboard Charlotte Maersk v0907

Cy Hung
On Behalf of Lucky Trading Co. Ltd

VI-33

Form Approved OMB No. 1651-0022
DEPARTMENT OF HOMELAND SECURITY
U.S. Customs and Border Protection
ENTRY SUMMARY

1. Filer Code/Entry No. MR4 1000880-6
2. Entry Type 01
3. Summary Date 1220XX
4. Surety No. 000
5. Bond Type 8
6. Port Code 3501
7. Entry Date 1210XX
8. Importing Carrier CHARLOTTE MAERSK
9. Mode of Transport
10. Country of Origin CN
11. Import Date 1210XX
12. B/L or AWB No. MAEU082930064/USG0859632
13. Manufacturer ID HKLUCTRA1234KOW
14. Exporting Country CN
15. Export Date 1110XX
16. I.T. No. V235698148
17. I.T. Date 1205XX
18. Missing Docs
19. Foreign Port of Lading 58201
20. U.S. Port of Unlading 2709
21. Location of Goods/G.O. No. 4824 MURPHY WAREHOUSE
22. Consignee No. SAME
23. Importer No. 41-197568200
24. Reference No.
25. Ultimate Consignee Name and Address SAME
City State Zip
26. Importer of Record Name and Address ENTERPRISING INDUSTRIES, INC
City CHICAGO State IL Zip 60610

27. Line No. | 28. Description of Merchandise | 29. A. HTSUS No. B. ADA/CVD No. | 30. A. Grossweight B. Manifest Qty. | 31. Net Quantity in HTSUS Units | 32. A. Entered Value B. CHGS C. Relationship | 33. A. HTSUS Rate B. ADA/CVD Rate C. IRC Rate D. Visa No. | 34. Duty and I.R. Tax Dollars Cents

ADDITION B.L.
082930064 /USG0859632 625 CTN
625.00 CTN
---- INV 04081223 ---
01 TOASTERS, ELECTRIC
8516.72.0000 702 1500NO 12500 5.30% 662.50
C2000
NOT RELATED
MERCHANDISE PROCESSING FEE .21% 26.25
HARBOR MAINTENANCE FEE 15.63

Other Fee Summary for Block 39
MPF FEE 499 49.88
HMF FEE 501 29.69
35. Total Entered Value $ 23,750.00
Total Other Fees $ 79.57
CBP USE ONLY
A. LIQ CODE
B. Ascertained Duty
REASON CODE
C. Ascertained Tax
D. Ascertained Other
E. Ascertained Total
TOTALS
37. Duty 1,078.75
38. Tax 0.00
39. Other 79.57
40. Total 1,158.32

36. DECLARATION OF IMPORTER OF RECORD (OWNER OR PURCHASER) OR AUTHORIZED AGENT
I declare that I am the ☐ importer of record and that the actual owner, purchaser, or consignee for CBP purposes is as shown above, OR ☒ owner or purchaser or agent thereof. I further declare that the merchandise ☒ was obtained pursuant to a purchase or agreement to purchase and that the prices set forth in the invoices are true, OR ☐ was not obtained pursuant to a purchase or agreement to purchase and the statements in the invoices as to value or price are true to the best of my knowledge and belief. I also declare that the statements in the documents herein filed fully disclose to the best of my knowledge and belief the true prices, values, quantities, rebates, drawbacks, fees, commissions, and royalties and are true and correct, and that all goods or services provided to the seller of the merchandise either free or at reduced cost are fully disclosed.
I will immediately furnish to the appropriate CBP officer any information showing a different statement of facts.

41. DECLARANT NAME IMA B. ROKER
TITLE ATTY-N-FACT
SIGNATURE Ima B. Roker
DATE 1220XX
42. Broker/Filer Information (Name, address, phone number)
US GROUP; 1325 EAGANDALE COURT; STE 120,
43. Broker/Importer File No. 000880/MLP
CBP Form 7501 (04/05)

Note: Exporter specific information and location information from the invoicing party is extracted from trade data.

