

NBER WORKING PAPER SERIES

A SEARCH AND LEARNING MODEL OF EXPORT DYNAMICS

Jonathan Eaton
Marcela Eslava
David Jinkins
C. J. Krizan
James R. Tybout

Working Paper 29100
<http://www.nber.org/papers/w29100>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2021

We gratefully acknowledge support from the National Science Foundation (Grant SES-0922358), the United States Census Bureau, and Banco de la Republica de Colombia. We also thank Monica Hernandez, Gustavo Caballero, and Camilo Acosta for excellent research assistance, as well as Enrique Montes for expert data advice. Finally we are grateful to the many seminar and conference participants who have commented on earlier drafts. This paper was written in part by Census Bureau staff. It has undergone a more limited review than official Census Bureau publications. All results were reviewed to ensure confidentiality. Any views, findings and opinions in the paper reflect the views of the authors and do not reflect the views of the National Science Foundation, the U.S. Census Bureau, the U.S. Department of Labor, or the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Jonathan Eaton, Marcela Eslava, David Jinkins, C. J. Krizan, and James R. Tybout. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

A Search and Learning Model of Export Dynamics

Jonathan Eaton, Marcela Eslava, David Jinkins, C. J. Krizan, and James R. Tybout

NBER Working Paper No. 29100

July 2021

JEL No. F12,F14

ABSTRACT

Exporting abroad is much harder than selling at home, and overcoming hurdles to exporting takes time. Our goal is to identify specific barriers to exporting and to measure their importance. We develop a model of firm-level export dynamics that features costly customer search, network effects in finding buyers, and learning about product appeal. Fitting the model to customs records of U.S. imports of manufactures from Colombia we replicate patterns of exporter maturation. A potentially valuable intangible asset of a firm is its customer base and knowledge of a market. Our model delivers some striking estimates of what such assets are worth. Averaging across active exporters, the loss from total market amnesia (losing its current U.S. customer base along with its accumulated knowledge of product appeal) is US\$ 3.4 million, about 34 percent of the value of exporting overall. About half is the loss of future sales to existing customers while the rest is the cost of relearning its appeal in the market and reestablishing visibility as an exporter. As finding buyers takes time, the 5-year response of total export sales to an exchange rate shock exceeds the 1-year response by about 40 percent, with the 1-year response reflecting mostly sales per exporter-importer match and the 5-year response, reflecting the number of matches almost as much as sales per match.

Jonathan Eaton
Department of Economics
Pennsylvania State University
303 Kern Graduate Building
University Park, PA 16801
and NBER
jxe22@psu.edu

Marcela Eslava
Department of Economics
Universidad de los Andes
Calle 19A # 1-37 Este. Bloque W of. 810
Bogotá, Colombia
meslava@uniandes.edu.co

David Jinkins
Copenhagen Business School
david.jinkins@gmail.com

C. J. Krizan
Center for Economic Studies
U.S. Census Bureau
4600 Silver Hill Road
Washington, DC 20233
c.j.krizan@fastmail.com

James R. Tybout
Department of Economics
Penn State University
517 Kern Graduate Building
University Park, PA 16802
and NBER
jtybout@psu.edu

1 Introduction

Quantitative models of global economic activity rely on trade costs to explain why trade flows between countries constitute only a fraction of total production. But aggregate data provide little insight into what these costs are. A recent literature has turned to firm-level evidence to dig deeper into the nature of barriers to trade and how they influence the evolution of trade flows.

Our contribution here uses customs records of Colombian manufacturing exports to the U.S. market for the period 1992-2009 to shed light on the specific nature of trade costs. These data allow us to identify individual buyers and sellers and their interactions over time. We can observe, for example, Colombian firms entering into exporting in any given year, and see how their relationships with their U.S. buyers evolve. We can thus connect aggregate trade flows to the evolution of the trade patterns of individual cohorts of Colombian exporters. The data reveal, for example, how the dynamics of individual firm entry jibe with aggregate magnitudes.

The data reveal a rich set of patterns. Most exporters abandon the foreign market after a single sale, while a few remain to accumulate a broad network of buyers. Sellers' relationships with individual buyers are also highly heterogeneous, both in terms of their longevity and frequency and size of sales. Finally, the frequency and success rates of firms' new matches vary systematically with their characteristics and histories.

The model: To identify and quantify specific trade barriers, we develop a dynamic model of buyer-seller relationships consistent with these patterns. A basic, quite standard, feature of the model is that, to connect with a possible buyer, a firm needs to engage in costly search. The more it spends on searching, the sooner, on average, it connects with a potential client.

We model a firm's relationships with its buyers in continuous time. The expected value of a relationship with a particular client depends on the frequency and size of the buyer's purchases, and the seller's willingness to invest in maintaining the relationship. The relationship may end either exogenously or because the seller deems expected future sales to the customer not worth further investment. The model can thus explain why some relationships end precipitously while others expire after a stretch of dwindling sales. The reward to search depends on both the expected value of a successful relationship once established and the probability of establishing one in the first place. The customs records allow us to quantify both the value of an ongoing relationship and a firm's ability to find new ones.

A firm enters a foreign market knowing its own efficiency but imperfectly informed about its product's popularity there. When the firm encounters a potential client, the client may or may not want to buy from the firm, informing the seller about its product's popularity

in the market. A string of successes signals a high level of buyer enthusiasm, encouraging the firm to search more intensely for new buyers, while a series of rejections indicates lack of buyer interest, leading the firm to reduce its search effort or quit the market altogether.

We model this “learning” channel as follows: A firm’s popularity in a market is simply the probability that a potential buyer there demands its product. The firm starts out with a belief about this prior that’s distributed beta, the conjugate prior of the binomial distribution. As the firm accumulates successes and failures in its encounters with potential clients it updates its beliefs and the beta distribution converges to the binomial with the true probability.

To this learning channel we also add a “visibility” channel. As an exporter accumulates clients, its increased visibility may reduce the cost of finding added buyers.¹

Our model thus incorporates three types of frictions, which we quantify individually, that drive the progress of an exporter in a foreign market: searching for buyers, learning about product appeal, and establishing visibility. The model replicates patterns in the customs data regarding, for example, the distribution of buyers across firms and transitions in the number of buyers for an individual firm.

Our model relates the heterogeneity in firms’ outcomes to their known fundamental efficiency, as in standard models, but also to the appeal of their product in the export market, which they learn only over time. Luck also plays a role in a firm’s success: a few successful matches early on establish visibility and encourage the firm to search more intensely.

Quantifying the model, we find that standard search frictions constitute the greatest cost to market penetration, with learning and establishing visibility playing only a modest role. But learning and visibility are important to understanding the behavior of new exporters. Their initial contribution to aggregate exports is small, but new exporters play a fundamental role in driving the overall dynamics of aggregate exports.

The various trade frictions in our model have implications for how aggregate exports respond to exchange rate shocks. The response to a permanent change in the exchange rate, for example, is slower than it is in models in the literature without learning.

Relationship to the literature: The body of work addressing firm dynamics in open economies is now large. Alessandria et al. (2020) provide a recent review. To suggest how our paper fits in, we classify existing work into partial and general equilibrium analysis.

The partial equilibrium literature treats individual firms’ as single agents in isolation, with wages, exchange rates, and other market-wide variables taken as exogenous. The payoff is

¹Additional clients may be also be harder to reach, so that having more existing customers means a higher cost of adding new ones. Our model allows for either possibility, but since we find that firms with more clients seem to have an easier time adding new ones, we use the term “visibility.”

greater flexibility to explore a broad set of features of exporting firms’ behavior. Particular features that the literature has addressed include investments in foreign customer accumulation (Ruhl and Willis, 2017, Fitzgerald et al., 2019, Rodrigue and Tan, 2019, and Pivetau, forthcoming), learning (Albornoz et al., 2012; Schmeiser, 2012; Nguyen, 2012; Aeberhardt et al., 2014; Timoshenko, 2015; Cebreros, 2016; Araujo et al., 2016; Ruhl and Willis, 2017; Arkolakis, et al., 2018; Li, 2018; Berman et al., 2019), and shipment size and timing (Kropf and Saure, 2014; Hornok and Koren, 2015; Bekes et al, 2017).

We follow in this single-agent tradition, adding to the literature in three respects. First, we combine the features enumerated above into an integrated framework in order to quantify their individual contributions. Second, we use our model to impute the the contribution of a foreign market to firm value, isolating the role of firms’ accumulated knowledge and market visibility. Finally, by incorporating learning by new exporters, we quantify the role of luck (early successes or failures) in establishing new exporters in a foreign market.

In the market equilibrium literature, aggregates respond to the collective behavior of individual firms. Examples include Alessandria and Choi (2007, 2014, 2019), Ruhl (2008), Atkeson and Burstein (2010), Drozd and Nozal (2012), Burstein and Melitz (2013), Alessandria et al. (2014), Impullitti et al. (2013), Arkolakis (2015), Eaton e al. (2016), Handley and Limao (2017), and Fajgelbaum (2020).² Endogenizing market-wide variables in this literature comes at the cost of sacrificing some of the richness in single-agent models.

While our analysis is partial equilibrium, we connect to the market-equilibrium literature by simulating transition paths for aggregate exports under alternative assumptions about the frictions exporters face. We can thus quantify the role of search, learning, and visibility in shaping short and long-run responses to real exchange rate shocks, distinguishing the extensive margin of new buyer-seller encounters, some of which evolve into successful matches, and the intensive margin among existing matches.³

2 Firm-Level Trade: Transaction-Level Evidence

Previous work has established a robust set of facts about the export behavior of firms.⁴ We design our model to capture these facts as revealed in our data on Colombian shipments to the United States. Before describing these patterns we first describe the data themselves.

²Blum et al. (2019) analyze shipment frequency and size in a market equilibrium model, but do not relate them to export dynamics.

³Piveteau (forthcoming) performs a related exercise.

⁴Early contributions include Brooks (2006), Besedes (2008), and Eaton et al. (2008). Bernard et al. (2017), Bernard and Moxnes (2018), and Alessandria et al. (2020) reference and discuss many of the more recent studies.

2.1 Data

We base our analysis on comprehensive data from the U.S. Census Bureau’s Longitudinal Foreign Trade Transactions Database (LFTTD), which covers all commercial shipments into and out of the United States, extracting shipments from Colombia during 1992-2009. Each transaction record includes a date, the US dollar value of the product shipped, a 6-digit harmonized system product code, a quantity index, and, critically, an ID for both seller and buyer.

These IDs allow us to identify the formation and dissolution of business relationships (“matches”) between an individual buyer in the U.S. and seller in Colombia. To identify the U.S. importer we use the buyer’s Employment Identification Number (EIN).⁵ To identify the Colombian exporter we used the manufacturer’s identification code.⁶

We limit ourselves to transactions between non-affiliated trade partners and consider only imports of manufactures.⁷ Our final data set, spanning the years 1992-2009, contains 26,625 unique Colombian exporters, 12,921 unique U.S. importers, and 42,767 unique trading pairs. Value data have been deflated to 1992 prices using the U.S. CPI.⁸

In addition to U.S. customs records, we use establishment-level survey data from Colombia’s national statistics agency, Departamento Administrativo Nacional de Estadística (DANE). These data provide annual information on the total sales, exports, and other characteristics of all Colombian manufacturing plants with at least 10 workers. We use these data to characterize the size distribution of Colombian plants, the fraction of Colombian plants that export, and, among exporting plants, the relationship between exports and domestic sales.⁹

⁵There are two ways to track U.S. importers in the LFTTD: EINs and the firm identifiers in the Longitudinal Business Database (“alphas”). Though an EIN does not necessarily identify a complete firm, it is unique to a firm, and there is an EIN associated with every import transaction. An alpha maps to an entire firm, but the match rate between trade transactions and alphas is only about 80 percent (Bernard, Jensen, and Schott, 2009). We use EIN’s to maximize coverage.

⁶This variable is based on Block 13 of CBP form 7501, the import declaration form. Customs brokers are required to input the data. This field is an amalgamation of the manufacturer’s country, company name, street address, and city. Anecdotal information from customs brokers indicates that commonly used software constructs the code automatically from the name and address information entered in other fields. So this variable is sensitive to differences in how exporters’ names and addresses are recorded as they pass through customs, and shipments from the same exporter can appear to originate from distinct Colombian firms. To gauge the importance of this problem, we’ve conducted various checks on the matches based on this variable. Appendix B explains these checks.

⁷We thus exclude oil and coffee, which constitute the bulk of Colombian exports to the U.S. The National Federation of Coffee Growers centralizes coffee exports. A few players also dominate oil exports.

⁸Because of disclosure restrictions, as well as our exclusion of non-manufactures and trade between affiliated parties, we cover only a fraction of the total value of Colombian exports to the U.S. Table 18 in Appendix B compares patterns in our sample to patterns in aggregates from both U.S. and Colombian official sources.

⁹Since these data have been used widely in other studies, we don’t provide further description here.

Table 1: Average aggregates by cohort age

Cohort age	Exporters	Total Exports	Average Exports
1 year	1	1	1
2 years	0.29	1.11	3.77
3 years	0.18	0.93	5.03
4 years	0.14	0.67	4.66
5 years	0.12	0.63	5.18
6 years	0.10	0.51	4.99
7 years	0.08	0.50	5.72
8 years	0.08	0.45	5.91
9 years	0.07	0.39	5.58
10 years	0.06	0.40	6.58

Notes: Based on LFTTD customs records, U.S. imports of manufactured goods from Colombia, 1992-2009. Figures for cohorts aged 2-10 are relative to the corresponding figure for one-year-old cohorts.

We now turn to some key patterns in the data that we seek to capture in our modeling and estimation.

2.2 Cohort maturation

Following Brooks (2006), Table 1 reports average patterns of maturation across cohorts of Colombian exporters of manufactured goods to the United States. Since maturation patterns vary little across individual cohorts, we've averaged across the seven cohorts entering each year between 1993 and 1999. We base the figures in the table on U.S. customs records from 1992 through 2009.

The second row of the Table implies that, on average, only 29 percent of the firms that entered initially (year one) continue exporting the next year (column 1), yet these survivors generated 11 percent more export revenue in year two than the entire cohort did in year one (column 2), because sales per survivor were 3.77 times as large in year two as sales per cohort member in year one (column 3). Subsequent rows apply to subsequent years of exporting by members of that cohort, all relative to the cohort's entry year.¹⁰

Column 1 of Table 1 shows the rate of decline in cohort membership is especially high between the first and second year, with more than 70 percent of firms dropping out. But conditional on making it to the second year, the survival probability is much higher, with an attrition rate around 40 percent the second year, with subsequent attrition lower still. Firms that were exporting to the United States in 1992 account for fewer than five percent

¹⁰Appendix tables A.1-A.3 provide a breakdown of the numbers appearing in Table 1. Similar tables for Colombian exports of all goods and to all destinations appear in Eaton, et al. (2008).

of the firms exporting to the United States towards the end of the sample.

Column 2 shows that the rapid initial decline in cohort exporters is not accompanied by a similar collapse in total cohort sales. The relative stability of total sales means that sales per firm are growing substantially.

