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

Ryan Monarch†
Federal Reserve Board

January 16, 2015

Abstract

This paper uses confidential U.S. Customs data on U.S. importers and their Chinese exporters to investigate the frictions from changing exporting partners. High costs from switching partners can affect the efficiency of buyer-supplier matches by impeding the movement of importers from high to lower cost exporters. I test the significance of this channel using U.S. import data, which identifies firms on both sides (U.S. and foreign) of an international trade relationship, the location of the foreign supplier, and values and quantities for the universe of U.S. import transactions. Using transactions with China from 2003-2008, I find evidence suggesting that barriers to switching exporters are considerable: 45% of arm's-length importers maintain their partner from one year to the next, and one-third of all switching importers remain in the same city as their original partner. In addition, importers paying the highest prices are the most likely to change their exporting partner. Guided by these empirical regularities, I propose and structurally estimate a dynamic discrete choice model of exporter choice, embedded in a heterogeneous firm model of international trade. In the model, importing firms choose a future partner using information for each choice, but are subject to partner and location-specific costs if they decide to switch their current partner. Structural estimates of switching costs are large, and heterogeneous across industries. For the random sample of 50 industries I use, halving switching costs shrinks the fraction of importers remaining with their partner from 57% to 18%, and this improvement in match efficiency leads to a 12.5% decrease in the U.S.-China Import Price Index. Similar gains from match efficiency result from estimation results for 50 of the largest product categories in U.S.-China trade.

JEL Codes: F23, F14, L14, D21;
Keywords: International Trade, International Business, Import Price, Transactional Relationships, Firm Behavior

*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, 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, Evan Starr, and Olga Timoshenko, and participants at 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. Andrew McCallum and Per Ruquist helped greatly with the use of Tomlab. The views expressed herein are mine, and should not be interpreted as reflecting the views of 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. Any opinions and conclusions expressed herein are mine and do not necessarily represent the views of 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.

†Contact: ryan.p.monarch@frb.gov

1 Introduction

An influential literature has studied the effects of misallocated resources on aggregate productivity: major improvements in productivity could theoretically be achieved by raising the marginal products of capital and labor of producers in developing countries to U.S. levels (Hsieh and Klenow 2009). However, a firm matching with a suboptimal supplier is another source of potential inefficiency, in particular if the supplier being used has high prices or poor quality. An anecdote illustrates that buyer-supplier mismatches are a significant bottleneck in international trade. In 2008, David Wei, CEO of e-commerce giant Alibaba.com announced a "Gold Supplier" identification program in order to clearly flag reliable exporters in China, with the explicit goal of making foreign sourcing decisions easier.[1] However, three years later, wide-scale fraud under this program came to light, resulting in Wei's resignation.[2] That a behemoth like Alibaba.com was both aware of the difficulties involved for buyers matching with suppliers and unable to resolve them in a satisfactory way demonstrates that finding the right supplier is not easy. If an importer is unaware of identical but lower-priced alternatives to its current supplier, or unwilling to bear the costs and uncertainty involved in such a change, then this creates the potential for inefficiency. Testing the importance of this channel is challenging, however, as data on relationships between final producers and their suppliers is sparse.

This paper utilizes confidential U.S. Customs and Border Protection data on U.S. importers and their Chinese exporting partners to explore the costs involved in changing partners and their impact on import prices. The database includes information for firms on both sides (U.S. and foreign) of an international trade transaction, allowing the study of these "switching costs" for U.S. importers. Empirical results indicate that such costs are likely to be substantial: from 2003-2008, 45% of arm's-length importers maintain their partner from one year to the next, even with an average of 35 available exporters selling the same product.[3] Nearly half of the total value of U.S. imports from China is concentrated among those importers who used the same exporter year-to-year. Furthermore, there is remarkable geographic inertia among importers who do change their partner: one-third of all switching importers remain in the same city as their original partner, even with an average of nine available cities to purchase their product. Finally, importer switching decisions are correlated with prices: those importers who paid the highest prices are much more likely to change their partner. Thus there is reason to believe in the potential for efficiency gains by reducing frictions to importer

[1] "[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 while the Gold Supplier Starter Pack will appeal to a wide range of potential new customers." – David Wei, CEO, Alibaba.com, 11/3/2008

[2] "Alibaba ... reported that 2,326 high volume sellers who pay a fee to the company to peddle their wares on the site - '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, and his deputy, COO Elvis Lee, both resigned yesterday." – Fortune Magazine, 2/22/2011

[3] I use the term importer to refer to a *firm-HS10 product pair*. Thus a firm that imports two HS10 products is considered to be two separate importers. HS10 is a ten-digit product code, the most disaggregated product code used by U.S. firms.

switching.

Guided by the above empirical regularities, I develop a dynamic discrete choice model of exporter choice. The importing firm decides which exporter to use by comparing partner-specific profits across all possible choices, including its current match. Considerations of which exporter to use depend not only on the price and quality of each exporting firm, but also per-unit switching costs, both at the partner and the city level. The key tradeoff that the firm faces is that switching to either a cheaper or a higher quality exporting partner raises profits, but changing one's current exporter or geographic location is costly. The model produces closed-form expressions of choice probabilities for each potential outcome, which allows computation of the switching cost parameters via maximum likelihood. I compute exporter quality using a procedure similar to the "control function" estimator of unobserved heterogeneity from Kim and Petrin (2010) and the quality ladder estimation of Khandewal (2010). I then use the Mathematical Programming with Equilibrium Constraints (MPEC) techniques developed by Dubé, Fox and Su (2012) and Su and Judd (2012) to solve the model. I estimate the parameters of the model industry-by-industry (at the HS6 level) and product-by-product (at the HS10 level).

The main quantitative results can be summarized as follows. First, model estimates of switching costs are large, heterogeneous across industries and products, and match the underlying data well. Based on model estimates, in order to be indifferent between its current partner and another partner in the same city charging the same price, the average importer would require a positive shock to profits of two standard deviations higher than average from the new partner. Industries with low switching costs have high amounts of switching, and products that are less substitutable tend to have higher switching costs. Second, the impact of switching frictions on importer prices is sizable. Using a randomly selected sample of 50 industries, I perform the following counterfactual: match importers with different exporters by altering switching cost parameter values, then construct the U.S.-China Import Price Index to determine aggregate price changes arising from these new matches. I find that reducing frictions by half reduces the Import Price Index by 12.5%, a reduction achieved by shrinking the percent of staying importers from 57% to 18%. Thus there is a substantial efficiency gain involved in lowering the cost of switching export partners. These results are qualitatively the same as estimating the same model on 50 of the largest HS10 product categories that the U.S. imports from China. On the other hand, tripling the barriers to switching results in a 7.62% increase in the Import Price Index, and 90% of importers remain with their partner year-to-year. Third, changing one friction without changing another has differential effects on prices. Eliminating geographic frictions while maintaining partner switching costs reduces the Import Price Index by 7.37%. However, keeping only the geographic switching cost reduces prices by 15.20%. Finally, I estimate the trade flow to a newly available supplier which is not subject to any geographic switching friction from the same random sample of

industries. If that supplier charges the median price among Chinese exporters within its product category and produces a high-quality variety, this collection of prospective suppliers (one per industry) would be able to attract approximately 4% of all imports from China. A supplier that can be switched to without geographic frictions is one way to consider a potential new U.S. supplier, but the median price among Chinese exporters is approximately 57% lower than the price charged by U.S. exporters for the same product mix. This demonstrates that there are substantial barriers to "re-shoring" Chinese imports back to U.S. suppliers.

The results described above demonstrate a significant effect of switching costs on exporter choice and import prices. There are a number of interpretations for what these switching costs represent. Allen (2012) shows the importance of information frictions, especially geographically, for Filipino farmers searching for buyers of their product. Thus one way to view the high cost of switching partners and additional cost of geographic switching is that importers are simply not aware of other low-price options that are available. A policy that reduces switching frictions is one that would reduce information asymmetry, such as a"gold standard" directory of all available exporters and prices put forth by both the U.S. and China, or the creation of an exporter marketplace for importers to utilize and select partners. A second interpretation is the existence of long-term trading contracts such as in Kleshchelski and Vincent (2009), that reduce uncertainty in product prices, quality, or lead time, but prevent more efficient matching. In this dimension, an overall improvement in contracting institutions leading to more widespread use of short-term contracts would be consistent with reductions in switching costs discussed above. A third explanation for high switching costs is related to the overall logistical difficulty in adjusting one's suppliers in response to short-term changes in purchasing prices 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. Indeed, intermediaries play an important in the Chinese export market, as described in Ahn, Khandewal, and Wei (2011) and Tang and Zhang (2012). In summary, the above quantitiative exercises have direct interpretations, and the result of policies that reduce matching frictions will be a significant improvement in productive efficiency at the firm level.

Although the field of international trade has focused on numerous aspects of firm-level participation in international activity, including especially the decision to export, import, engage in FDI, or use intermediaries, the study of individual exporter-importer relationships remains relatively sparse. 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 database used in this work, which provides information about U.S. importers and their foreign exporting partners. Eaton et. al. (2012) 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 Krizan (2012) use U.S. Census trade transaction data to document trends in the formation of importer-exporter relationships. 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. They find "importer-specific" spillovers are an important part of the general information spillovers that characterize exporting. Each of these puts the onus on the exporter to undertake searching behavior buyer by buyer, while 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 on trade: Bernard, Moxnes, and Ulltveit-Moe (2013) develop a model of relationhip-specific fixed costs to exporting using Norwegian buyer-supplier trade data, while Blum, Claro, and Horstmann (2010) 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 (2012) 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 Mathematical Programming with Equilibrium Constraints (MPEC) methodology for solving discrete choice problems found in Su and Judd (2012) and Dubé, Fox, and Su (2012). As in those studies, my model includes costs entering into a firm's profit function, where in my case, the costs are 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. The model also shares some similarities with Lincoln and McCallum (2012), who produce

estimates of exporting fixed costs that broadly measure the frictions involved in entering the export market that are comparable across industries. 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 (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, Hortacsu, and Syverson (2012) measure the extent to which firms rely on subsidiaries versus outside firms for intermediate input purchase. 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

The database I work with is 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. 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.[4] Through a variety of "external vailidity" checks outlined in Appendix A, I find substantial support for the use of the MID as a reliable, unique identifier, both over time and in cross-section. 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.[5]

[4]Specifically, the MID contains the first three letters of the producer's city, six characters taken from the producer's name, up to four numeric characters taken from its address, and the ISO2 code for the country of origin.

[5]Kamal, Krizan, and Monarch (2014) undertake a head-to-toe examination of the MID and its applicability towards studying buyer-seller relationships, including a description of the institutional background for the variable, various cleaning methods, and

At this stage, I perform an initial cleaning of the LFTTD, using methods outlined in Bernard, Jensen, and Schott (2009) and Pierce and Schott (2009). As in Bernard, Jensen and Schott (2009), I drop all transactions with imputed quantities or values (which are typically very low-value transactions) or converted quantities or values. I also eliminate all related-party transactions, as exporters who are importing from separate branches of the same firm will likely have very different relationship dynamics than arm's-length exporters. I concord HS codes over time according to the methodology in Pierce and Scott (2009). In addition, I clean up unreasonable values for the MID specifically related to U.S.-China trade. I restrict the sample to importers with a firm country identifier of China (meaning the producing firm is located in China). Due to the entrepôt nature of Hong Kong's international trade flows, I concentrate solely on Mainland China - deleting any observation that has any appearance of coming from Hong Kong, Macau, or Taiwan. For example, a city code of "HON", even with a country code for mainland China, is likely referring to Hong Kong and thus dropped.[6] Finally, I drop any firm that has a three-letter city code that is not in the top 300 cities of China by population.

2.2 Stylized Facts

The starting point of my analysis is to use the exporting partner MID to illustrate the extent of partner-switching. My unit of observation is a U.S. importer (firm-HS10 product combination). Some importers have more than one exporter, so I define an importer's *main exporter* in any time period to be the one from which the largest percentage of imports were delivered[7]. I define a firm as "staying" with one's partner if its main exporter remained the same over time[8]. I track the universe of U.S. importers from China in 2003-2008, and determine whether they (a) also imported one year previously, and if so, (b) whether they continued to import from the same exporting partner or geographic location as in the previous year. Figure 1 plots the fraction of importers staying with their partner, 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

comparisons of the number and composition of exporters from U.S. data with international data. Their findings give credence to use of the MID as a measure of foreign exporter identification in cross-section and in panel. I include a summary of the variable, including its construction and the relevant laws surrounding information provided in trade transactions in Appendix A.

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

[7]This simplification introduces the potential for "false switching", where an importer uses the same exporter in two periods, but changes the source of the plurality of its imports. Analysis of the LFTTD indicates that U.S. importers typically import a very large share of their total imports from only one partner. The average share of imports that come from a U.S. importer's main Chinese partner is 83.9%, with a standard deviation of 22%. Furthermore, multiple importers do not dominate in the data. Kamal and Krizan (2012) present some basic statistics on the number of exporting relationships that a U.S. importer may be in: across all U.S. importers, the average number of exporting partners for a U.S. importer is 1.8, and the average number of exporting partners for a "polygamous" U.S. importer is 4.

