

New Exporter Dynamics

Revised: March 2016
(First version: August 2005)

Kim J. Ruhl
New York University Stern School of Business

Jonathan L. Willis
Federal Reserve Bank of Kansas City

ABSTRACT

Models in which heterogeneous plants face sunk export entry costs are standard tools in the international trade literature. How well do these models account for the observed dynamics of new exporters? We document that new exporters initially export small amounts and — conditional on continuing in the export market — grow gradually over several years. New exporters are most likely to exit the export market in their first few years. We construct a dynamic discrete choice model of exporting and find that the standard sunk-cost model cannot replicate the behavior of new exporters: New exporters grow too large too quickly and live too long. We assess the quantitative importance of accounting for new exporter dynamics by extending the model to account for these facts. In this model, the present value of exporting falls relative to the sunk-cost model. As a result, the entry costs needed to account for the data are three times smaller than in the sunk-cost model. Our results suggest that dynamic models with richer plant-level heterogeneity are needed.

We would like to thank George Alessandria, Costas Arkolakis, Roc Armenter, and Timothy J. Kehoe for helpful comments. We thank Mark Roberts and Jim Tybout for allowing us access to the Colombian plant data. This work was undertaken with the support of the National Science Foundation under grant SES-0536970. The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Kansas City.

1 Introduction

Models in which heterogeneous plants face fixed export entry costs have become important tools for studying international trade patterns, measuring trade frictions, and evaluating policy.¹ These models, which are typically focused on steady-state analysis, have been successful in replicating several key features of the plant-level data in the cross section, such as the low export participation rate and the fact that exporters are larger than nonexporters.

While the steady-state properties of this class of models have provided insight into the plant's export decision, the presence of a sunk entry cost makes the dynamic content of these models extremely rich.² The key innovation in the dynamic models is that plants enter and exit the foreign market in response to changes in relative prices and productivity. In seminal work, Baldwin (1988), Baldwin and Krugman (1989), and Dixit (1989) demonstrated that the hysteresis implied by these models is important for understanding the effects of large and small exchange rate movements. More-recent work has incorporated this type of industrial structure into stochastic dynamic models disciplined by increasingly accessible plant-level data. For example, Melitz and Ghironi (2005) and Alessandria and Choi (2007a) use these types of models to study how the inclusion of plants' exporting decisions affects real exchange rate and net export dynamics over business cycles; and Ruhl (2008) demonstrates how export entry can produce asymmetric responses to temporary and permanent changes in expected export profits. Das, Roberts, and Tybout (2007) estimate a dynamic structural model of export entry and exit and use it to study the impact of trade policy. These models have typically focused on the aggregate implications of export entry and exit.

In this paper, we ask whether sunk cost models can reproduce the dynamics of new exporters. We begin by using plant-level data on Colombian manufacturers to document two key properties of new exporter behavior. First, we show that new exporters begin by exporting small amounts relative to total production and — conditional on continuing in the export market — gradually expand their export volumes over several years. This can be seen in figure 1a, in which we plot the average export to total sales ratio (the *export-sales ratio*) of the new exporters in our sample. Second, we show that, compared to plants that have exported for several years, new exporters are more likely to exit the export market. In figure 1b, we plot the share of plants that export $a + 1$ years after entry, conditional on having exported for a years. It takes three years for a new exporter's survival rate to level off. Our two findings have been subsequently confirmed in other data: Kohn, Leibovici, and Szkup (2014), for example, document similar patterns among Chilean manufacturers.

Next, we construct a dynamic stochastic model of a plant's export decision and calibrate it to the cross-sectional facts that are commonly used to parameterize fixed entry cost models. We find

¹Notable papers in this expansive literature include: Melitz (2003), Helpman, Melitz, and Rubinstein (2008), and Eaton, Kortum, and Kramarz (2011).

²In this class of models, a plant pays a fixed entry cost to enter the export market (f_0) and a fixed cost in each period thereafter to continue exporting (f_1). The difference between the two ($f_0 - f_1$) is the part of the entry cost that is irrecoverable — the part of the entry cost that is *sunk*.

that the model cannot — either qualitatively or quantitatively — reproduce the behavior of new exporters that we observe in the data. In the model, as opposed to the data, an exporter sells the most in the export market (relative to its total sales) in its first periods of exporting, and there is no gradual growth in export sales in subsequent years. The model is also unable to generate exit rates that decline with the number of years that a plant has exported. In the model, a plant is least likely to exit the foreign market in its first few years of exporting.

The failure of the standard sunk-cost model to capture the dynamics of new exporters is important. The present value of exporting — the key determinant of export entry — is fundamental to understanding both the plant-level and the aggregate responses to trade policy and to measuring export barriers, such as the costs of export entry. The present value of exporting is determined by both how long a plant expects to stay in the export market and when export profits are realized. By generating new exporters that live too long and export too much too soon, the model significantly overestimates the value of exporting. At the plant level, this implies that the estimated export entry costs will be overstated, which we quantify in section 6. Export entry costs fall by a factor of three when we augment the sunk-cost model so that export volumes grow as slowly as they do in the data.³ While the estimated entry cost falls when we account for new exporter dynamics, this does not mean that sunk costs are no longer important. We still estimate that a significant portion of the entry cost is sunk.

Why does the sunk-cost model fail to capture the dynamic properties of new exporters? With fixed entry costs and mean reverting shocks to productivity and the real exchange rate, the model generates selection into exporting when plants are strongest: New exporters are plants with good productivity or real exchange rate shocks, making them larger exporters and the least likely to exit. In the data, however, new exporters are small and face a low probability of success. Reconciling the strong entrants in the standard model with the weak entrants in the data is a fundamental challenge.

The sunk-cost model's failure to generate a gradual increase in export intensity is an implication of the fixed entry cost. Upon entry, a new exporter immediately adjusts its exports to the optimal level — since there is no other barrier to exporting, the plant immediately exports as much as it can. The optimal export-sales ratio is a function of the size of the export market, the elasticity of substitution, and the real exchange rate. The export market size and the elasticity are constant. If the real exchange rate were constant, too, then the export-sales ratio would be constant; a new exporter's export-sales ratio would jump from zero to its optimal level and stay at that level as long as the plant exported. With a variable real exchange rate, the average new exporter will enter when the exchange rate is above average. The exchange rate decays during the new exporter's tenure, generating an export-sales ratio that monotonically decreases. With these functional forms, the model cannot generate the gradual growth in a new exporter's export intensity.

The interaction between the sunk entry cost and the persistent nature of productivity and real

³In the model with stochastic entry costs, 60 percent of entering plants pay no entry cost and 40 percent pay an entry cost that is 61 percent as large as the one in the standard model. See section 6.2 for details.

exchange rate shocks is largely responsible for the model's failure to generate enough exit of new exporters. The sunk nature of the entry cost generates selection not only on a plant's current profitability, but also on its future profitability. Shocks are persistent, so a plant that receives a positive shock expects to be profitable in subsequent periods. This implies that the plants that enter the export market will be the most profitable and the least likely to exit in the few periods after entry, contrary to the patterns we observe in the data. Any parametrization of the standard model in which the entry cost is larger than the continuation cost will generate these counterfactual survival rates.

We make two changes to the standard sunk-cost model that allow it to account for the dynamics of new exporters. These modifications are not deep models of new exporter behavior; rather, our intention is to demonstrate the quantitative importance of getting new exporter behavior correct in the model. Below, we discuss theories that could potentially account for this behavior.

Our first modification is to export demand. We construct a foreign demand function that increases with the number of years a plant has been exporting. We calibrate the model so that the average export-sales ratio of new exporters matches that in the data. Our second modification is to the cost of entry, which we make stochastic. As discussed above, the sunk entry cost and persistent shock process imply that only plants that anticipate being profitable for several periods will enter the market — but we need a model in which some low-profit plants enter the export market. With a stochastic entry cost, some plants will draw very low (in fact, zero) entry costs, inducing entry of plants that otherwise would not have entered.⁴ Some of these plants will not find it optimal to continue exporting in the face of the ongoing exporting costs and will exit.

In this *extended* model, a new exporter starts small and grows over several periods, pushing the profits from exporting into the future. In the future, the favorable shocks that are realized currently will have decayed away, both decreasing profits and increasing the likelihood that the plant will exit the export market. Loading profits into the future decreases the present value of exporting, which is a risky endeavor that pays off only over long periods. The impact of this difference is most striking in the estimated export entry cost in the two models. To replicate the low export participation rate that we observe in the data, the sunk-cost model needs an entry cost more than three times as large as the one in the model with gradual adjustment. This type of result is robust to the underlying model of new exporter dynamics. While we have modified the standard model in a particular way, any model that lowers the present value of exporting for the marginal plant type will lead to lower estimated entry costs.

To further explore the implications of new exporter dynamics, we analyze the aggregate response to a transitory exchange rate depreciation and to a permanent decrease in trade costs. We carry out this analysis in the sunk-cost model and in the extended model that accounts for new exporter dynamics. When the exchange rate depreciates, both models generate an increase in the share of plants that export, but the sunk-cost model generates a ten-percent larger increase in ag-

⁴For plants that draw the low entry cost, their entry costs are smaller than their continuation costs — eliminating the sunk cost that drives the counterfactual behavior in the sunk-cost model.

aggregate export volumes because new exporters in the sunk-cost model are relatively large. In the extended model, the new exporters are relatively small, and, since the depreciation is temporary, most of the new exporters will exit before selling much.

Understanding new exporter dynamics has important policy implications. In contrast to an exchange rate depreciation, the response to a permanent decrease in trade costs is much larger in the extended model than in the standard sunk-cost model. In this case, the permanent decline in trade barriers allows exporters in the extended model to stay in the export market longer. The longer a plant can stay in the export market, the larger are its exports and the greater is its probability of continuing exporting. Thus, the aggregate response in the extended model is not only a function of new exporter entry, but also of the larger per-plant export volumes that come with, on average, more experienced exporters. This partial equilibrium result is further developed in general equilibrium in Alessandria, Choi, and Ruhl (2013).

