An Anatomy of International Trade: Evidence from French Firms

Jonathan Eaton,[1] Samuel Kortum,[2] and Francis Kramarz[3]

August 2008

[1]New York University (jonathan.eaton@nyu.edu).
[2]University of Chicago (kortum@uchicago.edu).
[3]CREST-INSEE and CEPR (kramarz@ensae.fr).

Abstract

We examine the cross section of sales of French manufacturing firms in 113 destinations, including France itself. Several regularities stand out: For example: (1) the number of French firms selling to a market, relative to French market share, increases systematically with market size; (2) sales distributions are very similar across markets of very different size and extent of French participation; (3) Average sales in France rise very systematically with selling to less popular markets and to more markets. We adopt a model of firm heterogeneity and export participation which we estimate to match moments of the French data using the method of simulated moments. The results imply that nearly half the variation across firms that we see in market entry can be attributed to a single dimension of underlying firm heterogeneity, efficiency. Conditional on entry underlying efficiency accounts for a much smaller variation in sales in any given market. Parameter estimates imply that fixed costs eat up a little more than half of gross profits. We use our results to simulate the effects of a counterfactual decline in bilateral trade barriers on French firms. The average firm in the top decile experiences significant expansion in total sales while average firms in lower deciles suffer losses or exit altogether.

1 Introduction

We exploit a detailed set of data on the exports of French firms to confront a new generation of trade theories. The data, from French customs, report the sales of over 200,000 individual firms to over 100 individual markets in a cross section.

We ask how well the model of the export behavior of heterogeneous firms introduced by Melitz (2003) and more concretely specified by Helpman, Melitz, and Yeaple (2004) and Chaney (2008) stands up to these data. Basic elements of the model are that firms' efficiencies follow a Pareto distribution, demand is Dixit-Stiglitz, and markets are separated by iceberg trade barriers and require a fixed cost of entry. The model is the simplest one we can think of that can square with the facts.

With this basic model in mind, we extract five relationships that underlie the data: (1) how entry varies with market size, (2) how the distribution of sales varies across markets, (3) how firms enter multiple markets, (4) how export participation abroad connects with sales at home, and (5) how sales abroad relate to sales at home. Through the haze of numbers we begin to see the outlines of the basic model, and even rough magnitudes of some parameters. The basic model fails to come to terms with some features of the data, however: Firms don't enter markets according to an exact hierarchy and their sales where they do enter deviate from the exact correlations the basic model would insist upon.

To reconcile the basic model with these failures we extend it by introducing market and firm-specific heterogeneity in demand and entry costs. We also incorporate a reduced form version of Arkolakis's (2008) market access cost. The extended model, while remaining very parsimonious and transparent, is one that we can connect more formally to the data. We describe how the model can be simulated and we estimate its parameters using the method of simulated moments.

With the parameter estimates in hand we find that the forces underlying the basic model remain powerful. Simply knowing a firm's efficiency improves our ability to explain the probability it sells in any market by nearly fifty percent. Conditional on a firm selling in a market, knowing its efficiency improves our ability to predict how much it sells there, but by much less. While these results leave much to be explained by the idiosyncratic interaction between individual firms and markets, they tell us that any theory that ignores features of the firm that are universal across markets misses much.

We embed our model into a general equilibrium framework with an arbitrary number of countries. Calibrating the framework to data on production and bilateral trade from 113 countries and the rest of the world, we can examine the implications of changes in exogenous parameters for income, wages, and prices in each country and for bilateral trade. We can use these counterfactual outcomes and our parameter estimates to simulate the implications for French firms. A striking finding is that lower trade barriers, while raising welfare in every country, leads to substantially more inequality in the distribution of firm size. Even though total output of French firms rises by 3.3 percent, all of the growth is accounted for by firms in the top decile with sales and sales per firm in every other decile falling. Import competition leads to the net exit of 26,212 firms, 11,323 accounted for by those in the bottom decile.