[8]All results are robust to different definitions of "staying", including staying with any one of the set of partners, or staying with the entire set of one's partners.

the same partner over time. Even though the number of potential exporting choices is increasing over this time period, the share of importers using the same supplier is 45.9% over this time. 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 switching partners would be 29/30, or 97%. Thus path dependence is far higher than would be expected if importers were choosing their partner randomly. These staying importers account for 44.9% of the total value of arm's length imports from China over this time period, which shows these are not simply very small importers who are switching. Secondly, among those firms who do choose to switch, approximately one-third of all importers remain in the same city as their original partner. Using the same benchmark as above, random exporter selection would imply an 86% chance of switching city.[9] Thus there is strong inertia keeping firms in their original city, even if they choose not to use the same exporting partner as before. There are many potential explanations for such a finding, including local network formation, efficient distribution channels centered on a particular geographic location, or agglomeration on the export side. In sum, the year-to-year figures show that supplier choices are highly correlated with previous supplier usage, and decisions of whom to switch to are highly 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.

The stylized facts about importer-exporter relationships described above are robust to a number of alternative checks and specifications. Firstly, Figure 2 shows that the tendency to stay with one's exporting partner is not concentrated among only small or large importers: even as the share of total imports from China increases dramatically from 2003-2008, the share of imports from importers keeping their same partner remained at about 40-45% of total imports, very similar to the overall share of importers remaining with their partner. Secondly, 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, including entry into the WTO, it is likely that many exporters are entering or exiting. To eliminate this channel, I recreate the results using only those matches where the exporting MID is found in both years. These results are found in Figure 3 Panel A. Mechanically, the percentage of firms staying with their partner must increase, but the two main stylized facts described above carry over: a significant share of (but not all) importers stay with their partners,

[9] There is an average of 9 cities for each HS10 product, but the number of exporters are not distributed equally across cities. I thus compute the probability of switching city for any one importer in the data if they were choosing their partner with the following formula:

$$Pr\left(CSwitch\right) = \frac{\sum_i \sum_c M_{ic}\left(1 - \frac{X_{ic}}{\sum_c X_{ic}}\right)}{\sum_i \sum_c M_{ic}}$$

where product i and city c have M_{ic} importers and X_{ic} exporters. The term in parentheses is the probability of an importer in city c switching city, which I then weight by the number of importers in that city. The denominator is the total number of U.S. importers from China.

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 my definition of an importer as a *firm-HS10 product* combination. Indeed, it is possible for one firm to appear multiple times in Figure 1. In particular, if one firm imports multiple products, then the counts may exaggerate or understate the effects of firm switching behavior. I thus perform the same decomposition considering a *firm* as the unit of analysis, rather than a *firm-HS10 product* combination. Even with this much more sparse assignment of switching firms, Figure 3 Panel B shows that the results remain very similar: now, close to two-thirds of continuing importers stay with their partners, with considerable geographic stickiness. These stylized facts are robust to a number of other specifications, including using only importing firms classified as manufacturing firms, an importer defined as a firm-HS6 product combination, individual year trends, different definitions of switching, and the share of importers not adjusting their supply network over a longer time frame than one year. These are described in greater detail in Appendix B.

2.3 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 model 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.

$$Stay^{j}_{i,t,t+1} = \alpha_0 + \sum_{k=1}^{10} \alpha_k \mathbb{1}\left[StandardPriceCat_{i,t} = k\right] + \gamma_1 ExpChar^{j}_{i,t} + \gamma_2 ImpChar^{j}_{i,t} + f_j + f_t + v_{i,t} \quad (1)$$

As above, I define importer i importing product j as staying ($Stay^{j}_{i,t,t+1} = 1$) with its export partner 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 "unit value" from the LFTTD, minus its HS10 product mean and divided by the standard deviation. Thus prices are comparable across industries. 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 constuct 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 Model (1) are reported in Table 1. Standard errors are clustered at the HS10 level.

I begin with the effect of price on the partner staying decision. 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 original (plurality) partner firm. 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 partner than the omitted group (5th decile). Those importers paying the lowest prices were also more likely to stay with their partner when accouting for importer and exporter characteristics, as can be seen in Columns (2)-(4). Figure 4 contains 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 partner than those in the middle and low price regions, while those paying the lowest prices are more likely to stay with their partners. Table 2 repeats the analysis simply using price entering the regression linearly, demonstrating again that higher prices make importers more likely to switch.[10]

The effects of the various importer and exporter characteristics are themselves of interest. Table 1 makes clear 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 partner. Thus there is substantial room for exporter and importer heterogeneity in explaining the staying decision for U.S. importers with their Chinese exporters.

In conclusion, 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

[10]The results are unchanged by estimating (1) using only one year of data, as I do in Appendix B.

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 the 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 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 to use to import its necessary inputs. At time t, final good producer m has a choice of which exporter $x_{j,t}$ to obtain its quantity of input j. 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 use their previously paid price to form expectations about the price from their original partner, and average price from each other exporter to form expectations about the price from that exporter.[11] Since the expectation differs depending on what partner was used, this expectation is both *importer-specific* (m) and exporter-specific ($x_{j,t}$), which allows the same exporter to charge different prices to different importers. I describe the calibration of this density in the next subsection.

[11] 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.

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 $\zeta^j_{c,j}$ 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 manner:

$$\bar{p}^m_{x,j,t} = \mathbb{E}\left[p^m_{x,j,t}\right] \exp\left\{\zeta_{x,j}\ \mathbb{1}\{x^m_{j,t} \neq x^m_{j,t-1}\} + \zeta_{c,j}\ \mathbb{1}\{c^m_{j,t} \neq c^m_{j,t-1}\}\right\} \tag{2}$$

where $\bar{p}^j_{x,t}$ 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 partner 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 $X^m_t = \left\{x^m_{j,t}\right\}_{j=1}^J$ be the vector of supplier choices made by importer m for each input $j = 1, ...J$ 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 and $\lambda_{x,j,t}$, the "quality" of the variety from exporter x.[12]

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

[12]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 4.1.1 below.

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

$$MC\left(X_t^m\right) = \frac{1}{\phi_m\left(X_t^m\right)} c_m\left(X_t^m\right) \tag{3}$$

Maximizing expected profits at time t means that importer m must set the price of their final good optimally and make its 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 (3) into the above profits equation gives the following equation:

$$\pi_t^m = \max_{X_t^m} \frac{1}{\sigma} B \left(\frac{\sigma}{\sigma-1}\right)^{1-\sigma} \left[\phi_m\left(X_t^m\right)\right]^{\sigma-1} c_m\left(X_t^m\right)^{1-\sigma} \tag{4}$$

Taking logs of (4) implies that the decision of where to obtain input j is additively separable from the choice of where to obtain all other inputs. Doing so, and defining the log expected profits attributable to using input j as $\ln \pi_{j,t}^m$ gives the following expression:

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

where

$$\ln \pi_{j,t}^m = \max_{x_{j,t}^m} \nu\left(\sigma-1\right) \ln \lambda_{x,j,t}$$
$$+ \left(1-\alpha\right)\left(1-\sigma\right)\gamma_j \left[\mathbb{E}\left[\ln p_{x,j,t}^m\right] + \zeta_{x,j} \mathbb{1}\{x_{j,t}^m \neq x_{j,t-1}^m\} + \zeta_{c,j} \mathbb{1}\{c_{j,t}^m \neq c_{j,t-1}^m\}\right] \tag{5}$$

and $A = \ln\left\{\frac{1}{\sigma} B \left(\frac{\sigma}{\sigma-1}\right)^{1-\sigma} w^{\alpha_m(1-\sigma)} \psi_m^{\sigma-1}\right\}$ captures all the terms not associated with the cost of an input bundle for the final good.[13] [14] Importer m will choose the supplier $x_{j,t}$ that gives the highest input-specific profits. 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

[13] In Equation (5), 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.

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

choice, so expected profits are calculated for each choice, and the partner with the highest expected profits will be chosen. Dividing Equation (5) through by $(\sigma - 1)$, I define the *exporter-specific* expected log profit term for any choice x_t^m as $\overline{\pi}_t^m\left(x_t^m, \boldsymbol{\beta}\right)$:

$$\overline{\pi}_t^m\left(x_t^m, \boldsymbol{\beta}\right) = \xi \ln \lambda_{x,t} + \beta_p \mathbb{E}\left[\ln p_{x,t}^m\right] - \beta_x \mathbb{1}\{x_t^m \neq x_{t-1}^m\} - \beta_c \mathbb{1}\{c_t^m \neq c_{t-1}^m\} \tag{6}$$

where $\xi = \nu$, $\beta_p = -(1-\alpha)\gamma$, $\beta_x = (1-\alpha)\gamma\zeta_x$, and $\beta_c = (1-\alpha)\gamma\zeta_c$. I summarize the vector of unknown parameters as $\boldsymbol{\beta} = \{\beta_p, \beta_x, \beta_c, \xi\}$.

Equation (6) is the cornerstone of my estimation strategy. Starting from a general model of importing behavior, I have derived a choice-specific profit equation that can be estimated using techniques of structural industrial organization[15]. 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 $\{\beta_p, \beta_x, \beta_c\}$ by industry 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 (6) can be achieved by calculating expected log profits from each exporter, 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. As in Rust (1987), I allow for a stochastic profit shock from choosing x_t that is observable to the importer, $\epsilon_{x,t}^m$, which helps the model match the data. This profit shock stands in for all of the other unobservable traits not included in (6) that may result in a match[16]. 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.

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 (2011), and continue to drop the j subscript. I assume monopolistic competition, fixed markups over marginal cost,

[15] This equation resembles the worker utility function from choosing different employers discussed in Fox (2010).
[16] Appendix E lays out a version of Equation (6) that accounts for the possibility of serial correlation.

and a random difference in prices across importers it serves $\rho^m_{x,t}$:

$$p^m_{x,t} = \mu MC_{x,t}\rho^m_t = \mu \frac{1}{z_x}\left(\lambda_{x,t}\right)^{\beta}\rho^m_{x,t} \tag{7}$$

where z is the idiosyncratic productivity of exporter x, and λ is quality. Exporters simply set a price and wait to be chosen by importers. This simple setup allows me to set up the expected price function importers use.

Taking logs, and assuming that the year-to-year changes in quality of an exporter is small over time, we obtain a transition rule for prices.

$$\ln p^m_{x,t} = \ln p^m_{x,t-1} + \left(\ln \rho^m_{x,t} - \ln \rho^m_{x,t-1}\right) \tag{8}$$

3.2.2 Price Evolution

Let $\mathbf{p_t}$ represent the vector comprising each exporter's price at time t, i.e. $\mathbf{p_t} = \{p_{1,t}, p_{2,t}, ...p_{X,t}\}$, for the total number of exporters X. In this section, I use Equation (8) to specify a simple one-step ahead process for prices, $f\left(p^m_{x,t} |\ \mathbf{p_{t-1}}, x^m_{t-1}, x^m_t\right)$ that is known to importers. Thus I can specify every exporter price that an importer faces.

Importers know both what prices were paid by other firms in previous periods, and the distribution of shocks to those prices in the next period. Guided by Equation (8), the basic formula I apply is as follows: firms expect to pay the price they paid last period if they keep their partner, while they expect to pay the average price all other importers paid from a given supplier if they switch to that supplier. In addition, importers are aware that prices in a given geographic area will change by a given amount from one year to the next, and know that change perfectly.[17] However, there is also a stochastic element to the price with a known mean and standard deviation.

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

$$\ln p^m_{x,t} = \ln p^m_{x,t-1} + \eta_{c,t} + u^m_{x,t}, \quad u^m_{x,t} \sim \mathcal{N}\left(0, \sigma^2_{x,t}\right) \tag{9}$$

Importing firms know there is a city-specific change in prices $\eta_{c,t}$, as well as the distibrution 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.

[17] In terms of Equation (8), one component of the $\left(\ln \rho^m_{x,t} - \ln \rho^m_{x,t-1}\right)$ is a known increase in average prices across the entire city. If Importer A chooses to stay with Supplier Z in city Q, they expect the price paid last period plus the increase in prices in city Q.

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

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

N is the number of firms who imported from firm $\widetilde{x}$ from the previous period, and they are indexed $n = 1, ...N$. Each price paid by importer n is $p^n_{\widetilde{x},t-1}$. 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 $f\left(p^m_{x,t} |\ \mathbf{p_{t-1}}, x^m_{t-1}, x^m_t\right)$. In other words, given state variables $\mathbf{p}$ and x, and 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 $\{\eta_{c,t}\}_{c=1}^{C}$ (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 $\{\{\eta_{c,t}\}_{c=1}^{C}, \{\sigma_{x,t}\}_{x=1}^{X}\}$ by using observed prices in both pre- and post-periods and estimating equations (9)-(10) for each industry[18].