A large literature establishes the relevance of sunk entry costs for export decisions. Early models, such as Baldwin (1988), Baldwin and Krugman (1989), and Dixit (1989) — and, more recently, Impullitti, Irarrazabal, and Oromolla (2013) — focuses on the hysteresis implied by the sunk nature of entry costs. Empirically, much of the evidence on sunk export entry costs came from reduced-form specifications such as Roberts and Tybout (1997) and Bernard and Jensen (2004), who establish that entry costs are important in accounting for the persistent nature of a plant's export status. Our work is closely related to that of Das, Roberts, and Tybout (2007), who also estimate a structural model of plant export decisions using Colombian data. Their model has many of the same characteristics as ours, but their focus is on the importance of plant-level decisions in shaping the aggregate response to trade policy.

A developing literature proposes structural models of exporting that can account for the dynamics of new exporters that we have documented here. In Nguyen (2012) and Albornoz et al. (2012), a plant's export demand is unknown but correlated across markets and time. Plants learn about their exporting profitability by trying out — possibly without success — new markets. Eaton et al. (2014) develop a model in which an exporter must be matched with an importer; once matched, the pair needs time to learn about the quality of their match. Rho and Rodrigue (2014) develop a model in which costs of adjusting the capital stock can lead to slow export expansion as the plant builds capacity. In Kohn, Leibovici, and Szkup (2014), plants need working capital loans and face collateral constraints. When a plant begins to export, it cannot immediately meet the increased need for working capital, so it takes time for the plant to grow to its optimal size.

The slow growth of entrants into a new market that we document here is not unique to exporters. There is a large literature on the entry of new plants into domestic markets (for example, Foster, Haltiwanger, and Syverson, 2008; 2012) that highlights the differences between new and established plants and finds that entrants are smaller than established industry competitors and that new plants grow slowly. Foster, Haltiwanger, and Syverson (2012) find that the differences between new and established plants are not driven by productivity differences, but by differences in

demand-side fundamentals. Their reading of the data is similar to ours. The productivity-driven sunk-cost model does a poor job of generating the behavior of new exporters, but our extended model — in which demand drives new exporter behavior — can account for the data.

In section 2, we document the two features of new exporter dynamics that are key to evaluating the plant’s benefit from exporting. We specify the standard sunk-cost model in section 3 and parameterize it in section 4. In section 5, we show how this standard framework cannot account for new exporter dynamics. We extend the sunk-cost model and discuss the plant-level implications of getting new exporter dynamics right in section 6, and in section 7, we consider the aggregate implications of accounting for new exporter dynamics. Section 8 concludes.

2 Data

In this section, we lay out two sets of facts. First, we summarize the features of the data that are frequently used to parameterize models based on their cross-sectional characteristics. These include the export participation rate, exporters’ entry and exit rates, and characteristics of the plant size distribution. Second, we highlight the transitional dynamics of new exporters, which is the focus of this paper.

We draw our data from an annual census of manufacturing plants in Colombia. The data were originally collected as a sequence of cross sections by the Departamento Administrativo Nacional de Estadística and were cleaned and linked into a panel, as described in Roberts (1996). The census covers all manufacturing plants with ten or more employees and includes variables about revenues, input costs, employment, and exporting revenue from 1981 to 1991.⁵ This time period, as well as the data we use, were previously studied in Roberts and Tybout (1996), Roberts and Tybout (1997), and Das, Roberts, and Tybout (2007). In this paper, we focus on the decision of an existing plant to enter the export market, so we balance the panel by dropping any plant that did not have at least 15 employees in each year of the sample. We follow Roberts and Tybout (1997) in using plants from 19 manufacturing industries. Finally, we exclude plants that experience large changes in real production or employment. The resulting sample contains 1,914 plants over 11 years. We provide more detail on the sample construction in the appendix.

2.1 Cross-sectional facts

We characterize a plant as an *exporter* in year t if export revenues for the plant are positive. A plant that enters the export market in year t , an *export starter*, is a plant that was not an exporter in year $t - 1$ and is an exporter in year t . An *export stopper* in year t is a plant that exported in $t - 1$ and does not export in year t . In our data, 25.5 percent of plants, on average, are exporters.⁶ That only

⁵Due to the focus on export entry, we drop the initial year, 1981, when computing some moments because of the need for a lag to determine the prior year’s export status of a plant.

⁶We report unweighted statistics. An alternative would be to weight plants by their employment size, so that the outcomes of larger plants are more important — an approach sometimes used in closed economy industry studies. Weighting by employment does not qualitatively change our results, and, quantitatively, the biggest difference is in the export participation rate, which is 48 percent when weighted by employment, reflecting the fact that exporters tend to be larger plants. In the productivity literature, whether or not to weight is still an unsettled matter; see, for example, Foster, Haltiwanger, and Syverson (2008).

26 percent of plants export anything has been an important fact in shaping the development of trade models. Discrete choice models, which have become workhorses in international trade, are ideally suited to generating this cross-sectional fact. If plants are heterogeneous, a large enough fixed exporting cost can shut out a fraction of the plants from the export market.

While many plants do not export, there is significant turnover in the export market: The average *export starter rate* (the number of export starters at t divided by the number of nonexporters at $t - 1$) is 5.2 percent, and the average *export stopper rate* (the number of export stoppers at t divided by the number of exporters at $t - 1$) is 10.6 percent. Modeling this turnover requires a dynamic model with some underlying process that drives entry and exit. Alessandria and Choi (2007b) use plant-level idiosyncratic productivity shocks, and Das, Roberts, and Tybout (2007) introduce idiosyncratic shocks to a plant's export profits. In these papers, as in this one, export entry is driven by the arrival of a favorable shock to a plant's productivity, to the real exchange rate, or to both.

Plants that export are, on average, larger than plants that do not export: This is known as the *exporter size premium*. In our sample, average employment in exporting plants is 2.6 times larger than average employment in nonexporting plants, and the average domestic sales of exporting plants is 3.8 times larger than that of nonexporting plants. This fact is often cited in support of models in which plants are heterogeneous in productivity, and more-productive plants are larger and more profitable. In these models, more-productive plants become exporters, generating an exporter size premium.

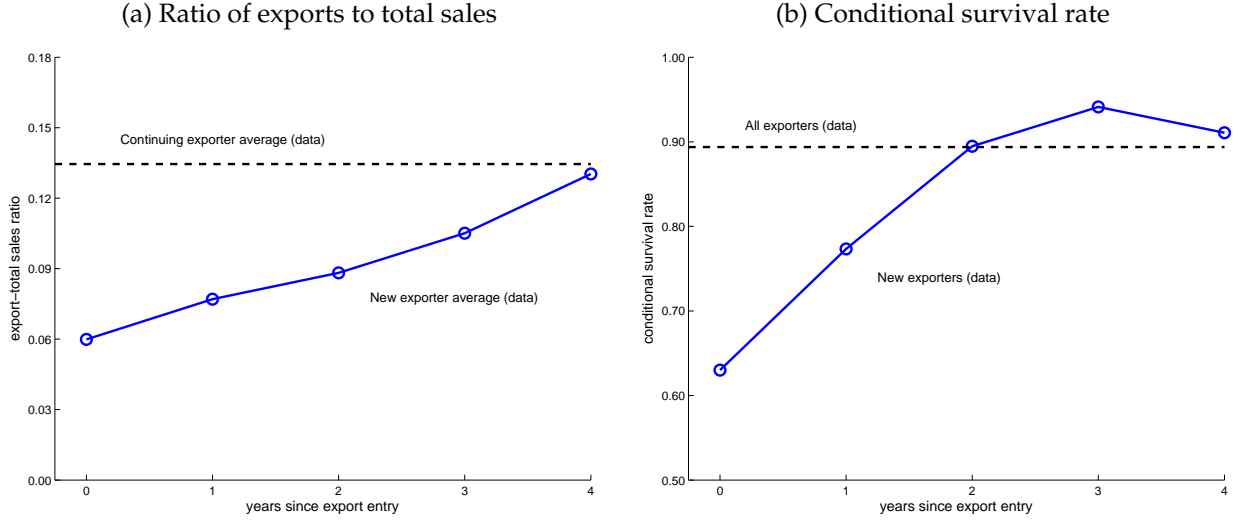
2.2 Dynamics of new exporter growth

The purpose of this study is to evaluate how well heterogeneous plant models replicate the dynamics of the export decision. Here, we focus on one aspect of the decision to export: the plant's decision about how much to produce for the export market. The now-standard models, such as Melitz (2003), feature fixed costs of exporting, which induce a discrete decision regarding entry into the export market. In this class of models, when a plant enters the export market, it immediately adjusts its export quantity to the optimal level. Do we see this pattern in the data?

In figure 1a, we plot the average export-sales ratio for new exporters in the Colombian data. For each plant in the panel that enters the export market, we compute the export-sales ratio of the plant for the year it entered — period zero on the x-axis — and the years following entry, conditional on a plant remaining in the export market for four years following entry. We restrict the sample to include only the plants that continued to export in the four years following entry, so that we are capturing the growth within a plant as it transitions from a new to an established exporter. Based on this criterion, the sample contains 136 entrants. From the figure, we see the discrete nature of the entry decision, as exports jump from zero (by definition of a nonexporter) to about six percent of total sales upon entry.