2.3 Patterns of buyer-seller matches

We now characterize buyer-seller matches during 1992-2009.

2.3.1 Monogamous and polygamous matches

The number of Colombian exporters in our sample grew at roughly 2 percent per year, from 2,232 in 1992 to 3,300 in 2009, while the number of U.S. importing firms grew by 3 percent per year, from 1,190 to 2,079 (Appendix A, Table 17). The number of Colombian exporter-U.S. importer pairs (representing at least one transaction between them in a year) also grew at an annual rate of 2 percent. Roughly 80 percent of matches are monogamous in the sense that the buyer deals with only one Colombian exporter and the exporter ships to only one buyer in the United States in a given year. However, since the remainder of the matches are polygamous, the average Colombian exporter sold to around 1.3 U.S. firms per year while the average U.S. buyer bought from around 2.3 Colombian firms per year. Both figures declined slightly over the period.

2.3.2 Transition probabilities

Like sellers' exporting stints (Table 1), most buyer-seller matches are short-lived. Even among those matches involving more than one shipment, the overall year-to-year death rate is roughly 40 percent, as we show later. There is a great deal of flux in an exporter's portfolio of clients.

Table 2 reports the probability with which a Colombian exporter (with the number of clients in the first column) transitions to the indicated number of clients (along the rest of the corresponding row) the following year. We classify a firm that stops exporting but re-appears as an exporter sometime later in our sample period as "dormant", in contrast with a firm that doesn't appear again in our sample, which we classify as "out". We treat the pool of potential entrants as firms that ever appear as exporters in our sample.

Among first-time exporters, roughly 93 percent sell to only one firm in their first year.¹¹ Of these, 62 percent don't export the next year, while only 6 percent go on to establish a larger number of relationships. For firms with 3 relationships in a year, 12 percent enter into

¹¹Many of these matches involve a single shipment. As we will show later, the overall match success rate (i.e., shipping to that buyer again) is roughly 41 percent.

a larger number the next year, but 67 percent lose clients. Firms starting with other client counts also, on average, move to a smaller number the following year. Hence, in addition to an enormous amount of churning among smaller exporters, we see a general tendency for firms to lose clients on net from one year to the next.

2.3.3 Ergodic degree distribution

What does this pattern of entry and growth imply about the ergodic distribution of relationships? The first row of Table 3 gives the ergodic distribution implied by the transition matrix in Table 2 under the assumption that the number of new entrants equals the number that exit.

The second row of Table 3 gives the corresponding actual distribution (averaged over 1992-2009). The ergodic and actual distributions are very close, suggesting that over our period the transition process has been quite stationary. Both distributions are very nearly Pareto, reflecting the coexistence of many small scale exporters with a few “super-exporters.”

2.3.4 Match maturation

A match is more likely to survive into the following year the larger sales in the initial year. Table 4 sorts matches into quartiles according to first-year sales, reporting year-to-year separation rates. In addition to the very low overall survival rates, two patterns stand out. First, the higher the quartile of initial sales, the lower the annual separation rate for the next four years. Second, survival probabilities rise year after year across the four quartiles.

Figure 1 shows average annual sales per match, broken down by initial sales quartile. The table further distinguishes matches according to their total life span: less than one year (life=0), 1 to 2 years (life=1), and so forth. For each cluster of bars, the left-most bar corresponds to sales in the initial year of the match, the next bar corresponds to sales during the second, and so forth.

The first message is that initial sales are a good predictor of sales in subsequent years, conditioning on survival. Annual sales in later years rise monotonically with sales in the first year across quartiles. (Note the different scales of the vertical axes in different panels of Figure 1.) Second, sales tend to jump from the first to the second year, in large part because observations on a match’s first year correspond to less than a full calendar year. (An analogous effect is at work in the final year of a match’s life.) Looking at complete-year observations reveals a tendency for annual sales to grow among matches that start small and survive, but no such tendency among matches that start in the largest quartile. Finally, looking across matches with different life spans, those that survive more years tend to have higher sales in all (full) years than matches that fail relatively quickly. This pattern is robust

Table 2: Transition Probabilities, Number of Clients

t \ t+1	Out	Dormant	1	2	3	4	5	6-10	11+
Out	.	.	0.932	0.055	0.009	0.002	0.001	0.001	0.000
Dormant	.	.	0.876	0.100	0.015	0.008	.	.	0.000
1	0.539	0.080	0.321	0.048	0.010	0.002	.	0.001	.
2	0.194	0.077	0.375	0.241	.	0.024	0.009	0.004	.
3	0.090	0.042	0.220	0.271	0.210	0.092	.	0.027	.
4	0.059	.	0.129	0.216	0.215	0.184	0.083	0.095	.
5	.	.	0.095	0.184	0.181	0.181	0.126	0.178	.
6-10	.	.	0.039	0.073	0.089	0.123	0.157	0.419	0.073
11+	.	0.000	0.000	0.000	.	.	.	0.432	0.526

Notes: Based on LFTTD customs records, U.S. imports of manufactured goods from Colombia, 1992-2009. Figures are cross-year averages of annual transition rates. Confidentiality restrictions prevent us from reporting numbers for cells that are too sparsely populated.

Table 3: Ergodic Client Distribution Implied by Transitions

	1	2	3	4	5	6-10	11+
Erg Distribution	0.792	0.112	0.031	0.016	0.009	0.022	0.016
Data	0.778	0.116	0.043	0.021	0.011	.	.


Notes: Based on transition probabilities reported in Table 2

Table 4: Separation Rates, by Age of Match and Initial Sales

	1 year	2 years	3 years	4 years	5+ years
Quartile 1	82.9	63.2	57.3	55.0	49.7
Quartile 2	75.6	58.4	49.4	46.8	43.7
Quartile 3	67.7	52.1	44.6	40.8	37.6
Quartile 4	52.1	44.5	40.3	39.2	36.7

Notes: Based on LFTTD customs records, U.S. imports of manufactured goods from Colombia, 1992-2009.

Figure 1: Average annual sales per match, by initial size quartile



Notes: Based on LFTTD customs records for manufactured goods imported from Colombia, 1992-2009.

across matches in the different quartiles of initial sales.

3 A Model of Exporting at the Transactions Level

Reflecting the data discussed in the previous section, our primary focus is understanding the dynamics of buyer-seller relationships between exporters from one country (in our case Colombia) and importers in a single foreign market (in our case the United States). Hence the model developed in sections 3.1, 3.2, and 3.3 applies to firms from a single source selling in a single foreign destination.

We show in section 3.4 how to extend the model to accommodate multiple foreign destinations. We don't have the data to pursue this extension here, but we are interested in connecting Colombian firms' activity in the U.S. market and at home. Hence in Section 4.1 we modify the the model developed in sections 3.1, 3.2, and 3.3 to apply to the home

market, giving us a two-market model with which to explore Colombian firms' activities in both markets.

With a single source and destination, our model provides a means of dissecting the dynamics of aggregate bilateral exports into, first, the sales of individual exporters and, then, into exporters' sales to individual clients.

Our model accounts for the irregular intervals at which buyer-seller relationships form and disband, and the fact that export shipments vary in size and are discrete events distributed unevenly through time. Formulating our model in continuous time facilitates capturing these features and more easily allows agents to update their behavior when their circumstances change. We treat exogenous random variables as Markov jump processes.

In presenting the model we first consider the relationship between a seller and an individual buyer. Having derived the seller's return from a relationship with an individual buyer, we turn to its learning about the popularity of its product in that market, i.e., the chance that a potential buyer there likes its product. Finally, we characterize its search for buyers.

3.1 A seller-buyer relationship

A relationship is a sequence of shipments from a seller to a buyer. We start with the seller's profit from an individual shipment, and then show how the dynamics of these shipments determine the overall value of the relationship.

3.1.1 Profit from a single shipment

Several features of our model are standard. At any time t seller j can hire workers at a wage w_t in real local currency units, each of whom can produce φ_j units of output, where φ_j is time-invariant and known by the seller. Hence seller j 's unit cost in local currency is w_t/φ_j . Selling at price p_{jt} in foreign currency unit profit in local currency is

$$p_{jt}/e_t - w_t/\varphi_j, \tag{1}$$

where e_t is the exchange rate.

Goods markets are monopolistically competitive with each producer supplying a unique product. Once buyer i has matched with seller j , the buyer periodically buys from j . Each shipment generates revenue:

$$X_{ijt} = \left(\frac{p_{jt}}{P_t} \right)^{1-\eta} y_{ijt} \bar{X}_t, \tag{2}$$

where $\eta > 1$ is buyers' elasticity of demand, p_{jt} is the price of seller j 's product, \bar{X}_t is the average spending level among all potential foreign buyers, P_t is the relevant price index for all

competing products in the foreign market, and y_{ijt} is a time-varying component of demand idiosyncratic to the ij relationship.¹²

We assume that the seller posts a non-negotiable price, charging the optimal markup over unit cost:¹³

$$p_{jt} = \frac{\eta}{\eta - 1} \frac{e_t w_t}{\varphi_j} \quad (3)$$

From (1), (2), and (3), the profit for seller j generated by an order from buyer i at time t is:

$$\pi_{ijt} = \frac{1}{\eta} \frac{\bar{X}_t}{e_t} \left(\frac{e_t w_t \eta / (\eta - 1)}{\varphi_j P_t} \right)^{1-\eta} y_{ijt}.$$

We can combine all the macroeconomic variables affecting the profit of any seller from this source selling in this destination, along with constants, as:

$$x_t = \frac{1}{\eta} \frac{\bar{X}_t}{e_t} \left(\frac{e_t w_t \eta / (\eta - 1)}{P_t} \right)^{1-\eta},$$

where x_t is common across all potential buyers in the foreign market. We can then write (dropping subscripts) the profit from a shipment as:

$$\pi_\varphi(x, y) = x \varphi^{\eta-1} y \quad (4)$$

Equation (4) is all we take from our specification of preferences and pricing behavior into the dynamics that follow. Any set of assumptions that deliver this simple multiplicative expression for a firm's profit from a shipment would serve us equally well. To facilitate our numerical implementation below we treat the set of possible values of x and y as discrete.

3.1.2 Relationship dynamics

A match can dissolve for two reasons. First, it can simply end exogenously with a constant hazard δ (due, say, to the demise of the buyer or the buyer's finding a more suitable or cheaper substitute). Second, after each sale to a particular buyer, the seller evaluates whether it's worth continuing the relationship. Doing so keeps the possibility of future sales to that buyer

¹²Since not all buyers necessarily face the same range of goods and hence the same aggregate price index P , we can treat i -specific components of the price index as P as embodied in y_{ijt} .

¹³Alternative specifications include bilateral bargaining between buyer and seller, as in Eaton et al. (2016), and pricing rules that recognize a link between current sales volume and future growth in customer base, as in Fitzgerald et al. (2019) and Piveteau (forthcoming). To keep our model tractable, and in view of Fitzgerald et al.'s (2019) finding that exporters' prices don't covary with market tenure, we opt for constant mark-up pricing.

alive, but requires paying a fixed cost F .¹⁴

When deciding whether to maintain a match, the seller knows its own efficiency φ , the macro state x , and profit from the current sale, $\pi_\varphi(x, y)$ to the buyer in question. It can thus infer this buyer's current y and calculate the value of the relationship as:

$$\tilde{\pi}_\varphi(x, y) = \pi_\varphi(x, y) + \max \{ \hat{\pi}_\varphi(x, y) - F, 0 \}$$

where $\hat{\pi}_\varphi(x, y)$ is the expected value of continuing a relationship that's currently in state (x, y) . The seller terminates this relationship if $\hat{\pi}_\varphi(x, y) < F$.

If the seller pays F to keep a match active one of several events will next affect it: with hazard δ the relationship is exogenously dissolved; with hazard λ^b , the buyer will place another order; with hazard $q_{xx'}^X$, x will jump to some new marketwide state $x' \neq x$; or, with hazard $q_{yy'}^Y$, y will jump to some new buyer-specific shock $y' \neq y$.¹⁵

Let τ_r be the random time that elapses until one of these (relationship-specific) events occurs. Given that x and y are independent Markov jump processes, τ_r is distributed exponentially with parameter $\lambda^b + \lambda_x^X + \lambda_y^Y$, where

$$\lambda_x^X = \sum_{x' \neq x} q_{xx'}^X \quad (5)$$

and

$$\lambda_y^Y = \sum_{y' \neq y} q_{yy'}^Y, \quad (6)$$

are the hazards of transiting from x to any $x' \neq x$, and from y to any $y' \neq y$, respectively. Then, assuming the seller has a discount factor ρ , the continuation value $\hat{\pi}_\varphi(x, y)$ solves the Bellman equation:

$$\begin{aligned} \hat{\pi}_\varphi(x, y) &= \mathbf{E}_{\tau_r} \left[e^{-(\rho+\delta)\tau_r} \frac{1}{\lambda^b + \lambda_x^X + \lambda_y^Y} \left(\sum_{x' \neq x} q_{xx'}^X \hat{\pi}_\varphi(x', y) + \sum_{y' \neq y} q_{yy'}^Y \hat{\pi}_\varphi(x, y') + \lambda^b \tilde{\pi}_\varphi(x, y) \right) \right] \\ &= \frac{1}{\rho + \delta + \lambda^b + \lambda_x^X + \lambda_y^Y} \left(\sum_{x' \neq x} q_{xx'}^X \hat{\pi}_\varphi(x', y) + \sum_{y' \neq y} q_{yy'}^Y \hat{\pi}_\varphi(x, y') + \lambda^b \tilde{\pi}_\varphi(x, y) \right) \end{aligned}$$

¹⁴The fixed cost could reflect maintenance of the account, technical support, or client-specific product adjustments. Colombian producers of construction materials interviewed for a related project (Domínguez et al, 2013) mentioned that a foreign buyer may request costly adjustments to a product or require special packaging.

¹⁵Since sales in the data are discrete events rather than flows, we model the buyer's purchases accordingly. We think of the buyer not as making use of the products continually but in discrete spurts. For example, the buyer might be a producer of a product that it makes in batches. At the completion of each batch it buys inputs for the next batch.

Before meeting a new buyer, the seller expects that the buyer will be in state y^s with probability $\Pr(y^s)$. The expected pay-off to forming a new match for a type- φ seller in market state x is thus:¹⁶

$$\tilde{\pi}_\varphi(x) = \sum_s \Pr(y^s) \tilde{\pi}_\varphi(x, y).$$

The term $\tilde{\pi}_\varphi(x)$, which is identified by data on match-specific revenue streams, determines a seller's search intensity.

3.2 Learning about product appeal

A seller searches for buyers in the market anticipating that that some fraction $\theta \in [0, 1]$ of them will be willing to do business with it. Given market state x , an encounter with a willing buyer generates the expected profit stream worth $\tilde{\pi}_\varphi(x)$ just derived, while an encounter with an unwilling buyer generates nothing then or subsequently.