Section 2 which follows explores five empirical regularities. With these in mind in Section 3 we turn to a model of exporting by heterogeneous firms. Section 4 explains how we estimate the parameters of the model while section 5 explores the implications of a lowering of trade barriers.

2 Empirical Regularities

Our data, described in Appendix A, are the sales, translated into U.S. dollars, of 229,900 French manufacturing firms to 113 markets in 1986.[1] Among them only 34,035 sell elsewhere than in France. The firm that exports most widely sells to 110 out of the 113 destinations.

We assemble our complex data in different ways that reveal sharp regularities. (1) We show how the number of firms selling in a market varies with the size of the market. (2) We look at features of the distribution of sales within individual markets. (3) We examine sets of destinations where firms sell. (4) We look at sales in France by firms selling (a) to more destinations and (b) to less popular destinations. (5) We compare what firms sell in export markets relative to their sales in France.

2.1 Market Entry

Figure 1a plots the number of French manufacturing firms N_{nF} selling to a market against total manufacturing absorption X_n in that market across our 113 markets.[2] While the number of firms selling to a market tends clearly to increase with the size of the market, the relationship is a cloudy one. Note in particular that more French firms sell to France than its market size

[1] Eaton, Kortum, and Kramarz (EKK, 2004), describe the data in detail, partitioning firms into 16 manufacturing sectors. While features vary across industries, enough similarity remains to lead us to ignore the industry dimension here.

[2] Manufacturing absorption is calculated as total production plus imports minus exports. Aggregate trade data are from Feenstra (2000) while production data are from UNIDO (2000). See EKK (2004) for details.

would suggest.

The relationship comes into focus, however, when the number of firms is normalized by the share of France in a market. Figure 1b continues to report market size across the 113 destinations along the x axis. The y axis replaces the number of French firms selling to a market with that number divided by French market share, π_{nF}, defined as total French exports to that market, X_{nF}, divided by the market's total absorption X_n, i.e.,

$$\pi_{nF} = \frac{X_{nF}}{X_n}.$$

Note that the relationship is not only very tight, but linear in logs. Correcting for market share pulls France from the position of a large positive outlier to a slightly negative one. A regression line has a slope of 0.65.

If we make the assumption that French firms don't vary systematically in size from other (non-French) firms selling in a market, the measure on the y axis indicates the total number of firms selling in a market. We can then interpret Figure 1b as telling us how the number of sellers varies with market size.

Models of perfect and Bertrand competition and the standard model of monopolistic competition without market-specific entry costs predict that the number of sellers in a market is invariant to market size. Figures 1a and 1b compel us to abandon these approaches.

The number of firms selling to a market increases with market size, but with an elasticity less than one. A mirror relationship (not shown) is that average sales per firm increase with market size as well, again with an elasticity less than one. We can get a fuller picture by asking how sales per firm rise with market size at different points in the sales distribution there. Figure 1c shows this relationship, reporting the 95th, 75th, 50th, and 25th percentiles of sales in each market (on the y axis) against market size (on the x axis). The upward drift is apparent across the board, although more weakly for the 25th percentile.

2.2 Sales Distributions

Our second exercise is to look at the distribution of sales within individual markets. We plot the sales of each firm in a particular market (relative to mean sales there) against the fraction of firms selling in the market who sell at least that much.[3] Doing so for all our 113 destinations a remarkable similarity emerges. Figure 2 plots the results for Belgium-Luxembourg, France, Ireland, and the United States, on common axes. Since there are many fewer firms exporting than selling in France the upper percentiles in the foreign destinations are empty. Nonetheless, the shape is about the same.