The discrete jump, however, is not the complete picture, as exports continue to grow over the years following entry. The dashed line in figure 1a is the average export-sales ratio for all

Figure 1: New exporter dynamics.



exporting plants. It takes a new exporter about five years to reach the average export-sales ratio.⁷ In table 1, we report the export-sales ratios that are the coefficients β from

$$\frac{ex_j}{sales_j} = \sum_{a=0}^4 \beta_a I_{ja} + \gamma_c + \gamma_i + \epsilon_j, \quad (1)$$

where I_{ja} is equal to one if plant j is exporting at age a and zero otherwise; γ_c are cohort fixed effects; and γ_i are industry fixed effects. We take the modal cohort (1984) and industry (textiles) as the reference groups. Controlling for the plant's cohort may be important; plants that enter in a better-than-average year may be able to initially sell more in export markets, changing the pattern described in the figure. When we control for an entrant's cohort, however, the export-sales ratio behaves almost identically to the simple averages. When we control for industry effects, the initial export-sales ratio falls to 1.5 percent, compared to the simple average of six percent.

Since our data contain information only on total export revenues, we do not know to how many countries a plant exports. While most exporters export to only one country, as Eaton, Kortum, and Kramarz (2004) document using French data, some of the growth we see in the new exporter export-sales ratio is a result of a plant adding new export markets, as documented in Eaton et al. (2007). Since we cannot separate out the growth in exports that arises from new markets, we will model the plant as choosing to export to a single global market.

Exporting is a persistent condition. In our data, if a plant exports in a given year, there is an 89-percent chance that it will still be exporting the next year. The persistence of export status is evidence in favor of models with sunk entry costs, such as those developed by Baldwin (1988) and Roberts and Tybout (1997). The unconditional survival rate, however, is much different from the

⁷Our data are reported annually, which raises an issue regarding partial-year effects at entry: If a plant begins exporting late in the year, its annual export-sales ratio will understate its export intensity. Bernard et al. (2014), using transaction-level data for Peru, find that the bias in the year of export entry can be substantial. While partial-year effects are likely present in our data, an entrant's export intensity remains below average for several years following entry.

Table 1: New exporter export dynamics: industry and cohort effects.

	Export-sales ratio [†]			Conditional survival rate [‡]		
Entry	5.99 (1.42)	5.52 (1.72)	1.49 (2.40)	0.63 (0.029)	0.66 (0.050)	0.68 (0.091)
Age=1	7.70 (1.42)	7.23 (1.72)	3.22 (2.40)	0.77 (0.025)	0.78 (0.033)	0.80 (0.038)
Age=2	8.82 (1.42)	8.35 (1.72)	4.41 (2.40)	0.89 (0.019)	0.90 (0.022)	0.89 (0.025)
Age=3	10.51 (1.42)	10.05 (1.72)	6.08 (2.40)	0.94 (0.014)	0.94 (0.017)	0.93 (0.019)
Age=4	13.03 (1.42)	12.56 (1.72)	8.47 (2.40)	0.91 (0.017)	0.91 (0.020)	0.89 (0.022)
Cohort Effects	No	Yes	No	No	Yes	No
Ind Effects	No	No	Yes	No	No	Yes
Adjusted R^2	0.24	0.24	0.35			

Notes: [†]OLS coefficients from (1). The sample consists of plants that entered in the years 1982-86 and continued to export for the next four years. [‡]OLS coefficients from the appropriate version of (2). Standard errors are reported in parentheses. The sample includes all exporters.

survival rates of new exporters.

In figure 1b, we plot the survivor rates for exporters. The survivor rate is the share of exporting plants of a given age that continue to export in the following period. We compute it as the coefficient β from

$$I_{ja} = \beta_a I_{j,a-1} + \gamma_c + \gamma_i + \epsilon_{ja} \quad a = 1, \dots, 4, \quad (2)$$

where I_{ja} is equal to one if plant j is exporting at age a and equal to zero otherwise.⁸ We allow for cohort fixed effects, γ_c , and industry fixed effects, γ_i . We report the coefficients from these regressions (one regression for each age) in table 1. The fixed effects regressions take the modal cohort (1984) and industry (textiles) as the reference group.

For new exporters, the probability of remaining in the export market for an additional period is 63 percent. This survivor rate increases to 77 percent for plants that have been exporting for two years and remains at or above that level for subsequent periods. This pattern indicates that new exporters have a higher level of uncertainty than established exporters regarding future participation in the export market. Eaton et al. (2007) also document the high rate of turnover among new exporters. The dashed line in the figure shows the unconditional survivor rate. A plant that has successfully exported for three or more years has about the same rate of export exit as the average exporter.

The two patterns apparent in figure 1 suggest that entering the export market is a slow process with an initially high level of failure. In what follows, we show how standard sunk entry cost models fail to account for these characteristics and, as a result, overstate the benefit from

⁸Note that β_a is the a th term in the product of the survival function (Kaplan and Meier, 1958).

exporting.

3 The sunk-cost model

In this section, we describe a model in which heterogeneous plants make decisions about export entry under uncertainty. Our focus is on the decisions that plants make in response to changes in relative prices and productivity. Thus, we abstract from general equilibrium effects by assuming a constant wage, interest rate, and domestic price level. We model Colombia as a small open economy that takes the real exchange rate as exogenous. In what follows, we suppress the time subscript on variables unless needed for clarity.

3.1 Demand

A representative agent in the domestic economy supplies labor inelastically and has preferences over an aggregate consumption good, C , that consists of many differentiated varieties. The individual varieties are aggregated according to

$$C = \left(\sum_{j=1}^J c_j^{\frac{\theta-1}{\theta}} \right)^{\frac{\theta}{\theta-1}}, \quad (3)$$

where θ is the elasticity of substitution between varieties, and J is the number of available varieties. The consumer chooses consumption of each variety to maximize utility subject to the budget constraint

$$\sum_{j=1}^J c_j p_j = I, \quad (4)$$

where I is the consumer's income. Taking prices as given, the representative agent's demand for variety j is

$$c_j = \left(\frac{p_j}{P} \right)^{-\theta} C, \quad (5)$$

where P is the price of a unit of the aggregate consumption,

$$P = \left(\sum_{j=1}^J p_j^{1-\theta} \right)^{\frac{1}{1-\theta}}. \quad (6)$$

The rest of the world is populated by a representative consumer with an analogous utility function and budget constraint. Foreign demand for variety j is

$$c_j^* = \left(\frac{p_j^*}{P^*} \right)^{-\theta} C^*, \quad (7)$$

where an asterisk denotes a foreign variable. Note that we have assumed that the representative agents in the domestic country and the rest of the world have the same elasticity of substitution,

θ . Countries differ, however, in the level of demand for varieties, C and C^* .

3.2 Plant's static problem

Let j index the plant that produces variety j . The market structure is monopolistic competition, with each plant producing a differentiated variety. A plant chooses how much to produce for the domestic market, y_j , and how much to produce and export to the rest of the world, y_j^* . Plants produce output using labor and capital. Our data do not allow us to reliably track the birth and death of a plant, so we model the export decision of a plant that already exists as a domestic producer. The plant's production function is

$$f(\tilde{\epsilon}_j, n_j, k_j) = \tilde{\epsilon}_j n_j^{\alpha_N} k_j^{\alpha_K}, \quad (8)$$

where $\tilde{\epsilon}_j$ is a plant-level idiosyncratic productivity shock; n_j is the amount of labor employed by plant j ; and k_j is the amount of a capital rented by the plant.

In each period, the plant chooses prices, production, input demand, and export status ($X'_j = 0$ if not exporting and $X'_j = 1$ if exporting) to maximize the plant's value. The plant's problem can be divided into two subproblems: a static problem in which the plant chooses prices, employment, capital, and output, given its export status; and a dynamic problem in which the plant chooses its export status. We lay out the static problem in this section and the dynamic problem in the next.

A plant's profits are measured relative to the domestic price level, P . Contemporaneous profits gross of exporting costs are based on revenue obtained from sales in the domestic market and the world market (if exporting) less input costs,

$$\Pi_j = \frac{p_j}{P} y_j + I(X'_j = 1) Q \frac{p_j^*}{P^*} y_j^* - w n_j - r k_j, \quad (9)$$

where $Q = \frac{eP^*}{P}$ is the real exchange rate; e is the nominal exchange rate (domestic currency relative to foreign currency); r is the rental rate of capital; and w is the price of labor, where both factor prices are relative to the price of consumption. The plant is subject to a feasibility constraint,

$$y_j + y_j^* = \tilde{\epsilon}_j n_j^{\alpha_N} k_j^{\alpha_K}. \quad (10)$$

We allow for decreasing returns to scale in production, $\alpha_N + \alpha_K \leq 1$. In models with linear production functions, the constant marginal cost of production separates the export decision from the domestic production decision. In the presence of plant-level decreasing returns to scale, the export and domestic production decisions are tied together.

We assume that plants satisfy the demand in the markets they choose to enter ($c_j = y_j$ and $c_j^* = y_j^*$ if $X'_j = 1$) and substitute the demand functions into the profit function. The plant's static maximization problem is

$$\Pi_j(X'_j, \epsilon_j, Q) = \max_{y_j, y_j^*} \left\{ C^{\frac{1}{\theta}} y_j^{\frac{\theta-1}{\theta}} + I(X'_j = 1) Q C^{*\frac{1}{\theta}} y_j^{*\frac{\theta-1}{\theta}} - w n_j - r k_j \right\}, \quad (11)$$

subject to (10). The optimal quantities shipped domestically and abroad are

$$\begin{aligned} y_j &= \frac{Q^{-\theta} \frac{C}{C^*}}{1 + Q^{-\theta} \frac{C}{C^*}} \tilde{\epsilon}_j n_j^{\alpha_N} k_j^{\alpha_K} \\ y_j^* &= \frac{1}{1 + Q^{-\theta} \frac{C}{C^*}} \tilde{\epsilon}_j n_j^{\alpha_N} k_j^{\alpha_K}. \end{aligned} \quad (12)$$

For an exporter, movements in the real exchange rate shift the fraction of production shipped domestically and abroad; the real exchange rate shifts the relative profitability of exporting. A change in the plant's productivity, in contrast, increases output but does not change the relative importance of the export market.