Each seller enters the market with an unknown θ drawn from the (common knowledge) beta distribution:

$$b(\theta|\alpha, \beta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \theta^{\alpha-1} (1 - \theta)^{\beta-1},$$

where $\Gamma(\phi) = \int_0^\infty z^{\phi-1} e^{-z} dz$ is the gamma function (needed to ensure that the distribution has the proper limits). Given its θ , the probability that a random sample of n potential buyers will yield a seller a interested customers is binomially distributed:

$$q[a|n, \theta] = \binom{n}{a} \theta^a (1 - \theta)^{n-a}.$$

Hence, after meeting n potential buyers, a of whom were interested in its product, a seller's posterior beliefs about its θ are distributed:

$$p(\theta|a, n) \propto q[a|n, \theta] \cdot b(\theta|\alpha, \beta),$$

where the factor of proportionality is the inverse of the integral of the right-hand side over the support of θ . A firm's expected success rate after a successes in n trials has the convenient closed-form representation:

$$\bar{\theta}_{a,n} = E[\theta|a, n] = \int_0^1 \theta p(\theta|a, n) d\theta = \frac{a + \alpha}{n + \alpha + \beta}. \quad (7)$$

¹⁶In our numerical analysis we take the probabilities $\Pr(y^m)$ to be the ergodic distribution of y implied by the transition hazards $q_{yy'}^Y$. We could assume that the distribution at the time of the first purchase is different from the ergodic one.

As the beta distribution is the conjugate prior of the binomial, this posterior mean converges to

$$plim \left(\frac{a}{n} \right) = \theta$$

as n gets large.

In our formulation a firm learns something about its demand in a market with each encounter with a new potential buyer, successful or otherwise. We thus depart from other models with learning in which there is only zero or one signal per period, depending upon the firm’s market participation (Timoshenko, 2015; Arkolakis et al., 2018; Fitzgerald et al., 2019). Our formulation creates an extra incentive for new entrants to search intensively, which we quantify in Section 6 below.

3.3 Searching for buyers

A seller continuously chooses a market-specific hazard s with which she encounters a potential buyer, incurring the instantaneous flow cost $c(s, a)$, which is increasing and convex in s .¹⁷ How $c(s, a)$ varies with the number of successful matches a depends on the relative strength of different forces. The cost might fall with a as successful matches increase the seller’s visibility with additional potential buyers. The cost might rise if the pool of easy-to-reach buyers becomes “fished out,” as in Arkolakis (2010). We leave it to the data to decide the direction and magnitude of the effect.

To derive the return to search, recall that when the foreign market state is x , a type- φ seller expects the value of a new successful match to be $\tilde{\pi}_\varphi(x)$, and the seller believes the next encounter will be successful with probability $\bar{\theta}_{a,n}$. Hence the expected value of an encounter is $\bar{\theta}_{a,n}\tilde{\pi}_\varphi(x)$

Let τ_s be the random time until the next search event, which could be either an encounter with a potential buyer or a change in the marketwide state x^f . Then the optimal search intensity s for a type- φ firm with foreign market search history (a, n) solves the the Bellman

¹⁷Interviews conducted with Colombian exporters revealed a variety of activities firms pursue to meet potential buyers in a foreign market (Domínguez, et al, 2013). Activities included maintaining a foreign sales office; paying the exports promotion office to organize visits with prospective clients, and sending their sales representatives to those visits; sending sales representatives abroad to visit potential clients on their own; attending trade fairs; paying a researcher to search the web for foreign firms that purchase products similar to their own; paying browsers to ensure that their site appear near the top of a search for their product type; maintaining a web site in English. Interviewees also reported that activities such as traveling to trade fairs or translating their websites to English led to relationships with one or two clients every few years. Establishing a larger network of clients required much more costly activities.

equation:

$$V_\varphi(a, n, x) = \max_s \mathbf{E}_{\tau_s} \left[-c(s, a) \int_0^{\tau_s} e^{-\rho t} dt + \frac{e^{-\rho \tau_s}}{s + \lambda_x^X} \cdot \left(\sum_{x' \neq x} q_{xx'}^X V_\varphi(a, n, x') \right) + s \left[\bar{\theta}_{a,n} (\tilde{\pi}_\varphi(x) + V_\varphi(a + 1, n + 1, x)) + (1 - \bar{\theta}_{a,n}) V_\varphi(a, n + 1, x) \right] \right]$$

(Recall that λ_x^X is given by (5).) Taking expectations over τ_s yields:

$$V_\varphi(a, n, x) = \max_s \frac{1}{\rho + s + \lambda_x^X} \left[-c(s, a) + \sum_{x' \neq x} q_{xx'}^X V_\varphi(a, n, x') + s \left\{ \bar{\theta}_{a,n} [\tilde{\pi}_\varphi(x) + V_\varphi(a + 1, n + 1, x)] + (1 - \bar{\theta}_{a,n}) V_\varphi(a, n + 1, x) \right\} \right] \quad (8)$$

Applying the multiplication rule for differentiation and using expression (8) for $V_\varphi(a, n, x)$, the optimal search intensity s^* satisfies:

$$\frac{\partial c(s^*, a)}{\partial s} = \bar{\theta}_{a,n} [\tilde{\pi}_\varphi(x) + V_\varphi(a + 1, n + 1, x)] + (1 - \bar{\theta}_{a,n}) V_\varphi(a, n + 1, x) - V_\varphi(a, n, x) \quad (9)$$

That is, the marginal cost of search equals the expected benefit of a match $\bar{\theta}_{a,n} \tilde{\pi}_\varphi(x)$ plus the expected value of the information and visibility it generates.

3.4 Entering multiple markets

So far we've focused on firm entry into a single foreign market. We can accommodate firms' activity across multiple markets, designating a particular market by m . We treat seller j from a particular source as having an efficiency φ_j that applies across markets. Prior to searching in a particular market m , a seller's prior about its product appeal there is $\theta_{j,0}^m$, which can vary across markets m , as can the seller's true product appeal θ_j^m . We assume that a seller needs to search independently in each market, and that all learning and visibility effects are market-specific.

If, as we assume in our application, that all sellers in the source country begin with the same prior $\theta_{j,0}^m = \theta_0^m$ in a particular market m , then the more efficient ones (with a higher φ_j) will initially search more intensely there. Less efficient firms will search less intensely or not at all. A feature of our specification is that uncertainty about product appeal might actually lead more firms to search in a market in hope of finding that their product is particularly appealing there.

Some of the firms actively searching will experience mostly failed encounters, lowering

their belief about θ_j^m and discouraging further search. Others, enjoying a string of successes, will revise their belief about θ_j^m upward, and look for new customers more intensely. The model thus delivers the prediction that firms' presence across different markets is positively correlated (due to φ_j), but imperfectly so (due to heterogeneity across markets in true product appeal θ_j^m and what a firm has learned about it in each market).¹⁸ How a firm fares in a market reflects its true product appeal θ_j^m only imperfectly. Some firms with strong product appeal may have had such low efficiency that they never bothered to search in the first place. Other unlucky ones may have been driven from the market, despite their high θ_j^m , by a series of unsuccessful encounters.

4 Specification for Estimation

To adapt our theoretical framework to the data at hand we make some specific assumptions about destination markets, search costs, and the stochastic processes that generate exogenous state variables.

4.1 Destinations

Our principal data are U.S. customs records reporting sales of manufactures by Colombian firms to unaffiliated U.S. buyers. Hence our source country is Colombia and our foreign destination is the United States. We also observe total sales by Colombian firms in Colombia itself. Hence we model relationships between Colombian firms with both their U.S. and domestic customers. In what follows we indicate magnitudes specific to the foreign (U.S.) market with $m = f$ and specific to the home (Colombian) market with $m = h$.

Colombian firms typically don't export to the United States until they've sold in the home market for several years. Since our model implies that learning effects would be exhausted by then, we treat firms as aware of their product appeal θ^h in the home market by the time they enter our window of reference. Since we treat θ^f and θ^h as uncorrelated, all Colombian firms enter the U.S. market equally naive about their product appeal there. The information that's relevant for their decision to embark on searching for buyers in the United States is their efficiency φ , which we treat as the same in each market. Hence Colombian firms that sell in the United States would tend to be larger in Colombia than those that don't, and firms selling in both markets that are larger in Colombia would tend to be larger in the

¹⁸Eaton et al. (2011)'s static model also treats firm efficiency as common across markets but demand shocks as market-specific. While our dynamic model implies positive correlation across destinations in the cross-section, since we treat φ_j as time invariant, our model doesn't deliver any ergodic correlation in sales across countries over time unless we introduce demand shocks y that are temporally correlated across destinations.

United States. As described below, we also assume that Colombian firms in the sample have exhausted network effects at home.

4.2 Matches and relationships

To connect our theory with data, we impose a particular interpretation of individual transactions. We treat a firm’s first sale to a new buyer as a “match,” which becomes a “successful match” or a “relationship” if and only if the firm sells to this buyer again. For each Colombian firm that ever sells in the U.S. market in our period of observation we can thus keep track of its number n of encounters with buyers, or “matches”, and the number $a \leq n$ that succeed.

4.3 Search costs

We generalize Arkolakis’s (2010) formulation of search costs to allow for network effects, specifying the cost of searching with intensity s^m in market m as:

$$c^m(s^m, a^m) = \kappa_0^m \frac{[(1 + s^m)]^{\kappa_1} - 1}{\kappa_1 [1 + \ln(1 + a^m)]^\gamma}. \quad (10)$$

Here a^m is the number of previous successful matches the seller has had in market m , κ_0^m is a market-specific cost parameter, while the parameters κ_1 and γ are common across markets.

Several properties of this function merit note. First, the parameter γ governs how the number of previous successes affects the current cost of search, with $\gamma > 0$ implying a benefit (say due to increased visibility) and $\gamma < 0$ consistent with a “fishing out” effect.¹⁹ Second, a seller who is not searching in a particular market incurs no search cost there: $c^m(0, a^m) = 0.$, while the marginal cost of increasing search at zero is strictly positive. Hence some firms will not search at all while others may search only minimally. Third, given the cumulative number of successful matches, a^m , the marginal cost of search increases with s at a rate determined by κ_1 . Finally, since a^m is the cumulative number of successes in market m , network effects endure, even after a particular match is severed or while a firm isn’t actively searching.

¹⁹To limit the dimensionality of our computational problem, we assume that firms with more than a^* buyers have both (i) exhausted their learning effects and (ii) reap no additional network effects from further matches. We choose a^* to exceed the observed maximum a for 99 percent of sellers in the U.S. market. Also, we set $a = a^*$ for all sellers in their home (Colombian) market.

4.4 Processes for exogenous state variables

Our exogenous state variables are φ , x^h , x^f , y^h and y^f . We normalize their logs to have zero mean, introducing scalars Π^m in the profit function (4) to accommodate level effects.

We assume that the distribution of efficiency φ across firms is log normal with variance σ_φ . We treat the Markov jump processes (x^h, y^h, x^f, y^f) as independent Ehrenfest diffusion processes, with the idiosyncratic match shocks y^f and y^h having the same distribution. We allow the overall market processes x^f and x^h to have different distributions, however, to accommodate, among other things, the different effects of exchange rate shocks across markets.

We specify an Ehrenfest process for z by discretizing its log into $2g + 1$ possible values a distance Δ apart, $g \in I^+ : z \in \{-g\Delta, -(g-1)\Delta, \dots, 0, \dots, (g-1)\Delta, g\Delta\}$. Given its current value z the process jumps to an adjacent value z' with hazard λ_z according to:

$$\ln z' = \begin{cases} \ln z + \Delta \\ \ln z - \Delta \\ \text{other} \end{cases} \text{ with probability } \begin{cases} \frac{1}{2} \left(1 - \frac{\ln z}{g\Delta}\right) \\ \frac{1}{2} \left(1 + \frac{\ln z}{g\Delta}\right) \\ 0 \end{cases} .$$

We allow for N^X possible values for the market state processes x^m , $m = f, h$, and N^Y possible values for the relationship-specific shocks y . Thus, given a grid size g , the intensity matrices $Q^X = \{q_{ij}^X\}_{i,j=1,N^X}$ and $Q^Y = \{q_{ij}^Y\}_{i,j=1,N^Y}$ implied by the transition hazards introduced in section 3.1 are each block-diagonal and characterized by a single parameter, Δ .

5 Estimation

Our estimation proceeds in two stages. We first estimate the processes for the aggregate market variables x^f and x^h using data on aggregate spending on manufactures in the United States and Colombia. We then use indirect inference to infer the remaining parameters.

5.1 Stage 1: estimating observable jump processes

We treat the annual manufacturing expenditure data for the two countries as discrete time observations of underlying independent Ornstein-Uhlenbeck processes of the form:

$$dz = -\mu z dt + \sigma dW$$

where $\mu = \lambda_z/g$, $\sigma = \sqrt{\lambda_z}\Delta$, and W follows a Weiner process.

Shimer (2005) shows that if z follows a continuous time Ehrenfest diffusion process, it asymptotes to an Ornstein-Uhlenbeck process with mean zero as the fineness of the grid

Table 5: Parameters of Market-wide Demand Processes

	<i>Parameter</i>	<i>value</i>
home macro state jump hazard	λ^{x_h}	1.200
foreign macro state jump hazard	λ^{x_f}	1.215
home macro state jump size	Δ^{x_h}	0.003
foreign macro state jump size	Δ^{x_f}	0.053

Notes: Our foreign market size measure is the OECD time series of US GDP in Industry, including imports and subtracting exports of manufactures. Our home market size measure is real Colombian expenditures on manufacturing goods, taken from DANE. We translated all of the data used for the estimation into real 1992 US dollars, deflating nominal US dollars with the consumer price index available on the US Bureau of Labor Statistic website. We used an official Colombian Peso-US Dollar exchange rate time series downloaded from the Central Bank of Colombia to translate Peso values to nominal US Dollar values.

increases.²⁰ We use our aggregate expenditure data to estimate, for $m = f, h$, the corresponding λ^{x_m} and σ^{x_m} . Given these estimates we infer Δ^{x_m} and λ^{x_m} .

Table 5 reports the estimates for λ^{x_m} and Δ^{x_m} . They imply that x^f and x^h both jump 1.2 times per year, on average. However, jumps in the U.S. market tend to be much larger, mainly because they reflect movements in the real exchange rate as well as movements in dollar-denominated expenditures.

5.2 Stage 2: indirect inference

Our data are relatively uninformative about the rate of time discount ρ and the demand elasticity η . Using values that are standard in the literature, we set $\rho = 0.05$ and $\eta = 5$. Also, to limit the size of the estimated parameter vector: (i) We set the exogenous match failure rate to be the observed match failure rate among matches at least 3 years old ($\delta = 0.326$); (ii) We specify the search cost function as quadratic in search intensity ($\kappa_1 = 2$); (iii) We assume that the hazard rate for the match-specific shock is once per quarter ($\lambda_y = 4$).²¹

All of the remaining parameters we estimate jointly using the transactions data summarized in Section 2.1 above. These parameters include the market size scalars (Π^h, Π^f), the fixed costs of maintaining a match (F^h, F^f), the parameters of the product appeal distributions (α, β), the dispersion of the productivity distribution (σ_φ), the jump size for the

²⁰Consider an Ehrenfest diffusion process with parameter vector (λ, g, Δ) . Rewriting the parameter vector as $(\lambda/\epsilon, g/\epsilon, \Delta\sqrt{\epsilon})$, $\epsilon > 0$, the autocorrelation parameter μ and the instantaneous variance parameter σ are invariant to ϵ . But as $\epsilon \rightarrow 0$, the innovation dW approaches normal.