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 (6). 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 $(M \times X)$ of next period's profit shocks $\boldsymbol{\epsilon_{t+1}}$. These states evolve according to the joint density $h\left(\mathbf{p_t}, \boldsymbol{\epsilon_{t+1}}\right)$.

[18] 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.

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 by Equation (6), the infinite-time problem for any importer (dropping the m superscript) is summarized by the following value function:

$$V\left(x_{t-1}, \mathbf{p_{t-1}}, \boldsymbol{\epsilon_t}\right) = \max_{\{x_t, x_{t+1}, \ldots\}} \mathbb{E}\left[\sum_{\tau=t}^{\infty} \delta^{\tau-t}\left(\overline{\pi}_\tau\left(x_\tau, \mathbf{p_{\tau-1}}, x_{\tau-1}, \boldsymbol{\beta}\right) + \epsilon_{x,\tau}\right)\right] \tag{11}$$

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 (11) can be rewritten as a Bellman Equation:

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

for

$$EV\left(x', \mathbf{p}, x, \boldsymbol{\epsilon'}\right) = \int_{\mathbf{p}'} \int_{\boldsymbol{\epsilon''}} V\left(x', \mathbf{p}', \boldsymbol{\epsilon''}\right) h\left(\mathbf{p}', \boldsymbol{\epsilon'}|\mathbf{p}, x, x', \boldsymbol{\epsilon'}\right) d\mathbf{p}' d\boldsymbol{\epsilon''}. \tag{12}$$

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 :

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 cost parameters $\boldsymbol{\beta}$, $P\left(x^C|s,\boldsymbol{\beta}\right)$, is:*

$$P\left(x^C|s,\boldsymbol{\beta}\right) = \frac{\exp\left[\overline{\pi}\left(x^C,s,\boldsymbol{\beta}\right) + \delta EV\left(x^C,s\right)\right]}{\sum_{\widehat{x}\in X}\exp\left[\overline{\pi}\left(\widehat{x},s,\boldsymbol{\beta}\right) + \delta EV\left(\widehat{x},s\right)\right]} \tag{13}$$

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

$$EV(x,s) = \int_{s'} \log\left\{\sum_{x'\in X}\exp\left[\overline{\pi}\left(x',s',\boldsymbol{\beta}\right) + \delta EV\left(x',s'\right)\right]\right\} f\left(s'|s,x\right) \tag{14}$$

Proof *See Appendix C.*

3.4 Maximum Likelihood Estimation

The parameters $\boldsymbol{\beta}$ can then be found via maximum likelihood estimation. Let x^m_{t-1} be the actual choice of exporter at time $t-1$ for importer m from data. Then the likelihood of observing importer m choosing exporter x^m_t is:

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

And thus the total likelihood function for the set of importer choices at time t is:

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

The constraints for the maximization problem are the system of fixed point equations defined by Equation (14). To solve this problem, I follow the MPEC approach as described in Su and Judd (2012) and Dubé, Fox, and Su (2012), namely an inner loop for solving the fixed point problem in (14) for the constraint vector $\mathbf{EV}$ 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 solve is:

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

$$\max_{\boldsymbol{\beta}} \sum_{m=1}^{M} \log \frac{\exp\left[\overline{\pi}_t^m\left(x_t^m, s_t^m, \boldsymbol{\beta}\right) + \delta EV\left(x_t^m, s_t^m\right)\right]}{\sum_{\widehat{x}_t^m \in X} \exp\left[\overline{\pi}_t^m\left(\widehat{x}_t^m, s_t^m, \boldsymbol{\beta}\right) + \delta EV\left(\widehat{x}_t^m, s_t^m\right)\right]} + \sum_{m=1}^{M} \log f\left(p_{x,t}^m | s_t^m, x_t^m\right) \tag{15}$$

$$s.t.$$

$$EV\left(x_t, s_t\right) = \int_{s_{t+1}} \log \left\{ \sum_{x_{t+1} \in X} \exp\left[\overline{\pi}_{t+1}\left(x_{t+1}, s_{t+1}, \boldsymbol{\beta}\right) + \delta EV\left(x_{t+1}, s_{t+1}\right)\right] \right\} f\left(s_{t+1} | s_t, x_t\right) \tag{16}$$

Solving this problem produces maximum likelihood estimates for the vector of parameters $\boldsymbol{\beta}$. The next section describes the particulars of how this model is taken to the data.

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 (15)-(16) and model performance on generated data with pre-set parameter values. 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 (16) 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 (16) as:

$$EV\left(x_t, \widehat{s}_t\right) = \sum_{\widehat{s}_{t+1}=1}^{N} \log \left\{ \sum_{x_{t+1} \in X} \exp\left[\overline{\pi}_{t+1}\left(x_{t+1}, \widehat{s}_{t+1}, \boldsymbol{\beta}\right) + \delta EV\left(x_{t+1}, \widehat{s}_{t+1}\right)\right] \right\} Pr\left(\widehat{s}_{t+1} | \widehat{s}_t, x_t\right) \tag{17}$$

where $\widehat{p}$ 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, Fox, and Su (2012). The MPEC maximization protocol uses values of the vector $\boldsymbol{\beta}$ that satisfy the fixed point equation (15), 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 $EV(x_t, \widehat{s}_t)$. Thus the constraint set in (17) 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 (15) to see which give the closest match between the estimated probabilities and the true data. Asymptotic normal standard errors are calculated following Rust (1994).

Before computing the model on U.S.-China trade data, I first set the parameters at fixed values and create 250 Monte Carlo replications of data based on these values. Every importer m is assigned an exporter x_{t-1} and price $p_{x,t-1}$ from a previous period, and predicts the expected price received from every potential exporter x_t, i.e. $\mathbb{E}[p_{x,t}|p_{x,t-1}, x_{t-1}, x_t]$. The importer matches with an exporter, given both these expected prices and the pre-set values of the parameters β_p, β_x, and β_c.[19] I then utilize the observed outcomes and prices from each dataset to run the maximum likelihood problem found in (15) and (17), extracting the cost parameters consistent with those choices. I set $\delta = 0.975$ and use $N = 5$ price states.

Since the total number of importers (M), exporters (X), and exporter cities (C) are free to choose, I set the number of importers at 30 and create three different samples: one small ($M = 30, X = 4, C = 3$), one with the average number of exporters in the data but few cities ($X = 33, C = 3$), and one matching both the average number of exporters and the average number of cities in the data ($X = 33, C = 9$).[20] I then run the estimation routine on each set of data, and report summary statistics for how well the procedure matches the pre-set values. The results of this procedure are presented in Table 3.

It is clear that the estimation routine matches the pre-set parameters poorly on the very small sample. Furthermore, mean estimates of the exporter switching effect β_x are extremely high. These results occur because with such a small sample space (four exporters and three cities), a number of Monte Carlo runs likely have very few cases of within-city switching, providing very high estimates for the base exporter switching cost. Yet even in this scenario, the median results preserve the ordering of the originally set parameters. In addition, the elasticity of the switching decision with respect to prices, β_p is of reasonable sign and size.

Once the sample size is extended to 33 exporters (the average number of exporters in an HS6 industry), some of the results improve. For Sample B, which contains fewer cities, we see a vast improvement in the measurement of the price coefficient β_p, and the exporter switching cost β_x. A greater number of exporter possibilities permits more partner switching observations, allowing for better estimation of this parameter.

[19] For this estimation, I do not include quality estimation terms, given the extra assumptions I would have to make to run the quality estimation protocol described below. The profit equation is the same as (3), only without the λ term.

[20] The average number of Chinese exporters in an HS6 industry is 33.58. The median number of exporters is eight. The industry at the 90th percentile of exporters contains 81 exporters. 15.3% of the 3000 or so HS6 codes found in the trade data contain only one exporter. These figures are based on computation of the "main" exporters, i.e. exporters who are found after assigning a "plurality exporter" to each importer.

Notably, the mean and median of the outside-city switching cost β_c is much higher than the pre-set value of the parameter, as the number of cities is small enough to make the estimation procedure assign higher costs of switching cities.

Finally, results for Sample C demonstrate that the procedure improves further when the number of cities is increased to the average number of exporting cities found in the LFTTD, $C = 9$. Not only do estimates of the city switching cost decrease in mean and median to levels much closer to the preassigned values, but the exporter switching cost and price elasticity similarly approach their values. My estimation procedure thus is expected to perform better in industries that have enough observations of within-city, out-of-city, and non-switching observations to estimate the parameters of interest, and in those industries, will deliver reasonable results.

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 (15) with the constraints in Equation (17). 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 (6). 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 calculating expected prices in Section 3.2. 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 given data in the LFTTD. The intuition follows from Kim and Petrin (2010) and Khandewal (2011): 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:

$$\ln p_{x,t} = \ln \mu - \log z_{x,t} + \ln \lambda_{x,t} \tag{18}$$

This equation follows directly from the specification of exporter prices described in Section 3.2.1, although the price is now the average price of exporter x_t charged to all importers and is not specific to an importer-exporter match.[21] From the trade data, I cannot observe the productivity of individual Chinese exporters. However, there are a number of variables in the data that I use to proxy for productivity, including total U.S. exports, number of HS products exported, number of years exporting to the U.S., number of import partners, and number of transactions. For each industry, I group these terms into a firm-specific vector of covariates Z_x, and together with time fixed effects, regress the exporter's average offered price (firm-level unit value) on these variables. I then take the residual from this regression and call it quality. The approach is similar to the "quality ladder" estimation procedure described in Khandewal (2010).[22]

Finally, I describe the final pieces of the puzzle necessary to run the MLE problem above. Firstly, with some industry variation, 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 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.[23]

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 broadly, consisting of a) any exporter used 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 majority supplier in time t.

The last step is to clean the LFTTD by eliminating unreasonable prices. 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. Before averaging prices across transactions, I eliminate any transactions with prices greater than the 90th percentile for an HS6 industry that are also greater than 10 times the median price in that industry. I then repeat the process, again eliminating prices that are greater than 10 times the new median price.

[21] I also assume the exponent on quality β is 1.

[22] 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.

[23] Analysis of the LFTTD indicates that U.S. importers typically import a very large share of their total imports from only one partner. The average share of imports that come from a U.S. importer's main Chinese partner is 83.9%, with a standard deviation of 22%.

I use the above procedure to estimate the model for a large number of industries, using data on U.S.-China trade from 2005-2006. I use TOMLAB / KNITRO to compute the Jacobian and the gradient for Equations (15) and (17) analytically, and then solve the above MLE problem.

4.3 Estimation Results

A key result from the estimation procedure is that the cost of switching is large. I run the model on 50 HS6 industries, chosen to represent industries across the spectrum of imported products from China that are also large enough to have both partner and city switching.[24] The value-weighted average of switching costs across the sample of industries is $\beta_x = 2.99$ and $\beta_c = 1.61$. The numeric results are interpreted in units of the Type I Extreme Value shock, which has mean 0 and standard deviation $\sqrt{\frac{\pi^2}{6}} \approx 1.29$, which gives 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 positive shock to profits that is approximately $2.99/1.29 = 2.3$ standard deviations from the mean. If that new partner is located in a separate city, then that new partner must provide a positive shock to profits that is $(2.99 + 1.61)/1.29 = 3.6$ standard deviations from the mean. I do the identical exercise for 50 of the largest HS10 products imported from China by volume, and find extremely similar results. Thus the estimated parameters confirm the reduced form results that frictions from switching are large.

Furthermore, the costs are highly heterogeneous across industries. Rather than presenting the full battery of results (available upon request), in this subsection I present estimates of the parameters in illustrative industries.[25] 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.[26] 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. Of all switching importers in HS6 841320 from 2005 to 2006, only 30% switched cities. For HS6 820310, the respective figure is 15%. 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,

[24] The list of industries and their trade shares are listed in Table A2.

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

[26] Asymptotically normal standard errors are calculated, and are available upon request. All results discussed in this subsection are significantly different from zero according to the 99% threshold.

given that calculations of these parameters take into account how much switching is actually occurring. 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 a high proportion of the total exporter 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 provided in Broda and Weinstein (2005), 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). The estimates are based on import data from the United States in the 1990s. 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 substition (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, Van Biesebroeck, Wang, and Zhang (2012)

between HS codes and the China Industry Code (CIC) system to analyze different types of industries based on their domestic characteristics. Using China National Bureau of Statistics firm-level data for 2005, I isolate CIC industries where exporters have particular characteristics relevant to the importer-exporter partnership decision. For example, I compare industries with highly skilled workers to unskilled workers, and industries composed of large firms to industries composed of small firms. 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.

The first set of results relates to the labor productivity of workers in different exporting industries. Chinese exporters in the industry "Arms and Ammunition" (CIC 3663) are in the lower tail of value added per worker relative to other industries. On the other hand, exporters in the industries "Rolling and Processing of Rare Earths" and "Tungsten and Molybdenum Smelting" have very high levels of value added per worker. As can be seen from the top panel of Table 5, industries with lower worker productivity tend to be characterized by lower exporter switching costs, while those industries with high levels of worker productivity have much higher exporter switching costs. This result is intuitive, as it implies importing firms who are importing products with highly productive workers receive greater relationship-specific benefits, and breaking up is more costly. On the other hand, firms with unproductive workers have little to distinguish themselves from competing firms, and thus have lower costs of switching from one to another.