To interpret these figures as distributions, let x_n^q be the q'th percentile of French sales in market n normalized by mean sales in that market. We can write:

$$\Pr\left[x_n \leq x_n^q\right] = q$$

where x_n is sales of a firm in market n relative to the mean. Suppose the sales distribution is Pareto with parameter $a > 1$ (so that the minimum sales relative to the mean is $(a-1)/a$). We could then write:

$$1 - \left(\frac{a x_n^q}{a-1}\right)^{-a} = q$$

[3] Following Gabaix and Ibragimov (2008) we construct the x axis as follows. Denote the rank in terms of sales of French firm j in market n, among the N_{nF} French firms selling there, as $r_n(j)$, with the firm with the largest sales having rank 1. For each firm j the point on the x axis is $(r_n(j) - .5)/N_{nF}$.

or:

$$\ln\left(x_n^q\right) = \ln\left(\frac{a-1}{a}\right) - \frac{1}{a}\ln(1-q),$$

implying a straight line with slope $-1/a$. At the top percentiles the slope does appear nearly constant and below -1 but at the lower tails it is much steeper, reflecting the presence of suppliers selling very small amounts.[4] This shape is well known in the industrial organization literature looking at various size measures in the home market.[5] What we find here is that this shape is inherited across markets looking at the same set of potential sellers.

2.3 Entry Strings

We now examine entry into different markets by individual firms. As a starting point for this examination, suppose firms obey a hierarchy in the sense that any firm selling to the $k+1$st most popular destination necessarily sells to the kth most popular destination as well. Not surprisingly firms are less orderly in their choice of destinations. A good metric of how far they depart from a hierarchy is elusive. We can get some sense, however, by looking simply at exporters to the top seven foreign destinations. Table 1 reports these destinations and the number of firms selling to each. It also reports the total number that export to at least one of these destinations and the total number of exporters. Note that 4,106 of the 34,035 exporters, constituting only 12 percent, don't sell in the top 7. The last column of the table reports, for each top 7 destination, the marginal probability of selling there conditional on selling somewhere among the top 7.

[4] Considering only sales by the top 1 percent of French firms selling in the four destinations depicted in Figure 2, regressions yield slopes of -0.74 (Belgium), -0.87 (France), -0.69 (Ireland) and -0.82 (United States).

[5] See Simon and Bonini (1958) and Luttmer (2006), among many, for a discussion and explanations.

We use these marginal probabilities to calculate the probability of selling to the markets in the order prescribed by a hierarchy if the probabilities of selling in any market were independent across markets. The first column of Table 2 lists each of the strings of destinations that obey a hierarchical structure while column 2 reports the number of firms exporting to that string, irrespective of their export activity outside the string. The third column reports the probability that a firm would export to that string if the probabilities of exporting to each destination were independent across markets. Independence implies that only 13.5 percent of exporters would obey the required ordering (for example, selling to Belgium and Germany but not the other five), so that 86.5 would deviate from it (e.g., by selling to Belgium, Germany, and the United Kingdom but nowhere else). In fact more than twice that number, 30.9 percent, adhere to the hierarchy. Column 4 reports the implied number of firms that would sell to each string under independence. Note that many more sell to the short strings and fewer to the long strings than independence would imply. We conclude that a model needs to recognize both a tendency for firms to export according to a hierarchy while allowing them significant latitude to depart from it.

2.4 Export Participation and Size in France

How does a firm's participation in export markets relate to its sales in France? We organize our firms in two different ways based on our examination of their entry behavior above.

First, we group firms according to the minimum number of destinations where they sell. All of our firms, of course, sell to at least one market while none sell to all 113 destinations. Figure 3a depicts average sales in France on the y axis for the group of firms that sell to at least k markets with k on the x axis. Note the near monotonicity with which sales in France rise with the number of foreign markets served.

Figure 3b reports average sales in France of firms selling to k or more markets against the number of firms selling to k or more markets. The highly linear, in logs, negative relationship between the number of firms that export to a group of countries and their sales in France is highly suggestive of a power law. The regression slope is -0.66.