Using (12), we can define the maximized value of profits for the plant, given its export status, idiosyncratic productivity, and the real exchange rate, as

$$\Pi(X'_j, \epsilon_j, Q) = \max_{n_j, k_j} \left(1 + I(X'_j = 1) Q^\theta \frac{C^*}{C} \right)^{\frac{1}{\theta}} C^{\frac{1}{\theta}} \tilde{\epsilon}_j^{\frac{\theta-1}{\theta}} n_j^{\alpha_N \frac{(\theta-1)}{\theta}} k_j^{\alpha_K \frac{(\theta-1)}{\theta}} - w n_j - r k_j. \quad (13)$$

We normalize the size of domestic aggregate demand, $C = 1$, so that C^* is the size of world aggregate demand relative to domestic aggregate demand. Since the idiosyncratic shocks, ϵ , are stationary, we define $\epsilon_j = \tilde{\epsilon}_j^{\frac{\theta-1}{\theta}}$, and normalize the mean of the ϵ process to one.

3.3 Dynamic programming problem

The presence of sunk export entry costs makes the plant's export decision a dynamic one. The plant faces costs of entering and maintaining an export operation. When a plant enters the export market having not exported in the previous period — export entry — it must pay f_0 . This cost represents the initial outlays required to set up exporting operations. If the plant has exported in the previous period and wishes to continue to export, it must pay f_1 . The cost of maintaining exporting operations will induce some plants to exit the export market when the present value from exporting becomes low enough. The cost of entry above the continuation cost, $f_0 - f_1$, is the part of the entry cost that is “sunk.” The exporting cost function is given by

$$f_X(X_j, X'_j) = f_0 I(X'_j = 1 | X_j = 0) + f_1 I(X'_j = 1 | X_j = 1). \quad (14)$$

The plant's state variables are the individual state variables (ϵ_j, X_j) and the aggregate state variable Q . The random variables Q and ϵ_j are modeled as time-invariant AR(1) processes,

$$\ln \epsilon_t = \rho_\epsilon \ln \epsilon_{t-1} + \omega_{\epsilon,t}, \quad \omega_\epsilon \sim N(0, \sigma_\epsilon^2) \quad (15)$$

$$\ln Q_t = \rho_Q \ln Q_{t-1} + \omega_{Q,t}, \quad \omega_Q \sim N(0, \sigma_Q^2). \quad (16)$$

In each period, the plant makes a discrete choice about its export status. Plants discount future profits at the rate $R = 1/(1+r)$. A plant's dynamic decision problem is given by the Bellman

equation,

$$V(X_j, \epsilon_j, Q) = \max_{X'_j} \left\{ \Pi(X'_j, \epsilon_j, Q) - f_X(X_j, X'_j) + R \mathbb{E}_{\epsilon'_j, Q'} V(X'_j, \epsilon'_j, Q') \right\}. \quad (17)$$

The fixed costs of exporting generate a value function with a kink at the productivity level of the marginal exporter, ruling out most approximation techniques. To solve (17), we discretize the state space and iterate on the value function directly. The solution to the plant's dynamic programming problem generates a policy function for export entry that is characterized by

$$X'_j(0, \epsilon_j, Q) = \begin{cases} 1 & \text{if } \Pi(1, \epsilon_j, Q) + R \mathbb{E}_{\epsilon'_j, Q'} [V(1, \epsilon'_j, Q') - V(0, \epsilon'_j, Q')] - f_0 \geq 0 \\ 0 & \text{otherwise.} \end{cases} \quad (18)$$

The first term on the right-hand side of (18) is the profit earned from exporting. The second term is the difference in the expected value of the plant if it chooses to export rather than remain a nonexporter. The plant enters the export market whenever the sum of these two terms (the present value of entry) is greater than the cost of entry. This optimality condition ties together new exporter dynamics — which strongly influence the present value of exporting — and the export entry cost.

Note that (18) will provide much (though not all) of the identification of the entry cost, f_0 . With a continuum of plant types, the marginal export entrant has a present value of entry exactly equal to the entry cost. In section 6, we will extend the model in a particular way that accounts for new exporter dynamics, but any model that decreases the present value of exporting for the marginal plant will lead to lower estimated entry costs.

4 Estimation

Our intent is to parameterize the model so that it replicates the cross-sectional features of the data — in particular (see table 3), the export starter and stopper rates, the average export-sales ratio, the size distribution of plants, and the autocorrelation of plant output — and then to ask how well the model accounts for new exporter dynamics. To do so, we need to jointly estimate the costs associated with exporting and the parameters that govern the production process.

We use an indirect inference estimation strategy, as described in [Gourieroux and Monfort \(1996\)](#). In that approach, an auxiliary model is specified to capture key features of the data. For expositional convenience, we refer to our auxiliary model as a collection of moments from the data. We estimate the structural parameters by matching the moments from the simulated data to moments from the observed data. Within the estimation procedure, the dynamic programming problem in (17) is solved for each guess of the parameter vector; a corresponding dataset is then simulated, from which moments are computed for the moment-matching exercise. The identification of the structural parameters comes from the selection of appropriate moments, which is discussed below. The next section discusses the parameterization of the model.

4.1 Parameterization

We can divide the set of parameters in the model into two subsets: a set that can be estimated without having to solve for the model’s equilibrium, and a set that requires the computation of the model’s equilibrium to estimate. We discuss the former set of parameters first.

A period in our model is one quarter. We set the annual interest rate, r , to be 10.9 percent, which is the average ex ante real lending rate in Colombia during our sample period. We use the producer price index to deflate the nominal lending rate — using the consumer price index, or the GDP deflator, yields similar rates. The parameters that describe the process for the real exchange rate in (16), ρ_Q and σ_Q , are estimated from the data. We estimate an AR(1) process for Q using quarterly data on the real effective exchange rate for Colombia from 1980–2005. We find that $\rho_Q = 0.83$, with a standard error of 0.053, and that $\sigma_Q = 0.036$.

The value of the real wage is normalized so that the size of the median plant in our model matches the median plant size in the data. In the Colombian data, the median plant (which is a nonexporter) has 63 employees. We set the wage so that the model’s median plant, which also does not export, has 63 employees. The average value of labor’s income share in our data implies that $\alpha_N = 0.45$. This value is very similar to the labor share in manufacturing, computed from the Colombian national accounts (0.42) and the labor share in the entire economy (0.54). In the sunk-cost model, we assume constant returns to scale, which implies that $\alpha_K = 0.55$. In section 5.2, we explore the robustness of our estimates to different assumptions regarding plant-level returns to scale.

Our final parameter in this set is the elasticity of substitution between varieties, θ . We are not able to estimate this parameter from our data, so we choose $\theta = 5$, as is common in the literature. We check our model’s sensitivity to this parameter in section 5.2. The parameter values are summarized in table 2.

4.2 Estimation results

The remaining parameters in the model are estimated using an indirect inference method that chooses the model’s parameters so that key moments generated by the model match those in the Colombian data. The parameters to be estimated, $\phi = (f_0, f_1, C^*, \rho_\epsilon, \sigma_\epsilon)$, govern the costs of exporting, the level of foreign demand, and the plant-level idiosyncratic shock process.

For a given vector of parameters, we solve the Bellman equation in (17) and find the associated policy functions of the plant. We then simulate a panel of 1,914 plants for 420 quarters. We drop the first 20 quarters to mitigate initial conditions and use the remaining periods to compute the moments in our simulated data exactly as we have done in the Colombian data. Note that we parameterize the model at a quarterly frequency, but the data are collected annually. We aggregate the quarterly data from our model to an annual panel before computing the moments. This allows plants in the model to make decisions at a quarterly frequency, yet keeps the model comparable to the data. This procedure also generates a time aggregation effect, as is likely present in the data (Bernard et al., 2014).

Table 2: Parameters common across models.

Parameter	Value	Target
r (annual)	0.109	Average observed interest rate
ρ_Q	0.826	Real effective exchange rate
σ_Q	0.036	Real effective exchange rate
α_N	0.450	Labor share of income
α_K	0.550	Plant-level returns to scale
θ	5.0	Elasticity of substitution

To obtain our estimates, we search over the parameter space to solve

$$L(\phi) = \min_{\phi} (m_s(\phi) - m_d)' W (m_s(\phi) - m_d), \quad (19)$$

where ϕ is the vector of parameters; $m_s(\phi)$ is the vector of moments from the simulated model data; m_d is the vector of moments from the Colombian data; and W is the weighting matrix. The weighting matrix is the inverse of the covariance matrix of the moments in the data.⁹ The function $L(\phi)$ is not analytically tractable, so we use computational methods to solve (19).

To identify the five parameters of interest, we choose five moments from the data that are informative about the parameters. Based on the analysis in section 2, we choose: 1) the fraction of nonexporting plants that begin exporting (the starter rate); 2) the fraction of exporting plants that stop exporting (the stopper rate); 3) the average export-sales ratio of exporting plants; 4) the coefficient of variation of plant size, as measured by domestic sales; and 5) the coefficient β from

$$\log y_{i,t} = \gamma_i + \delta_t + \beta \log y_{i,t-1} + \nu_{i,t}, \quad (20)$$

where $y_{i,t}$ are the domestic sales of plant i at time t . The regressors γ and δ are plant and time fixed effects. For all of the moments except the coefficient in (20), we construct their values as the average across the sample, 1981-1991.

Our aim is to choose moments that characterize the stationary equilibrium of the model. The export participation rate,¹⁰ the export-sales ratio, and the size distribution of plants are commonly used to parameterize “static” models in which plant productivity is constant and the export decision becomes a once-and-for-all choice (e.g., Melitz, 2003). The domestic sales regression (20) allows us to identify the persistence of productivity shocks using all the plants in our sample — most of which are not exporters.

The model does not admit a closed-form mapping of each parameter to each particular moment; the parameters jointly determine the moments. The cost of entering the export market (f_0) affects the starter rate but also influences the stopper rate, as the higher barrier to entry implies

⁹We calculate the bootstrapped estimate of the covariance matrix from 1,000 panels of 1,914 plants — the same size as our sample. Each bootstrapped sample is drawn with replacement from the Colombian data.

