A Search and Learning Model of Export Dynamics

(Preliminary and Incomplete)

Jonathan Eaton, Marcela Eslava, C. J. Krizan, Maurice Kugler, and James Tybout[1]

February 2010

[1]We gratefully acknowledge support from the National Science Foundation (Grant SES-0922358), the United States Census Bureau, and Banco de la República de Colombia. We also thank Monica Hernández, Yanping Liu and Van Voung for excellent research assistance, as well as Enrique Montes for expert data advice. 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 U.S. Census Bureau.

1 Introduction

Economists have yet to develop models that reliably explain export dynamics at the micro level. Traditional gravity models focus on long run determinants of aggregate bilateral export flows, and are poorly-suited for the analysis of firm-level export fluctuations.[1] Sunk-cost hysteresis models—which emphasize the start-up costs that new exporters face—do help us understand patterns of foreign market entry and exit by individual firms (Dixit, 1989; Baldwin and Krugman, 1989; Das, et al, 2007). But they provide little guidance as to why new exporters either exit or rapidly expand, while established exporters' sales are stable. Nor do they convincingly reconcile the sustantial market entry costs that they posit with the fact that many firms export for short periods on a very small scale. Finally, while recent work by Arkolakis (2007, 2009) accounts for small-scale exporters and the age-dependence of export growth rates, it lacks the market frictions needed to explain the lags and irreversibilities observed in firms' exporting behavior

This paper develops a model that explains small-scale exporting, age-dependent export growth, and lags and irreversibilites. It is based on the conjecture that firms' exporting behavior reflects search and learning processes in foreign markets. That is, producers who are interested in a particular foreign market devote resources to identifying potential buyers there. When they find one, they learn something (receive a noisy signal) about the appeal of their products in this market. They also learn about foreign demand for their product from their experiences in their home markets. Taking stock of the available information, these firms update their beliefs concerning the scope for export profits, and they adjust the

[1] Recent contributions to the gravity literature include Helpman et al (2008) and Anderson and van Wincoop (2003). Deardorff (1998) provides a survey of the earlier literature.

intensity of their search efforts accordingly, attempting to maximize their net expected profit streams. Export surges take place when home-market firms receive positive early signals about the scope for profits—both from their own experiences and from the experiences of rivals—and they intensify their search/marketing efforts, adding quickly to their foreign client base. Export collapses occur when firms allow their portfolio of buyers to shrink.

The motivation for this paper comes from descriptive analysis of a decade's worth of individual merchandise shipments from Colombia to the United States. We begin by reviewing the stylized facts that come out of this analysis, including a number of finding that we have not reported in our earlier work (Eaton et al, 2008). Then we introduce our model, discuss its calibration, and demonstrate that, by adopting the assumptions mentioned in the previous paragraph, we are able to explain the basic features of the shipments data.

2 Firm-Level Trade: Transaction Level Evidence

The emprical motivation for our model comes from two sources. The first is a comprehensive data set that describes all shipments from Colombia to the United States (and elsewhere) that passed through Colombian customs during the period 1996-2005. Each customs record includes a date, the US dollar value of the product shipped, a 6-digit harmonized system product code (augmented by addition product information), a quantity index, a seller ID code, and the location of the buyer.[2] The second data base provides analogous information for the period 1992-2005. However it is based on U.S. Customs records, and it describes

[2] Because we use the same data that are used for official statistics, the merchandise exports in our data set aggregate to within one percent of total merchandise exports reported by the Colombian Bureau of Statistics (Departamento Administrativo Nacional de Estadística or DANE). The deviation is due to mistakes in the records of tax identifiers. Since following firms over time is central to our analysis, our database includes only records of transactions in which the tax identifier has the appropriate format. Not satisfying this requirement is a clear indication that the firm is not correctly identified in the record.

imports by buyers in the United States from Colombian exporters (as well as other sources). Critically, in addition to providing all of the information contained in the Colombian records, the U.S. customs data include ID codes for both sellers and buyers. It therefore allows us to identify the formation and dissolution of business relationships between individual buyers in the U.S. and sellers in Colombia, hereafter referred to as "matches."

2.1 Evidence from Colombian Customs Records

Following Brooks (2006) and Eaton et al. (2008), Table 1 provides various annual measures of Colombian exports to the United States for the years 1996-2007.[3] Each column follows an exporting cohort—i.e., a group of firms that began exporting in a particular year, after at least one year of no exporting—from the year of its appearance through time. (Since we don't know the history of firms before 1996, the 1996 "cohort" consists of all firms present that year regardless of when they began exporting.) The panels of the Table report number of exporters, total exports, and exports per firm, respectively.

[3] Similar tables for Colombian exports to all destinations combined appear in Eaton, et al, 2008.

Table 1a: Number of Exporting Firms, by Entry Cohort

	Year of entry into U.S. market									
Year	**1996**	**1997**	**1998**	**1999**	**2000**	**2001**	**2002**	**2003**	**2004**	**2005**
1996	10,517	0	0	0	0	0	0	0	0	0
1997	4,414	6,049	0	0	0	0	0	0	0	0
1998	3,306	1,002	3,389	0	0	0	0	0	0	0
1999	2,718	617	938	2,492	0	0	0	0	0	0
2000	2,539	552	761	938	2,847	0	0	0	0	0
2001	2,418	523	700	735	1,113	3,348	0	0	0	0
2002	2,260	484	632	621	833	1,156	3,116	0	0	0
2003	2,200	465	578	553	697	903	1,048	3,655	0	0
2004	2,089	435	528	519	637	759	859	1,131	4,377	0
2005	2,051	420	362	407	505	568	578	769	1,000	5,060

Table 1b: Value of Exports, by Entry Cohort (millions of $US)

	Year of entry into U.S. market									
Year	**1996**	**1997**	**1998**	**1999**	**2000**	**2001**	**2002**	**2003**	**2004**	**2005**
1996	10,651	0	0	0	0	0	0	0	0	0
1997	11,182	369	0	0	0	0	0	0	0	0
1998	10,053	361	477	0	0	0	0	0	0	0
1999	10,514	421	392	241	0	0	0	0	0	0
2000	11,723	475	335	377	207	0	0	0	0	0
2001	10,373	483	296	395	525	233	0	0	0	0
2002	10,049	422	286	362	406	240	136	0	0	0
2003	10,651	490	358	381	546	228	222	251	0	0
2004	13,547	442	409	342	600	366	269	329	427	0
2005	16,207	725	451	588	891	435	295	349	585	665

Table 1c: Exports per Firm, by Entry Cohort (thousands of $US)

	Year of entry into U.S. Market									
Year	**1996**	**1997**	**1998**	**1999**	**2000**	**2001**	**2002**	**2003**	**2004**	**2005**
1996	1013	0	0	0	0	0	0	0	0	0
1997	2533	61	0	0	0	0	0	0	0	0
1998	3041	360	141	0	0	0	0	0	0	0
1999	3868	683	418	97	0	0	0	0	0	0
2000	4617	861	440	402	73	0	0	0	0	0
2001	4290	923	423	537	471	70	0	0	0	0
2002	4446	872	452	584	487	208	44	0	0	0
2003	4841	1053	620	689	783	252	212	69	0	0
2004	6485	1016	776	658	942	482	313	291	98	0
2005	7902	1725	1247	1444	1764	766	510	454	585	131

Consider panel 1a first. Naturally, each cohort's membership falls as it matures. Note that there is very high attrition the first year, with at least half and up to three-fourths of firms dropping out. Conditional on making it to the second year, the survival probability is much higher, however, with an average attrition rate around 20 percent. Thus, in terms of numbers, the most recent cohort is always larger than any previous one (excepting the 1996 cohort, which is a special case). Note that firms that were exporting to the United States in 1996 account for only about one seventh of the firms exporting to the United States in 2005.

Panels 1b shows that, despite the rapid initial decline in its membership, the total sales of a cohort tends to rise over time, although quite unevenly. By the end of the period the 1996 cohort contributes about 76 percent of total sales, despite their relatively small number. The 2005 cohort contributes the second largest share.