²¹While we could in principle have estimated these last three parameters, we found them not well-identified in our data.

match-specific shocks (Δ_y), the hazard rate for shipments (λ_b), the network parameter (γ), and the market-specific cost function scaling parameters (κ_0^h, κ_0^f). We collect these parameters into the vector Λ :

$$\Lambda = \left\{ \Pi^h, \Pi^f, F^h, F^f, \alpha, \beta, \sigma_\varphi, \Delta_y, \lambda_b, \gamma, \kappa_0^h, \kappa_0^f \right\}$$

We estimate Λ using the method of indirect inference (Gouriéroux and Monfort, 1996). For each candidate Λ , we first use the model to simulate the foreign and domestic transactions of an artificial sample of producers. Then, using these simulated data, we estimate a set of reduced-form regressions that summarize the relationships we want our model to capture. Finally, searching the support of Λ , we choose the one that makes the regression coefficients from simulated data as close as possible to the corresponding regression coefficients from the sample data. Algebraically, our estimator is

$$\hat{\Lambda} = \min_{\Lambda} [\bar{m} - m(\Lambda)]' W [\bar{m} - m(\Lambda)],$$

where \bar{m} is a column vector of regression coefficients from sample data, $m(\Lambda)$ is the analogous vector of regression coefficients from data simulated at Λ , and W is a compatible non-singular weighting matrix. Setting $W^{-1} = \text{var}(\bar{m} - m(\Lambda))$ maximizes the efficiency of this estimator, but any non-singular W yields consistent estimates. We use a block-diagonal version of $\text{var}(\bar{m} - m(\Lambda))$, with each block corresponding to the moments from a particular regression.

Tables 6, 7, and 8 report the coefficients from the data-based regressions (with standard errors in parentheses below).²² We now describe the individual regressions, our reasoning in choosing them, and the parameters they help most to identify. We base our discussion about identification partly on our calculation of Andrews et al.’s (2017) sensitivity matrix, which we report in full in Appendix C.²³

Search. Equation (*i*) in Table 6 summarizes the effects of a firm’s market experience on its search intensity (*s*). Recall that our definition of a match allows us to infer both the number of matches *n* of a firm (a sale to a new buyer) and its number of successes $a \leq n$ (a sale to a new buyer followed by at least one subsequent sale to that buyer). The dependent variable is the inverse of the time between firm *j*’s n^{th} and $n + 1^{\text{st}}$ matches, regardless of

²²To facilitate comparison between the data-based and simulation-based regressions, where possible and with no loss of information, we’ve replaced the intercept of each regression with the mean value of the dependent variable. Several regressions use real peso values as reported by DANE. We’re not confident that they’re strictly comparable to the real dollar units used in U.S. customs records.

²³Andrews et al. (2017) propose using the sample analog of the matrix $(G'WG)^{-1}G'W$, where $G = -\partial[m(\Lambda)]/\partial\Lambda'$ is the Jacobian of the vector of simulated moments. “Intuitively, this matrix is a local approximation to the mapping from moments to estimated parameters.” (p. 1555) We report the results of this calculation at our benchmark estimates (discussed below) in elasticity form in Appendix C.

whether either is successful. We treat this observation as a proxy for firm j 's search intensity in a market. The right-hand side is a second-order translog function of firm j 's cumulative number of successes (a_{nj}) after n matches and cumulative success rate (a_{nj}/n_j) in the market. (To deal with firms that have had no successes, we add 1 to a_{nj} and to a_{nj}/n_j before taking logs.).²⁴

Since the coefficients in this equation reflect all of the elements of Λ , there's no obvious mapping between them. Our sensitivity matrix suggests that the regression helps most in identifying the fixed costs of maintaining a relationship, F^h and F^f , the parameters of the success rate (θ) distribution, α and β , and the network parameter, γ .

Because equation (i) involves second-order terms, it's difficult to interpret its coefficients. Evaluating its predicated values on a grid of success rates and cumulative successes shows that search intensity is only mildly sensitive to success rates, but strongly increases with cumulative successes.

Separation. Equation (ii) captures a second basic feature of a firm's exporting behavior: termination of matches. Here the unit of observation is seller j 's i^{th} match in year t and the dependent variable, $D^{exit\ match}$, takes a value of one when this match is in its final year.²⁵ In our model the seller endogenously terminates a match when $\hat{\pi}_{\varphi_j}(x_t, y_{ijt}) < F^f$, which is more likely when the firm's productivity φ_j or the demand shock y_{ijt} is low. Since we don't observe these variables we use several of their correlates as predictors: current match sales, X_{ijt}^f , age of the match, A_{ijt} , and export market tenure, Δ_{ijt} (all in logs). Reflecting the patterns in Table 4, we allow a firm in its first year of exporting ($D^{new\ to\ mkt} = 1$) to experience a different failure rate.²⁶

The results in Table 6 reflect patterns we saw in Table 4: matches in their first year are relatively likely to fail, as are matches that start with relatively small sales; more experienced exporters tend to have longer-lived relationships, reflecting cross-firm variation in productivity levels φ .

The sensitivity matrix (Appendix C) implies that equation (ii) helps most to identify the fixed costs of maintaining an established match, F^h and F^f , and the jump size, Δ_y , which affects the option value of keeping a match active.

Match success rates. The remaining regressions in Table 6 concern the distribution of success rates. Equation (iii) relates the average success rate of an active exporter to its cumulative number of previous meetings (n). Equation (iv) relates the dispersion in success

²⁴This equation is roughly a second-order approximation to the foreign market policy function (9), ignoring, for instance, the nonlinear firm effects generated by φ and θ^f .

²⁵We include only active matches, in which there is a sale in year t , in the sample.

²⁶Note, however, that Table 4 includes single-shipment matches while our estimation of equation (ii) doesn't drop them, since we don't consider the match "successful".

rates (the squared residuals from equation (iii)) to n . The unconditional mean and variance in success rates are 0.41 and 0.09, respectively. Among experienced (high- n) firms, the mean is significantly higher and the variance is significantly lower. Both regressions are informative about α and β , as well as selection due to learning.

Client distributions and shipment frequencies. Equation (v) of Table 7 relates to the information on client distributions in Table 3. With $\Phi(\ell)$ representing the fraction of exporters with no more than ℓ active clients, column (v) reports a regression of $\ln(1 - \Phi(\ell))$ on $\ln \ell$ and $(\ln \ell)^2$.²⁷ This functional form nests a linear relationship between $\ln(1 - \Phi(\ell))$ and $\ln \ell$ implied by a Pareto shape of client distributions found in earlier studies. The small coefficient on the quadratic term (-0.055) confirms a Pareto shape. Equation (v) helps identify the parameters of the theta distribution (α and β) as well as the network parameter γ , the key determinants of the relative prevalence of large versus small firms.

Equation (vi) in Table 7 establishes the mean log number of shipments per year in a continuing match. It serves as a target for the shipment arrival hazard and helps identify λ_b .

Match-level and firm-level sales. Table 8 collects regressions on the time series properties of firms' exports, cross-firm dispersion in exports, and patterns of correlation between exports and domestic sales. These equations are particularly informative about the parameters $\{\Pi^h, \Pi^f, F^h, F^f, \sigma_\varphi, \Delta_y\}$.

Equation (vii) is an AR1 in log match revenues, conditioned on match age and a dummy to control for first-year effects. Following the discussion in Section 5.1 above, the root (0.826) and root mean square error (1.208) help identify the jump size Δ_y and the cross-firm variance in productivity, σ_φ . Also, together with equation (ii), the mean log annual revenue per match (10.67) essentially pins down the profit function scalar Π_f and the fixed cost of maintaining a foreign match F^f .

The last four equations in Table 8 concern domestic sales. Since we don't observe firms' individual matches in the domestic market, these regressions describe establishment-level panel data merged with Colombian customs records.²⁸

Equation (viii) is an AR1 for home sales, informative about how much firms adjust their domestic connections and their associated match-specific sales in response to idiosyncratic shocks. The equation is particularly helpful in identifying κ_0^h and F^h , and the mean squared error helps identify σ_φ and $\alpha/(\alpha + \beta)$.

Equation (ix) projects Colombian firm-level exports to the United States on firm-level domestic sales. The coefficient speaks to the variance of productivity shocks (σ_φ^2), which are

²⁷By construction, the intercept of the (non-parametric version of) this regression must be zero.

²⁸Regressions (viii) through (x) in Table 8 use a combination of the Colombian Annual Manufacturing Survey (AMS) and Colombian administrative records of exports transactions involving the U.S. market over 1993-2007, merged using firm identifiers.

common to both markets, relative to the variance of market-specific appeal draws, θ^h and θ^f . It also reflects the size of match-specific (and thus market-specific) idiosyncratic shocks, Δ_y .²⁹

Finally, equations (x) and (xi) describe the relative importance of home versus foreign sales. Equation (x) reveals the share of firms that participate in the foreign market. It speaks to the relative return to maintaining foreign versus domestic business connections, as implied by the sets of parameters (Π^f, F^f, κ_0^f) versus (Π^h, F^h, κ_0^h) . Equation (xi) gives the average share of exports to the U.S. market in total sales of exporting firms. It largely reflects the number of clients in each market, and thus responds especially to differences between κ_0^f and κ_0^h .

5.3 Parameter estimates for the benchmark model

The first two columns of Table 9 report estimates of the parameter vector Λ for the benchmark model. We normalize moments about month-to-month customer transitions from the customs records to one year. Our estimate of δ , for example, implies that, on average, matches last roughly a third of a year before separating for exogenous reasons.

Our estimates of the fixed costs ($F^f = \$US 0.30$, $F^h = \$US 0.03$) are both very small. While their effect on major exporters is negligible, they keep out fringe players that would otherwise sell tiny amounts.

The profit and cost function scalars are more important. We estimate much lower search costs in the home market ($\kappa_0^h = 859$ versus $\kappa_0^f = 3,080$) and much higher profit per sale ($\Pi^h/\Pi^f = \exp(-3.88 + 6.14) = 9.77$). Both help explain the small share of output exported to the U.S. (Table 8, regression *xi*). The difference between the two sets of scalars is identified by their different effects on match arrival rates (Table 6, regression *i*) versus revenues from ongoing matches (Table 8, regressions *vii* and *viii*).

Other parameters are hard to interpret individually. More enlightening are their collective implications explored in Section 6 below.

5.4 Model fit

In general, our model replicates the patterns in the data, though not all of the model-based equation estimates correspond closely to their data-based analogs. The model captures average exporting rates, match-specific sales dynamics, and the client distribution well, as well as mean values of dependent variables. The model fails, however, to generate the association

²⁹Given the average success rate, $\alpha/(\alpha + \beta)$, the variances of θ^h and θ^f depend only on $\alpha + \beta$.

Table 6: Match hazards, success rates, and endurance

	(i) $\ln(s_{ij})$	(ii) $D_{ijt}^{exit\ match}$	(iii) $\frac{a_{ij}}{n_{ij}}$	(iv) $u_{a_{ij}/n_{ij}}^2$
mean, dep. variable	-0.719 (0.621e-2)	0.395 (0.319e-2)	0.413 (0.153e-2)	0.091 (0.26e-3)
$\ln(1 + n_{ij})$	—	—	0.093 (0.003)	-0.056 (0.000)
$\ln(1 + a_{ij})$	-0.818 (0.113)	—	—	—
$\ln(1 + a_{ij})^2$	0.312 (0.017)	—	—	—
$\ln(1 + \frac{a_{ij}}{n_{ij}})$	-1.132 (0.296)	—	—	—
$[\ln(1 + \frac{a}{n})]^2$	2.451 (0.396)	—	—	—
$\ln(1 + a_{ij}) \cdot \ln(1 + \frac{a_{ij}}{n_{ij}})$	-0.708 (0.134)	—	—	—
$D_{ijt}^{new\ to\ mkt}$	—	0.034 (0.011)	—	—
$\ln X_{ijt}^f$	—	-0.031 (0.002)	—	—
$\ln A_{ijt}$	—	-0.054 (0.009)	—	—
$\ln \Delta_{jt}$	—	-0.028 (0.007)	—	—
observations (rounded)	38,500	23,500	35,800	35,800

Notes: Unit of observation, columns *i*, *iii* and *iv*: seller *j*'s *i*th match. Unit of observation, column *ii*: seller *j*'s *i*th match in its *t*th year. s_{ij} = inverse of time interval between commencement of match *i* and commencement of the next one for exporter *j*. $D_{ijt}^{exit\ match} = 1$ if exporter *j*'s *i*th match dies in year *t*. a_{ij} = cumulative number of successes for exporter *j* at time of match *i*. $D_{ijt}^{new\ to\ mkt} = 1$ if exporter *j*'s *i*th match is in its first year. $\ln A_{ijt}$ = log age of exporter *j*'s *i*th match. $\ln \Delta_{jt}$ = log age of exporter *j* in year *t*. X_{ijt}^f = foreign sales volume generated by exporter *j*'s *i*th match.

Table 7: Client distribution and shipment frequency

	(v) $\ln(1 - \Phi(\ell))$	(vi) $\ln(s_{ijt})$
mean, dep. variable	-5.973 (2.173)	0.971 (0.004)
$\ln(\ell)$	-1.8813 (0.1123)	-
$(\ln \ell)^2$	-0.0545 (0.0211)	-
sample restrictions	$\ell > 0$	$s_{ijt} > 0$
observations	43	87,000

Notes: ℓ : number of active clients; $\Phi(\cdot)$ = cumulative distribution of exporters in terms of ℓ ; s_{ijt} = number of shipments per year to client i by exporter j in year t .

Table 8: Home and foreign sales regressions

	(vii) $\ln X_{ijt}^f$	(viii) $\ln X_{jt}^h$	(ix) $\ln X_{jt}^f$	(x) D_{jt}^f	(xi) $\frac{X_{jt}^f}{X_{jt}^f + X_{jt}^h}$
mean, dep. variable	10.665 (0.002)	-	-	0.102 (0.003)	0.127 (0.002)
R_{ijt-1}	0.328 (0.018)	-	-	-	-
$\ln X_{ijt-1}^f$	0.826 (0.004)	-	-	-	-
$\ln X_{jt-1}^h$	-	0.976 (0.029)	-	-	-
$\ln X_{jt}^h$	-	-	0.323 (0.110)	-	-
$\ln \Delta_t$	0.063 (0.014)	-	-	-	-
root mse	1.2079	0.4621	2.1665	0.303	0.243
sample restrictions	$X_{ijt}^f, X_{ijt-1}^f > 0$	$X_{jt}^h, X_{jt-1}^h > 0$	$X_{jt}^f, X_{jt}^h > 0$	$X_{jt}^h > 0$	$X_{jt}^f, X_{jt}^h > 0$
observations	25,400	99,300	11,600	119,800	12,500

Notes: $R_{ijt} = 1$ if exporter j 's i^{th} match is in its first year. $\ln \Delta_{jt} = \log$ age of exporter j . X_{ijt}^f = foreign sales volume generated by exporter j 's i^{th} match. X_{jt}^f = total foreign sales volume generated by firm j . X_{jt}^h = total home sales volume generated by firm j . $D_{jt}^f = 1$ if firm j is an exporter.