I also compare results for firms of different employment sizes. Exporters in the industry "Other Ward Care and Medical Equipment" (CIC 3689) are of very small size, compared to exporters in the industry "Arms and Ammunition" (CIC 3663). The bottom panel shows importers importing in a product category that is dominated by small firms tend to value their relationship more, while an industry dominated by large firms is characterized by smaller exporter switching costs and more relationship breakups. 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 (2010).

In summary, the results are broadly what we would expect of the estimates *ex ante*: higher exporter switching costs relative to city switching costs appear in 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 whole 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 matching U.S. import prices from China, would change in response to falling switching costs. Conversely, I also examine how prices would be affected by increases in switching costs, such that switching occurred far less frequently than is seen in the data.

I follow the procedure outlined by the BLS Handbook of Methods to calculate the Import Price Index for my sample (U.S. Department of Labor, 2013).[27] Using the same sample of 50 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), I generate outcomes given randomly drawn extreme-value shocks and the estimated parameters.

Specifically, for each industry j in the set of industries I use J, the industry price index P_j sums together firm-level prices, weighted by the share of one firm's imports in total industry imports:

$$P_j = \sum_{i \in I_j} \omega_i p_i \tag{19}$$

In the above, p_i is a summary measure-the mean or median- of firm i's received price across 1000 replications.[28] 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 w_j across the industries in my sample:

$$P = \sum_{j \in J} w_j P_j \tag{20}$$

The result is an price index that accounts for firm size and industry size. I create the same index for

[27] I make one deviation from the BLS methodology, as I compute the index for each 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.

[28] 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.

each different simulation and compare it to the generated data according to the original parameters.[29]The results are in Table 6.

The first thought experiment is to reduce both β_x (the partner-switching cost) and β_c (the city-switching cost) by half for all industries, and determine the size of the efficiency gain when more importers can separate and/or find better matches. I find that the U.S.-China Import Price Index decreases by 12.5% in response to such a change, as seen in Table 6 Column 2. Since the fixed switching cost is measured in units of the Type I Extreme Value ϵ shock, and the average β_x is approximately 3 (about 2.5 standard deviations of the shock), 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 sample of industries I examine) approximately 57% of firms stayed with their partner. This reduction results in only about 18% of all total importers now staying with their partners. I calculate the price index for each of 1000 Monte Carlo simulations under both the original and the adjusted parameter set, and present the kernel densities in Figure 5.

One can see the same pattern in individual industries as well. Figure 6 shows the distribution of industry price indices (Equation 19) 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 infrastructure and its national development strategy focusing on 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, 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 short-term contracts while still remaining confident in the ability to find quality inputs at acceptable prices over

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

the long-term.

A second thought experiment is to weigh the relative effects of β_x vs. β_c. These parameters are broadly interpretable as the overall difficulty in leaving one's trading partner, for reasons of information frictions, long-term contracts, etc. versus city-specific effects, such as geographic agglomeration, better distribution networks, or how rapidly developing cities compare to well-known exporting hubs. Table 6 Column 3 shows then when the partner-switching cost β_x is reduced to zero, the amount of importers staying reduces even further, such that the total number of staying importers drops to 8%. Importers are also far more likely to switch city under this scenario, as the total cost of switching cities $\beta_x + \beta_c$ reduces substantially (β_x is typically 2 times higher than β_c in the average industry). The effects on prices are even stronger than the original case: if almost all importers can break up from their partners, then the new matches have a 15% lower Import Price Index, as compared to the original frictions. On the other hand, reducing the city friction to zero means that, as expected, importers do switch city more often, as can be seen in Table 6 Column 4. However, the partner friction is important enough to still leave 30% of all importers staying with their original partner, smaller than either of the previous two counterfactuals. The effects on prices is also smaller, with the Import Price Index reducing by 7.37%.

Finally, I consider the case where both the frictions are tripled, thus shutting down many of the original switches under the original parameters. Table 6 Column 5 demonstrates that such a change increases the number of importers staying with their partner to 90%, meaning the vast majority of firms are now unwilling to leave their partner. Indeed, the effect of importers being unable to move increases the prices received by 7.62%. Such a finding confirms the importance of supplier switching in overall price changes.

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 the size of efficiency gains through lower prices are a robust prediction of the model I estimate, and are large in scope.

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 use my model to estimate how low prices would have to be 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 one firm, and assign it a different price to create separate scenarios. 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 "quality" (residual of the regression of price on observable exporter characteristics). I then re-solve the fixed point equation in (17) 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 even farther 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 use the techniques of industrial organization to estimate the parameters of interest. Switching costs are large, and heterogeneous across 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 12.5% 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. The geographic link between importers and exporters gives us a new way to understand how shocks in a specific area move through international trade, a field of study that has thus far been limited to industry-to-industry linkages. 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 will asses the impact of specific regional policies on importer behavior, such as the formation of special export zones in cities such as Shenzhen. Being able track exporter dropouts from the U.S. import data presents a reasonable degree of exogenous variation that can be used to determine U.S. final good producer behavior in response to an unexpected loss of members of its supply chain. 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

[1] ACE Entry Summary, U. S. Customs and Border Protection Business Rules and Process Document (Trade), Version 3.3. U.S. Customs and Border Protection. http://www.cbp.gov/linkhandler/cgov/trade/trade_programs/entry_summary/ace/ace_es_business.ctt/ace_es_business.pdf (Accessed May 10, 2013).

[2] Ackerberg, Daniel (2009) "A new use of importance sampling to reduce computational burden in simulation estimation" Quantitative Marketing and Economics, December 2009, Volume 7, Issue 4, pp 343-376.

[3] Ahn, JaeBin, Amit K. Khandewal, and Shang-Jin Wei (2011), "The Role of Intermediates in Facilitating Trade". Journal of International Economics, vol. 84 (1), 73-85.

[4] Alessandria, George (2009), "Consumer Search, Price Dispersion, And International Relative Price Fluctuations". International Economic Review, vol. 50(3), pages 803-829, 08.

[5] "Alibaba.com Announces Major Initiatives for Buyers and Suppliers" (2008). Alibaba.com Limited Press Release, November 3, 2008. Accessed October 29, 2013, http://www.alibaba.com/aboutalibaba/releases_081103.html

[6] Allen, Treb (2012), "Information Frictions in Trade". 2012 Meeting Papers 125, Society for Economic Dynamics.

[7] Antràs, Pol and Arnaud Costinot (2011). "Intermediated trade". The Quarterly Journal of Economics 126 (3), 13191374.

[8] Armenter, Roc, and Miklós Koren, (2014). "A Balls-and-Bins Model of Trade." American Economic Review, 104(7): 2127-51.

[9] Bernard, Andrew B., J. Bradford Jensen, and Peter K. Schott (2009), "Importers, Exporters, and Multinationals: A Portrait of Firms in the U.S. that Trade Goods". In *Producer Dynamics: New Evidence from Micro Data*, University of Chicago Press, p. 513-552.

[10] Bernard, Andrew B., Andreas Moxnes, and Karen Helene Ulltveit-Moe (2013), "Two-Sided Heterogeneity and Trade" CEPR Discussion Paper No. DP9681, October 2013.

[11] Brandt, Loren, Johannes Van Biesebroeck, Luhang Wang, and Yifan Zhang (2012), "WTO Accession and Performance of Chinese Manufacturing Firms". CEPR Discussion Papers 9166, C.E.P.R. Discussion Papers.

[12] Broda, Christian and David Weinstein (2008), "Globalization and the Gains from Variety". Quarterly Journal of Economics Volume 121, Issue 2.

[13] Campa, Jose Manuel, Linda S. Goldberg, and Jose M. Gonzalez-Minguez (2005). "Exchange rate pass-through to import prices in the Euro area". Staff Reports 219, Federal Reserve Bank of New York.

[14] CBP Form 7501 Instructions, Department of Homeland Security, U.S. Customs and Border Protection. http://forms.cbp.gov/pdf/7501_instructions.pdf (Accessed May 10, 2013).

[15] Drozd, Lukasz A., and Jaromir B. Nosal (2012). "Understanding International Prices: Customers as Capital". American Economic Review, 102(1), 364-95.

[16] Dubé, Jean-Perre, Jeremy T. Fox, and Che-Lin Su (2012), "Improving the Numerical Performance of Static and Dynamic Aggregate Discrete Choice Random Coefficients Demand Estimation". Econometrica, Vol. 80 (5), 22312267

[17] Eaton, Jonathan and Samuel Kortum (2002), "Technology, Geography, and Trade". Econometrica vol. 70(5), pages 1741-1779, September.

[18] Eaton, Jonathon, Marcus Eslava, David Jinkins, C.J. Krizan, Maurice Kugler, and James Tybout, (2012), "A Search and Learning Model of Export Dynamics". Mimeo.

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

[20] Feenstra, Robert C., Zhiyuan Li and Miaojie Yu, (2011), "Exports and Credit Constraints Under Incomplete Information: Theory and Evidence from China". NBER Working Papers 16940, National Bureau of Economic Research, Inc.

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

[22] Hallak, Juan Carlos and Jagadeesh Sivadasan (2011), "Firms' Exporting Behavior under Quality Constraints. Research Seminar in International Economics No. 628, September 2011.

[23] Hsieh, Chang-Tai and Peter J. Klenow (2009), "Misallocation and Manufacturing TFP in China and India". Quarterly Journal of Economics, Volume 124, Issue 4, Pages 1403-1448.

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

[25] Kamal, Fariha and C.J. Krizan (2012), "Decomposing Aggregate Trade Flows: New Evidence from U.S. Traders". Working Papers 12-17, Center for Economic Studies, U.S. Census Bureau.

[26] Kamal, Fariha C.J. Krizan, and Ryan Monarch (2014), "Identifying Foreign Suppliers in U.S. Merchandise Import Transactions". Mimeo.

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

[28] Khandewal, Amit (2010) "The Long and Short (of) Quality Ladders". Review of Economic Studies, 77(4), 1450-1476.

[29] Kleshchelski, Isaac and Nicolas Vincent (2009), "Market Share and Price Rigidity". Journal of Monetary Economics, Volume 56, Issue 3, April 2009, Pages 344352.

[30] Krizan, C.J. (2012), "The Nuts and Bolts of International Trade Transactions: A Data User's Primer". U.S. Census Center for Economic Studies.

[31] Kim, Kyoo-il and Amil Petrin (2010), "Control Function Corrections for Omitted Attributes in Differentiated Product Models'.' Working Paper, University of Minnesota.

[32] Kranton, Rachel E. and Deborah F. Mineheart (2001), "A Theory of Buyer-Seller Networks". American Economic Review, Vol. 91(3), 385-508.

[33] Lincoln, William and Andrew McCallum (2011), "Entry Costs and Increasing Trade". Working Papers 11-38, Center for Economic Studies, U.S. Census Bureau.

[34] McFadden, Daniel (1981), "Econometric Models of Probabilistic Choice", in *Structural Analysis of Discrete Data*. Cambridge, MIT Press.

[35] Petropoulou, Dimitra (2011), "Information Costs, Networks and Intermediation in International Trade", Globalization and Monetary Policy Institute Working Paper 76, Federal Reserve Bank of Dallas.

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

[37] Powell, Bill (2011), "Why Alibaba's CEO had to go". *Fortune*, February 22, 2011. Accessed October 29, 2013, `http://tech.fortune.cnn.com/2011/02/22/why-alibabas-ceo-had-to-go/`

[38] Rust, John (1987), "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher". Econometrica Vol. 55 (5), 999-1033.

[39] Rauch, James (2001), "Business and Social Networks in International Trade". Journal of Economic Literature, Vol. 39, 1177-1203.

[40] Rauch, James and Joel Watson (2004), "Network Intermediates in International Trade". Journal of Economics and Management Strategy, vol. 13(1), 69-93.