The decline in number of firms per cohort along with their increasing contribution to total sales means, of course, that sales per firm are growing substantially (panel 1c). In fact, export sales for young survivors in each cohort tend to grow far more rapidly than total export sales, so that cohorts' market shares tend to rise *despite* rapid attrition during their early years. Finally, note that cohort size and success (in terms of survival and sales) vary substantially across entry years. For example, the 2000 cohort appears very robust both in terms of number of exporters and exports per firm, with 1998 weak by comparison. This suggests that entry selection mechanisms vary over time in response to market-wide forces.

2.2 Evidence from U.S. Customs records

Individual buyers and sellers are identified in the transaction level data collected by the United States Census Bureau. Accordingly, this data set allows us to keep track of how many buyers each Colombian exporter is shipping to, and to see when buyers are dropped or added. We next use these data to characterize the buyer-seller matchings that took place during our sample period of 1992-2005.

2.2.1 Monogamous and polygamous matches

The number of Colombian exporters appearing in the sample grew from 3,742 in 1992 to 5,297 in 2005, a growth of 3.5 per annum, while the number of U.S. importing firms grew by 4.4 percent (Table 2). The number of Colombian exporter-U.S. importer pairs (representing at least one transaction between them in a year) grew at an annual rate of 3.3 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. However, since the remainder of the matches are polygamous, the average Colombian exporter was involved in around 1.4 relationships with U.S. firms while the average U.S. buyer was involved in around 4 relationships with Colombian firms. Both figures declined slightly over the period.

Table 2

	Colombian Exporters	**U.S. Importers**	**Exporter-Importer Pairs**
1992	3,742	1,265	5,297
2005	5,297	2,214	8,046

2.2.2 Transition Probabilities

Most matches are very short-lived. Of the buyer-seller matches that existed at the beginning of the period, 47 percent didn't make it to 1993. But of those that survived into that year, almost 70 percent made it into the next year. Similarly, of the relationships that existed in 2005, 48 percent started that year, but of those that started the previous year, 65 percent had been around at least 3 years before. Of the 5,297 matches identified in 1992, only 85 endure (are present every year) throughout the period.

Table 3 reports the probability with which a Colombian firm participating in certain number of relationships with buyers transits into different number of relationships the following year. This table reports the annual average for 1992-1997 across all industries. Numbers for later periods are very similar. Thus, of firms not exporting to the United States in year t but exporting in year $t+1$, 92.5 percent sell to only one U.S. firm, etc. Of those that sell to one U.S. buyer in a year, 63 percent don't export the next year,while only about 6 percent go on to establish a larger number of relationships. For firms with two relationships in a year, about 14 percent enter into a larger number of relationships, etc. Hence there is an enormous amount of churning at the lower end. Even for firms with a large number of relationships the most likely outcome is to have fewer the next year.

Table 3: Transition Probabilites, Number of Clients

t+1\t	**0**	**1**	**2**	**3**	**4**	**5**	**6-10**	**11-25**
0	**0.000**	0.630	0.265	0.153	0.050	0.024	0.039	0.000
1	0.925	**0.310**	0.344	0.246	0.131	0.079	0.039	0.000
2	0.056	0.046	**0.244**	0.222	0.202	0.211	0.087	0.000
3	0.012	0.010	0.096	**0.186**	0.223	0.168	0.082	0.000
4	0.004	0.003	0.031	0.116	**0.165**	0.184	0.117	0.000
5	0.002	0.000	0.012	0.045	0.108	**0.105**	0.169	0.380
6-10	0.002	0.000	0.004	0.016	0.113	0.205	**0.429**	0.620
11-25	0.000	0.000	0.004	0.016	0.009	0.024	0.039	**0 000**
sum	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000

We can ask what this pattern of entry and growth implies about the ergodic distribution of relationships. If we assume that the number of entrants in a year replace exiters to the extent that the overall number of firms rises by 3.5 percent a year, the ergodic distribution implied by this transition matrix is given by:

Table 4

	1	**2**	**3**	**4**	**5**	**6 to 10**	**11 to 25**
ergodic	.809	.109	.039	.019	.010	.013	.002
period average	.800	.114	.041	.020	.010	.012	.003

For purposes of comparison, the year-specific average share of Colombian firms in each group is reported as well. Note that the ergodic distribution implied by the transition matrix is very close to the distribution in the data.

2.2.3 Match maturation

The survival probability of new matches increases with initial sales volume. Table 5 sorts observations on matches according to their size in their first year of existence and reports year-to-year separation rates. In addition to the very low survival rates, two patterns stand out. First, those matches that begin with sales in the top quartile among all new matches are more likely to survive than matches that begin with smaller sales volumes. Second, survival probabilities improve after the intial year, especially for the surviving matches from the smallest quartile.

Table 5: Separation Rates, Age of Match, and Initial Sales*

	Age of match (in years)				
Inital sales volume	**1**	**2**	**3**	**4**	**5+**
1st quartile	62.0	50.7	48.4	51.1	45.6
2nd quartile	57.3	47.7	42.6	43.5	46.3
3rd quartile	50.6	44.0	41.4	41.3	44.9
4th quartile	47.1	38.2	45.1	44.5	42.4

*Figures exclude matches between affiliated parties, as well as those that involve trade in minerals, unprocessed agricultural products, and services.

Further features of the match maturation process are evident in Figure 1, which shows the log of annual sales per match, broken down by initial size quartile. For each size quartile, matches are further distinguished according to their life span: less than one 1, 1 to 2 years, and so forth. And for each cluster of bars, the left-most bar corresponds to sales in the initial year of the match's existence, the next bar corresponds to sales during the second year of the match's existence, and so forth.

The first message of these graphs is that initial sales are good predictor for sales in subsequent years, conditioned on survival. Those matches with first-year sales in the smallest quartile systematically generated the lowest annual sales in subsequent years, and more generally, first-year sales are monotonically related to annual sales in subsequent years. Second, there is a strong tendency for first-year sales to fall below sales in subsequent years, partly because observations on a match's first year correspond to less than a full calendar year. (There is an analogous effect at work in the final year of a match's life.) Third, however, after the first year there is no general tendency for annual sales to grow in any of the size quantiles. 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.

3 A Model of Exporting at the Transactions Level

We propose a model that is consistent with the patterns documented in the previous section, and that provides new micro foundations for export booms. It explains firm-specific export adjustments on three margins: clients (buyers) per destination market, per-period sales per client, and duration of the buyer-seller relationship.The model captures four key patterns documented above: (1) many new exporting firms appear each period; (2) most new exporters sell tiny amounts and disappear from export markets in the following period; (3) those exporters who survive expand their export volume very rapidly over the following period, often accumulating additional buyers; and (4) firms that sell more initially are more likely to survive into the following period. It also explains (5) inter-temporal fluctuations in the size of the entering cohort, and (6) market-wide and relationship-specific fluctuations in per-period sales volumes.

The model builds on existing models of firm heterogeneity and exporting. As in Melitz (2003) and Bernard et al. (2003), firms are heterogeneous in terms of their underlying efficiency, with more efficient firms having greater incentive to overcome trade costs to sell in foreign markets. As in Das et al. (2007) and Irarrazabal and Opromolla (2007) firms experience shocks to their efficiency that lead them to switch into or out of exporting. As in Arkolakis (2008), by incurring a larger fixed cost a firm can increase the number of buyers it can reach. Finally, as in Rauch and Watson (2003), learning takes place after matches are made.

What we add to these models is a characterization of decision-making and learning by *exporters*. Before it enters an export market a firm is unsure of the appeal that its product has to buyers there. However, the firm can invest in activities that bring its product to the attention of individual buyers, such as advertising, participation in trade fairs, and maintenance of a foreign sales office. The more a firm spends on these activities, the more likely it will encounter a foreign buyer per unit of time. And when a match does occur, its sales not only generate a profit for the firm, it conveys information to the firm about its product's appeal in that market. On the basis of this information the firm updates its beliefs about its product's ultimate chances for success in that market. Good news means that future matches are likely to be more profitable, so it strengthens its efforts to encounter buyers, while bad news discourages the firm from putting in so much effort.