Table 9: Structural parameter estimates

	<i>Parameter</i>	Benchmark model		Known- θ^f variant	
		<i>value</i>	<i>std. error</i>	<i>value</i>	<i>std. error</i>
log of domestic profit scalar	$\ln \Pi^h$	-3.879	(0.1364)	-3.460	(0.0725)
log of foreign profit scalar	$\ln \Pi^f$	-6.135	(0.1993)	-6.273	(0.0759)
fixed cost, domestic	F^h	0.027	(0.0047)	0.037	(0.0064)
fixed cost, foreign	F^f	0.296	(0.0428)	0.301	(0.0359)
First θ distribution parameter	α	0.571	(0.0454)	0.581	(0.0703)
Second θ distribution parameter	β	1.894	(0.2320)	4.661	(0.2107)
demand shock jump size	Δ^y	1.882	(0.2222)	1.951	(0.1810)
shipment order arrival hazard	λ_b	15.426	(0.1991)	15.431	(0.1428)
std. deviation, log firm type	σ_φ	1.386	(0.0095)	1.401	(0.0051)
network effect parameter	γ	0.383	(0.0485)	0.508	(0.0479)
log of home search cost scalar	$\ln \kappa_0^h$	11.722	(0.1486)	12.480	(0.0850)
log of foreign search cost scalar	$\ln \kappa_0^f$	13.002	(0.0095)	13.666	(0.1373)
log of fit metric	$\ln(\Lambda)$	10.806		11.346	

Notes: Both models were fit using the method of simulated moments, targeting the statistics in Tables 6-8 using a block-diagonal weighting matrix based on the covariance matrices of the targeted regressions. Standard errors were constructed using the Delta method.

between success rates and firms’ search intensities that we observe in the data.³⁰ Appendix D provides a more detailed summary of the fit, juxtaposing the data-based moments, \bar{m} , with their simulated counterparts, $m(\Lambda)$, from the benchmark model.

We haven’t targeted the patterns described in Section 2.1, since we treat a successful match in our data as beginning with the first sale and in our model only with the second.³¹ It’s nevertheless instructive to ask how our model replicates them. Table 10 repeats the calculations reported in Table 1 using data simulated with the benchmark model. Qualitatively, the patterns match up: The largest drops in the number of exporters occur during a cohort’s first two years, with cohort size dropping gradually thereafter. Total exports rise early in a cohort’s life, declining thereafter. Finally, exports per surviving firm grow rapidly over time, reflecting both the exit of small-scale firms and client accumulation among survivors.

On the other hand, “Average exports” and “total exports” vary less dramatically with cohort age in the actual data than in the simulated data. Also, the drop in cohort membership is more dramatic in the simulated data during the first year. The difference in the definition

³⁰One reason is that this relationship is statistically weak, and it therefore doesn’t receive much weight in the fit metric. (Note the large standard errors for the coefficients on $\ln(1 + a/n)$ and $[\ln(1 + a/n)]^2$ in column 1 of Table 6.)

³¹Thus, relative to our model-based definition, the tables in Section 2.1 inflate the one-year-old firm and total export counts, while they depress mean exports among one-year olds. Restrictions on data access have temporarily prevented us from making these tables fully compatible. We plan to address the issue in a future draft.

Table 10: Cohort evolutions: simulated data

Cohort age	Exporters	Total Exports	Average Exports
1 year	1.00	1.00	1.00
2 year	0.61	1.73	2.84
3 years	0.35	1.34	3.81
4 years	0.19	1.81	9.50
5 years	0.10	2.29	22.74
6 years	0.06	2.12	34.43
7 years	0.05	1.89	39.69
8 years	0.04	1.69	43.23
9 years	0.03	1.89	63.69
10 years	0.02	1.46	65.17

Notes: Figures for cohorts aged 2-10 are expressed relative to corresponding figures for one-year-old cohorts.

Table 11: Exporter distribution: simulated data

Number of buyers	share of exporters
1	0.77
2	0.10
3	0.05
4	0.03
5	0.02
6-10	0.03
11+	0.01

Notes: Figures give the ergodic distribution of current buyer counts across exporting firms.

of matches between Table 1 and 10 may explain these discrepancies.

Table 11 reports the distribution of client counts across exporters implied by our model. They match quite closely the actual distributions reported in Table 3, although the Table 3 figures show more exporters with two clients (and fewer with more than two clients) than the model predicts.

Finally, Table 12 revisits the analysis of match exit rates in Table 4 with simulated data. (As with the figures in Table 6, the two datasets define matches differently.) The model replicates the higher failure rates among first-year matches, and the tendency for matches that begin in the largest sales quartile to fail less frequently than others. However, the high failure rates are concentrated among one-year-old matches in the simulated data, while they decline more more gradually with age in Table 4. Also, unlike in Table 4, the simulated exporters that begin in the smallest size quartile exhibit failure rates as low as those of the largest exporters. (Here too, since single-shipment encounters are concentrated in the

Table 12: Match separation rates: simulated data

Match age	Quartile 1	Quartile 2	Quartile 3	Quartile 4
1 year	0.60	0.88	0.89	0.63
2 years	0.27	0.29	0.31	0.27
3 years	0.30	0.32	0.33	0.30
4 years	0.31	0.28	0.20	0.32
5+ years	0.28	0.30	0.36	0.36

Notes: Figures are percentages of the exporters in each age-initial size category that do not export during the following year.

smallest size category, differences in the treatment of these encounters in the two tables contribute to the discrepancy.)

5.5 Results for the known product appeal variant

To understand the role of learning in the foreign market we estimate a variant of the model in which a firm knows its true product appeal θ^f from the outset. In this world a firm with low appeal doesn't bother to invest much or at all in searching in the foreign market. A fully-informed firm has less incentive to search since there's no information from a match.

The last two columns of Table 9 report parameter estimates for this variant. Most are similar to those from the benchmark model. But the known- θ^f implies a larger network effect ($\gamma = 0.50$ versus $\gamma = 0.38$) and search cost ($\kappa_0^h = 859$ and $\kappa_0^f = 3,079$ versus $\kappa_0^h = 1,826$ and $\kappa_0^f = 5,982$). These higher values help the known θ^f model explain the observed pattern of small entry, gradual growth, and eventual dominance by high- θ entrants without relying on learning. The known- θ^f model does substantially worse according to Rivers and Vuong's (2002) test statistic for non-nested comparisons.³²

³²The Rivers and Vuong (2002) statistic takes the form $T_n = (\sqrt{n}/\hat{\sigma}_n) [\hat{\Lambda}^1 - \hat{\Lambda}^2]$, where $\hat{\Lambda}^1$ and $\hat{\Lambda}^2$ are the MSM fit metrics for the two models, and $\hat{\sigma}_n^2$ approximates $var [\hat{\Lambda}^1 - \hat{\Lambda}^2]$. This statistic has a standard normal distribution under the null $E(\hat{\Lambda}^1) = E(\hat{\Lambda}^2)$. With model 1 the benchmark and model 2 the known- θ^f variant, we get $T_n = -1,583.2$ (treating the weighting matrix W as nonstochastic). Two caveats apply. First, since the targeted regression coefficients are based on a variety of samples, it's not obvious what sample size n we should use for this statistic. We use a very conservative approximation to the number of firms we base our inferences on ($n = 1000$). Second, this test statistic doesn't recognize randomness in the fit statistics due to simulation.

6 Implications of the Results

Our results allow us to quantify the value of participating in the domestic and foreign markets. Ingredients to this calculation are the value of a successful relationship, the cost of search, and the probability that a match is successful.

6.1 The value of a relationship

In our model a seller benefits from matching with a potential client for three reasons. First is the obvious one that the seller earns profits from sales to the new client. Second, since we find a positive network effect ($\gamma = 0.383$), a successful match lowers the cost of finding additional buyers. Third, the success or failure of the match informs the seller about the popularity of its product, so it can adjust its subsequent search effort accordingly. What do our estimates imply about the magnitudes of these benefits?

Profits per shipment vary widely with macro conditions, the seller's efficiency, and match-specific shocks. When macro conditions are average, foreign shipments generate an average profit of \$US 3,413 with a standard deviation of \$US 3,458. (The analogous figures in the home market are \$US 32,567 and \$US 32,993.)³³ And matches active in either market for an entire year generate an average of $\lambda_b = 15.4$ shipments.

These profits are generated within the match itself. The network and learning benefits from a match derive from subsequent successful relationships. Quantifying these benefits requires not only assessing the return from such relationships, but the cost of searching for them and the probability of success.


6.2 The cost of search

For a firm with no prior success in the foreign market, a search intensity sufficient to yield an average of one new match per year costs $c^f(1, 0) = \$US 1,539$ while an expected yield of four new matches, about one successful match for a firm with average product appeal, costs $c^f(4, 0) = \$US 24,637$. (The analogous figures in the home market are $c^h(1, 0) = \$US 428$ and $c^h(4, 0) = \$US 6,848$.)

Our estimated network parameter $\gamma = 0.383$ implies that these costs fall substantially as a firm racks up successes: A firm with 2 successful foreign matches pays an expected $c^f(4, 2) = \$US 20,142$ for the next match, roughly 20 percent less than the cost of the first. (The analogous figure in the home market is $c^h(4, 2) = \$US 5,598$.)

³³All figures are in 1992 US dollars.

Figure 2: Log continuation values conditioned on match history



Notes: Continuation value trajectories for firms with productivity in the 10th, 50th, and 90th percentiles of the simulated productivity distribution of exporters. For each productivity type, we plot values for all successful matches, alternating success and failure, and all failures.

6.3 The probability of success

A firm’s true probability of success in the foreign market is drawn from a beta distribution, which we estimate to have mean $\alpha/(\alpha + \beta) = 0.23$ and variance $\alpha\beta/[(\alpha + \beta)^2(\alpha + \beta + 1)] = 0.23^2$. Hence, before acquiring export market experience, a firm expects that roughly 1 in 4 encounters with a potential buyer will lead to a business relationship. It also expects to learn a good deal from the outcomes of its early matches.


6.4 Network and learning effects

Combining these ingredients we can assess the combined importance of network and matching effects. Figure 2 shows the perceived continuation value from each additional meeting for firms drawn from the 10th, 50th, and 90th percentiles among simulated exporters. These values depend on the firm’s belief at each moment about its popularity $\bar{\theta}^f$, which in turn depends on the number of matches (n) it’s had already. We show three extreme histories: an unbroken string of successes ($a = n$), an unbroken string of failures ($a = 0$), and alternating success and failure ($n \approx 2a$).

Differences in firms’ efficiency imply big differences in their initial perception of the value of their participation at the outset: the high-productivity firm perceives a value of US\$ 53,800; the median a value of US\$ 452; and the low-productivity firm only US\$ 5.

Both network and learning effects are most powerful for neophyte exporters who haven’t yet formed networks or learned anything about their appeal. The first match has the biggest

Figure 3: Evolution of success probability belief



Notes: Beliefs of a firm with productivity in the 90th percentile of exporters over success probability. Top line is five success followed by five failures. Bottom line is five failures followed by five successes.

impact on continuation values, and most of the impact of additional information is gone by the twentieth match. For example, if its first match is a success, the highest productivity firm’s value jumps to US\$ 165,000. But failures quickly erase firm value. The continuation value of the median productivity firm with four successful matches is almost the same as the value of the high productivity firm with four failed matches, at US\$ 5,669.


Learning effects can cause two *ex ante* identical firms to have very different long term experiences in an export market, depending on whether their early matches succeed. Because match histories affect continuation values, they also affect the intensity with which a firm searches for new clients.

Figure 3 plots the evolution of beliefs about θ^f for two firms both in the 90th productivity percentile (where, for concreteness, searching with intensity s means waiting exactly $1/s$ for the next match). The top line plots perceived appeal from a sequence of five successes followed by five failures, and the bottom line perceived appeal from five failures followed by five successes. They end at the same point, but if success comes early, it takes 10.5 years to get 10 matches while, with initial failure, it takes more than 43 years. The discouraged failure-first firm takes four times longer to get to 10 meetings because it searches less intensively.

6.5 Value dynamics when product appeal is known

The patterns we’ve depicted so far reflect both network and learning effects. To gauge the relative importance of each we redo Figure 2 under the assumption that firms know their true product appeal from the start. We set $\theta^f = 0.43$ for all firms, corresponding to the

Figure 4: Log continuation values conditioned on match history, no learning



Notes: Continuation value trajectories for firms with productivity in the 10th, 50th, and 90th percentiles of the simulated productivity distribution of exporters in the known- θ^f version of the model. For each productivity type, we plot values for all successful matches, alternating success and failure, and all failures.

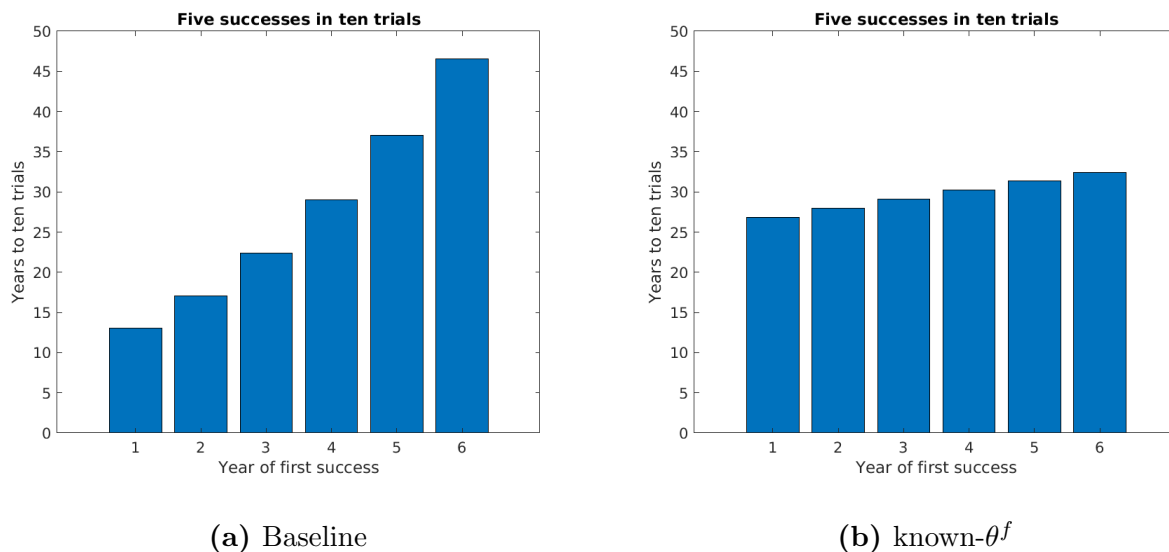
65th percentile of success probabilities among active exporters in our simulated data. Using the estimated “known- θ^f ” policy function in Table 9, we simulate the continuation values of firms at the 10th, 50th, and 90th productivity percentiles.