¹⁰In the stationary distribution, the starter rate and the stopper rate will give rise to a constant export participation rate. We could have chosen to target any two of these three moments in our estimation procedure.

Table 3: Moments in the data and in the models.

Moment	Data	Sunk cost	Gradual demand	Extended
Starter rate	0.0517	0.0517	0.0517	0.0517
Stopper rate	0.1062	0.1062	0.1062	0.1062
Average export-sales ratio	0.1346	0.1346	0.1345	0.1345
Coef. of variation, domestic sales	0.2090	0.2090	0.2090	0.2090
Slope, domestic sales reg.	0.6482	0.6482	0.6482	0.6482
Constant, export growth reg.	0.3078	0.9574*	0.3078	0.3078
Slope, export growth reg.	0.1255	0.0179*	0.1255	0.1255
Survival rate, export entrant	0.6300	0.9978*	0.9304*	0.6300
Non-targeted moments				
Exporter size premium, employment	2.6	3.1	3.0	1.9
Exporter size premium, domestic sales	3.8	2.7	2.6	1.7

Notes: The sample covers 1981–1991. The domestic sales slope is the coefficient from the regression in (20). The export growth constant and slope coefficients are from a regression of the export-sales ratio (as a fraction of the average export-sales ratio) on the number of years that a plant has been exporting in (23). The exporter size premium is the ratio of the geometric mean of exporter size to the geometric mean of nonexporter size. *Moments not targeted.

that, on average, more-productive plants will choose to be exporters, making them less likely to exit.

The serial correlation in the idiosyncratic shocks (ρ_ϵ) directly affects the autocorrelation of domestic sales, but it also influences the starter and stopper rates, as plants have stronger reactions to more-persistent shocks. Combined with the standard deviation of the idiosyncratic shocks (σ_ϵ), the serial correlation of plant-level shocks also determines much of the plant size distribution, as summarized by the coefficient of variation of domestic sales. The relative size of aggregate demand (C^*) determines the average exports-sales ratio, which affects the profitability of export opportunities.

5 Results

The second column of table 3 reports the moments from the simulated model data and shows that the sunk-cost model fits the chosen moments well. As an additional check, we report the exporter size premium, which is not a target moment in our estimation procedure. The exporter size premium is an informative data moment, as it reflects the fundamental mechanism in this class of models: More-productive plants select into exporting, and these more-productive plants are larger. In the sunk-cost model, the exporter size premium, as measured by employment, is 2.6, compared to 3.1 in the data. The model’s export size premium, when measured by domestic sales, is 2.7, compared to 3.8 in the data.

The estimated parameter values are shown in table 4. The persistence in export status that we observe in the data is commonly attributed to two causes: sunk investments made in entering the export market and persistent unobservable productivity. The estimation places weight on both of these factors. The idiosyncratic shock is strongly autocorrelated, although the standard deviation of the innovations to the shocks is also large. There is also a significant sunk aspect of the exporting

cost structure. The entry cost is almost 20 times as large as the continuation cost of exporting.

One advantage of our structural model is that we can recover estimates of the export fixed costs. The parametric costs are denominated in units of the median plant’s total sales. The estimates in table 4 imply that the export entry cost is 96 percent of the median plant’s annual sales, and the export continuation cost is about five percent of the median plant’s sales. Note that the median plant is a nonexporter: This is not surprising given that, in the sunk-cost model, exporting requires an entry cost of almost one year’s sales. Given the selection of large, productive plants into exporting, however, the estimated entry cost is only 31 percent of the average entrant’s first-year exports and 173 percent of the entrant’s first-year gross export profits. Below, we show that the size of the entry costs falls significantly when we account for new exporter dynamics.

Table 4: Estimated parameter values.

	f_0^\dagger	f_0^\ddagger	f_1^\dagger	f_1^\ddagger	C^*	σ_ϵ	ρ_ϵ	γ_0	γ_1	ζ_L
Sunk cost	0.961 (0.102)	0.305	0.047 (0.005)	0.015	0.146 (0.010)	0.116 (0.011)	0.873 (0.023)			
High elasticity	0.736 (0.095)	0.207	0.034 (0.003)	0.010	0.135 (0.010)	0.087 (0.009)	0.873 (0.023)			
Low elasticity	1.312 (0.136)	0.537	0.075 (0.005)	0.031	0.157 (0.010)	0.170 (0.017)	0.872 (0.024)			
Decreasing returns	0.897 (0.112)	0.314	0.046 (0.004)	0.016	0.147 (0.010)	0.134 (0.013)	0.874 (0.023)			
Gradual demand	0.286 (0.126)	0.219	0.064 (0.008)	0.049	0.198 (0.019)	0.116 (0.011)	0.873 (0.023)	0.258 (0.082)	0.024 (0.006)	
Extended model	0.590 (0.479)	0.799	0.057 (0.006)	0.077	0.185 (0.017)	0.116 (0.011)	0.873 (0.023)	0.278 (0.146)	0.026 (0.009)	0.009 (0.003)

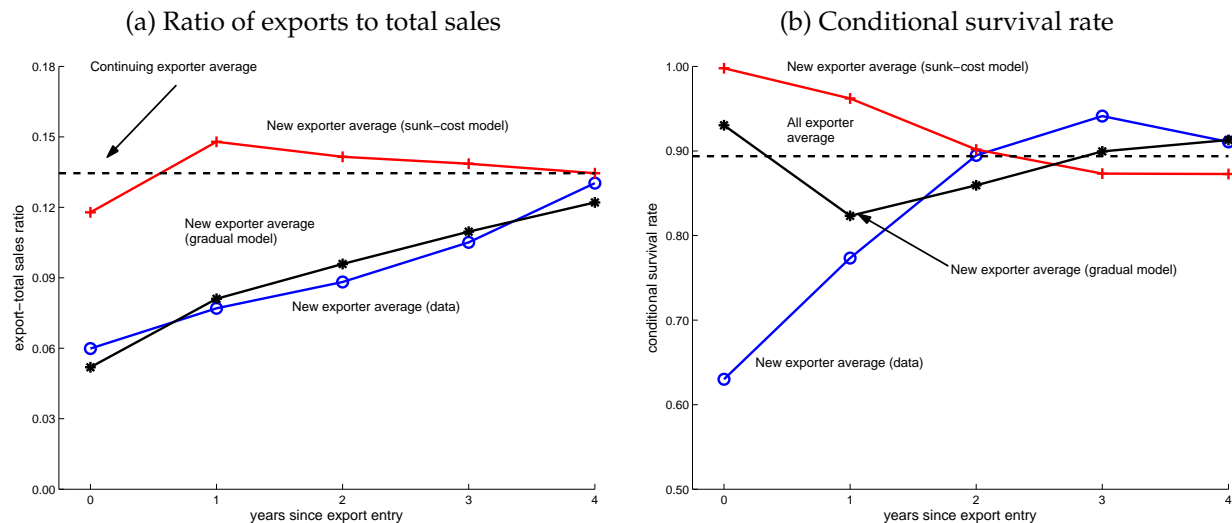
Notes: Standard errors are reported in parentheses. [†]Measured as a fraction of the median plant’s sales. [‡]Measured as a fraction of the average entrant’s first-year export revenue.

5.1 New exporter dynamics

Having parameterized the model, we can now ask: How well does it account for the dynamic behavior of new exporters? We plot the export-sales ratio for new exporters in figure 2a. Relative to the data, the export-sales ratio for new exporters displays a large jump: In its first year as an exporter, the plant exports, on average, almost 12 percent of output. This increases slightly in the second year and remains roughly constant thereafter. The dynamics from the model are in sharp contrast to the data, in which exports grow over a period of several years. The pattern of export behavior in the model is largely driven by the structure of the fixed entry costs. Once a plant enters the export market, there are no further costs for adjusting the export-sales ratio, so it immediately adjusts exports to the optimal level.

Plants enter the export market for two reasons: favorable movements in the real exchange rate and positive shocks to productivity. Both of these shocks are persistent, which influences the plant’s behavior. The declining export-sales ratio of one- to four-year old exporters is driven by the behavior of the real exchange rate, as is evident in (12). A plant is more likely to enter the export market when the real exchange rate is high. As the innovation to the real exchange

Figure 2: The sunk-cost model and the gradual demand model.



rate deteriorates, the export-sales ratio returns to its average level. The non-monotonicity in the year of entry is caused by a partial-year effect. We aggregate our quarterly model's data to a yearly frequency, so plants that enter later in a calendar year will have artificially lower export-sales ratios in the first year. Without this partial-year effect, the export-sales ratio would decline monotonically.

Implications of the shock process can be seen further in figure 2b, which plots the conditional survivor rates for entrants. The conditional survival rates in the model inherit the persistence of the underlying shocks. Plants enter the export market in response to a good shock; as the shock dies out, plants are more likely to exit the market. Thus, the rate of survival is high initially and falls through time. This is in contrast to the data, in which we see plants more likely to drop out of the export market initially.

The behavior of the model in these two dimensions is sharply at odds with the data and contributes to the large entry cost we have estimated. In the model, an entrant's probability of staying in the export market is initially high, and the plant can immediately access the export market. The plant is able to immediately realize export profits and has a high probability of staying in the market for the initial periods, so the value of being an exporter is large. From (18), we see that the large present value of exporting requires a large entry cost to keep the export participation rate as low as it is in the data. In section 6, we consider an extension of the model that forces exports to grow as they do in the data.

5.2 Robustness

In this section, we show that our results so far — i.e., compared to the data, new exporters export too much too soon and survive too long — are robust to our choice of substitution elasticities and assumptions about the scale of production.

We re-estimate the model under three alternative scenarios. In the *high elasticity* model, we set

Table 5: New exporter characteristics.