3.1 Profits

To characterize the profit flow, consider firm j with an efficiency φ_{jt} (taking into account transport costs) at time t. This efficiency is known to the firm and evolves over time with idiosyncratic shocks. Given that it pays a wage (or more generally, unit input price) w_t it can produce at cost w_t/φ_{jt} in terms of local currency. If the exchange rate is e_t, its unit cost in the foreign market is $e_t w_t/\varphi_{jt}$. So assuming that all foreign buyers have Dixit-Stiglitz preferences with known demand elasticity η, seller j offers price:

$$p_{jt} = \frac{\eta}{\eta - 1} \frac{e_t w_t}{\varphi_{jt}} \tag{1}$$

to any foreign buyer i with whom it matches.[4]

If potential buyer i is confronted with an opportunity to purchase firm j's product in period $\tau \leq t$, that is, if j matches with i in τ, its period t sales to i conditioned on match survival are:

$$X_{ijt} = \exp(z_j^f + \epsilon_{ij}) \left(\frac{p_{jt}}{P_t^f} \right)^{1-\eta} \overline{B}_t^f. \tag{2}$$

Here we introduce the market-wide spending levels among potential buyers, $\overline{B}_t^f$, a price index for all competing products in the destination market, P_t^f.[5] Also we measure the appeal of firm j's product to buyer i with $z_j^f + \epsilon_{ij}$, which includes a component that is general to all

[4]For simplicity we assume that the firm makes a take-it-or-leave-it price offer. An alternative specification would introduce bilateral bargaining between buyer and seller, although the seller's uncertainty about the buyer's evaluation of the product renders this second approach substantially more complicated. Drozd and Nozal (2008) incorporat this type of bargaining in their model.

[5]Not all buyers necessarily face the same range of goods and hence the same aggregate price index for competing goods. We treat idiosyncratic components of the price index as P_t^f as reflected in ϵ_{ijt}.

buyers and a component that is idiosyncratic to buyer i.

The flow profit in home currency implied by (1) and (2) is:

$$\begin{aligned} & \pi(P_t^f, P_t^d, e_t, z_j^f, \epsilon_{ij}, \varphi_{jt}) \\ = \; & \frac{1}{\eta}\frac{\overline{B}_t^f}{e_t P_t^d}\exp(z_j^f + \epsilon_{ij})\left(\frac{e_t w_t \eta/(\eta-1)}{\varphi_{jt} P_t^f}\right)^{1-\eta}, \end{aligned} \tag{3}$$

where P_t^d is the price level in the domestic (home) market. Or, combining all the aggregate variables and constants:

$$\pi(B_t^f, z_j^f, \epsilon_{ij}, \varphi_{jt}) = B_t^f \exp(z_j^f + \epsilon_{ij})\varphi_{jt}^{\eta-1} \tag{4}$$

where $B_t^f = \frac{1}{\eta}\frac{\overline{B}_t}{e_t P_t^h}\left(\frac{e_t w_t \eta/(\eta-1)}{P_t}\right)^{1-\eta}$ captures the market-wide forces that influence the payoff to all matches in the foreign market. We assume that B_t^f and φ_{jt} evolve over time according to a Markov process, so that given (B_t, φ_{jt}) in period t, the period $t+1$ values have a joint distribution $G(B^{f\prime}, \varphi' | B_t^f, \varphi_{jt})$.

For purposes of the dynamic optimization problem to be introduced below, it will be convenient to define $\widetilde{\pi}_0(B_t^f, z_j^f, \varphi_{jt})$ as the expected present value of firm j's entire profit stream associated with a new match as perceived at time t, conditional on $(B_t^f, z_j^f, \varphi_{jt})$. That is, $\widetilde{\pi}_0(B_t^f, z_j^f, \varphi_{jt})$ is the discounted expected value of the $\pi(B_t^f, z_j^f, \epsilon_{ij}, \varphi_{jt})$ trajectory from period t forward, with expectations taken over ϵ_{ij} and the future trajectory of (B_t^f, φ_{jt}).

In addition to its arguments, the value of $\widetilde{\pi}_0(B_t^f, z_j^f, \varphi_{jt})$ depends on the firm's discount rate r, the rate at which matches terminate for exogenous reasons, δ, and the per-period fixed cost F that firms must pay to maintain each existing client relationship. More precisely, in period t the present value of a relationship that began in period $\tau \leq t$ is:

$$\begin{aligned} & \widetilde{\pi}_{t-\tau}(B_t^f, z_j^f, \varphi_{jt}) \qquad (5) \\ = \quad & B_t^f \exp(z_j^f + \sigma_\epsilon^2/2)\varphi_{jt}^{\eta-1} \\ & + \frac{1-\delta}{1+r} \max\left\{ \int_{B'.} \int_{\varphi'} \widetilde{\pi}_{t+1-\tau}(B^{f\prime}, z_j^f, \varphi') dG(B', \varphi' | B_t, \varphi_{jt}) - F, 0 \right\}. \end{aligned}$$

Thus $\widetilde{\pi}_0$ can be recovered by evaluating (5) at $\tau = t$.

This formulation of match pay-offs has several desirable features. First, once matches are formed, sales continue to fluctuate in response to market-wide shocks B_t and shocks to the exporter's efficiency φ_{jt}. Second, as these fluctuations occur, matches disolve endogenously if their continuation value falls below the fixed costs of maintaining them, F. Finally, matches that generate low sales volumes are relatively likely to fail after a single period, and those low sales matches that survive the first period are relatively likely to fail in the future. So the model is capable of capturing both the association between initial sales and match survival and the rising survival rates documented in table 5.

At the same time that firms are matching with buyers in foreign markets, they are doing so at home. We assume that (5) characterizes the payoff to these home market matches as well, replacing B_t^f and z_j^f with their domestic market counterparts, B_t^d and z_j^d. Given that tastes are correlated across countries we expect that $cov(z_j^f, z_j^d) \neq 0$. Also, since Colombian factor prices and the real exchange rate affect profits for all firms in both markets, $cov(B_t^f, B_t^d) \neq 0$.

3.2 Information about product appeal

In addition to generating profits, each match conveys information to an exporting firm about its product's appeal to foreign consumers, and thereby affects its efforts to locate more buyers abroad. We consider two ways that this might occur. Our first treatment of learning presumes that firm j does not know the average level of demand for its product in the destination market, z_j^f, but it uses information from its previous matches and its home market experiences to formulate beliefs about this variable. Our alternative treatment presumes that some fraction of matches are rejected by buyers after receiving a sample shipment, and that firms learn about this fraction through experience. This learning is based on their history of successes with matches in the destination market, as well as their success rates in establishing business relationships in their home market. We now develop the details of each type of learning.

3.2.1 learning about the scale of demand per buyer

Under our first representation of learning, we assume that exporting firms are able to observe the market-wide aggregate B_t^f and their own productivity, φ_{jt}. Hence, after making its first sale to buyer i, firm j can use (2) to infer the associated demand shifter $s_{ij} = z_j^f + \epsilon_{ij}$. This statistic serves as a noisy signal of its product appeal z_j^f in the foreign market, and thereby affects its search intensity. [6]

More precisely, before it has met any foreign buyers, firm j's beliefs concerning z_j^f are based solely on its home market product appeal index, z_j^d, which we assume has been revealed to it through many matches with domestic buyers. Given that z_j^f and z_j^d are jointly normally distributed with zero means across the population of firms, these prior beliefs are distributed $N(\alpha z^d, \sigma_\nu^2)$ where $\alpha = cov(z^f, z^d)/var(z^d)$ and $\sigma_\nu^2 = var(z^f - \alpha z^d)$. However, each time a firm matches with a foreign buyer it learns something about its product's appeal to foreign consumers. Let the buyer-specific component of foreign product appeal, ϵ_{ij}, be distributed

[6] It would be possible to treat ξ as unobserved to the exporter. This would complicate the learning process, however, since exporters when then continue to extract information from all active matches. We feel the benefits of this extra complexity do not warrant the computational costs.