Parallel to Figure 2, Figure 4 shows the histories of only successes, only failures, and alternating success and failure. With firms knowing their θ^f , continuation values move much less with experience. A successful match does (more modestly) raise the continuation value through the network effect of lowering the cost of search, with unsuccessful matches having no effect.

How do match arrival times depend on successes and failures in the known- θ^f variant? Since firms know their success probabilities, the known- θ^f version of Figure 3 (not pictured) is simply two horizontal lines with height θ^f . But the lengths of these lines still depend on match histories through the network effect. Figure 5 demonstrates this dependence by plotting the expected time to ten meetings when five consecutive meetings succeed and the others fail. The x -axis is the number of meetings before the first success. (For example, if it’s 3, the first 3 meetings fail, the next 5 succeed, and the last 2 fail.)

Comparing again the expected time to ten meetings with five consecutive successes and five consecutive failures, Figure 5 presents results for both the baseline model (panel a) and the known- θ^f variant (panel b) for a firm in the 90th percentile of exporter productivity (again with $\theta^f = 0.43$ in the known- θ^f variant). As in Figure 3, the time it takes a learning firm to reach ten meetings depends heavily on the placement of the successes (panel a), taking 12

Figure 5: Time to ten meetings by placement of five consecutive successes



years if the successes come first but 45 years if the failures come first. For a known- θ^f firm (panel b), it takes 27 years to reach ten meetings if the successes come first and 32 years if they come last.³⁴ Hence learning and network effects both matter, but learning contributes much more to heterogeneity in experience.

6.6 Foreign-market amnesia

Our analysis allows us to quantify the value of a firm’s experience in the export market, which reflects a combination of the expected future profits generated by current business relationships, the benefits of client networks in reducing the cost of finding new clients, and, in the benchmark with learning, the knowledge a firm has acquired about its market appeal.

To assess the value of these intangible assets we generate 2000 different 100-year foreign demand trajectories using the process reported in Table 5. We calculate the value of access the foreign market for our firms in years 25 to 100, after burn in. The value to an average exporter is around US\$ 3.4 million. About half represents future sales to existing customers. The rest reflects mostly the value of knowing market appeal θ^f rather than the visibility generated by its existing customer network.

Multiplying by the average number of exporting firms observed during our sample period implies a total value of US\$ 9.6 billion, about 34 percent of the total value of export revenues,

³⁴The reason it takes so long is that the firm knows that only around half of its meetings will succeed. With θ^f near one it would only take a few years to reach 10 meetings.

US\$ 28.3 billion. In other words, a hypothetical experiment of total amnesia about the foreign market (losing contact with clients and forgetting past successes and failures) would wipe out about a third of the value of the U.S. market for our Colombian firms.

Since the distribution of export sales is highly skewed across firms, the largest exporters dominate these averages. The median exporter would stand to lose only about US \$ 370,000 from total U.S. market amnesia, hardly more than a tenth of the average. This loss reflects mostly lost sales to existing customers. To summarize, the value of a large exporter's experience in a market is nearly as much in its enhanced ability to find new customers as in its expected future sales to existing ones. The main benefit of experience for a smaller exporter is its expected future sales to existing customers.

7 Exchange rate dynamics


A natural task for our model is characterizing the response of exports to an exchange rate shock. We consider a 20 percent devaluation by scaling the 100-year foreign market demand trajectories generated by the process reported in Table 5 by 1.2 after the 50th year. Since we treat the shock as unanticipated, the first 50 years replicate the experiment just discussed in Section 6.6. We first look at what happens to an average firm in our simulation. We then turn to aggregate trade dynamics, breaking down the response into different margins of adjustment. We conclude with a discussion of the implied short and long-run trade elasticities.

Our results for individual exporters are based on the 2000 simulations described in Section 6.6. In this section, we consider the total value of the firm, including both its value from experience and its option value of reentering the foreign market with the same efficiency φ but in a state of foreign-market amnesia. Calculated this way, the mean value of an active exporter before the shock is about 5 million dollars, and the median is 540 thousand dollars. The value of the mean exporter increases 34 percent just after the shock. Average exporter value jumps more than the mechanical 20 percent from the devaluation, because after the shock exporters search harder and expect to both learn and become visible more quickly. As time passes, the value of the average exporter decreases. This pattern reflects the role of matching frictions, which delay the entry of marginally profitable firms.

How do these changes in value translate into aggregate export dynamics? Figure 6 summarizes the results of simulating the aggregate export trajectories associated with the exchange rate devaluation.³⁵ The permanent 20 percent real peso devaluation occurs at the

³⁵Our single-agent model misses interactions between exporters in the foreign market. But since Colombia constitutes a small share of the U.S. spending, such general equilibrium effects are likely negligible.

Figure 6: Baseline response to a permanent devaluation: export aggregates



Notes: Figures depict aggregate responses to a permanent 20 percent real devaluation at time 0. Shaded areas in panels a and b reflect contributions of matches that existed at time 0 (yellow), matches formed after time 0 by exporters that were active at time 0 (red), and matches formed after time 0 by exporters that entered the foreign market after time 0 (blue). Thin lines show patterns in the absence of the shock. Panel c depicts incumbent exporters active before time zero (red), and exporters entering after time zero (blue). All series are averages across 2000 simulations.

end of the 50th year (marked period 0 in the figure). Panels a and b break down the total value of a particular aggregate export sales and matches into three segments: contributions from matches created before period 0 (yellow area), contributions from matches created after period 0 with exporters that existed in period 0 (red area), and matches formed after period 0 with exporters that entered after period 0 (blue area). Panel c breaks down the aggregate number of active exporters into those active before period 0 (red area) and those who entered after period 0 (blue area). The lines show how the boundaries between the shaded areas would have differed if there had been no permanent devaluation.³⁶

Panel *a* describes total export sales. The rapid turnover in matches is striking. Within several years, incumbent matches have lost about three-quarters of their market share, regardless of whether the exchange rate depreciates or fluctuates around a stationary mean (the latter case shown by the thin superimposed lines). Nonetheless, given their persistent productivity (φ) and product appeal (θ), incumbent exporters retain more than 50 percent of the market after 25 years by regularly replenishing their client portfolio.

Despite rapid match turnover, adjustments to the new exchange rate take time to play out. During the first year, sales within matches account for almost all of the movement in

³⁶Piveteau (forthcoming) provides similar graphs that inspired Figure 6. To highlight the role of learning and endogenous match separations, we use a decomposition that distinguishes matches to new exporters from others. Piveteau (forthcoming) distinguishes the consumer margin, the extensive margin, and an aggregate valuation effect.

Table 13: Simulated Trade Elasticities

Time since shock	1 year	5 years	25 years
Sales	1.20 (0.05)	1.69 (0.05)	1.88 (0.05)
Matches	0.26 (0.03)	0.80 (0.03)	0.94 (0.03)
Exporters	0.14 (0.02)	0.50 (0.02)	0.60 (0.02)

Notes: All elasticities are based on 2000 simulations of a 20 percent real devaluation of the Colombian peso. Standard errors based on cross-simulation standard deviations are in parentheses.

aggregate export sales (Panels *a*). The total effect of the exchange rate shock grows over time. The effect grows both because the total number of exporters adjusts, and also because the number of matches per exporter increases as is implied by Panels *b* and *c*. Over a period of roughly 15 years, these effects add an extra 40-50 percent to the initial response.

While it is not immediately apparent from Figure 6, the percentage contribution of each type of match to the corresponding aggregate is nearly invariant to the devaluation shock. Hence, if we were only interested in the rate at which new exporters displace incumbent exporters, or the rate at which new matches displace incumbent matches, it would matter very little whether we were analyzing the aftermath of a permanent devaluation or a period without any regime switching.³⁷

Table 13 reports the short, medium, and long-run trade elasticities implied by the permanent 20% real peso devaluation with standard errors based on cross-simulation standard deviations in parentheses. Our long run sales elasticities resemble Piveteau’s (forthcoming) and Boehm et al.’s (2020), but the learning effects in our framework imply a somewhat longer transition period. Our elasticities are substantially lower than the long run elasticities typically generated by calibrated general equilibrium models (e.g., Alessandria and Choi, 2014; Alessandria, et al., 2018).³⁸

³⁷This is true across exchange rate shocks both in the baseline model and in the Known- θ^f model described in Table 9. The contributions of types of match do differ, however, when we compare across models. In particular, we find that in the Known- θ^f model the share of new matches going to new exporters is higher than in the baseline. We discuss these and related comparisons across models in Appendix Section E.

³⁸Alessandria and Choi (2014) use a symmetric 2-country dynamic model with endogenous firm creation, capital accumulation, fixed exporting costs, and iceberg costs. Analyzing movement from a global 8 percent tariff to free trade, they find the a trade elasticity rises of about 5 in the short run and 8 in the long run, which is reached in 5-8 years. In a similar model, but with firms’ exporting costs depending upon their incumbency, Alessandria et al. (2018) estimate a short-run trade elasticity of 4 and a long-run elasticity of 11.55. Their model generates transition dynamics over 10-15 years.

8 Summary

Research exploiting customs records has generated a robust set of stylized facts regarding firm-to-firm trade dynamics: First, most exporters are inexperienced, ship small amounts, and have few foreign clients. Second, the typical buyer-seller relationship lasts only a year or two, so business connections evolve rapidly, and it's common to see firms with only a few clients cease exporting entirely, giving way to the next entering cohort of inexperienced exporters. Third, however, each new cohort contains a small number of firms that survive and grow many times faster than aggregate exports. They do so not by selling more to the same clients, but by finding new customers.

We confirm these patterns for Colombian manufacturers shipping to the United States, and develop a continuous-time model to account for them. Firms wishing to export must engage in costly search to find potential buyers, who may either reject their products or form finite-lived business relationships with them. Buyers who form business relationships with exporters send them favorable signals about the appeal of their products and, in doing so, encourage them to search more intensely for additional buyers (learning effects). Successful business relationships also reduce sellers' search costs by improving their visibility (network effects). Finally, sellers' search intensities depend both on their permanent idiosyncratic characteristics and on market conditions.

Fit using the method of simulated moments, the model replicates these patterns in customs records and allows us to quantify several types of trade costs, including the cost of searching for potential clients and the cost of maintaining business relationships with existing clients. It also allows us to estimate the network effect of previous successes on the costs of meeting new clients, and to characterize the cumulative effects of learning on firms' search intensities and intangible capital stocks.

While our model delivers similar long-run elasticities to other one-sided search models, the presence of learning means that it takes longer to reach the long-run. The reason the long-run effect of learning and visibility is modest is that they are most important among newer exporters, which account for a small share of total export volume. Much more important for aggregate export dynamics are the search frictions that limit the ability of exporters to connect with new potential buyers.

References

- Aeberhardt, R., I. Buono, and H. Fadinger (2014): “Learning, Incomplete Contracts and Export Dynamics: Theory and Evidence from French Firms.” *European Economic Review* 68: 219–249
- Albornoz, Facundo, Hector Calvo Pardo, Gregory Corcos, and Emanuel Ornelas (2012) ”Sequential Exporting.” *Journal of International Economics* 88: 17-31.
- Alessandria, George, Costas Arkolakis, and Kim Ruhl (2020) ”Firm Dynamics and Trade” NBER Working Paper 27934.
- Alessandria, George and Horag Choi (2007) ”Do Sunk Costs of Exporting Matter for Net Export Dynamics?” *Quarterly Journal of Economics* 122(1): 289-336.
- Alessandria, George and Horag Choi (2014) ”Establishment Heterogeneity, Exporter Dynamics, and the Effects of Trade Liberalization.” *Journal of International Economics* 94: 207-233.
- Alessandria, George and Horag Choi (2019) ”The Dynamics of the U.S. Trade Balance and Real Exchange Rate: The J Curve and Trade Costs?” NBER Working Paper 25563.
- Alessandria, George, Sangeeta Pratap, and Vivian Yue (2014) ”Export Dynamics in Large Devaluations.” Working Paper, Federal Reserve Bank of Philadelphia.
- Alessandria, George, Horag Choi, and Kim Ruhl (2018) ”Trade Adjustment Dynamics and the Welfare Gains from Trade.” Working Paper, The University of Rochester.
- Andrews, Isiah, Matthew Gentzkow, and Jesse Shapiro (2017) “Measuring the Sensitivity of Estimated Parameters to Estimation Moments.” *Quarterly Journal of Economics* 132(4): 1151-1199.
- Araujo, Luis, Emanuel Ornelas and Giordano Mion (2016) ”Institutions and Export Dynamics.” *Journal of International Economics* 98: 2-20.
- Arkolakis, Konstantinos (2010) “Market Access Costs and the New Consumers Margin in International Trade.” *Journal of Political Economy* 118(6): 1151-1199.
- Arkolakis, Konstantinos (2015) “A Unified Theory of Firm Selection and Growth.” *Quarterly Journal of Economics* 131(1): 89-155.
- Arkolakis, Konstantinos, Theodore Papageorgiou and Olga Timoshenko (2018). ”Firm Learning and Growth.” *Review of Economic Dynamics* 27: 146-168.

- Atkeson, Andrew and Ariel Burstein (2010) "Innovation, Firm Dynamics, and International Trade." *Journal of Political Economy* 118(3): 433-484.
- Békés, Gábor, Lionel Fontagné, Balázs Murakozy, and Vincent Vicard (2017). "Shipment Frequency of Exporters and Demand Uncertainty." *Review of World Economics* 153(4): 779-807.
- Berman, N., V. Rebeyrol, and V. Vicard (2019) "Demand learning and Firm Dynamics: Evidence from Exporters." *Review of Economics and Statistics* 101(1): 91-106.
- Bernard, Andrew, J. Bradford Jensen, and Peter K. Schott (2009) "Importers, Exporters, and Multinationals: A Portrait of Firms in the U.S. that Trade Goods," in Timothy Dunne, J. Bradford Jensen and Mark J. Roberts eds. *Producer Dynamics*, University of Chicago Press.
- Bernard, Andrew, Esther Ann Boler, Renzo Massari, Jose-Daniel Reyes, and Daria Taglioni (2017) "Exporter Dynamics and Partial-Year Effects." *American Economic Review* 107(10): 3211-3228.
- Bernard, Andrew and Andreas Moxnes (2018) "Networks and Trade." *Annual Review of Economics* 10(65): 65-85.
- Besedes, Tibor (2008). "A Search Cost Perspective on the Formation and Duration of Trade." *Review of International Economics* 16(5): 835-849.
- Blum, Bernardo S., Sebastian Claro, Kunal Dasgupta, and Ignatius J. Horstmann (2019). "Inventory Management, Product Quality, and Cross-country Income differences." *American Economic Journal: Macroeconomics* 11(1): 338-388.
- Boehm, Christoph, Andrei Levchenko and Nitya Pandalai-Nayar (2020). "The Long and Short (Run) of Trade Elasticities." Working Paper, The University of Michigan.
- Brooks, Eileen (2006) "Why Don't Firms Export More? Product Quality and Colombian Plants" *Journal of Development Economics* 80: 160-178.
- Burstein, Ariel and Marc Melitz (2013) "Trade Liberalization and Firm Dynamics," in *Advances in Economics and Econometrics Tenth World Congress*. Applied Economics, Econometric Society Monographs. Vol. 2. Cambridge, UK: Cambridge University Press.
- Cebreros, Alfonso (2016). "The Rewards of Self-discovery: Learning and Firm Exporter Dynamics." Banco de México Working Paper 2016-08.