Years since export entry	Data	Sunk cost	High elas.	Low elas.	DRS	Gradual demand	Extended model
Export-sales ratio							
0	6.0	11.8	12.7	11.0	11.7	5.2	5.2
2	8.8	14.2	14.1	14.2	14.2	9.6	9.5
4	13.0	13.5	13.2	13.7	13.5	12.2	12.4
Survivor rates							
0	63.0	99.8	99.8	99.8	99.8	93.0	63.0
2	89.5	89.9	89.9	90.5	90.1	85.9	91.3
4	91.1	87.7	87.3	87.1	87.4	91.3	94.1
Exporter size premium							
0	2.21	6.74	8.80	3.19	5.18	6.66	1.98
2	3.34	2.70	3.49	1.39	2.14	3.27	2.25
4	4.14	2.17	2.80	1.14	1.73	2.22	1.65

Notes: Statistics are based on a simulation of 1914 plants for 400 quarters. The export-sales ratio is the per-plant export-sales ratio, averaged over all exporting spells that begin with entry into the export market and are followed by at least five years of continuous exporting. The survivor rate is computed as the fraction of all plants of age a that continue to export at age $a + 1$. The size premium is the (geometric) average domestic sales of exporters divided by the average domestic sales of nonexporters.

$\theta = 7$. In the *low elasticity* model, we set $\theta = 3$. In the *decreasing returns to scale* (DRS) model, we scale α_N and α_K so that they sum to 0.95. Table 4 reports the estimated parameters from these models. In the interest of brevity, we do not report the moments from these three models — the moments match the data exactly, in the same way that the sunk-cost model does in table 3.

Changing the elasticity of substitution has a direct effect on a plant’s profitability. A lower elasticity allows the plant to charge a higher markup, increasing its profits.¹¹ With greater potential profits from exporting, the model needs a larger export entry cost ($f_0 = 1.3$), compared to ($f_0 = 0.96$) the sunk-cost model. The greater entry cost offsets some of the greater expected profit from exporting. Increasing the entry cost strengthens the selection on productivity, so the continuation cost must rise to keep the stopper rate from falling. The variance of the plant-size distribution depends on both the variance of the idiosyncratic shock process and the elasticity of substitution. A lower elasticity works to decrease the variance of the plant-size distribution (the plant’s sales are less responsive to the plant’s price), so the model needs more-volatile idiosyncratic productivity shocks, $\sigma_\epsilon = 0.17$. Increasing the elasticity works in an analogous way, leading to a smaller export entry cost and a less-volatile idiosyncratic shock process compared to the sunk-cost model.

In the decreasing returns to scale model, the exporting decision is no longer independent of the plant’s domestic production decisions. If entering the export market increases total production, then the marginal cost of production for both the domestic and export markets will increase. This

¹¹The markup in this model is the standard monopolistic competition markup, $\theta/(\theta - 1)$.

characteristic makes exporting less desirable than in models with constant returns to scale, which is reflected in a smaller entry cost than in the sunk-cost model.¹² Decreasing returns to scale decreases the variance in the plant size distribution, so the volatility of the plant productivity process increases to offset this force.

While the different models generate parameter estimates that differ from the sunk-cost model's, changing the elasticity of substitution or the returns to scale in production does not help the model account for the behavior of new exporters. We report summary statistics about the export-sales ratio and the survival rates from the three models in table 5. Along these two dimensions, the three models are virtually identical.

6 Model extensions

The sunk-cost model does a poor job of matching the dynamics of the export entrant's export-sales ratio and survival rates. In this section, we extend the model in two ways that allow the model to generate realistic new exporter dynamics. Our goal in this section is not to offer a deep model of new exporter dynamics, but to quantitatively assess the importance of getting the dynamics of new exporters right.

6.1 Gradual demand

To generate slow export growth after entry, we modify the demand presented to an entrant. A plant that has been an exporter for a periods faces the demand curve,

$$c_j^*(a) = \gamma(a) \left(\frac{p_j^*(a)}{P^*} \right)^{-\theta} C^* \quad (21)$$

$$\gamma(a) = \begin{cases} \gamma_0 + \gamma_1 \times a & \text{if } a = 0, \dots, 21 \\ 1 & \text{if } a > 21. \end{cases} \quad (22)$$

This specification implies that a new exporter's demand in the export market grows over 20 quarters and remains at "full" demand for the rest of the time it spends in the export market. The rest of the model remains as in the sunk-cost model.

Our specification of the demand function is meant to capture the forces that might lead a new exporter to slowly increase its exports. For example, learning has been suggested, both from the demand side, as in Rauch and Watson (2003) and Eaton et al. (2014), and on the supply side, as in Jovanovic (1982). We discuss structural models that may be able to generate the slow growth in new exporter sales in section 8.

The *gradual demand model* adds two new parameters to those in the sunk-cost model, γ_0 and γ_1 . To estimate these parameters, we take as moments from the data those from the sunk-cost model,

¹²Blum, Claro, and Horstmann (2013) construct a model in which increasing marginal costs cause plants to enter and exit the export market in response to stochastic demand.

plus the coefficients from the regression

$$\frac{ex_a}{sales_a} = g_0 + g_1 \times a + \epsilon_a. \quad (23)$$

The left-hand variable is the average export-sales ratio for plants of age a — the values in the first column of table 1. Before fitting a line to the data, we normalize the export-sales ratios by dividing them by the average export-sales ratio of all plants. These two data moments are reported in table 3.¹³ As figure 1a makes clear, the new exporter export-sales path is close to linear, so (23) captures the pattern in the data well.

We estimate the gradual demand model adding γ_0 and γ_1 to the parameter vector from the sunk-cost model. Table 3 summarizes the model’s fit, and table 4 lists the estimated parameter values. The estimated parameters from this model are very similar to the ones in the sunk-cost model, except for the cost of entry, which is much smaller.

In figure 2a, we plot the export-sales ratio in the sunk-cost model, the gradual demand model, and the data. The gradual demand model has very different implications regarding the value of exporting. By forcing exports in the model to grow as they do in the data, a plant’s export profits are pushed out into the future. Since plants discount future profits, and shocks are persistent but not permanent, this redistribution of profits into the future lowers the present value of export entry. This is reflected in the estimated parameters: The export entry cost, f_0 , is only 30 percent of that in the sunk-cost model. When export sales grow slowly, as in the data, the cost of entering the export market falls to 29 percent of the median plant’s sales. Note that, although the size of the entry cost has fallen, so, too, has the size of the entrant. The entry cost in the gradual demand model is 23 percent of a new exporter’s first-year exports (174 percent of first-year gross export profits), compared to 31 percent (173 percent) in the sunk-cost model. We analyze the dynamic properties of new exporter profits — which are quite different across models — in section 6.2.

Introducing gradual export demand growth to the model changes the exit behavior of new exporters, but only in a small way. As figure 2b shows, the gradual demand model cannot account for the pattern of new exporter survival that we find in the data. The smaller entry cost that is estimated in the gradual demand model allows smaller, less productive plants to enter the export market. These plants will be closer to the exit threshold and are more likely to exit in response to lower productivity or appreciating real exchange rates. The effect of this can be seen in the figure: The survival rate of an entrant is lower than that in the sunk-cost model, but it is still much higher than in the data.

6.2 Gradual demand and stochastic entry costs

The model with gradual export demand expansion is able to capture the dynamics of new exporter sales, but it cannot account for the pattern of exit among new exporters. The problem is that plants tend to enter the export market in response to favorable shocks, which, in our models, are

¹³Note that (22) is specified at our model’s quarterly frequency, but the model output is aggregated to an annual frequency to be comparable with the data in (23).

autocorrelated: The few periods after entry tend to be very good periods for the plant, so it has little incentive to exit early in its life cycle.

To generate a low survival rate for new exporters, the model needs to allow for some “bad” plants to enter the export market. These plants try exporting for a period, do not earn enough profit to justify paying the fixed costs of continuing to export (f_1), and exit. In the two models considered above, the entry cost is large enough that unproductive plants do not find it worthwhile to enter because they do not expect to earn large enough profits going forward. We extend the gradual demand model by making the export entry cost random. Occasionally, plants with low productivity will face a low entry cost. These plants will enter, but many will not find it profitable to continue exporting and exit quickly.


Following Das, Roberts, and Tybout (2007), we modify our exporting technology so that the cost of entry, f_0 , is random. We choose the simplest specification for this technology: With probability $1 - \zeta_L$, the entry cost is *high*, $f_0 = f_H$, and with probability ζ_L , the entry cost is *low*, $f_0 = 0$. This implies that a nonexporting plant will randomly be offered free entry into the export market.

We estimate this model, adding ζ_L to the parameter vector from the gradual demand model. To the vector of moments to match, we add the initial survival rate of a new exporter: Survival rates beyond the first year of entry are not directly targeted. We report the estimated parameters for this *extended model* in table 4 and the model’s fit in table 3. The probability of receiving free entry into the export market is small, $\zeta_L = 0.009$. Most of the model’s estimated parameters are similar to the ones in the gradual demand model, except for the export entry cost. The entry cost in the extended model is twice as large as in the gradual demand model (about 60 percent of the median plant’s annual sales), but this cost is not paid often. Indeed, of all the new exporters in our simulated data, 68 percent of them entered when they drew the low entry cost. Having fewer plants entering when faced with the high entry cost weakens the identification of this cost, which is reflected in the larger standard error for this parameter.

Comparing the fixed exporting costs across models, we find that incorporating slow-growing export demand significantly lowers the size of the estimated entry cost. As we decrease the present value of exporting, we decrease the entry cost needed to limit export entry. This result does not mean, however, that sunk entry costs are no longer important once we have corrected for new exporter dynamics. As table 4 indicates, the entry cost is still about ten times larger than the continuation cost in the extended model. While this is half as large as the ratio of 20 in the sunk-cost model, the extended model still needs a significant sunk cost component to match the average stopper rate (0.11) in the data. While sunk costs remain important, our analysis demonstrates the need for additional sources of plant heterogeneity in order to explain exporter behavior.