$N(0,\sigma_\epsilon^2)$ across the population of possible matches Then after meeting n foreign buyers, a firm's posterior beliefs concerning z^f are distributed $N(\widehat{z}^{f,n},\sigma_n^2)$ where:

$$\widehat{z}^{f,n} = \alpha z^d \frac{\sigma_\nu^{-2}}{\sigma_\nu^{-2}+n\sigma_\varepsilon^{-2}} + \overline{s}^n \frac{n\sigma_\epsilon^{-2}}{\sigma_\nu^{-2}+n\sigma_\epsilon^{-2}}, \tag{6}$$

$$\sigma_n = \left(\sigma_\nu^{-2}+n\sigma_\epsilon^{-2}\right)^{-1/2}, \tag{7}$$

and for firm j, $\overline{s}^n = n^{-1}\sum_{i=1}^n s_{ij}$.

3.2.2 Learning the probability of product acceptance

Our alternative representation of learning more closely resembles Rauch and Watson's (2003) formulation, in which buyers place small trial orders with matched exporters of questionable quality. Many matches fail thereafter as the buyers examine the sample shipments and learn more about the sellers, but those that survive move on to larger shipments and are less likely to fail in future periods.

To keep the model tractable, we presume that firms know their z^f values by virture of their experiences in their home market, so shipment sizes do not induce learning among sellers. Instead, sellers learn about the fraction of the buyer pool that finds their products appealing. That is, each time j matches with a foreign client, an initial sample shipment occurs and, depending upon the buyer's opinion of this shipment, the match leads to subsequent orders with probability θ_j. Conditioning on the macro state and firm j's productivity, a match at time t that succeeds (i.e., leads to subsequent orders) generates expected period t sales of $B_t \exp(z_j+\sigma_\epsilon^2/2)\varphi_{jt}^{\eta-1}$. As in the previous formulation, matches eventually disolve in response to idiosyncratic buyer or seller shocks.

Learning amounts to discovering the fraction of the pool of potential buyers who are interested in one's product, θ_j. Given n meetings with buyers, the likelihood of k successes is binomial:

$$\Pr(k|\theta,n) = \binom{n}{k}\theta^k(1-\theta)^{n-k}$$

Further, if the seller has priors concerning θ that are Beta-distributed with parameters (a,b), the posterior distribution of θ also has a Beta distribution

$$L(\theta|n,k,a,b) = \frac{\theta^{k+a-1}(1-\theta)^{n-k+b-1}}{B(k+a,n-k+b)}$$

where $B(\cdot)$ is the scaler that makes the area under this density unity.[7] The mean of this posterior distribution is

$$\hat{\theta}(k,n,a,b) = \frac{a+k}{a+b+n} \tag{8}$$

Accordingly, with $\theta-$learning, the *ex ante* expected payoff to matching becomes

$$\hat{\theta}\cdot\widetilde{\pi}_1(B_t,z_j^f,\varphi_{jt}),$$

presuming that there are no profits or losses associated with shipping a sample. (This assumption is easily relaxed.)

3.3 Search intensity

It remains to characterize the optimal search policy for each type of learning. Let firm j experience new matches with hazard λ when it spends $c(\lambda)$ on search activities, where $c(\cdot)$ is

[7] The parameters (α,β) can be based upon firms' experiences in their domestic market, or the distribution of success rates observed in the foreign market, or some combination these two sources of information.

increasing and convex.[8] The optimization problems for our two versions of learning are then as follows.

3.3.1 Learning the level of demand per buyer

Consider first the case in which firms learn about the average level of demand for their product. If a firm receives an average signal of $\overline{s}_n$ after n encounters and calculates the associated $\widehat{z}^{f,n}$ using (6), its value of continued searching in the foreign market is:

$$
\begin{aligned}
& V(\widehat{z}^{f,n}, n, B^d, \varphi) \qquad (9) \\
= \; & \max_{\lambda} \Big\{ -c(\lambda) + \lambda \int_z \widetilde{\pi}_0(B^f, z^f, \varphi) dF(z^f | \widehat{z}^{f,n}, n) + \frac{1-\lambda}{1+r} \int_{B'.} \int_{\varphi'} V(\widehat{z}^{f,n}, n, B^{f\prime}, \varphi') dG(B^{f\prime}, \varphi' | B^f, \varphi) \\
& + \frac{\lambda}{1+r} \int_{B'.} \int_{\varphi'} \int_{\widehat{z}'} V(\widehat{z}^{f,n\prime}, n+1, B^{f\prime}, \varphi') d\Phi(\widehat{z}^{f,n\prime} | \widehat{z}^{f,n}) dG(B^{f\prime}, \varphi' | B^f, \varphi) \Big\}
\end{aligned}
$$

Here r is the discount rate, $F(z^f | \widehat{z}^{f,n}, n)$ is the posterior distribution for z after the n^{th} match, and $\Phi(\widehat{z}^{f,n\prime} | \widehat{z}^{f,n}) = N(\widehat{z}^n, \sigma_{n+1})$ is the posterior distribution for z that the firm expects to prevail after the $n+1^{st}$ match, given $\widehat{z}^{f,n}$.

A simplified version of (9) characterizes firms' search behavior in their home market, since they have already learned their products' appeal to domestic consumers:

$$
\begin{aligned}
& V^d(z^d, B^d, \varphi) \qquad (10) \\
= \; & \max_{\lambda^h} \left[-c(\lambda^d) + \lambda \widetilde{\pi}_0(B^d, z^d, \varphi) \right] + \frac{1}{1+r} \int_{B^{h\prime}.} \int_{\varphi'} V^d(z^d, B^{d\prime}, \varphi') dG(B^{d\prime}, \varphi' | B^d, \varphi)
\end{aligned}
$$

[8] Following Arkolakis (2008), if we think that the market has M potential buyers and sampling occurs without replacement we can generalize the hazard rate to be $\widetilde{\lambda} = \lambda \cdot h(n)$ where $h(n)$ is decreasing in n, bounded on [0,1], and $h(M) = 0$. For example, if the probability of a match is proportional to the pool of potential buyers who have not yet been visited, this function might take the form: $h(n) = \frac{M-n}{M}$. Working against this effect is the possibility that as matches accumulate, a firm's reputation grows, and it becomes *less* costly to reach new custumers. Hence a general expression for $h(n)$ that does not impose a sign on its derivative may be the most appropriate formulation. If this function is identified, it provides a test of Arkolakis (2008).

Two margins of firm-level export response to idioysncratic and market-wide shocks are characterized by these value functions: the present value of sales per buyer, and the number of buyers per firm (which is governed by λ and λ^h).

3.3.2 Learning the probability of post-match product acceptance

For the case in which firms learn about their success rates,the value of search for a firm that has encountered k successes in n foreign matches is:

$$\begin{aligned} & V(n,k,\theta^d,B^f,z^f,\varphi) \\ = \; & \max_{\lambda}\Big\{-c(\lambda)+\lambda\widehat{\theta}\cdot\widetilde{\pi}_1(B^f,z^f,\varphi)+\frac{1-\lambda}{1+r}\int_{B^{f\prime}.}\int_{\varphi'}V(n,k,B^{f\prime},z^f,\varphi')dG(B^{f\prime},\varphi'|B^f,\varphi) \\ & +\frac{\lambda}{1+r}\int_{B'.}\int_{\varphi'}\Big[\widehat{\theta}V(n+1,k+1,B^{f\prime},z^f,\varphi')+(1-\widehat{\theta})V(n+1,k,B^{f\prime},z^f,\varphi')\Big]\,dG(B^{f\prime},\varphi'|B^f,\varphi)\Big\} \end{aligned} \tag{11}$$