- Domínguez, Juan Camilo, Jonathan Eaton, Marcela Eslava, and James Tybout. (2013) "Search and Learning in Export Markets: Case Studies for Colombia." Pennsylvania State University, Working Paper.
- Drozd, Lukasz A. and Jaromir B. Nosal (2012) "Understanding International Prices: Customers as Capital." *American Economic Review* 102(1): 364-395.
- Eaton, Jonathan, Marcela Eslava, Maurice Kugler and James Tybout (2008). "Export Dynamics in Colombia: Firm-Level Evidence," in Elhanan Helpman, Dalia Marin and Thierry Verdier, eds., *The Organization of Firms in a Global Economy*, Cambridge, MA: Harvard U. Press.
- Eaton, Jonathan, Samuel Kortum and Francis Kramarz (2011). "An Anatomy of International Trade: Evidence From French Firms" *Econometrica* 79(5): 1453-1498.
- Eaton, Jonathan, David Jinkins, James Tybout, and Daniel Xu (2016) "Two-sided Search in International Markets." Working paper, Pennsylvania State U.
- Fajgelbaum, Pablo D. (2020). "Labor Market Frictions, Firm Growth, and International Trade." *Review of Economic Studies* 87(3): 1213-1260.
- Fitzgerald, Doireann, Stefanie Hallerz, and Yaniv Yedid-Levi (2019) "How Exporters Grow." Working Paper, Federal Reserve Bank of Minneapolis.
- Gouriéroux and Monfort, 1996. *Simulation-Based Econometric Methods*. New York: Oxford U. Press.
- Handley, Kyle and Nuno Limão (2017). "Policy Uncertainty, Trade, and Welfare: Theory and Evidence for China and the United States." *American Economic Review* 107(9): 2731-2783.
- Hornok, Cecília and Miklós Koren (2015). "Per-Shipment Costs and the Lumpiness of International Trade." *The Review of Economics and Statistics* 97(2): 525-530.
- Impullitti, Giammario, Alfonso Irarrazabal, and Luca Oromolla (2013) "A Theory of Entry into and Exit From Export Markets." *Journal of International Economics* 90: 75-90.
- Kropf, Andreas and Philip Sauré (2014). "Fixed Costs per Shipment." *Journal of International Economics* 92(1): 166-184.
- Li, Shengyu (2018) "A Structural Model of Productivity, Uncertain Demand, and Export Dynamics." *Journal of International Economics* 115: 1-15.

- Nguyen, Daniel (2012) "Demand Uncertainty, Exporting Delays and Exporting Failures." *Journal of International Economics* 86: 336-344.
- Piveteau, Paul (forthcoming) "An Empirical Dynamic Model of Trade with Consumer Accumulation." *American Economic Journal: Macroeconomics*.
- Rivers, Douglas and Quang Vuong (2002). "Model Selection for Nonlinear Dynamic Models." *Econometrics Journal* 5: 1-19.
- Rodrigue, Joel and Yong Tan (2019). "Price, Product Quality, and Exporter Dynamics: Evidence from China." *International Economic Review* 60(4): 1911-1955.
- Ruhl, Kim (2008) "The International Elasticity Puzzle." Working Paper, The University of Wisconsin.
- Ruhl, Kim and Jonathan Willis (2017) "New Exporter Dynamics." *International Economic Review* 58(3): 703-725.
- Schmeiser, Katherine N (2012). "Learning to Export: Export Growth and the Destination Decision of Firms." *Journal of International Economics* 87(1): 89-97.
- Shimer, Robert (2005) "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *The American Economic Review* 95(1): 25-49.
- Timoshenko, O. A. (2015): "Learning Versus Sunk Costs Explanations of Export Persistence." *European Economic Review* 79: 113-128

FOR ONLINE APPENDIX

A data tables

year	1992	1993	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	total
1992	2,232																		2,232
1993	823	1,235																	2,058
1994	583	330	1,160																2,073
1995	440	213	339	953															1,945
1996	372	163	178	255	899														1,867
1997	321	128	133	170	248	877													1,877
1998	268	104	124	132	153	256	893												1,930
1999	232	85	87	114	117	187	262	1,026											2,110
2000	203	85	79	91	103	136	170	344	1,372										2,583
2001	187	70	65	79	85	109	145	229	389	1,251									2,609
2002	173	64	62	72	68	88	112	171	242	399	1,373								2,824
2003	165	51	58	62	62	77	86	140	185	301	440	1,719							3,346
2004	150	52	41	53	63	76	80	132	164	223	327	616	1,768						3,745
2005	140	52	47	39	54	77	69	115	145	196	235	398	661	1,902					4,130
2006	122	46	44	39	44	71	65	110	131	157	168	308	410	564	1,896				4,175
2007	113	37	39	31	42	55	48	91	101	132	156	240	305	365	548	1,681			3,984
2008	93	29	30	24	38	50	45	74	90	117	130	184	198	230	331	447	1,455		3,565
2009	80	25	28	24	28	40	39	60	72	88	97	145	175	157	230	248	386	1,378	3,300

Table 14: Number of Exporting Firms, by Entry Cohort

year	1992	1993	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	total
1992	469																		469
1993	352	83																	435
1994	336	83	92																510
1995	313	75	102	58															549
1996	256	67	62	40	60														484
1997	247	84	43	41	48	119													581
1998	225	49	42	36	45	131	63												590
1999	207	51	49	41	39	197	74	81											739
2000	180	53	55	37	51	102	53	158	109										799
2001	150	22	51	41	28	57	36	80	101	111									677
2002	124	23	47	34	27	28	23	45	65	83	40								538
2003	147	42	51	31	42	24	22	37	71	107	50	78							702
2004	156	43	53	19	57	21	23	42	78	106	60	107	90						855
2005	150	22	75	17	52	18	23	43	78	80	58	81	75	84					855
2006	117	31	52	14	64	43	17	38	61	79	32	51	52	112	78				838
2007	103	7	18	11	67	58	19	30	28	64	22	35	33	66	67	62			689
2008	95	6	9	8	33	37	17	33	26	34	20	31	37	54	42	53	57		591
2009	68	22	7	6	13	24	10	23	16	16	14	22	41	25	39	37	36	64	485

Table 15: Value of Exports, by Entry Cohort (millions of \$US)

year	1992	1993	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	pooled
1992	210																		210
1993	428	67																	211
1994	576	251	79																246
1995	712	353	300	61															282
1996	687	411	346	158	67														259
1997	771	652	321	241	192	136													310
1998	839	468	339	269	297	510	71												306
1999	893	601	561	361	336	1,054	281	79											350
2000	885	623	697	407	496	750	313	460	80										309
2001	801	316	783	519	329	521	251	350	259	89									260
2002	716	353	757	473	399	318	207	260	268	207	29								191
2003	891	827	870	493	677	315	257	260	385	355	114	46							210
2004	1,039	828	1,281	358	900	281	291	318	478	476	183	174	51						228
2005	1,071	413	1,593	444	967	231	326	375	535	408	248	204	113	44					207
2006	958	675	1,177	356	1,448	605	256	341	464	505	188	165	126	198	41				201
2007	915	175	466	357	1,606	1,048	391	327	278	481	140	145	108	181	123	37			173
2008	1,023	208	283	341	860	747	379	443	289	287	153	166	186	236	125	120	39		166
2009	855	864	262	266	478	607	255	389	221	176	143	152	235	162	169	151	93	47	147

Table 16: Exports per Firm, by Entry Cohort (thousands of \$US)

Year	Colombian Sellers	U.S. Importers	Pairs
1992	2,232	1,190	3,087
1993	2,058	1,183	2,824
1994	2,073	1,212	2,810
1995	1,945	1,173	2,588
1996	1,867	1,191	2,490
1997	1,877	1,208	2,480
1998	1,930	1,191	2,495
1999	2,110	1,386	2,793
2000	2,583	1,661	3,411
2001	2,609	1,698	3,483
2002	2,824	1,826	3,733
2003	3,346	2,110	4,483
2004	3,745	2,296	5,071
2005	4,130	2,457	5,552
2006	4,175	2,471	5,607
2007	3,984	2,343	5,307
2008	3,565	2,221	4,751
2009	3,300	2,079	4,467

Table 17: Exporters and importers by year

B Data checks

To investigate the quality of the exporter id (`manuf_id`) in the U.S. import records, we ran a series of robustness checks. The Colombian and U.S. data overlap for the years 2000-2008 and both contain measures of the value of exports as well as the number of exporting firms. If the `manuf_id` variable is error-prone and noisy, we would expect the U.S. data to over-report the number of Colombian firms exporting to the U.S. That is, each time a customs broker wrongly enters the data in the field, a new firm would be created. Table 18 below summarizes the total value of exports to the U.S. and the number of Colombian firms, by year, for each data set.

The datasets align much more closely on value than they do on firm counts. The difference in value is never more than 10% while the firm count difference ranges from 18% to 74%. The differences are stable over time.

To look more closely at the cause of the difference in firm counts, we compared the number of firms across sources by HS2 categories. The counts in the LFTTD were higher than the Colombian data in only 28 of the 82 codes and by far the biggest differences are in HS codes 61 and 62: textiles. In these two product classes the U.S. data identify 4025 more firms than the Colombian data. If we remove these two sectors from the list, the difference in firm counts flips and the Colombian data contain 1001 more firms than the LFTTD.

Title 19 of U.S. code requires that the `manuf_id` variable for textile products represent the manufacturer of the textile products, not an intermediary. That is, for this sector, in particular, the CBP 7501 form must report the manufacturer, not an intermediary. By contrast, prior work by several authors of this paper has shown that the Colombian data report the exporter, which may or may not be the manufacturer. Given that previous research (Tybout, 2000 JEL) has shown that developing countries tend to have a disproportionately large share of small manufacturing firms, it's reasonable to assume that a large part of the

Year	Colombia		United States		% difference	
	# exporters	value	# exporters	value	# exporters	value
2000	1775	1038	2721	1140	53%	10%
2001	2026	995	2744	1019	35%	2%
2002	2230	870	2986	855	34%	-2%
2003	2800	1113	3579	1119	28%	1%
2004	3035	1379	4002	1415	32%	3%
2005	2861	1554	4288	1438	50%	-7%
2006	2689	1665	4361	1552	62%	-7%
2007	2420	1540	4175	1496	73%	-3%
2008	2161	1570	3758	1474	74%	-6%

Table 18: Colombian versus U.S. Customs Records

reason the U.S. data report so many more firms in the textile sector is that the U.S. data count many small manufacturers while the Colombian data are, in many cases, reporting aggregators and intermediaries.

As a final check of the integrity of the `manuf_id` variable, and the robustness of our main results, we experimented with a “fuzzy” version of the `manuf_id` variable that did not contain any street numbers in the string (a likely source of input errors). The effect is to reduce the number of Colombian firms in the data, an approximation of fixing any extraneous noise from data entry errors. Next we re-ran Table 4 with the fuzzy data and compared the results to the original version.

One of the key findings from Table 4 is the high match separation rates ranging from about 40% to 66%. Using the fuzzy version did not reduce the separation rates substantially and left the patterns intact. The fuzzy separation rates ranged from 26% to 62%, a drop of 6% on average. It does not appear that our results are sensitive to a modest reduction in data entry errors.

C Identification

Following Andrews, et al. (2017), Table 19 below reports estimates of $(G'WG)^{-1}G'W$, where $G = -\partial[m(\Lambda)]/\partial\Lambda'$ is the Jacobian for the vector of simulated moments. It was constructed using our benchmark parameter estimates with all elements converted to elasticity form to make them unit-free.

The text discusses results associated with particular equations. Some general observations are as follows. Most parameters respond to many moments rather than one or a few. Among parameters with elasticities having absolute value greater than 0.1, most respond significantly to at least 5 moments and several (F^f, F^h, γ) respond to more than 15. All parameters respond to at least 2. The moments affecting the most parameters are those generated by the match sales autogression (equation *vii*), the shipping rate regression (equation *vi*), the domestic sales autoregression (equation *viii*), the regression explaining the variance in success rates (equation *iv*), and the fraction of firms that export (equation *x*).