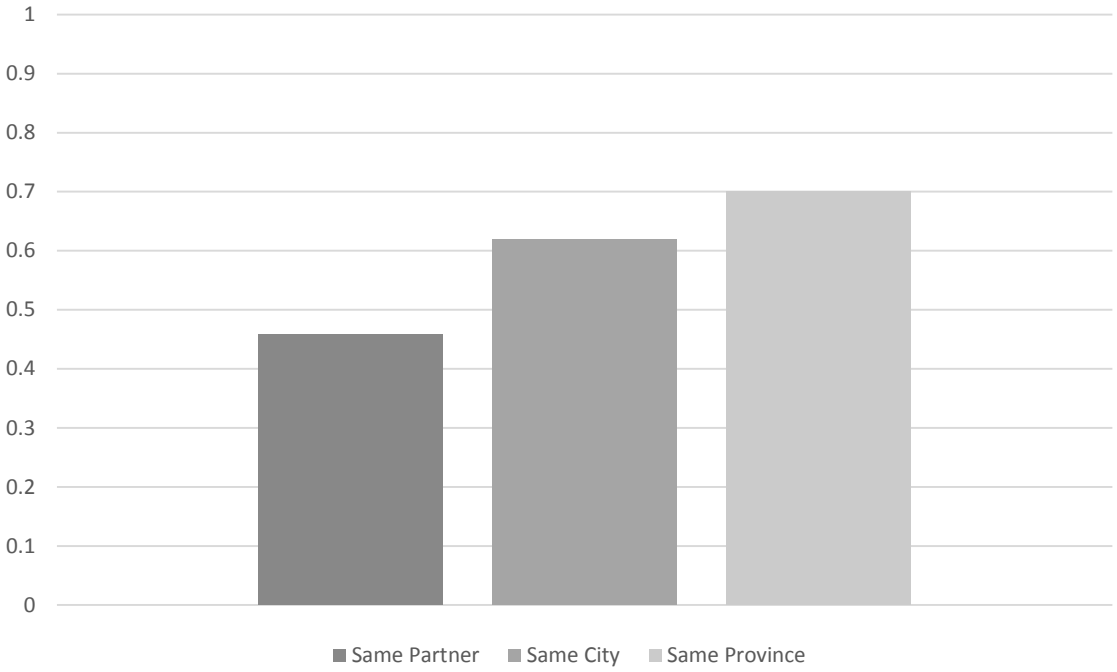
[41] Su, Che-Lin and Kenneth L. Judd (2012), "Constrained Optimization Approaches to Estimation of Structural Models". Econometrica, vol. 80(5), 2213-2230.

[42] Tang, Heiwai and Yifan Zhang (2012), "Quality Differentiation and Trade Intermediation". Centro Studi Luca d'Agliano Development Studies Working Paper No. 340.

[43] U.S. Department of Labor, Bureau of Labor Statistics (2013). *BLS Handbook of Methods*, Chapter 15. `http://www.bls.gov/opub/hom/homch15_f.htm` Last modified July 10, 2013, accessed October 23, 2013.

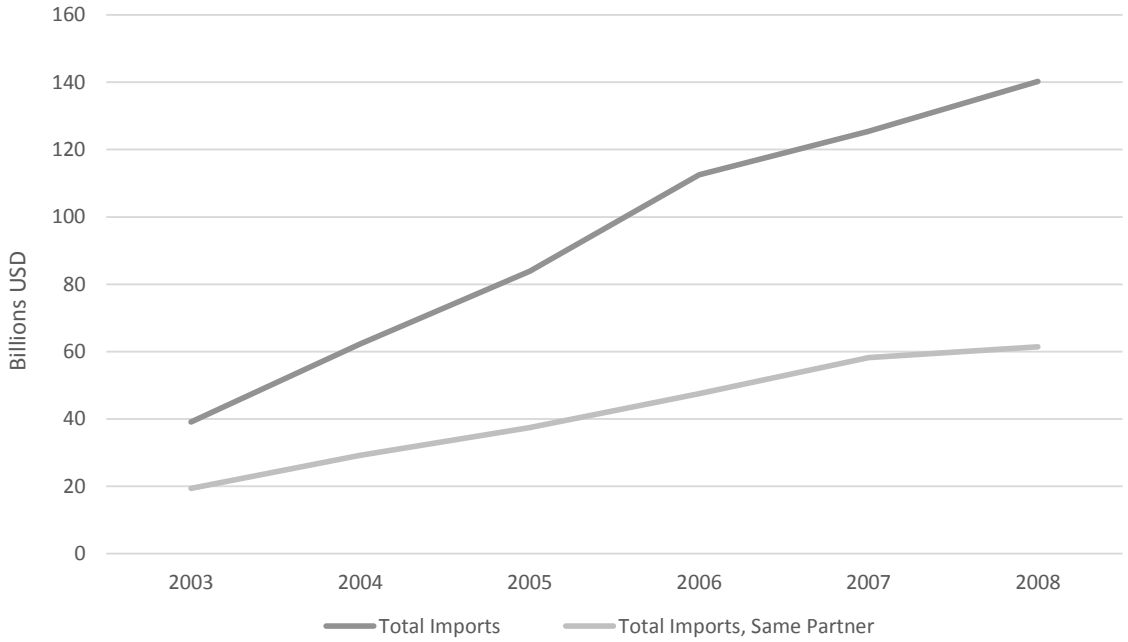
Tables and Figures

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



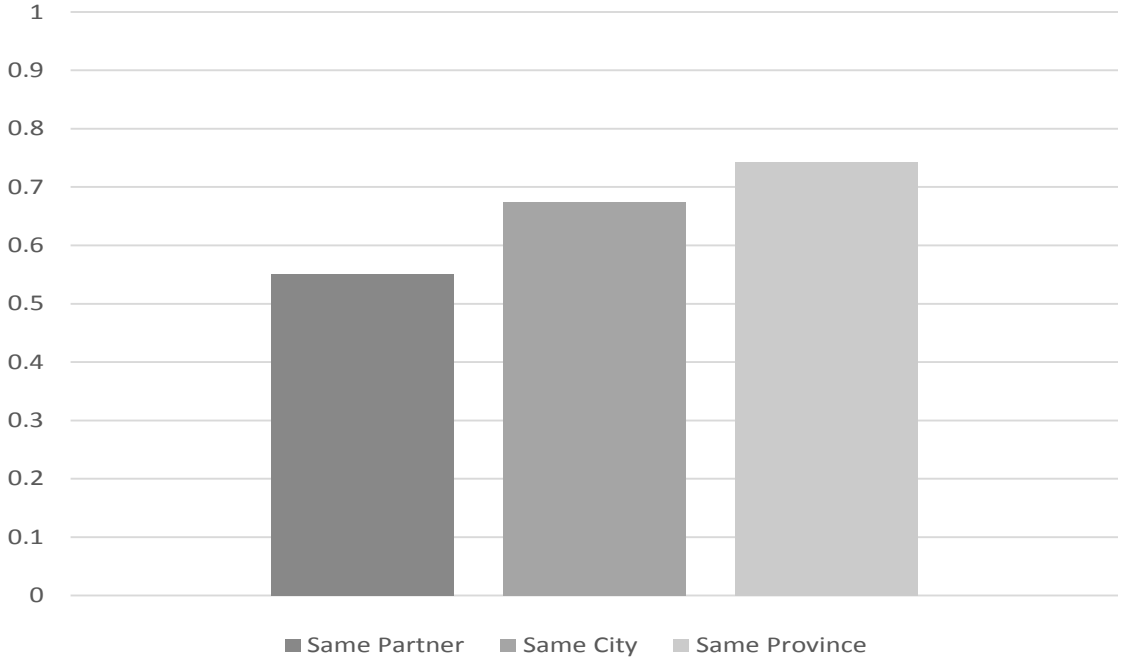
Notes: To determine if a U.S. importer (firm + HS10 product) kept the same exporting partner from one year to the next, I calculate the majority partner for each importer in each year, using the value imported from each manufacturing ID in the Longitudinal Foreign Trade Transaction Database (LFTTD). If this majority partner remained the same from year-to-year, then the importer "stayed" with its partner. If the city of the majority partner remained the same, then the importer stayed in its city, and if the majority 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: U.S. Imports and "Same Partner" Imports from China, 2003-2008



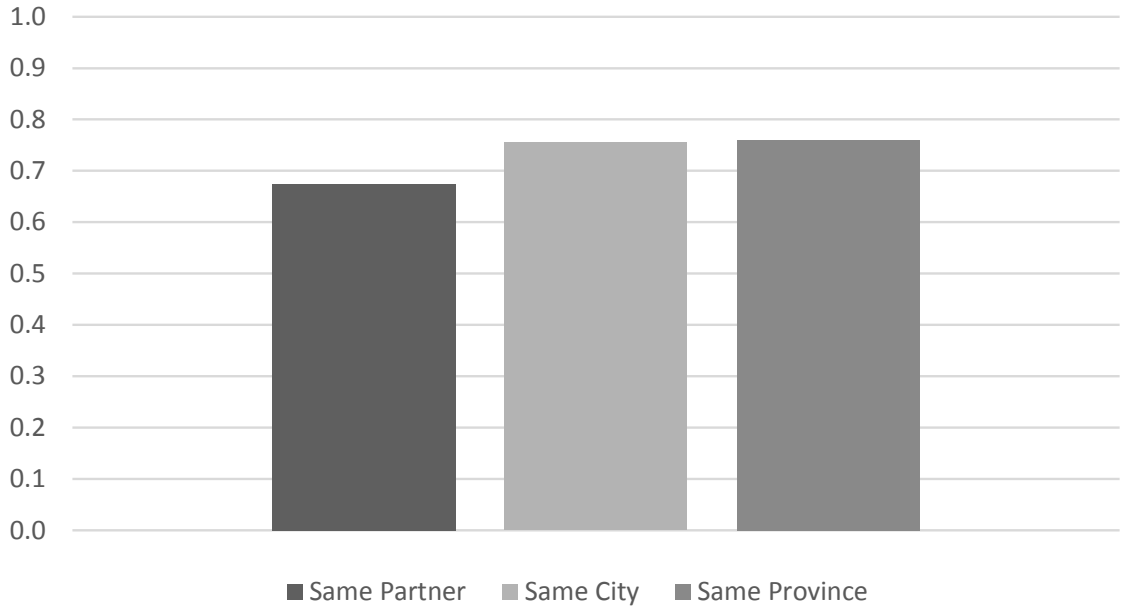
Notes: Total U.S. imports from China comes is the sum of all arm's length import value from the Longitudinal Foreign Trade Transaction Database (LFTTD). To determine if a U.S. importer (firm + HS10 product) kept the same exporting partner from one year to the next, I calculate the majority partner for each importer in each year. If this majority partner remained the same from year-to-year, then the importer "stayed" with its partner. Total imports among "Stayers" is the sum of all arm's length import value from these importers. I apply the panel concordance for HS10 products developed by Pierce and Schott (2012). Note that "total imports" also includes importers who began importing the year in question (the "extensive margin") while total imports from "stayers", by definition, cannot. Imports from non-entrants (the "intensive margin") are typically 85-90% of total imports from China over this time period.

Figure 3 Panel A: Year-to-Year "Staying" Percentages, Surviving Exporters, 2003-2008



Notes: For this figure, an importer is considered to have switched (not stayed with) its exporter only if two conditions are met: (a) the majority partner changed from one year to the next, and (b) the majority partner in the original year is still found exporting to someone else. The same procedure as above is followed to determine whether an importer stayed with its partner, city, or province.

Figure 3 Panel B: Year-to-Year "Staying" Percentages of U.S. Firms, 2003



Notes: For this figure, 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. The same procedure as above is followed to determine whether an importer stayed with its partner, city, or province.

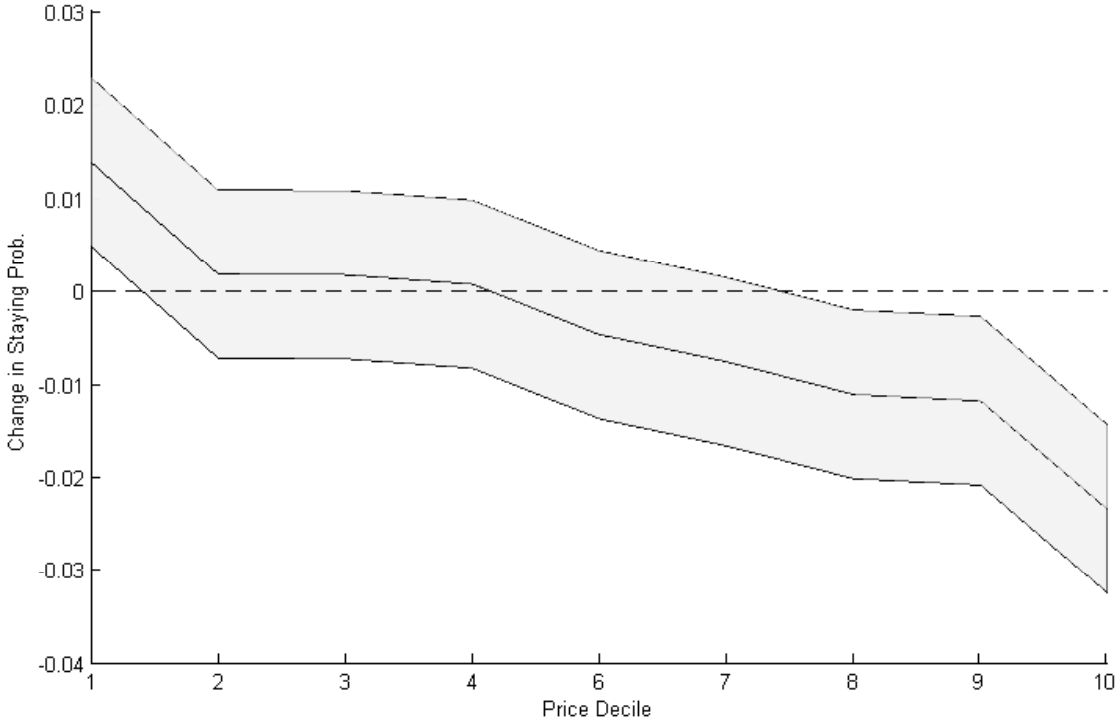
Table 1: Determinants of Supplier Stay/Switch Decision