Similarly, firm j's value of search in its home market is:

$$\begin{aligned} & V(\theta^d,B^d,z^d,\varphi) \\ = \; & \max_{\lambda}\left[-c(\lambda)+\lambda\theta^d\cdot\widetilde{\pi}_1(B^d,z^d,\varphi)\right]+\frac{1}{1+r}\int_{B^{d\prime}.}\int_{\varphi'}V(\theta^d,B^{d\prime},z^d,\varphi')dG(B^{d\prime},\varphi'|B^d,\varphi) \end{aligned} \tag{12}$$

3.4 Specification for Numerical Solution

To solve our model numerically we parameterize the cost of matching as:

$$c(\lambda)=b\left(\frac{\lambda}{1-\lambda}\right)+f\cdot 1\left[\lambda>0\right],\quad \lambda\epsilon[0,1) \tag{13}$$

where f is the fixed cost of maintaining positive levels of search. Also, we treat shocks to efficiency and macroeconomic shocks as following independent first-order autoregressive processes, so that:

$$\ln \varphi_{jt} = \psi^{\varphi} \ln \varphi_{jt-1} + \nu_t^{\varphi} \tag{14}$$

$$\ln B_t^f = \psi^B \ln B_{t-1}^f + \nu_t^{Bf} \tag{15}$$

$$\ln B_t^d = \psi^{B^h} \ln B_{t-1}^d + \nu_t^{B^d} \tag{16}$$

where:

$$\nu_t^{\varphi} \tilde{} i.i.d.\ N(0, \sigma_{\varphi}^2) \tag{17}$$

$$\left(\nu_t^{B^d}, \nu_t^{Bf}\right) \tilde{} i.i.d.\ N(\mathbf{0}, \Sigma_B) \tag{18}$$

To summarize, the model incorporates seven types of random shocks: cross-firm variation in foreign product appeal z^f, cross-firm variation in home market product appeal, z^d, noise around true product appeal associated with each match, ϵ, shocks to productivity, ν_t^{φ}, and the market-wide shocks, ν_t^{Bf} and $\nu_t^{B^d}$. It is fully described by the search cost function (13), the expression for profit (3), and the Bayesian updating function (6) or (8), from which we can calculate the expected value of a relationship (5), the value function (11) or (12) and a policy function that give the optimal level of search.

3.5 Fitting the z-learning model to data

3.5.1 Efficiency process

To implement the version of our model with learning about z, we require values for a variety of parameters. First, we need to estimate the AR process that governs firm-level efficiency trajectories (14). The Colombian data are unusually well-suited to this task, since they allow us to construct firm-level quantity indices for both inputs and outputs (e.g., Eslava et al, 2004). Nonetheless, these indices are measured in different units at different firms, since each produces its own variety of output with its own input varieties. We therefore sweep out units of measurement by estimating (14) in growth terms:

$$\Delta \ln \varphi_{jt} = \psi^{\varphi} \Delta \ln \varphi_{jt-1} + \Delta \nu_{jt}^{\varphi}. \tag{19}$$

Several econometric issues arise here. One is that the error in a differenced AR(1) model is correlated with the lagged dependent variable, since ν_{jt-1}^{φ} helps determine $\ln \varphi_{jt-1}$. We handle this problem by using Blundell and Bond's (1998) GMM estimator, with $\ln \varphi_{jt-2}$, $\ln \varphi_{jt-3}$ and other twice-lagged plant characteristics (like output and capital stocks) serving as instruments. There is also a selection problem, since disproportionate exit occurs among low-productivity firms. This we handle with Mills ratios based on survival probabilities.. Finally, to recover $var(\nu^{\varphi})$, we note that by (17), $var(\nu^{\varphi}) = \frac{1}{2} var(\Delta \nu^{\varphi})$.

Preliminary estimates are reported in Table 6 below for all plants and for subsets of plants based on exporting status. As is typically the case, differencing the data removes much of the persistence in measured efficiency, but lagged efficiency remains significant. In the present context we interpret this result to imply that permanent differences in product appeal and/or units of measurement account for the observed strong persistence in domestic sales. It should be noted that the reported specfication fails a Wald specification test for the exogeneity of the instrument (twice lagged productivity).

Table 6: GMM estimates of productivity process

	All plants	Non-exporters	Exporters	Continuing Exporters
$\hat{\psi}^{\varphi}$	0.155	0.109	0.112	0.224
$\widehat{se}(\hat{\psi}^{\varphi})$	(0.026)	(0.045)	(0.053)	(0.074)
$se(\Delta\nu^{\varphi})$				
fixed effects	yes	yes	yes	yes
observations	27,608	14,859	8,468	6416
Wald test (p value)	0.000	0.016	0.034	0.016

3.5.2 Market-wide shocks

We also require estimates for the processes that generate $\ln B_t^f$ and $\ln B_t^d$. These we obtain (up to the intercept) using aggregate real consumption of manufactured goods in the United States and Colombia, respectively. Both are expressed in real pesos, so $\ln B_t^f$ incorporates the effects of exchange rate fluctuations.[9] When industries are pooled we allow for industry-specific intercepts and we estimate a differenced form of (15) and (16), as with the efficiency process. When we focus on an individual industry this is of course not necessary.

Estimates based on pooled industry-level data are reported in Table 7 below.

Table 7: GMM estimates of Industry-wide expenditure Processes

	Col. dom. sales	U.S. Sales	U.S. sales $\times$ real exch. rate
$\hat{\psi}^B$	0.300	0.953	0.827
$\widehat{se}(\hat{\psi}^B)$	(0.089)	(0.024)	(0.100)
$se(\nu_t^B)$		0.052	0.135
fixed effects	yes	no	no
observations	149	44	44
Wald test (p value)	0.001	–	–

3.5.3 Remaining parameters

Four types of parameters remain, all of which we identify using the simulated method of moments. First, there are those that characterize the joint distribution of product appeal determinants: z_j^f, z_j^d and ϵ_j. These are all normalized to zero, so the variances of each $(\sigma^2_{z^d}, \sigma^2_{z^f}, \sigma^2_\epsilon)$, plus the coefficient α are sufficient to identify their joint density. Since they govern the varation in sales across buyers for a given seller, the variation in exports across sellers, and covariance bewteen foreign and domestic sales among exporters, key moments are: $var(\ln X_{jt}^d)$,

[9] If both the log exchange rate and the log of domestic expenditures on manfactured goods in the U.S. follow AR1 processes, then $\ln B_t$ is the sum of two AR1's, which is generally an ARMA(2,1) process. To avoid introducing another state variable in our model we treat $\ln B_t$ as a simple AR1, implicitly assuming that both the log exchange rate and the log of domestic expenditures have the same root. This assumption could easily be relaxed at the expense of computational speed.

$var(\ln X_{jt}^f)$, $var(\ln X_{ijt}^f | X_{jt}^f/n_{jt})$ and $cov(\ln X_{jt}^d, \ln X_{jt}^f)$ where X_{jt}^d is firm i's sales in its home market, and X_{jt}^f/n_{jt} is its average sales per foreign client.

Second, there are parameters that govern matching and separation processes: f, F, δ. The fixed costs of maintaining a relationship and the exogenous match destruction rate δ determine observed rates of match separations. Their effects are distinctive in that δ affects all firms equally, while the effect of f declines as $\tilde{\pi}_0$ increases. Hence key moments for identifying these parameters are the level of failure rates and the covariance bewteen failure rates and $\ln X_{jt}$. The fixed costs of searching, F, determine which firms abstain from exploring export markets altogether, so the fraction of firms that never export helps to identify this parameter.

Finally, there are several nuicance parameters: the rate of time preference, r, and the profit function scale parameter. Since rates of time preference are typically poorly identified in dynamic structural models, we follow convention and simply fix r at a plausible value. Given other parameters, intercepts for the $\ln B^f$ and $\ln B^d$ processes are chosen to replicate observed industry-level sales volumes as closely as possible.

3.6 Implications of the z-learning model

Estimation of the parameters is in progress. To give a preliminary sense for the behavior of our model, we use a simulated method of moments estimator to fit the z-learning model to the moments listed in table 8 below. (Model-based simulated moments are reported next to their data-based counter-parts.) The resulting parameter estimates are reported in Table 9, where figures with asterisks were chosen *a priori* rather than estimated.