Table C1: List of Products Used in Counterfactuals (HS10)

HS10	Description	Trade Share
8521900000	Video Recording Or Playing Equip,Exc Tape	0.10
8525209070	Cellular Radiotelephones For Pcrs, 1 Kg And Under	0.09
8520900080	Other Magnetic Sound Recording Or Reproducing Equi	0.05
8471500085	Digta Proc Unit W Storage/Input/Output Units,Nesoi	0.04
9405106010	Househld Chandelierelec Ceiling Lgt Base Mt,Nesoi	0.04
7321116000	Nnprtbl Cookng Applncs A Plte Wmrs Nesoi Ios Gas	0.04
8528215501	Video Mnitors,Clr,Flat Pnel Scr,W/ Rec/Rep,<=34.29	0.04
8414513000	Fans For Permanent Installation	0.03
8471604580	Display Units W/Cathode Ray Tube Exc Color, Nesdoi N	0.03
6403996040	Ftwr Sol R/P Up Lthr Exc Pigskin Tenis-Gym Shoe Me	0.03
9401790005	Hshld Outdoor Seats W Mtl Frame, W Text Covrd Cush	0.02
8525404000	Digital Still Image Video Cameras	0.02
8471801000	Control Or Adapter Units For Adp Machines	0.02
6403999031	Ftwr S R/P U-L Exc Pgskn Gt $2.50Pr Ten For Women	0.02
4011201015	Radial Tire Use Bus/Truck, Exc Lt Truck, On Hwy	0.02
8708704545	Road Wheels, Of Aluminum, For Vehicles, Nesoi	0.02
8467210010	Electric Hand Drills, Rotary, Battery Powered	0.02
6403916040	Ftwr:So R/P:Up Lthr:Cv Ank:Tennis Basket Shoes Men	0.02
3925301000	Blinds(Including Venetian Blinds)Of Plastic	0.02
8528127201	Rec Tv,Color,Flat Panel Screen,Display Exc 34.29Cm	0.02
8415103040	Air-Condtnrs,Wind/Wall, Self-Contain Lt 2.93 Kw/Hr	0.02
6402999030	Oth Ftwr Rub/Plas Valued Over $12/Pair For Men	0.02
8472909080	Office Machines, Nesoi	0.01
8519990045	Audio Laser Disc Players (Compact Disc)	0.01
8467220070	Electric Hand Saws, Reciprocating And Jig Types	0.01
6402999060	Oth Ftwr Rub/Plas Valued Over $12/Pair For Women	0.01
8414519090	Fans, Not Perm Inst,Slf-Cont Elec M Tr N/E 125 W	0.01
7202700000	Ferromolybdenum	0.01
3926201010	Gloves,Seamless,Surgical & Medical,Of Plastic	0.01
8509100080	Electric Domestic Vacuum Cleaners, Wgt> 5Kg, Nesoi	0.01
6302319020	Sheets, Cotton, Not Printed/Knitted/Napped/Trim	0.01
6110121060	Women'S Or Girls' Pullovers, Sweatshirts, And Si	0.01
9401308030	Swivel Seats With Variable Height Adjustment, Neso	0.01
8516290030	Electric Portable Space Heaters, Fan-Forced Air	0.01
8712003500	Bicycles Having Both Wheels Excd 65Cm Diam, Nesoi	0.01
8516500090	Microwave Ovens Having Capacity >31.0 Liters	0.01
8539310060	Dschrge Lmps,(Ex Ultrvio.),Flrscent,Sngle Screw-In	0.01
8708395030	Brake Drums And Rotors Of Heading 8701 To 8705	0.01
6111206010	Babies' Sunsuits & Similar Apparel Of Cotton, Knit	0.01
8527136040	Radiobroadcast Rec Comb Inc Optical Disc Play Rec	0.01
8527316040	Radio-Recorder Comb, Inc Optical Disc Player/Recrd	0.01
3926201020	Gloves,Seamless,Excp Surg & Medical,Disposabl,Plas	0.01
8418210010	Refrigerators, Household, Comp Typ, Vol ¡184 Liter	0.01
3913902000	Polysaccharides & Deriv, Cellulose In Granular Etc	0.01
8516710020	Automatic Drip & Pump Type Electric Coffee Makers	0.01
6402991815	Oth Ftwr R/P Upper>90% R/P Tennis/Basketball Shoes	0.01
4412140540	Plywd Birch Face/Hdwd Outr Only Wd Sht Nt Surf Cov	0.01
9401790015	Outdoor Seats W Mtl Frame, W Text Covered Cushions	0.01
6306229030	Tents, Except Screen Houses, Of Synthetic Fibers	0.01
8471704065	Hard Disk Drive Unt, Nesoi, W/Out Extnl Powr Suply	0.01

These shares are the percent of import value compared to the total among these 50 products. Together the total import value of these products is 15% of U.S. import value from China in 2005.