Second, we rank countries according to their popularity as destinations for exports. The most popular destination is of course France itself, where all of our firms sell, followed by Belgium-Luxembourg with 17,699 exporters. The least popular is Nepal, where only 43 French firms sell (followed in unpopularity by Afghanistan and Uganda, with 52 each). Figure 3c depicts average sales in France on the y axis plotted against the number of firms selling to the kth most popular market on the x axis. The relationship is tight and linear in logs as in Figure 3b, although slightly flatter, with a slope of -0.57. Selling to less popular markets has a very similar positive association with sales in France as selling to more markets.

We conclude that firms that sell to less popular markets and sell to more markets systematically sell more in France. Delving further into the French sales of exporters to markets of varying popularity, Figure 3d reports the 95th, 75th, 50th, and 25th percentile of sales in France (on the y axis) against the number of firms selling to each market. Note the tendency of sales in France to rise with the unpopularity of a destination across all percentiles (less systematically so for the 25th percentile). A challenge for modeling is reconciling the stark linear (in logs) relationships in Figures 3b, 3c, and 3d with the more nuanced size distributions in Figure 2.[6]

[6]We were able to observe the relationship between market popularity and sales in France for the 1992

2.5 Export Intensity

Having looked separately at what exporters sell abroad and what they sell in the French market, we now examine the ratio of the two. We introduce the concept of firm's j's normalized export intensity in market n which we define as:

$$\frac{X_{nF}(j)/\overline{X}_{nF}}{X_{FF}(j)/\overline{X}_{FF}}.$$

Here $X_{nF}(j)$ is French firm j's sales in market n and $\overline{X}_{nF}$ are average sales by French firms in market n ($X_{FF}(j)$ and $\overline{X}_{FF}$ are the corresponding magnitudes in France). Scaling by $\overline{X}_{nF}$ removes any effect of market n as it applies to sales of all French firms there. Scaling by $X_{FF}(j)$ removes any direct effect of firm size.

Figure 4 plots median normalized export intensity for each foreign market n (on the y axis) against the number of firms selling to that market (on the x axis) on log scales. Two aspects stand out.

First, while we have excluded France from the figure, its y coordinate would be 1. Note that the y coordinates in the figure are at least an order of magnitude below one. Hence, a typical exporter's sales are more oriented toward the domestic market.

Second, as a destination becomes more popular, normalized export intensity rises. The slope is 0.38. Hence if the number of sellers to a market rises by 10 percent, normalized export intensity rises by around 4 percent.

cross-section as well. The analog (not shown) of Figure 3c is nearly identical. Furthermore, the changes between 1986 and 1992 in the number of French firms selling in a market correlates as it should with changes in the mean sales in France of these firms. The only glaring discrepancy is Iraq, where the number of French exporters plummeted between the two years, while average sales in France did not skyrocket to the extent that the relationship would dictate.

3 Theory

In seeking to explain these relationships we turn to a parsimonious model that delivers predictions about where firms sell and how much they sell there. We infer the parameter values of the model from our observations on French firms. To this end we build on Melitz (2003), Helpman, Melitz, and Yeaple (2004), Chaney (2008), and Arkolakis (2008). The basic structure is monopolistic competition: goods are differentiated with each one corresponding to a firm; selling in a market requires a fixed cost while moving goods from country to country incurs iceberg transport costs; firms are heterogeneous in efficiency as well as in other characteristics while countries vary in size, location, and fixed cost of entry.

We begin with the determination of unit costs of different products in different countries around the world (whether or not these products are produced or supplied in equilibrium). Unit costs depend on input costs, trade barriers, and underlying heterogeneity in the efficiency of potential producers in different countries.