Table 8: Simulated versus data moments

Data-based versus simulated statistics

Among exporting plants	Data	Model	**Among exporting firms, number of U.S. clients**	Data	Model
E[(exports)/(dom. sales)]	0.120	0.019	one client	0.809	0.800
$corr$[ln(dom. sales), ln(exports)]	0.260	0.258	two clients	0.109	0.159
$corr$[ln(dom. sales+exports), φ]	0.210	0.212	three clients	0.039	0.031
$corr$[ln(exports), φ]	0.063	0.072	four clients	0.019	0.001
			five or more clients	0.020	0.001
For all plants			**Hazard rates by match age**		
$corr$[ln(dom. sales), φ]	0.210	0.199	first year	0.543	0.568
			second year	0.452	0.520
			third year	0.444	0.447
			fourth year	0.451	0.413
			fifth year and older	0.448	0.397

Table 9: Parameters for Simulations

Parameter		*value*
rate of time preference	r	0.05*
rate of exogenous separation	δ	0.081
profit function scale parameter, foreign	s^f	0.502
profit function scale parameter, domestic	s^d	1.50*
fixed cost of searching	f	0.125
fixed cost of sustaining a match	F	1.878
standard deviation of noise in signal	σ_ϵ^2	0.646
standard deviation of product appeal abroad	σ_{z^f}	1.221
standard deviation of product appeal at home	σ_{z^d}	0.939
correlation, home and foreign mkt. appeal	α	0.573
root of efficiency process	ψ^φ	0.155
root of foreign mkt. process	ψ^{B^f}	0.953
root of home mkt. process	ψ^{B^d}	0.700*
standard deviation of efficiency innovation	σ_{ν^φ}	0.100*
standard deviation of foreign mkt. shock	$\sigma_{\nu^{B^f}}$	0.053
standard deviation of home mkt. shock	$\sigma_{\nu^{B^d}}$	0.100*

Overall, the model does a good job of replicating the data moments, and the parameter estimates seem sensible. Among other things, they imply that home market success helps firms to predict their product appeal in foreign markets, and that there are small costs associated with maintaining search activities abroad. It is substantially more costly to maintain a relationship with a client, once it is established, however, and these costs are the main reason that small-scale sellers let their relationships expire. (Match seperation rates range between 0.4 and 0.7, but the probability of an exogenous separation is only $\delta = 0.081$.)

3.6.1 Policy functions

Some implications of our model are summarized by figures 2 through 6. The first panel of figure 2 shows the value of access to foreign buyers that firms perceive after one signal, as a function of the signal they have received. (Expectations are taken over productivity and the macro state, B^f.) Not surprisingly, there is a positive relationship, and firms that receive better signals choose to search more intensively. The second panel of this figure shows how values and search intensities have changed after 5 signals have accrued. (The horizontal axis is the posterior mean after 5 signals, $\widehat{z}^5$.) Note that the value of search has fallen relative to its value after one signal for those firms with low average signals that were initially searching because these signals become increasing precise as experience accumulates. (When five buyers tell you they don't care for your product, there is a good chance that your product has poor market potential.) The last two panels of figure 2 translate search values into match probabilities, and tell the same qualitative story. Below some threshold signal, the return to search is less than the associated fixed cost (f), and so no search takes place. If f were to increase, this cutoff would shift to the right (not pictured).

Figure 3 shows how the policy function characterized in figure 2 translates into behavior for a simulated set of 1,000 firms. Here the horizonal axis is the true z^f value rather than signal. The first panel describes match hazards for a new cohort of firms, none of which has received any signals yet. Since all firms share the same z^h in this simulation, and thus have the same priors, there is no relationship between z^f values and search intensity. However, some firms don't search very intensively because their current productivity level is low. After 5 periods, a relationship between z^f and search intensity emerges, with many low-z^f firms dropping out of foreign markets. This replicates patterns seen in Tables 1 and 4. Note that considerable heterogeneity in behavior remains, given z^f. This reflects both productivity differences and differences in the idiosyncratic features of the buyers (ϵ's) that the exporters have randomly matched with. It also reflects the magnitude of fixed search costs.

3.6.2 Foreign and domestic sales

Figure 4 aggregates the domestic sales and exports of 1000 simulated firms over a 70 year time horizon. All firms are assumed to begin with no matches at home or abroad, so the early years in figure 4 correspond to a transition period during which customer bases are being developed in both markets. One immediate implication is that it can take 10 to 20 years to reach the ergodic distribution for total export sales in a new destination market.

Limiting our attention to the later years, another interesting implication emerges. Export sales amount to 10-20 percent of domestic sales, which is consistent with plant-level survey data from Colombia spanning the period 1986-1996. This is far lower than the ratio of foreign to home profit functions scalers, which is 0.33 ($s^f = 0.50$, $s^d = 1.50$). The reason is that search intensities move in sympathy with expected profits per match, making the ratio of clients abroad to clients at home substantially less than unity. That is, our model explains the border puzzle as partly due to lower search intensities in foreign markets.

3.6.3 Export trajectories

The search frictions in our model lead to a new kind of inertia and hysteresis in export markets, especially for higher quality exporters, who tend to form more durable relationships. Unlike the earlier sunk cost hysteresis literature, which emphasized substantial market entry costs, our formulation accomdates the stylized fact that many small-scale exporters appear and disappear each period. The other formulation which does this, Arkolakis (2009), does so in a frictionless environment and thus does not speak to response lags or irreversibilities in exporter behavior.

Are these features of our model important? Figure 5-6 aggregate the simulated firm-level export trajectories and numbers of foreign clients used to construct figure 4, thereby imputing economy-wide series for each. Both series are in logs and normalized to zero in the initial year. Figure 5 depicts the log of total exports and the log market-wide shifter (B_t^f) through time; the latter can be thought of as mainly reflecting movements in the real exchange rate. Clearly, exports are responsive to the exchange rate, and much more volatile. Partly this reflects the fact that the elasticityof demand is $\eta = 5$, causing small changes in export prices to trigger large changes in demand. But this does not explain why responses are larger in some periods than in others–something that would not occur in a frictionless model. Note in particular the large drop in exports that occurs around period 40, despite the mild appreciation of the exchange rate.

The reason for the relatively dramatic response here is clear in figure 6, which shows the log of the total number of clients, again with the log of total exports. This series drops dramatically around period 40, reflecting the fact that many firms reduced their search efforts and allowed their existing matches to expire. Thus it appears that there are threshold values of expected profitability per match below which many producers dramatically curtail their foreign business relationships. And, unlike most trade models with heterogeneous firms, the fixed costs that create these threshold values influence both large and small-scale exports, since all have matches that are marginally profitable.

For the same reason, exports respond asymmetrically to expansions and contractions. A sufficiently large drop in demand can induce firms to let their customer relationships expire, as occured around period 40. But when foreign demand increases, they must respond mainly by increasinng their sales per customer; they cannot instantaneously increase their clientele. (Note that spikes in demand occuring in periods 15 and 20 led to large percentage-wise expansion in total exports, but much smaller expansions in clientele.)

3.7 Fitting the $\theta-$Learning Model to Data (to come)

3.8 Implications of the $\theta-$Learning Model (to come)

4 Conclusions (to come)

References

Anderson, James and Eric van Wincoop. (2003). “Gravity with Gravitas: a Solution to the Border Puzzle,” *American Economic Review*, 93, pp. 170-92.

Arkolakis, Konstantinos (2007) “Market Access Costs and the New Consumers Margin in International Trade,” University of Minnesota, Department of Economics, Working Paper.

Arkolakis, Konstantinos (2009) “A Unified Theory of Firm Selection and Growth,” Yale University, Department of Economics, Working Paper.

Baldwin, Richard. E. and Paul Krugman (1989): "Persistent Trade Effects of Large Exchange Rate Changes." *Quarterly Journal of Economics*, 104, 635-654.