Table 19: Sensitivity matrix, elasticity form

	$\ln \Pi^h$	F^h	F^f	$\ln \Pi^f$	α	β	Δ_y	λ_b	γ	$\ln \kappa_0^h$	$\ln \kappa_0^f$	σ_φ
avg. mat death	0.008	0.064	0.010	0.005	-0.009	0.018	0.001	0.001	-0.031	-0.002	-0.002	0.000
new to mkt	-0.029	0.167	-0.189	-0.014	0.066	-0.054	-0.085	0.003	0.168	-0.004	-0.001	-0.004
current sales	-0.019	-0.360	-0.234	0.016	0.000	-0.008	-0.195	-0.009	0.180	0.035	0.009	0.008
exporter age	0.020	-0.259	0.193	0.013	-0.055	0.046	0.050	-0.005	-0.118	0.006	0.004	0.005
match age	0.005	0.086	-0.015	0.002	0.002	0.009	-0.004	0.002	-0.007	-0.004	-0.002	-0.001
avg. match sales	-0.536	-4.896	-1.569	-0.134	0.762	-0.259	-1.329	-0.125	2.067	0.239	0.181	0.084
1st yr dummy	-0.353	-2.996	1.325	-0.745	0.622	-1.855	2.991	0.058	1.111	-0.172	0.161	-0.107
match sales, t-1	0.107	1.303	-0.200	0.112	-0.194	0.414	0.055	0.029	-0.745	-0.004	-0.050	0.005
exporter age	-0.007	-0.105	0.040	-0.010	0.018	-0.038	-0.003	-0.002	0.065	-0.002	0.004	-0.001
MSE, match AR1	0.065	0.392	-0.325	0.014	-0.054	0.040	0.035	-0.002	-0.033	-0.001	-0.011	-0.002
degree dist. slope	0.010	0.091	-0.004	0.027	-0.150	-0.479	0.013	0.039	0.120	-0.024	0.013	0.014
degree dist. curv.	0.020	0.094	-0.031	0.021	-0.078	-0.109	-0.020	0.009	-0.003	-0.005	0.005	0.008
avg. ln #shipments	-0.079	-1.086	1.195	-0.043	0.120	-0.050	-0.636	1.034	-0.560	0.046	0.000	-0.023
export/dom coef.	0.079	0.546	0.163	-0.081	0.011	-0.111	0.639	0.030	-0.122	-0.078	-0.007	-0.023
dom. sales AR1	-0.705	1.813	-13.736	0.457	-1.018	1.636	0.672	-0.067	-2.912	1.209	-0.074	0.121
avg. match hazard	0.000	-0.245	-0.663	0.001	0.010	-0.019	-0.028	0.005	0.558	0.034	-0.001	-0.073
$\ln(1+a)$	-0.004	-0.025	-0.005	-0.004	0.002	-0.018	0.009	0.000	0.015	0.001	0.001	-0.001
$\ln(1+a)^2$	0.000	-0.003	0.017	-0.003	0.001	-0.010	0.010	0.001	0.005	-0.002	0.000	-0.001
$\ln(1+\frac{1}{n})$	-0.018	-0.107	-0.037	-0.018	-0.002	-0.072	0.048	0.000	0.035	0.008	0.005	-0.002
$\ln(1+\frac{1}{n})^2$	-0.039	-0.216	-0.149	-0.024	0.009	-0.108	0.039	-0.004	0.097	0.020	0.009	-0.002
$\ln(1+\frac{1}{n}) \cdot \ln(1+a)$	-0.013	-0.069	-0.067	-0.004	0.006	-0.019	-0.001	-0.002	0.030	0.008	0.002	0.000
avg. $\frac{a}{n}$	0.022	-0.603	-0.870	0.206	1.399	-0.090	0.440	-0.082	-0.693	0.053	0.001	0.032
coef., $\ln n$	-0.003	-0.006	-0.061	-0.005	0.010	0.009	0.037	0.000	-0.001	0.000	0.001	0.000
$var(\frac{a}{n} n)$	-0.238	-0.435	1.257	-0.361	1.001	-1.035	-0.447	-0.015	-0.974	-0.009	0.077	-0.053
coef., $\ln n$	0.100	1.217	-1.121	0.096	0.117	0.749	-0.350	-0.020	-0.675	0.012	-0.025	0.022
$\frac{foreign\ sales}{total\ sales}$	0.157	-0.155	0.419	-0.016	0.048	-0.072	-0.039	0.010	0.172	-0.055	0.006	-0.002
$\frac{\#exporters}{\#firms}$	0.118	0.572	-1.517	0.048	-0.282	0.014	0.229	0.021	-0.354	0.117	-0.024	0.003

D Model Fit

The tables in this appendix compare data-based moments (top row) with their model-based counterparts (middle row). The bottom row reports standard errors for the data-based moments in parentheses. The last two rows repeat Tables 6, 7 and 8 in the test.

Table 20, column 1, shows that the model understates monthly log match hazards. The quadratic relationship between match hazards and cumulative successes in the data appears in the model-based simulations, albeit somewhat dampened. And the relation between success rates and match hazards changes curvature. Column 2 shows that the model underpredicts match death rates, though it picks up their negative relationship to match sales and age. (The first year effect seems to be entirely absorbed by this age variable.) As for success rates, the model comes reasonably close to the data. It misses the positive association between this variable and number of matches, but does replicate the reduction in success rate dispersion as the cumulative number of matches grows.

Turning to Table 21, we see that model gets the nearly-Pareto distribution of client counts across firms, as the coefficient on $\ln(\ell)^2$ is negative but close to zero, just as in the data. However, the slope of regression v is less negative in the simulated data than in the actual data, implying that the model predicts relatively more exporters have high-client counts. As for equation (vi) , the estimated model generates more shipments per month among active matches than we find in the data.

Finally, Table 22 shows that the model does a good job of explaining match-level sales dynamics (equation vii), including the dependence of sales on exporters' market tenure. It also gets the persistence in home market sales almost exactly right (equation $viii$). It's less successful at explaining the weak correlation between domestic and foreign sales, perhaps because the dependent variable is exports destined for the U.S. alone, and not exports to other destinations, which are not in our model.

Table 20: Match hazards, success rates, and endurance: Model vs. Data

	(i) $\ln(s_{ij})$	(ii) $D_{ijt}^{exit\ match}$	(iii) $\frac{a_{ij}}{n_{ij}}$	(iv) $u_{a_{ij}/n_{ij}}^2$
mean, dep. variable	-0.719 1.527 (0.621 E-2)	0.395 0.267 (0.319 E-2)	0.413 0.470 (0.153 E-2)	0.091 0.066 (0.265 E-3)
$\ln(1 + a_{ij})$	–	–	0.093 –0.009 (0.003)	–0.060 –0.033 (0.000)
$\ln(1 + a_{ij})$	–0.818 –0.371 (0.113)	–	–	–
$[\ln(1 + a_{ij})]^2$	0.312 0.024 (0.017)	–	–	–
$\ln(1 + \frac{a_{ij}}{n_{ij}})$	–1.132 3.774 (0.296)	–	–	–
$[\ln(1 + \frac{a_{ij}}{n_{ij}})]^2$	2.451 –5.555 (0.396)	–	–	–
$\ln(1 + a_{ij}) \cdot \ln(1 + \frac{a_{ij}}{n_{ij}})$	–0.708 0.564 (0.134)	–	–	–
$D_{ijt}^{new\ to\ mkt}$	–	0.034 –0.133 (0.012)	–	–
$\ln X_{ijt}^f$	–	–0.032 –0.033 (0.002)	–	–
$\ln A_{ijt}$	–	–0.054 –0.077 (0.009)	–	–
$\ln \Delta_{jt}$	–	–0.028 0.020 (0.007)	–	–

Notes: Unit of observation, columns *i*, *iii* and *iv*: seller *j*'s *i*th match. Unit of observation, column *ii*: seller *j*'s *i*th match in its *t*th year. s_{ij} = inverse of time interval between commencement of match *i* and commencement of the next one for exporter *j* $D_{ijt}^{exitmatch} = 1$ if exporter *j*'s *i*th match dies in year *t*. a_{ij} = cumulative number of successes for exporter *j* at time of match *i*. $D_{ijt}^{newtomkt} = 1$ if exporter *j*'s *i*th match is in its first year. $\ln A_{ijt}$ = log age of exporter *j*'s *i*th match. $\ln \Delta_{jt}$ = log age of exporter *j* in year *t*. X_{ijt}^f = foreign sales volume generated by exporter *j*'s *i*th match.

Table 21: Client distribution and shipment frequency, model vs. data

	(v)	(vi)
	$\ln(1 - \Phi(\ell))$	$\ln(\lambda_b)$
mean, dep. variable	-	0.971 1.489 ()
$\ln(\ell)$	-1.881 -1.199 (0.112)	-
$(\ln \ell)^2$	-0.056 -0.115 (0.021)	-
sample restrictions	$\ell > 0$	$\lambda_b > 0$
observations	43	87,000

Notes: ℓ : number of active clients; $\Phi(\cdot)$ = cumulative distribution of exporters in terms of ℓ ; s_{ijt} = number of shipments per year to client i by exporter j in year t

Table 22: Home and foreign sales regressions

	(vii)	(viii)	(ix)	(x)	(xi)
	$\ln X_{ijt}^f$	$\ln X_{jt}^h$	$\ln X_{jt}^f$	D_{jt}^f	$\frac{X_{jt}^f}{X_{jt}^h + X_{jt}^f}$
mean, dep. variable	10.665 10.957 (0.002)	-	-	0.102 0.141 (0.003)	0.127 0.062 (0.002)
R_{ijt-1}	0.328 0.607 (0.018)	-	-		
$\ln X_{ijt-1}^f$	0.826 0.848 (0.004)	-	-		
$\ln X_{jt-1}^h$	-	0.976 0.964 (0.001)	-		
$\ln X_{jt}^h$	-	-	0.323 0.811 (0.012)		
$\ln \Delta_{jt}$	0.063 0.060 (0.014)	-	-		
sample restrictions	$X_{ijt}^f > 0$ $X_{ijt-1}^f > 0$	$X_{jt}^h > 0$ $X_{jt-1}^h > 0$	$X_{jt}^f > 0$ $X_{jt}^h > 0$	$X_{jt}^h > 0$	$X_{jt}^f, X_{jt}^h > 0$
observatiaons	25,400	99,300	11,600	119,800	12,500

Notes: $R_{ijt} = 1$ if exporter j 's i^{th} match is in its first year. $\ln \Delta_{jt} = \log$ age of exporter j . X_{ijt}^f = foreign sales volume generated by exporter j 's i^{th} match. X_{jt}^f = total foreign sales volume generated by firm j . X_{jt}^h = total home sales volume generated by firm j . $D_{jt}^f = 1$ if firm j is an exporter.

E Learning, networks, and aggregate export dynamics

The relative contributions of the different types of matches depend on the learning and visibility effects in our model. To demonstrate how, we contrast the post-devaluation evolution of aggregates implied by our benchmark model (column 1 of Table 9) with the pattern implied by the known- θ^f variant (column 3 of Table 9). And we contrast the dynamic implications of the known- θ^f version of the model with those of a model in which θ^f 's are known *and* all firms have maximum visibility, regardless of their exporting history.³⁹

Figure 7 summarizes our findings. Panels a and b describe the market share of incumbent matches (yellow region), new matches with incumbent exporters (red region), and new matches with new exporters (blue region) after an exchange regime shock. Panel c describes the number of incumbent exporters who were active before time 0 (red region), and those who enter after time 0 (blue region). We have superimposed lines that show how the shaded areas would have shifted if firms had known their true θ^f 's with certainty (dashed red lines), and if they not only knew their θ^f 's, but also had maximum visibility (dashed green lines).

Several messages emerge. First, when firms know their true θ^f from the start, the share of new matches that goes to new exporters is substantially higher. Why? firms with high θ^f 's dominate entering cohorts, and these firms search more intensely. Firms that were incumbent in period 0 also tend to have high θ^f 's and invest heavily in search, but since many of these firms were already well established, revealing their θ^f draws to them doesn't change their behavior as much.

Second, when firms start out knowing their θ^f draws, giving everyone maximum visibility tends to increase share of new matches with new exporters. The simple reason is that in the benchmark model new exporters have less visibility than incumbents. This experiment eliminates their disadvantage. Visibility effects aren't as important as learning effects.

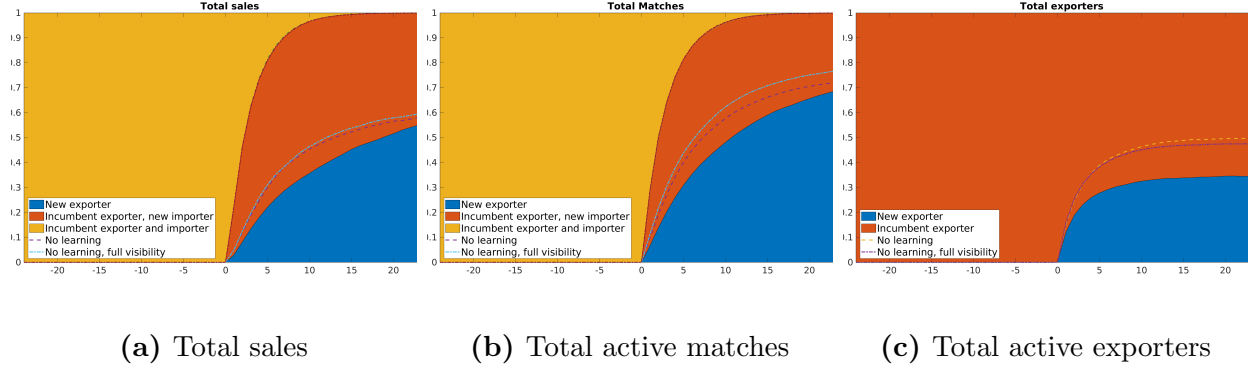
Finally, regardless of whether all firms have maximum visibility or full knowledge of their θ^f draws, the matches of entering cohorts displace incumbent matches at almost exactly the same pace. That is, the border between the orange and yellow regions nearly coincides with the dashed red and green lines in all panels. This sameness reflects the fact that incumbent matches are unaffected by either beliefs about θ^f or search costs.

In Table 23, we compare the short-, medium-, and long-run trade elasticities across models. The first three columns present results for the benchmark model identical to those

³⁹In our simulations of the benchmark model, the maximum number of successful matches is approximately $a = 40$ successful matches. So to characterize full visibility, we replace equation (10) with

$$c^m(s^m, 40) = \kappa_0^m \frac{[(1 + s^m)]^{\kappa_1} - 1}{\kappa_1 [1 + \ln(1 + 40)]^\gamma}. \quad (11)$$

Figure 7: Responses to a permanent shock: baseline vs. other specifications



Notes: Figures depict aggregate responses to a permanent 20 percent real devaluation at time 0. Shaded areas in panels a and b reflect contributions of matches that existed at time 0 (yellow), matches formed after time 0 by exporters that were active at time 0 (red), and matches formed after time 0 by exporters that entered the foreign market after time 0 (blue). Thin dashed lines show patterns that would have obtained if learning effects (red) or both learning and network effects (green) had been absent. All series are averages across 2000 simulations of the exchange rate process.

Table 23: Simulated Trade Elasticities: Comparing across models

Favorable Time since shock	Baseline			Known θ^f			Known θ^f & Full visibility		
	1 year	5 years	25 years	1 year	5 years	25 years	1 year	5 years	25 years
Sales	1.20 (0.05)	1.69 (0.05)	1.88 (0.05)	1.18 (0.04)	1.66 (0.04)	1.65 (0.04)	1.22 (0.03)	1.63 (0.03)	1.67 (0.03)
Matches	0.26 (0.03)	0.80 (0.03)	0.94 (0.03)	0.27 (0.02)	0.76 (0.02)	0.85 (0.02)	0.29 (0.02)	0.72 (0.02)	0.76 (0.02)
Exporters	0.14 (0.02)	0.50 (0.02)	0.60 (0.02)	0.19 (0.02)	0.51 (0.02)	0.61 (0.02)	0.16 (0.02)	0.49 (0.02)	0.55 (0.02)

Notes: All elasticities are based on 2000 simulations of favorable and unfavorable 20 percent changes in the mean real exchange rate. Standard errors based on cross-simulation standard deviations are in parentheses.

in Table 13, the middle three columns present results for the known- θ^f variant, and the last three present results for the known- θ^f variant with full-visibility (network effects exhausted).

The results aren't very sensitive to our assumptions regarding learning or visibility, though the known- θ^f model generates a bit *less* long run responsiveness to positive shocks, mainly because of smaller match elasticities. When potential firms are ignorant of their potential, they search more intensely in response to favorable long-run changes in demand since higher potential profits increase the incentive to learn one's market appeal. As for negative shocks, the associated reduction in entry makes learning less important.

But visibility does matter over the long run. Giving all firms full visibility eliminates the search-cost disadvantage of young firms, moving them away from their exit threshold and dampening extensive margin responses to the shock. Both learning and visibility matter mainly for new, small-scale exporters, so their impact on aggregate export fluctuations is quantitatively small.