In figure 3, we plot the export-sales ratio and the conditional survivor rates that characterize new exporters for the extended model, the sunk-cost model, and the data. Adding stochastic entry costs allows the extended model to capture the conditional survivor rate without diminishing its ability to account for the export intensity dynamics. We have calibrated the model to the entrant’s first-year survival rate because that is when the most turnover occurs. To be comparable with

Figure 3: The sunk-cost model and the extended model.



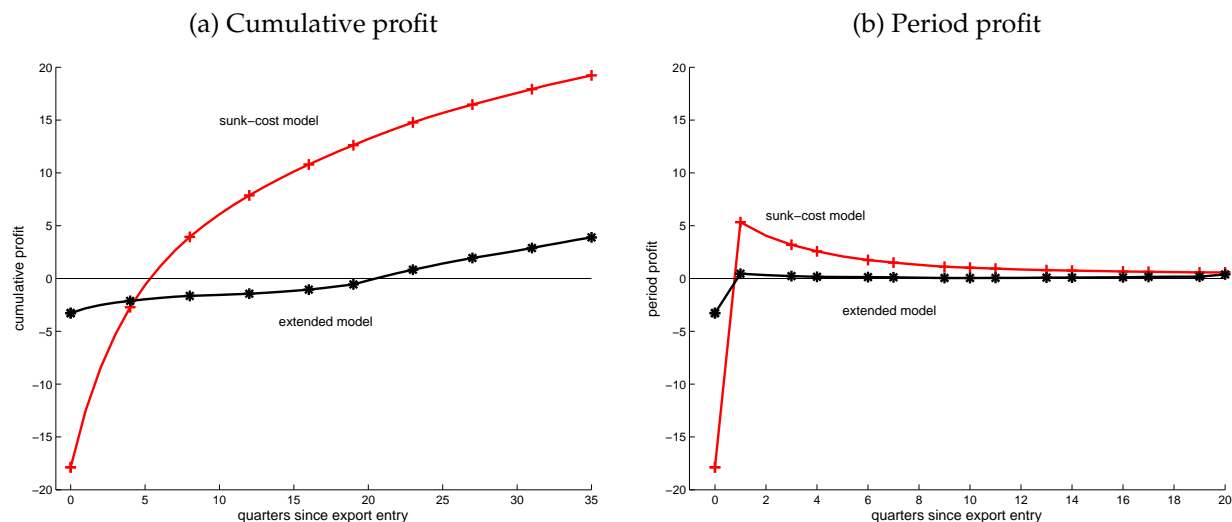
the export-sales ratio (which evolves over five years), however, we can measure the fraction of entrants that are still exporting after five years. In the extended model, this five-year survival rate is 44 percent, compared with 37 percent in the data and 65 percent in the sunk-cost model. It may be possible to achieve a better fit by introducing age-indexed export continuation costs, but given the reduced-form nature of this exercise, we leave it for future research to develop economic models of the survival rate.

A new exporter's profits following entry are useful in understanding how the extended model works. In figure 4a, we plot the average discounted cumulative export profit for entrants. In the sunk-cost model, the average entrant incurs a large negative profit on entry, reflecting the large entry cost. The quarters following entry, though, are highly profitable. We plot the average per-period profit of an export entrant in figure 4b. Entry into the export market occurs because the plant receives a favorable shock to the plant's productivity, to the real exchange rate, or to both. The profit of the entrant inherits the persistence of the underlying shock process, so the new exporter's profits are front-loaded. It takes seven quarters for the average export entrant to break even.

The average cumulative profit of an entrant in the extended model is qualitatively different from that in the sunk-cost model. In the extended model, the initial quarter's profit is negative, but much less so than in the sunk-cost model. This is the result of the smaller export entry cost estimated in the extended model and the fact that many plants are entering the export market when the entry cost is zero. Following entry, the plant's cumulative profit grows slowly, and it takes 22 quarters for the average entrant to break even — three times longer than in the sunk-cost model. This is driven primarily by the slow increase in demand for the new exporter's variety. While the favorable shock that led to entry is as persistent as it is in the sunk-cost model (table 4), it is largely offset by the slow growth in the demand for the entrant's variety.

Even though the extended model can account for the export volume and persistence of new

Figure 4: Average export profit of export entrants.



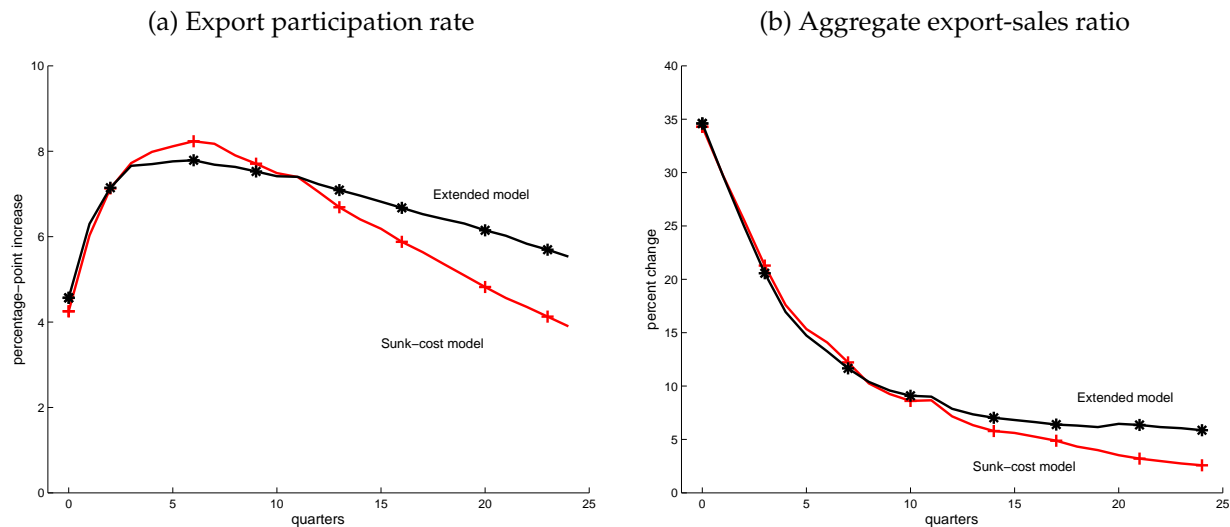
exporters, there are dynamics that it cannot capture. In table 5, we report the average size — as measured by domestic sales — of a new exporter relative to an average nonexporter. In the data, the average new exporter sells about twice as much to the domestic market as the average nonexporter, and this exporter premium doubles if the exporter continues for five years. In the sunk-cost and gradual demand models, the selection induced by the large entry costs generates an initial size premium that is too large, highlighting the relatively large productivity shock needed to generate entry. In the extended model, entry of smaller plants that received the low entry cost draw drives down the initial premium to levels close to the data, but none of the models can generate a size premium that grows with exporting tenure. The shrinking exporter premia are driven by the selection of plants into the export market in response to a positive shock, which decays in the periods following entry. This discrepancy between the model and the data is further evidence that a richer model of plant heterogeneity is needed — and likely one that does not hinge on autoregressive shocks.

7 Aggregate implications

In this section, we explore the extent to which new exporter dynamics influence the aggregate response of the economy to a transitory real exchange rate depreciation and a permanent demand shock. This is a partial equilibrium model, so the aggregate responses shown here are likely to be muted by general equilibrium effects on input costs and increased competition. Alessandria, Choi, and Ruhl (2013) further explore the aggregate impact of new exporter dynamics in general equilibrium.

Figure 5 reports the share of plants that export (the export participation rate) and the aggregate export-sales ratio following a one-time, one-standard-deviation shock to the real exchange rate. In figure 5a, we plot the increase in the export participation rate. The models display significant hysteresis: After 20 quarters, the real exchange rate has returned to its steady-state value, but

Figure 5: Aggregate response to a real exchange rate depreciation.



Notes: The figures show the impulse responses to a one-standard-deviation shock to the real exchange rate. In panel a, the vertical axis measures the percentage-point increase in the export participation rate relative to the model in which the real exchange rate is one. In panel b, the vertical axis measures the percentage change in the aggregate export-sales ratio relative to the model in which the real exchange rate is one.

the export participation rate is still significantly above its steady-state level. In both models, the hysteresis is a result of the sunk part of the export cost. As the real exchange rate returns to its steady-state value, plants that have already paid the entry cost can still profitably export because they have to pay only the smaller continuation cost.

The extended model includes a second mechanism that generates hysteresis. Plants that have entered and successfully increased their export demand are also reluctant to exit because reentry involves not only repaying the entry cost, but also having to rebuild export demand. These two forces generate hysteresis in the extended model that is stronger than in the sunk-cost model. After 14 quarters, the export participation rate in the extended model rises above that in the sunk-cost model.

In figure 5b, we plot total exports relative to total sales — the aggregate export-sales ratio. The behavior of the aggregate export-sales ratio in the first 10 quarters following the exchange rate depreciation is similar in the two models. In both models, most of the change in aggregate exports comes from incumbent exporters responding to the change in the exchange rate. In the following quarters, aggregate exports in the extended model are greater than those in the sunk-cost model. The larger response in the extended model is the result of the stronger response in export participation; by the twentieth quarter, the exchange rate has returned to its steady-state value.