Bernard, Andrew and J. Bradford Jensen (1999) “Exceptional Exporter Performance: Cause, Effect, or Both?” *Journal of International Economics* , 47: 1-25.

Bernard, Andrew, J. Bradford Jensen, Samuel Kortum and Jonathan Eaton (2003) “Plants and Productivity in International Trade,” *American Economic Review* 93(4), 1268-1290

Bernard, Andrew, J. Bradford Jensen, J. Stephen J. Reading, and Peter K. Schott (1999) “Firms in International Trade,” forthcoming, *Journal of Economic Perspectives*.

Besedes, Tibor. 2006. “A Search Cost Perspective on Duration of Trade,” working Paper, Louisiana State University.

Blum, Bernardo S., Sebastian Claro, and Ignatius Horstmann (2009). "Intermediation and the Nature of Trade Costs: Theory and Evidence." Working Paper, The University of

Toronto.

Brooks, Eileen (2006) “Why don’t firms export more? Product Quality and Colombian Plants” *Journal of Development Economics,* 80: 160-178.

Clerides, Sofronis, Saul Lach and James Tybout (1998) “Is Learning-by-Exporting Important? Micro-dynamic Evidence from Colombia, Mexico and Morocco,” *Quarterly Journal of Economics*, pp. 903-947.

Das, Mita, Mark Roberts and James Tybout (2007) “Market Entry Costs, Producer Heterogeneity and Export Dynamics,” *Econometrica* 75(3), pp. 837-874.

Deardorff, Alan (1998). “Determinants of Bilateral Trade: Does Gravity Work in a Neoclassical World?” in Jeffrey A. Frankel, ed., *The Regionalization of the World Economy*, University of Chicago Press, pp. 7-28.

Drozd, Lukasz A. and Jaromir B. Nosal (2008) "Understanding International Prices: Customers as Capital," Working Paper, The University of Wisconsin.

Dixit, Avinish (1989), "Hysteresis, Import Penetration, and Exchange Rate Pass-Through," *Quarterly Journal of Economics*, Vol. 104, No. 2 (May), pp. 205-228.

Eaton, Jonathan, Samuel Kortum, and Francis Kramarz (2004) “Dissecting Trade: Firms, Industries, and Export Destinations,” *American Economic Review Papers and Proceedings,* 94: 150-154.

Eaton, Jonathan, Samuel Kortum, and Francis Kramarz (2007) “An Anatomy of International Trade: Evidence from French Firms,” Working Paper, New York University, Department of Economics.

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., T*he Organization of Firms in a Global Economy*, Cambridge, MA: Harvard U. Press.

Eslava, Marcela, John Haltiwanger, Adriana Kugler, and Maurice Kugler (2004) "The Effects of Structural Reforms on Productivity and Profitability Enhancing Reallocation: Evidence from Colombia," *Journal of Development Economics,* 75: 333-371.

Helpman, E., M. Melitz and A. Rubenstein (2008). "Estimating Trade Flows: Trading Partners and Trading Volumes." *Quarterly Journal of Economics* (May) 123(2), pp. 441-488.

Irarrazabal, Alfonso A. and Luca David Opromolla (2006) "Hysteresis in Export Markets," New York University, Working Paper.

Kugler, Maurice (2006) "Spillovers from foreign direct investment: within or between industries?" *Journal of Development Economics,* 80(2): 444-477.

Jovanovic, Boyan (1982) "Selection and the Evolution of Industry," *Econometrica,* 50: 649-670.

Luttmer, Erzo (2008) "Selection, Growth, and the Size Distribution of Firms," *Quarterly Journal of Economics.*

Melitz, Marc (2003) "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica* 71, 1695-1725.

Rauch, James and Joel Watson (2003) "Starting Small in an Unfamiliar Environment," *International Journal of Industrial Organization* 21: 1021-1042.

Roberts, Mark and James Tybout (1997a) "The Decision to Export in Colombia: An Empirical Model of Entry with Sunk Costs," *American Economic Review* 87(4), pp. 545-563.

Roberts, Mark and James Tybout (1997b) *What Makes Exports Boom*? Directions in Development Monograph Series, The World Bank, Washington, DC.

Tauchen, George (1986) "Finite State Marko-Chaid Approximation to Univariate and Vector Autoregressions," *Economics Letters,* 20.2. 177-181.