Dependent Variable: Stayed with Chinese Exporter Year-to-Year, 2002-2008

	(1)	(2)	(3)	(4)
Log Price				
1st Decile	-0.0085**	0.0171***	0.0143***	0.0138***
	(0.004)	(0.003)	(0.003)	(0.003)
2nd Decile	-0.0052	0.0026	0.0020	0.0018
	(0.003)	(0.003)	(0.003)	(0.003)
3rd Decile	-0.0015	0.0015	0.0015	0.0017
	(0.003)	(0.003)	(0.003)	(0.003)
4th Decile	0.0005	0.0008	0.0007	0.0007
	(0.003)	(0.003)	(0.003)	(0.003)
6th Decile	-0.0038	-0.0043	-0.0046	-0.0047
	(0.003)	(0.003)	(0.003)	(0.003)
7th Decile	-0.0064**	-0.0069**	-0.0073**	-0.0076**
	(0.003)	(0.003)	(0.003)	(0.003)
8th Decile	-0.0109***	-0.0097***	-0.0110***	-0.0111***
	(0.003)	(0.003)	(0.003)	(0.003)
9th Decile	-0.0138***	-0.0091***	-0.0112***	-0.0118***
	(0.003)	(0.003)	(0.003)	(0.003)
10th Decile	-0.0357***	-0.0189***	-0.0225***	-0.0234***
	(0.003)	(0.003)	(0.003)	(0.003)
Log Supplier Size		0.0400***	0.0643***	0.0643***
		(0.001)	(0.001)	(0.001)
Supplier Age		-0.0017***	-0.0031***	-0.0029***
		(0.000)	(0.000)	(0.000)
Importer Size			-0.0322***	-0.0306***
			(0.001)	(0.001)
Constant	0.4438***	0.0211***	0.1047***	0.0863***
	(0.003)	(0.007)	(0.008)	(0.008)
Entry Year FE	No	No	No	Yes
N	510,485	510,485	510,485	510,485
R^2	0.07	0.09	0.09	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 equal to 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 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. Importer Entry Year is the first year a U.S. importers is found importing from China. Any importer that has the same share of imports from two separate Chinese suppliers is dropped.

Figure 4: Role of Price in Supplier Stay/Switch Decision



Notes: 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. 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 (HS10 product code and fim combination) 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 decile of price is excluded.

Table 2: Determinants of Supplier Stay/Switch Decision

Dependent Variable: Stayed with Chinese Exporter Year-to-Year, 2002-2008

	(1)	(2)	(3)	(4)
Log Price	-0.0084***	-0.0099***	-0.0104***	-0.0106***
	(0.001)	(0.001)	(0.001)	(0.001)
Log Supplier Size		0.0399***	0.0643***	0.0643***
		(0.001)	(0.001)	(0.001)
Supplier Age		-0.0024***	-0.0031***	-0.0029***
		(0.000)	(0.000)	(0.000)
Importer Size			-0.0322***	-0.0306***
			(0.001)	(0.001)
Constant	0.4353***	0.0193***	0.1011***	0.0823***
	(0.003)	(0.007)	(0.007)	(0.007)
Entry Year FE	No	No	No	Yes
N	510,485	510,485	510,485	510,485
R^2	0.07	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 equal to 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 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. Importer Entry Year is the first year a U.S. importers is found importing from China. Any importer that has the same share of imports from two separate Chinese suppliers is dropped.

Table 3: Monte Carlo Replication Results, based on 250 Replications

	β_p	β_x	β_c
Pre-Set Values	-0.5	-1	-3
Sample A: $M = 30, X = 4, C = 3$			
Mean	-0.626	1.269	4.728
Median	-0.564	0.954	3.690
Sample B: $M = 30, X = 33, C = 3$			
Mean	-0.543	0.843	8.209
Median	-0.540	0.837	5.017
Sample C: $M = 30, X = 33, C = 9$			
Mean	-0.538	0.985	3.427
Median	-0.512	1.055	3.081

Table 4: Selected Quantitative Estimates, HS Industrial Classification

HS6 Industry	β_p	β_x	β_c	β_c/β_x
Geographic Characteristics				
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
Market Size Characteristics				
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
Substitutability of Product				
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

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

CIC Industry	β_p	β_x	β_c	β_c/β_x
Worker Characteristics				
Low Skilled Workers				
Arms and Ammunition	-0.03	1.67	2.66	1.58
High Skilled Workers				
Rolling and Processing of Rare Earths	-0.01	2.63	1.86	0.70
Tungsten and Molybdenum Smelting	0.04	2.64	1.73	0.71
Firm Size Characteristics				
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

Table 6: Counterfactual Results (I)

	Original Sample	$\beta_x \downarrow 50\%$ $\beta_c \downarrow 50\%$	$\beta_x \downarrow 100\%$ $\beta_c =$	$\beta_x =$ $\beta_c \downarrow 100\%$	$\beta_x \uparrow 200\%$ $\beta_c \uparrow 200\%$
Price Index	-	-12.50%	-15.20%	-7.37%	+7.62%
Staying	*57%*	18%	8%	31%	90%
City Staying	*75%*	43%	47%	46%	93%

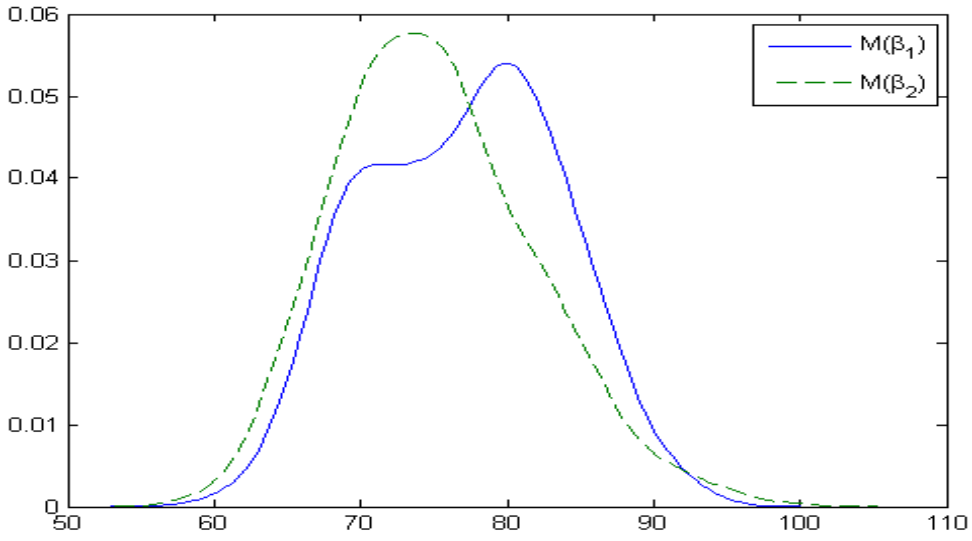
Notes: Objects computed by the model simulated with the originally estimated parameters are compared to the same objects in each of four counterfactual experiments: partner cost and city cost each reduced by half; partner cost reduced to zero, city cost unchanged; partner cost unchanged, city cost reduced to zero; and partner cost and city cost increased by three times. 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	3.86%	1.75%	56.27%
Mean	3.76%	1.53%	47.47%
75th Pct.	3.18%	0.77%	34.04%

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.

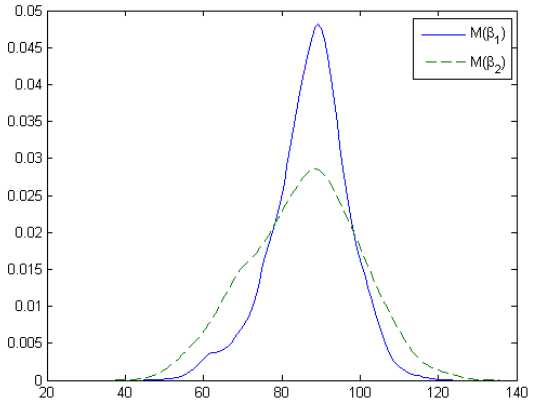
Figure 5: Kernel Density Plots, Original $\boldsymbol{\beta}$ vs. Divided by Half



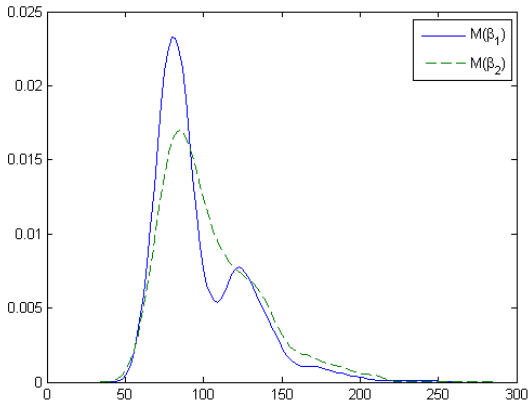
Notes: This figure presents kernel densities for the weighted average Import Price Index for importer-exporter matches predicted under the original parameters (solid line) and reducing switching costs by half (dashed line). I run 1000 replications of the model under each parameter set, and calculate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

Figure 6: Kernel Density Plots, Original $\boldsymbol{\beta}$ vs. Divided by Half, Selected Industries

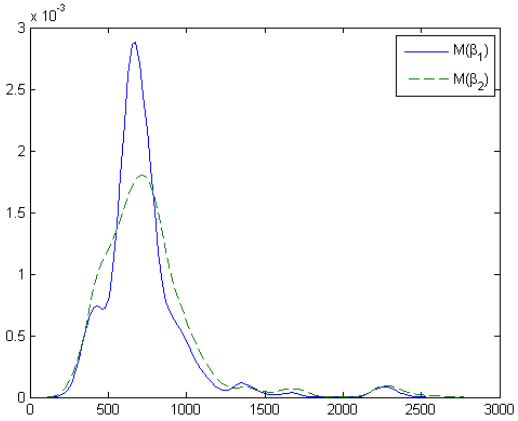
Panel A: HS 401120, Rubber Tires



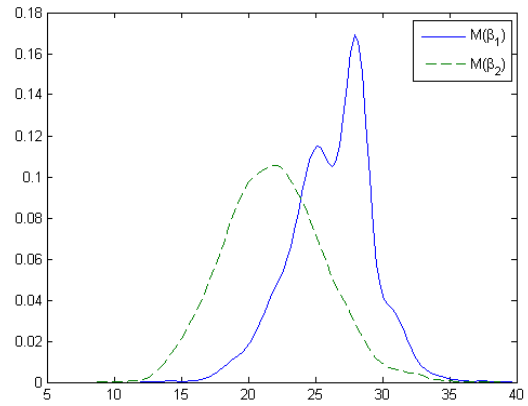
Panel D: HS 610432, Women/Girls' Cotton Jackets



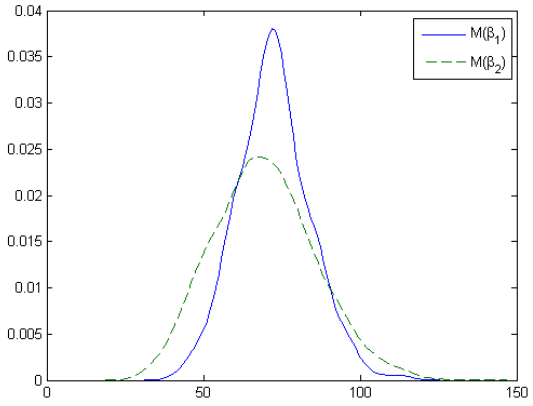
Panel B: HS 847130, Laptop Computers



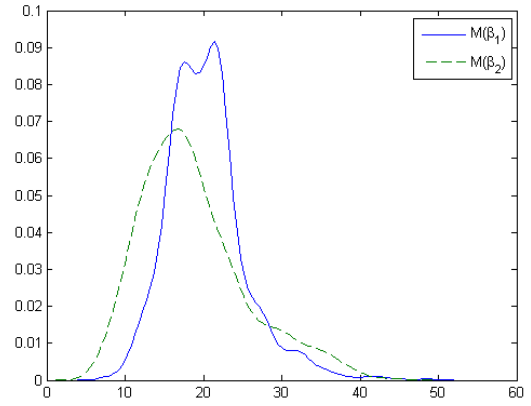
Panel E: HS 640340, Metal Toe-Cap Footwear



Panel C: HS 852520, Cell Phones



Panel F: HS 850940, Mixers/Blenders



Notes: This figure is 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. I run 1000 replications of the model under each parameter set, and calculate the price index for the matches predicted in each replication.

Appendix

Appendix A Robustness and External Validity of the MID

At this point, 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[30]. 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. I therefore first present some 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[31]. 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

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

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

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.[32]. 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 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 gurantee 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 from Table 1 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 have undergo such a change.

[32]For more information on this database, see Feenstra, Li, and Yu (2011).