Next, we consider the aggregate implications of a permanent, unforeseen increase in demand. In table 6, we compare the initial steady state of the calibrated model to a new steady state in which the level of foreign demand, C^* , has risen by four percent.¹⁴ The starter rate rises and the stopper

¹⁴We have abstracted from proportional trade costs in our model. Given our elasticity, the four-percent change in

Table 6: Aggregate response to 4 percent decrease in trade costs.

	percent change		levels		
	aggregate export/sales	export part.	starter rate	stopper rate	plant-level export/sales
Original			0.052	0.106	0.135
Liberalized:					
Sunk cost	7	8	0.055	0.100	0.138
Gradual	15	15	0.060	0.092	0.142
Extended	23	21	0.054	0.080	0.144

Notes: The variable aggregate export/sales is total exports of all plants divided by total sales of all plants. The variable plant-level export sales is the export-sales ratio of each exporting plant, averaged across all exporting plants.

rate falls in all of the models, resulting in a greater fraction of plants exporting. The aggregate export-sales ratio rises by seven percent in the sunk-cost model and 23 percent in the extended model. In a general equilibrium model, export entry and the expansion of existing exporters would drive wages up and the aggregate price level down, dampening the aggregate response to trade liberalization that we find in our partial equilibrium model.¹⁵

The extended model's larger response is again driven by the entry of new exporters. After the decrease in trade costs, young exporters are more profitable. This means that some exporters who would have exited are now able to stay in the export market, grow their export demand, and increase the probability that they stay in the export market. This can be seen in the last two columns of table 6: compared to the sunk-cost model, stopper rates fall by more, and average export intensity rises by more in the extended model.

8 Conclusion

In this paper, we assess the ability of sunk cost models of trade to account for the plant-level dynamics of new exporters. We document the experience of new exporters with data from Colombian manufacturing plants and find that a new exporter's export-sales ratio grows slowly following entry. The average new exporter takes four years to reach the export intensity of the (unconditional) average exporter. In addition, we show that new exporters experience a period of shakeout upon entry. The probability that an entrant will still be exporting in the next year is only 63 percent, and this survival rate is increasing with the exporter's tenure. These findings add to the growing body of empirical knowledge on plant-level behavior.

We construct a model in which plants are subject to persistent shocks to productivity and the real exchange rate and face fixed costs of export entry. We calibrate the model so that it replicates key cross-sectional facts that are frequently used in parameterizing this class of models. We find that the workhorse models — those featuring fixed entry costs and plant-level heterogeneity —

demand is approximately equal to a one-percent change in a proportional trade cost.

¹⁵Alessandria, Choi, and Ruhl (2013) analyze a similar model in general equilibrium. They also find a long-run trade elasticity of about 7 in a model similar to our sunk-cost model. They find a trade elasticity in a model similar to the our extended model of 12, compared to our 23.

do not replicate the gradual growth of new exporters and the pattern of plant export survival.

The failure of the model to generate these features has important implications for the value of exporting. The standard model front-loads the profits from exporting: An entrant is able to immediately adjust its export sales, and the autocorrelation of the shocks ensures that the good state of the world will likely continue for the first few periods. Exporting has a large present value, so large entry costs are needed to keep most plants from exporting.

We modify the standard sunk-cost model so that the entrants' behavior matches the data. In this model, plants enter the export market small, and they grow as they continue as exporters. This pushes the profits from exporting further into the future. With discounting and uncertainty, the present value of these future profits is significantly lower than in the sunk-cost model. The difference between the two models is most striking with regard to the calibrated entry costs. The model in which new exporters are forced to grow slowly has an entry cost that is three times smaller than that in the standard sunk-cost model.

Understanding the dynamics of new exporters is important. In the data, an exporter's first few years are typically spent exporting small amounts, and the probability of leaving the export market is high. When models do not generate this behavior, the value of exporting is large, and large entry costs are needed to produce the right export participation rate.

References

- Albornoz, Facundo et al. (2012). "Sequential exporting." *Journal of International Economics* 88 (1), pp. 17–31.
- Alessandria, George and Horag Choi (2007a). "Do Sunk Costs Of Exporting Matter for Net Export Dynamics?" *Quarterly Journal of Economics* 289.
- (2007b). "Establishment Heterogeneity, Exporter Dynamics, and the Effects of Trade Liberalization." *FRB of Philadelphia Working Paper No. 07-17*.
- Alessandria, George, Horag Choi, and Kim J. Ruhl (2013). "Trade Adjustment Dynamics and the Welfare Gains from Trade." Unpublished Manuscript.
- Baldwin, Richard E. (1988). "Hysteresis in Import Prices: The Beachhead Effect." *American Economic Review* 78 (4), pp. 773–85.
- Baldwin, Richard E. and Paul R. Krugman (1989). "Persistent Trade Effects of Large Exchange Rate Shocks." *Quarterly Journal of Economics* 104 (4), pp. 635–54.
- Bernard, Andrew B. and Bradford J. Jensen (2004). "Entry, Expansion and Intensity in the U.S. Export Boom, 1987-1992." *Review of International Economics* 12 (4), pp. 662–675.
- Bernard, Andrew B. et al. (2014). "Exporter Dynamics, Firm Size and Growth, and Partial Year Effects." Unpublished Manuscript.
- Blum, Bernardo S., Sebastian Claro, and Ignatius J. Horstmann (2013). "Occasional and Perennial Exporters." *Journal of International Economics* 90 (1), pp. 65–74.
- Das, Sanghamitra, Mark J. Roberts, and James R. Tybout (2007). "Market Entry Costs, Producer Heterogeneity, and Export Dynamics." *Econometrica* 75 (3), pp. 837–873.
- Dixit, Avinash K. (1989). "Hysteresis, Import Penetration, and Exchange Rate Pass-Through." *Quarterly Journal of Economics* 104 (2), pp. 205–228.
- Eaton, Jonathan, Samuel S. Kortum, and Francis Kramarz (2004). "Dissecting Trade: Firms, Industries, and Export Destinations." *American Economic Review* 94 (2), pp. 150–154.
- (2011). "An Anatomy of International Trade: Evidence from French Firms." *Econometrica* 79 (5), pp. 1453–1498.
- Eaton, Jonathan et al. (2007). "Export Dynamics in Colombia: Firm-Level Evidence." *KSG Working Paper No. RWP07-050*.
- Eaton, Jonathan et al. (2014). "A Search and Learning Model of Export Dynamics." Unpublished Manuscript.

- Foster, Lucia, John Haltiwanger, and Chad Syverson (2008). "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *The American Economic Review* 98 (1), pp. 394–425.
- Foster, Lucia, John C. Haltiwanger, and Chad Syverson (2012). "The Slow Growth of New Plants: Learning about Demand?" NBER Working Paper 17853.
- Gourieroux, Christian and Alain Monfort (1996). *Simulation-Based Econometric Methods*. New York: Oxford University Press.
- Helpman, Elhanan, Marc Melitz, and Yona Rubinstein (2008). "Estimating Trade Flows: Trading Partners and Trading Volumes." *Quarterly Journal of Economics* 123 (2), pp. 441–487.
- Impullitti, Giammarco, Alfonso A. Irarrazabal, and Luca David Opromolla (2013). "A Theory of Entry into and Exit from Export Markets." *Journal of International Economics* 90 (1), pp. 75–90.
- Jovanovic, Boyan (1982). "Selection and the Evolution of Industry." *Econometrica* 50 (3), pp. 649–670.
- Kaplan, Edward L and Paul Meier (1958). "Nonparametric Estimation from Encomplete Observations." *Journal of the American Statistical Association* 53 (282), pp. 457–481.
- Kohn, David, Fernando Leibovici, and Michal Szkup (2014). "Financial Frictions and New Exporter Dynamics." Unpublished Manuscript.
- Melitz, Marc J. (2003). "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica* 71 (6), pp. 1695–1725.
- Melitz, Marc J. and Fabio Ghironi (2005). "International Trade and Macroeconomic Dynamics with Heterogeneous Firms." *Quarterly Journal of Economics* 120 (3), pp. 865–915.
- Nguyen, Daniel X. (2012). "Demand Uncertainty: Exporting Delays and Exporting Failures." *Journal of International Economics* 86 (2), pp. 336–344.
- Rauch, James E. and Joel Watson (2003). "Starting Small in an Unfamiliar Environment." *International Journal of Industrial Organization* 21 (7), pp. 1021–1042.
- Rho, Young-Woo and Joel Rodrigue (2014). "Firm-level Investment and Export Dynamics." Unpublished Manuscript.
- Roberts, Mark J. (1996). "Colombia, 1977–85: Producer Turnover, Margins, and Trade Exposure." *Industrial Evolution in Developing Countries: Micro Patterns of Turnover, Productivity, and Market Structure*. Ed. by Mark J. Roberts and Jamer R. Tybout. New York: Oxford University Press, pp. 227–59.
- Roberts, Mark J. and James R. Tybout (1996). *Industrial Evolution in Developing Countries: Micro Patterns of Turnover, Productivity, and Market Structure*. New York: Oxford University Press.

- Roberts, Mark J. and James R. Tybout (1997). "The Decision to Export in Colombia: An Empirical Model of Entry with Sunk Costs." *The American Economic Review* 87 (4), pp. 545–564.
- Ruhl, Kim J. (2008). "The International Elasticity Puzzle." Unpublished manuscript.

A Appendix

A.1 Sample construction

- Our sample covers the years 1981-1991. The construction of the underlying data is described in Roberts (1996).
- Following Roberts and Tybout (1997), we use data from 19 industries: 311 (Food Products), 312 (Other Food Products), 321 (Textiles), 322 (Clothing and Apparel), 323 (Leather Products excl. Clothing and Shoes), 324 (Leather Shoes), 341 (Paper), 342 (Printing and Publishing), 351 (Industrial Chemicals), 352 (Other Chemicals), 356 (Plastic Products), 362 (Glass Products), 369 (Other Products of Non-metallic Minerals), 371 (Iron and Steel), 381 (Metal Products excl. Machinery and Equipment), 382 (Machinery excl. Electrical Machinery), 383 (Electronic Machinery and Equipment), 384 (Transportation Equipment), and 390 (Miscellaneous Manufacturing Industries).
- We balance the panel by dropping any plant whose employment falls below 15 employees in any year of the sample.
- We exclude plants that experience an annual absolute change in real production or employment greater than 1.5, where the change is measured by

$$\Delta = \left| \frac{x_t - x_{t-1}}{0.5 \times (x_t + x_{t-1})} \right|.$$