Figure 1: Annual Log Sales by Match Age, Match Duration, and Initial Log Sales

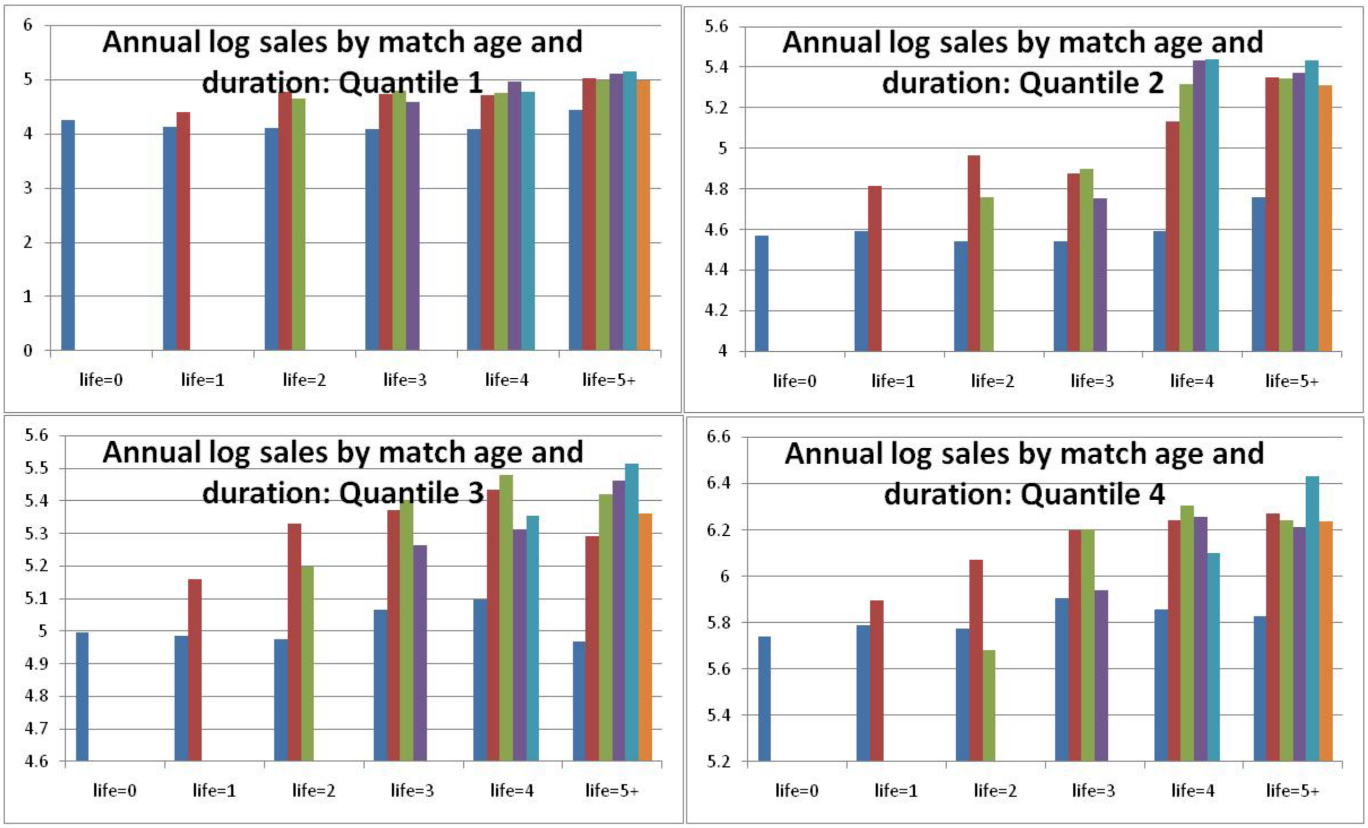

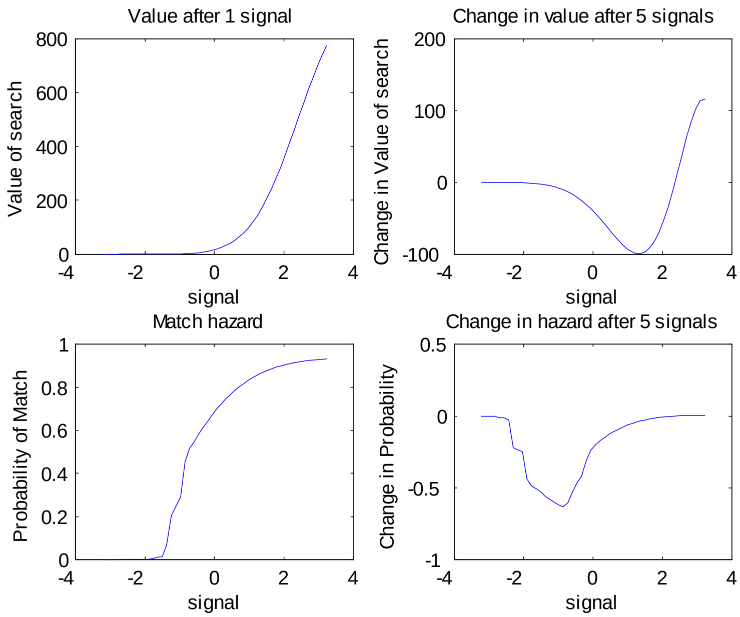

Figure 2:

Signal, value, match hazard, and learning

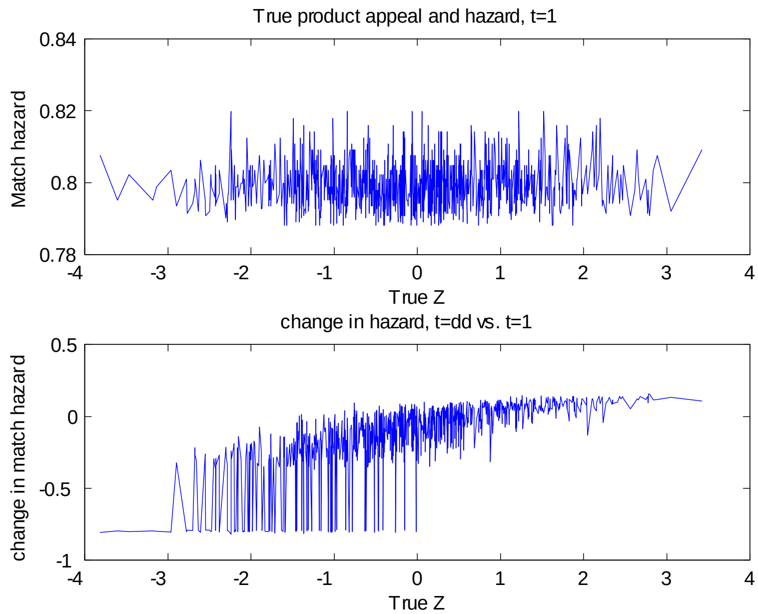

Figure 3:

True product appeal and match hazard: initial and change after 5 signals

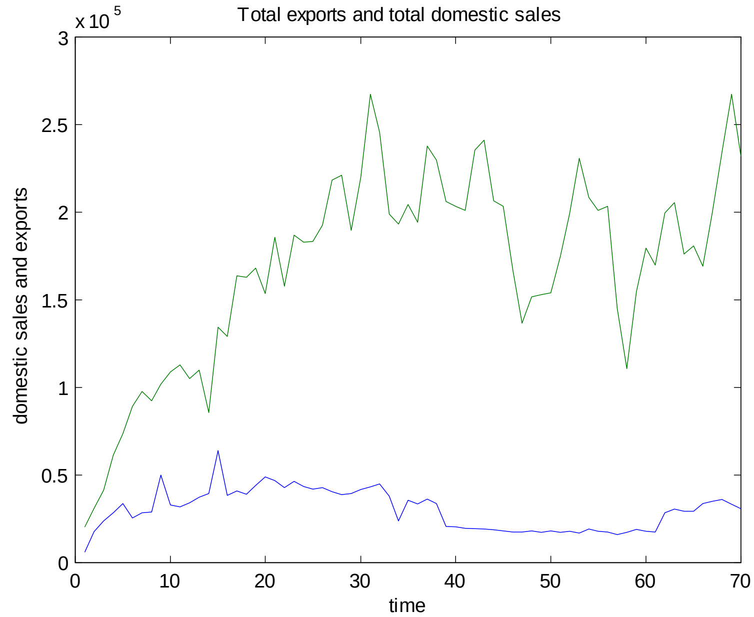

Figure 4:

Total Exports and Total Domestic Sales

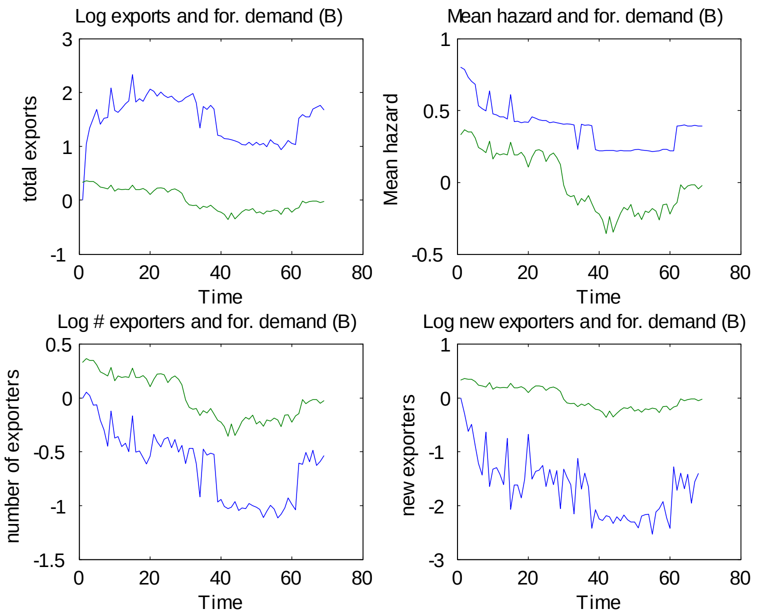

Figure 5:

Total log exports and log foreign market shifter (B)

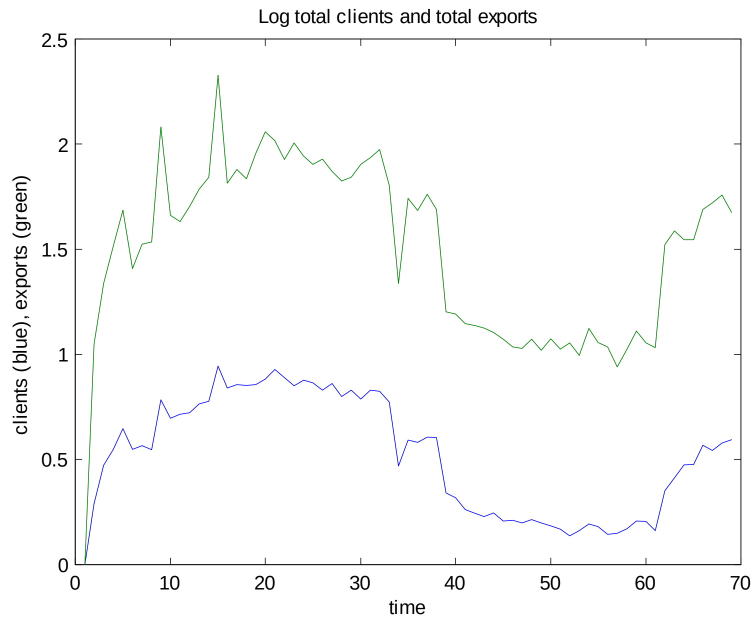

Figure 6:

Total foreign clients and total exports