Appendix B Robustness Checks for the Stylized Facts

In Section 2, I demonstrated that among U.S. continuing importers from China, partner switching is pervasive, there is a strong geographic component to the switching decision, and that these two trends become more pronounced over time. My baseline specification is to define a "switch" as a U.S. importing having a completely different set of partners for a product from one year to the next. Below I lay out a number of other specifications.

Figure B1 shows that the same stylized facts carry through even when using an alternate of switching: defining an importer as "staying" if it stayed with at least one or more of its partners: importer-exporter relationships are highly volatile, and geography matters a great deal in the switching decision. The same results come through if I analyze only U.S. manufacturing firms (as identified in the 2002 Census of Manufacturers) as in Figure B2, or if I use firm-HS6 product as the unit of analysis, as shown in Figure B3. Figure B4 presents the individual years that make up the switching data in Figure 1. Ultimately, it is clear that the stylized facts described above are consistent across a variety of different specifications.

I also check the estimation of Equation (1) using data from 2005-2006. As can be seen in Tables B1 and B2, there is still a strong correlation between high prices and the decision to switch, with those importers paying the highest price close to 4% more likely to switch than the omitted decile of prices.

Appendix C Proof of Proposition 1

Assumption 1 (Conditional Indepdendence) *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:*

$$P\left(x^C|s,\beta\right)=\frac{\exp\left[\overline{\pi}\left(x^C,s,\beta\right)+\delta EV\left(x^C,s\right)\right]}{\sum_{\widehat{x}\in X}\exp\left[\overline{\pi}\left(\widehat{x},s,\beta\right)+\delta EV\left(\widehat{x},s\right)\right]} \tag{21}$$

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

$$EV\left(x,s\right)=\int_{s'=0}^{\infty}\log\left\{\sum_{x'\in X}\exp\left[\overline{\pi}\left(x',s',\beta\right)+\delta EV\left(x',s'\right)\right]\right\}f\left(s'|s,x\right) \tag{22}$$

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

$$\begin{aligned}&S\left(\left[\overline{\pi}\left(s,\beta\right)+\delta EV\left(s\right)\right]\right)\\&\equiv\int_{\epsilon}\max_{x}\left[\overline{\pi}\left(x,s,\beta\right)+\delta EV\left(x,s\right)\right]g\left(\epsilon\right)\end{aligned} \tag{23}$$

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 $\overline{\pi}\left(x,s,\beta\right)$. Furthermore, the function $EV\left(x,s\right)$ 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 distibution 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 distibution 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 (20) as:

$$\int_{\epsilon}\max_{x}\left[\overline{\pi}\left(x,s,\beta\right)+\delta EV\left(x,s\right)+\epsilon\left(x\right)\right]g\left(\epsilon\right)=\mathbb{E}_{\epsilon}\left\{\max_{x}\mu_{x}+\epsilon\left(x\right)\right\} \tag{24}$$

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

For any n indepdent 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 dis-

tubtion 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 $\{\epsilon_i\}_{i=1}^n$ distibuted 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 (21) is therefore distibuted multivariate extreme value with location parameter $-\gamma + \log \sum_{x \in X} \exp\left(\mu_x\right)$. Since the expectation of any random variable distibuted multivariate extreme value with location parameter μ is $\mu + \gamma$, the social surplus function from (24) 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 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 (10) 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.

Figures D1 and D2 present a separate summary measure: rather than summarizing 1000 outcomes for each firm, I can alternatively create the price index P across all firms and industries for each Monte Carlo run, and compare them. By either taking the weighted average of the price across firms in an industry (Figure D1), or the median price across firms in an industry (Figure D2), I can generate density plots. Again, as the above results also show, the model generally tends to underpredict the price index.

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 swiching 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 E Potential for Serial Correlation

In Section 3, I model the importer's decision to choose a particular supplier as a dynamic discrete choice model with switching costs, where one-period profits- including an i.i.d. error term- take the following form:

$$\overline{\pi}_t^m\left(x_t^m, \boldsymbol{\beta}\right) + \epsilon_{x,t}^m = \xi \ln \lambda_{x,t} + \beta_p \mathbb{E}\left[\ln p_{x,t}^m\right] - \beta_x \mathbb{1}\{x_t^m \neq x_{t-1}^m\} - \beta_c \mathbb{1}\{c_t^m \neq c_{t-1}^m\} + \epsilon_{x,t}^m \tag{25}$$

One problem with the above equation is it excludes the possibility of serial correlation in the error term. For example, if an importer chooses exporter x two periods in a row, the model would interpret that as evidence for state dependence, when it could be the case that importer has some characteristic- constant over time- that makes them prefer exporter x in both periods. If the results are being driven by this heterogeneity, then the switching cost estimates might be overestimated.

In order to account for such bias, it is possible to allow an importer-exporter specific term to enter into the profit equation:

$$\overline{\pi}_t^m\left(x_t^m, \boldsymbol{\beta}\right) + \epsilon_{x,t}^m = \xi \ln \lambda_{x,t} + \beta_p \mathbb{E}\left[\ln p_{x,t}^m\right] - \beta_x \mathbb{1}\{x_t^m \neq x_{t-1}^m\} - \beta_c \mathbb{1}\{c_t^m \neq c_{t-1}^m\} + \alpha_x^m + \nu_{x,t}^m \tag{26}$$

A reduced form way to account for this issue is to include *lagged quality* $\lambda_{x,t-1}$ in the profit equation. The idea is since lagged quality does not affect current profits (something that is straightforward to test empirically), but is both exporter and importer specific, it could be included as the α term in the above equation.

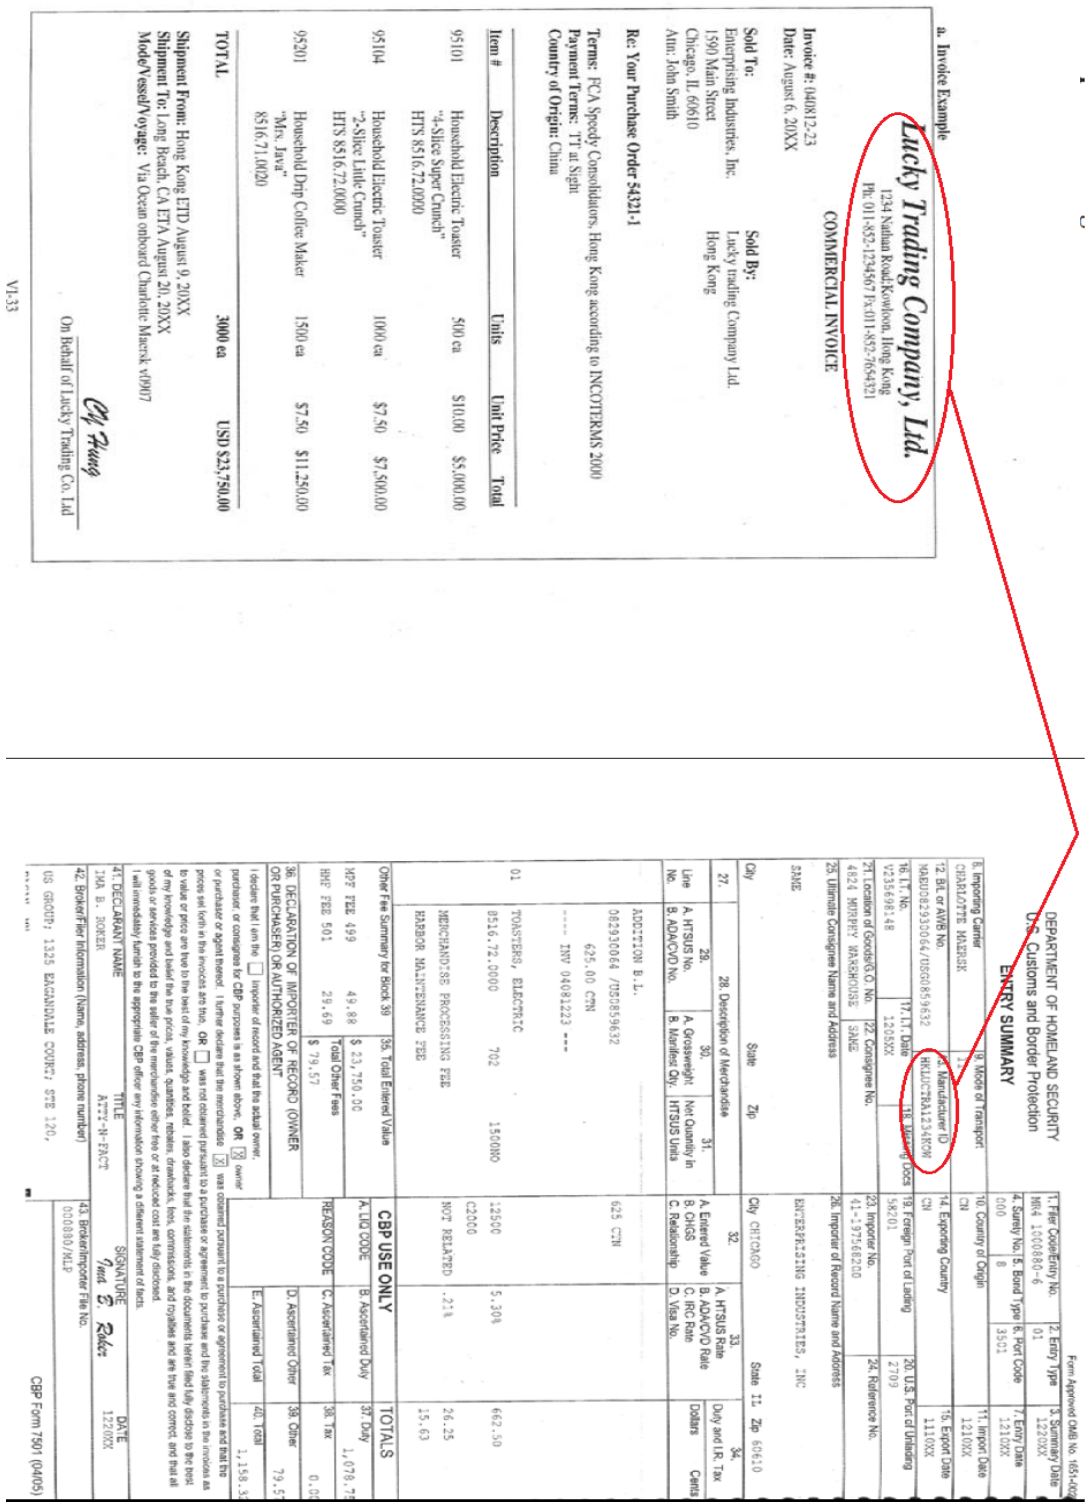
Alternatively, the α term in Equation (26) can be estimated as an additional set of parameters via the maximum likelihood process. Each importer would have a particular realization from the distribution of the exporter's α. With an average of 35 exporters per product, this would mean 35 new state variables to include in the dynamic programming problem (a 35-dimensional random effect). The brute force method would be to calculate the distibutions of these 35 separate state variables (for example α_{MEAN} and α_{SD} for each exporter $x = 1, ...35$) as an additional loop in the maximum likelihood problem, as the dynamic programming problem would be different for each vector of α realization. Ackerberg (2009) presents a method to simplify the problem through use of importance sampling to reduce the computation time for such a problem.

A simplifying solution would be to assume that the individual realization of α_x^m is not observed until the match is actually made- a "limited memory". This would increase the state space by only one variable, and

we would be using the data to estimate what those realizations must have been for the observed choices to have been made. In this case, the α term in Equation (26) would be importer-exporter-relationship-time specific. This would also account for the possiblity of serial correlation in the error term of Equation (25). Estimation results of the model under these extra conditions is will be made available shortly.

Appendix Tables and Figures

Figure A1: Sample Invoice



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

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

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.