Table C2: List of Industries Used in Counterfactuals (HS6)

HS6	Description	Trade Share
291560	Butyric Acid, Valeric Acid, Their Salts and Esters	0.00
291631	Benzoic Acid, Its Salts and Esters	0.00
293629	Other Vitamins and Their Derivatives (Unmixed)	0.01
340120	Soap in other forms	0.01
392020	Other Plates, Sheets, Film, Foil, Tape, Strip of Propylene Polymers (Non-cellular)	0.01
481810	Toilet paper	0.01
481960	Box files, letter trays, storage boxes and similar articles, used in offices, shops	0.01
490300	Children's picture, drawing or coloring books	0.01
520831	Plain Woven Fabrics, Cotton (Cotton 85% or More; Dyed; Not >100g/m2)	0.01
560312	Nonwovens of man-made filament,>25g/m2	0.01
570210	Kelem, Schumacks, Karamanie and Similar Hand-woven Rugs	0.01
580639	Other Narrow Woven Fabrics of Other Textile Materials	0.01
591190	Other Textile Products and Articles, for Technical Use	0.01
610432	Women's or Girls' Jackets of Cotton, Knitted or Crocheted	0.01
610791	Men's or Boys' Bathrobes, Dressing Gowns, of Cotton, Knitted or Crocheted	0.02
620339	Men's or Boys' Jackets, Blazers, of Other Textile Materials	0.01
621230	Corsets	0.01
621490	Shawls, Scarves, Mufflers, Mantillas, Veils, of Other Textile Materials	0.00
640219	Other Sports Footwear, Outer Soles and Uppers of Rubber or Plastics	0.08
640340	Other Footwear, Incorporating Protective Metal Toe-cap	0.11
650699	Headgear of Other Materials	0.00
650700	Headbands, Linings, Covers, Hat Foundations, Hat Frames, for Headgear	0.00
670411	Complete Wigs of Synthetic Textile Materials	0.02
730722	Threaded elbows, bends and sleeves, of Stainless Steel	0.00
730830	Doors, windows and their frames and thresholds for doors, of Iron or Steel	0.01
731814	Self-tapping screws of Iron or Steel	0.03
731930	Other pins of Iron or Steel	0.00
820310	Files, rasps, and similar tools	0.01
820890	Other (including parts) (Knives and Blades for machines and appliances)	X
830300	Armored/ reinforced safes, strong-boxes, safe deposit lockers, of base metal	0.04
830990	Stoppers, Caps, Lids, Seals, Other Packing Accessories, of Base Metal	0.01
841320	Hand Pumps for Liquids	0.00
841360	Other Positive Rotary Displacement Pumps	0.00
841370	Other Centrifugal Pumps	0.02
841420	Hand or Foot Operated Air Pumps	0.01
841850	Refrigerating, Freezing Chests, Cabinets, Display Counters, Show-cases & Similar	0.00
848110	Pressure-reducing Valves	X
850650	Lithium primary cells and primary batteries	0.01
850910	Vacuum Cleaners, With Self-contained Electric Motor	0.14
850940	Food Grinders and Mixers; Fruit or Vegetable Juice Extractors	0.09
853641	Relays, for a Voltage Not Exceeding 60v	0.02
870893	Clutches and parts thereof	0.02
871110	Motorcycles, Side-cars, Reciprocating Engine, cylinder capacity not >50 cc	0.03
871120	Motorcycles, Side-cars, Reciprocating Engine, cylinder capacity >50 cc not 250 cc	0.07
900580	Monoculars, Other Optical Telescopes; Other Astronomical Instruments	0.02
902910	Revolution counters, production counters, taximeters, odometers, pedometers etc	0.01
920590	Other wind musical instruments	0.01
950631	Golf Clubs, Complete	0.05
960321	Tooth Brushes	0.01
960910	Pencils and Crayons, With Leads Encased in a Rigid Sheath	0.02

These shares are the percent of import value compared to the total among these 50 industries. The combined value share of HS 848110 and 820890 (redacted for disclosure purposes) is 0.37%.

Table D1: Model Fit

	Data	Median over 1000 runs	%
Price Index			
Weighted Average	84.6239	76.4979	90.4
Median	66.1725	61.7019	93.2

	Data	Industry Median	%	Industry Mean	%
Total Switching Partner	714	711	99.6	708.85	99.3
Total Switching City	416	469	112.7	469.76	112.9

Notes: Objects computed by the model simulated with the estimated parameters are compared to the same objects in the data. To compute the Price Index, I first take the median received price across 1000 simulations for each importer. I then either weight each importer by its industry share, and sum up ("Weighted Average") or I simply compute the median across importers in an industry. I then apply industry weights based on total trade among along simulated industries to make an aggregate price index. The switching and city switching figures are the number of importers switching partner or city in the data compared to either the mean number of firm switching/city switching for each industry, or the median number of firms switching/city switching for each industry.