Table A2: List of Industries Used in Counterfactuals

HS6 Code	Description	Trade Share
291560	Butyric Acid, Valeric Acid, Their Salts and Esters	0.31%
291631	Benzoic Acid, Its Salts and Esters	0.25%
293629	Other Vitamins and Their Derivatives (Unmixed)	0.86%
340120	Soap in other forms	0.66%
392020	Other Plates, Sheets, Film, Foil, Tape, Strip of Propylene Polymers (Non-cellular)	0.83%
481810	Toilet paper	1.02%
481960	Box files, letter trays, storage boxes and similar articles, used in offices, shops	0.55%
490300	Children's picture, drawing or coloring books	0.86%
520831	Plain Woven Fabrics, Cotton (Cotton 85% or More; Dyed; Not >100g/m2)	0.72%
560312	Nonwovens of man-made filament,>25g/m2	0.87%
570210	Kelem, Schumacks, Karamanie and Similar Hand-woven Rugs	0.65%
580639	Other Narrow Woven Fabrics of Other Textile Materials	0.62%
591190	Other Textile Products and Articles, for Technical Use	1.07%
610432	Women's or Girls' Jackets of Cotton, Knitted or Crocheted	0.66%
610791	Men's or Boys' Bathrobes, Dressing Gowns, of Cotton, Knitted or Crocheted	1.84%
620339	Men's or Boys' Jackets, Blazers, of Other Textile Materials	1.23%
621230	Corsets	0.59%
621490	Shawls, Scarves, Mufflers, Mantillas, Veils, of Other Textile Materials	0.15%
640219	Other Sports Footwear, Outer Soles and Uppers of Rubber or Plastics	7.93%
640340	Other Footwear, Incorporating Protective Metal Toe-cap	10.65%
650699	Headgear of Other Materials	0.23%
650700	Headbands, Linings, Covers, Hat Foundations, Hat Frames, for Headgear	0.13%
670411	Complete Wigs of Synthetic Textile Materials	1.65%
730722	Threaded elbows, bends and sleeves, of Stainless Steel	0.18%
730830	Doors, windows and their frames and thresholds for doors, of Iron or Steel	0.87%
731814	Self-tapping screws of Iron or Steel	3.11%
731930	Other pins of Iron or Steel	0.21%
820310	Files, rasps, and similar tools	0.11%
820890	Other (including parts) (Knives and Blades for machines and appliances)	X
830300	Armored/ reinforced safes, strong-boxes, safe deposit lockers, of base metal	4.44%
830990	Stoppers, Caps, Lids, Seals, Other Packing Accessories, of Base Metal	0.61%
841320	Hand Pumps for Liquids	0.18%
841360	Other Positive Rotary Displacement Pumps	0.42%
841370	Other Centrifugal Pumps	2.13%
841420	Hand or Foot Operated Air Pumps	0.46%
841850	Refrigerating, Freezing Chests, Cabinets, Display Counters, Show-cases & Similar	2.46%
848110	Pressure-reducing Valves	X
850650	Lithium primary cells and primary batteries	0.82%
850910	Vacuum Cleaners, With Self-contained Electric Motor	14.30%
850940	Food Grinders and Mixers; Fruit or Vegetable Juice Extractors	9.81%
853641	Relays, for a Voltage Not Exceeding 60v	2.01%
870893	Clutches and parts thereof	1.74%
871110	Motorcycles, Side-cars, Reciprocating Engine, cylinder capacity not >50 cc	3.33%
871120	Motorcycles, Side-cars, Reciprocating Engine, cylinder capacity >50 cc not 250 cc	7.10%
900580	Monoculars, Other Optical Telescopes; Other Astronomical Instruments	1.99%
902910	Revolution counters, production counters, taximeters, odometers, pedometers etc	0.60%
920590	Other wind musical instruments	0.60%
950631	Golf Clubs, Complete	4.60%
960321	Tooth Brushes	1.10%
960910	Pencils and Crayons, With Leads Encased in a Rigid Sheath	2.10%

These shares are the percent of import value compared to the total among these 50 industries. The number of importing firms in HS 848110 is too few to report importing information: the combined value share of HS 848110 and 820890 is 0.37%.

Figure B1: Year-to-Year "Staying" (New Definition) Percentages of U.S. Importers from China

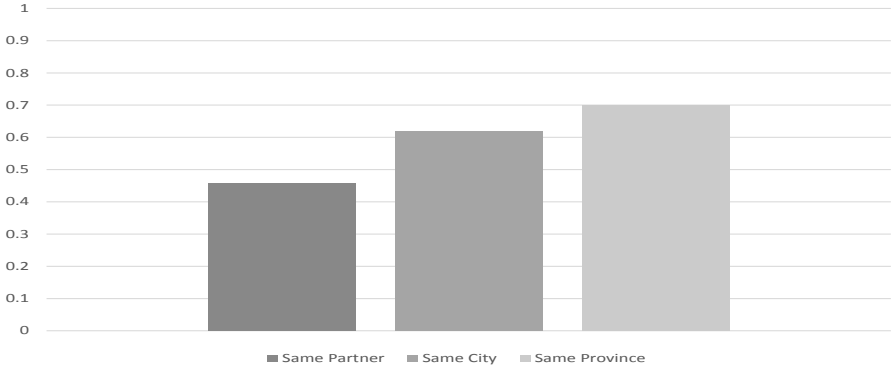


Figure B2: Year-to-Year "Staying" Percentages of U.S. Importers from China, Manufacturers Only

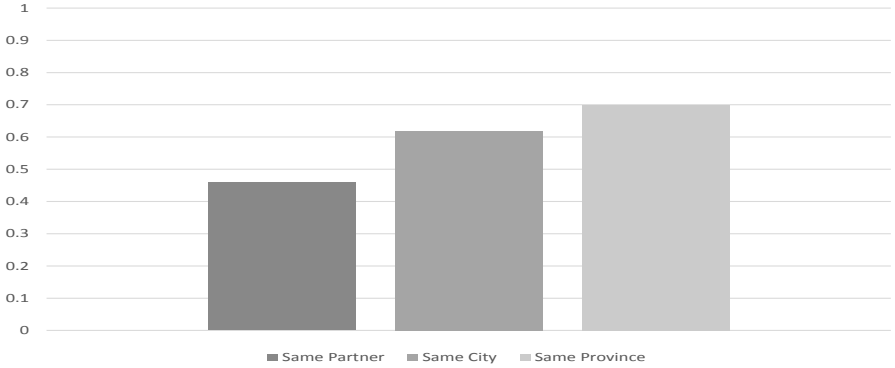


Figure B3: Year-to-Year "Staying" Percentages of U.S. Importers from China, Firm-HS6

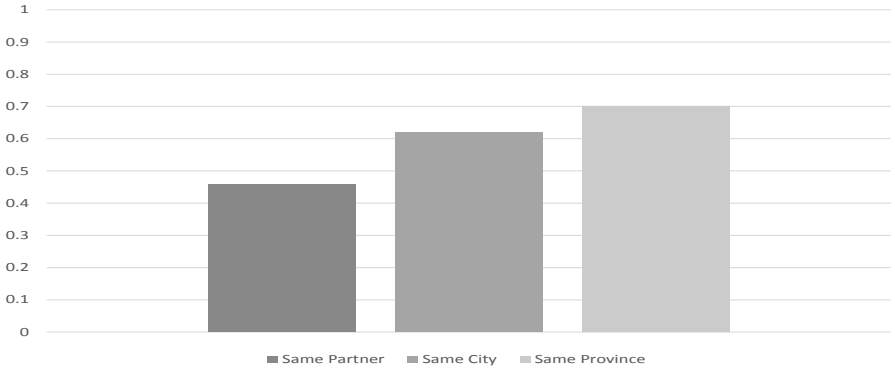


Figure B4: Year-to-Year "Staying" Percentages of U.S. Importers from China, Individual Years

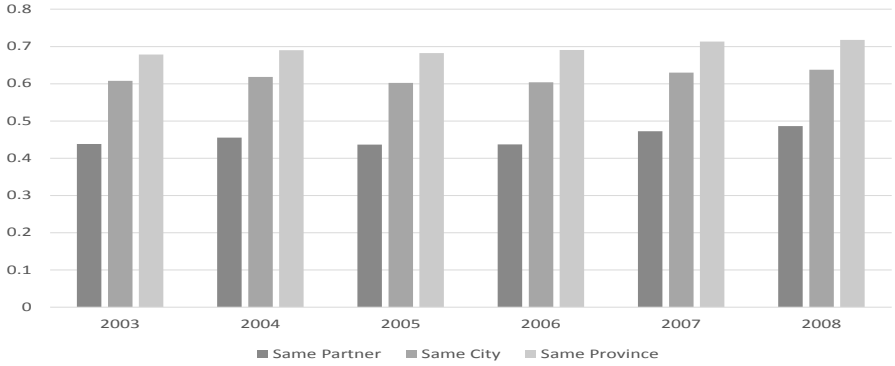


Table B1: Determinants of Supplier Stay/Switch Decision

Dependent Variable: Stayed with Chinese Exporter Year-to-Year, 2005-2006

	(1)	(2)	(3)	(4)
Log Price				
1st Decile	-0.0065	0.0176**	0.0143*	0.0137*
	(0.007)	(0.007)	(0.007)	(0.007)
2nd Decile	-0.0115	0.0041	0.0044	0.0050
	(0.008)	(0.008)	(0.008)	(0.008)
3rd Decile	-0.0024	0.0002	0.0003	0.0001
	(0.008)	(0.008)	(0.008)	(0.008)
4th Decile	-0.0018	-0.0007	-0.0004	-0.0007
	(0.007)	(0.007)	(0.007)	(0.007)
6th Decile	-0.0128*	-0.0131*	-0.0132*	-0.0132*
	(0.007)	(0.007)	(0.007)	(0.007)
7th Decile	-0.0097	-0.0088	-0.0086	-0.0088
	(0.008)	(0.008)	(0.008)	(0.008)
8th Decile	-0.0141*	-0.0126*	-0.0135*	-0.0136*
	(0.008)	(0.008)	(0.007)	(0.007)
9th Decile	-0.0218**	-0.0158**	-0.0170**	-0.0175**
	(0.007)	(0.007)	(0.007)	(0.007)
10th Decile	-0.0414***	-0.0256***	-0.0285***	-0.0293***
	(0.007)	(0.008)	(0.007)	(0.007)
Log Supplier Size		0.0360***	0.0671***	0.0671***
		(0.001)	(0.001)	(0.001)
Supplier Age		-0.0020***	-0.0028***	-0.0026***
		(0.000)	(0.000)	(0.000)
Importer Size			-0.0410***	-0.0396***
			(0.002)	(0.002)
Constant	0.4547***	0.0699***	0.1768***	0.1530***
	(0.005)	(0.012)	(0.013)	(0.014)
Entry Year FE	No	No	No	Yes
N	93,530	93,530	93,530	93,530
R^2	0.13	0.14	0.15	0.15

Notes: Robust standard errors clustered at the HS10 level in brackets. *** significant at the 1% level, ** significant at the 5% level. HS10 fixed effects are included. The sample is the universe of U.S. importers (HS10 product code and fim combination) from China who are found 2005-2006. 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 equal to 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 is the total estimated exports of a Chinese supplier in the HS10 product code in 2005, 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 2005. Importer size is the total size of imports in that HS10 product code in 2005 for any U.S. firm. Importer Entry Year is the first year a U.S. importers is found importing from China. Any importer that has the same share of imports from two separate Chinese suppliers is dropped.

Table B2: Determinants of Supplier Stay/Switch Decision

Dependent Variable: Stayed with Chinese Exporter Year-to-Year, 2005-2006

	(1)	(2)	(3)	(4)
Log Price	-0.0108***	-0.0116***	-0.0121***	-0.0121***
	(0.002)	(0.002)	(0.002)	(0.002)
Log Supplier Size		0.0290***	0.0518***	0.0514***
		(0.001)	(0.002)	(0.002)
Supplier Age		-0.0017***	-0.0024***	-0.0022***
		(0.000)	(0.000)	(0.000)
Importer Size			-0.0269***	-0.0251***
			(0.002)	(0.002)
Constant	0.4425***	0.1392***	0.1823***	0.1560***
	(0.000)	(0.011)	(0.011)	(0.013)
Entry Year FE	No	No	No	Yes
N	93,530	93,530	93,530	93,530
R^2	0.13	0.14	0.14	0.14

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 equal to 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 is the total estimated exports of a Chinese supplier in the HS10 product code in 2005, 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 2005. Importer size is the total size of imports in that HS10 product code in 2005 for any U.S. firm. Importer Entry Year is the first year a U.S. importers is found importing from China. Any importer that has the same share of imports from two separate Chinese suppliers is dropped.

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.

Figure D1: Price (Weighted Average) Kernel Density Plot

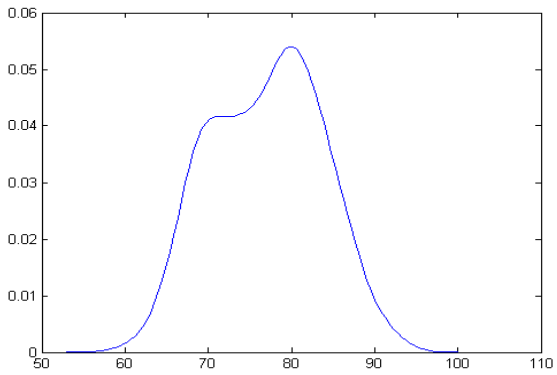
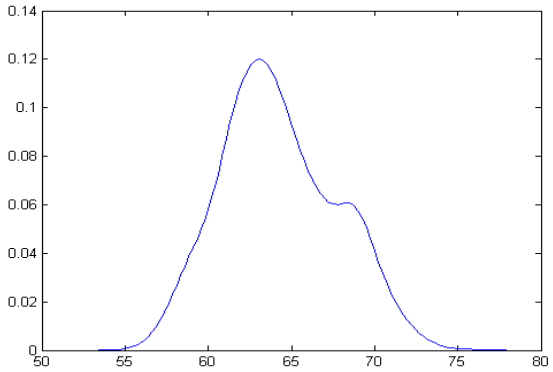


Figure D2: Price (Median) Kernel Density Plot



Notes: These figures are the analogue of the weighted average and median price kernel density plots described above.