3.1 Producer Heterogeneity

A potential producer of good j in country i has efficiency $z_i(j)$. A bundle of inputs there costs w_i, so that the unit cost of producing good j is $w_i/z_i(j)$. Countries are separated by iceberg trade costs, so that delivering one unit of a good to country n from country i requires shipping $d_{ni} \geq 1$ units, where we set $d_{ii} = 1$ for all i. Combining these terms, the unit cost to this producer of delivering one unit of good j to country n from country i is:

$$c_{ni}(j) = \frac{w_i d_{ni}}{z_i(j)}. \tag{1}$$

The measure of potential producers in country i who can produce their good with efficiency at least z is:

$$\mu_i^z(z) = T_i z^{-\theta} \quad z > 0, \tag{2}$$

where $\theta > 0$ is a parameter.[7] Using (1), the measure of goods that can be delivered from country i to country n at unit cost below c is $\mu_{ni}(c)$ defined as:

$$\mu_{ni}(c) = \mu_i^z\left(\frac{w_i d_{ni}}{c}\right) = T_i(w_i d_{ni})^{-\theta} c^{\theta}.$$

The measure of goods that can be delivered to country n from anywhere at unit cost c or less is therefore:

$$\mu_n(c) = \sum_{i=1}^{N} \mu_{ni}(c) = \Phi_n c^{\theta}, \tag{3}$$

where $\Phi_n = \sum_{i=1}^{N} T_i(w_i d_{ni})^{-\theta}$.

Within this measure, the fraction originating from country i is:

$$\frac{\mu_{ni}(c)}{\mu_n(c)} = \frac{T_i(w_i d_{ni})^{-\theta}}{\Phi_n} = \pi_{ni}. \tag{4}$$

[7] We follow Helpman, Melitz, and Yeaple (2004) and Chaney (2007) in treating the underlying heterogeneity in efficiency as Pareto. Our observations above on patterns of sales by French firms in different markets are very suggestive of an underlying Pareto distribution. A Pareto distribution of efficiencies can arise naturally from a dynamic process that is a history of independent shocks, as shown by Simon (1956), Gabaix (1999), and Luttmer (2006). The Pareto distribution is closely linked to the type II extreme value (Fréchet) distribution used in Kortum (1997), Eaton and Kortum (1999), Eaton and Kortum (2002), and Bernard, Eaton, Kortum, and Jensen (2003). Say that the range of goods is limited to the interval $j \in [0, J]$ with the measure of goods produced with efficiency at least z given by: $\mu_i^Z(z; J) = J\left\{1 - \exp\left[-(T/J)z^{-\theta}\right]\right\}$ (where $J = 1$ in these previous papers). This generalization allows us to stretch the range of goods while compressing the distribution of efficiencies for any given good. Taking the limit as $J \to \infty$ gives (2). (To take the limit rewrite the expression as $\left\{1 - \exp\left[-(T/J)z^{-\theta}\right]\right\}/J^{-1}$ and apply L'Hôpital's rule.)

where π_{ni}, which arises frequently in what follows, is invariant to c.

We now turn to demand and market structure in a typical destination.

3.2 Demand, Market Structure, and Entry

A market n contains a measure of potential buyers. In order to sell to a fraction f of them a producer selling good j must incur a fixed cost:

$$E_n(j) = \varepsilon_n(j) E_n M(f). \tag{5}$$

Here $\varepsilon_n(j)$ is a fixed-cost shock specific to good j in market n and E_n is the component of the cost shock faced by all who sell there, regardless of where they come from. The function $M(f)$, the same across destinations, relates a seller's fixed cost of entering a market to the share of consumers it reaches there. Any given buyer in the market has a chance f of accessing the good while f is the fraction of buyers reached.

In what follows we use the specification for $M(f)$ derived by Arkolakis (2008) from a model of the microfoundations of marketing:

$$M(f) = \frac{1 - (1 - f)^{1-1/\lambda}}{1 - 1/\lambda},$$

where the parameter $\lambda \geq 0$ reflects the increasing cost of reaching a larger fraction of potential buyers.[8] This function has the desirable properties that the cost of reaching 0 buyers in a market is 0 and that the total cost is increasing (and the marginal cost weakly increasing) in the fraction f of buyers reached. Taking the limit $\lambda \to \infty$ implies a constant marginal cost of reaching an added buyer. Since buyers in a market turn out to be identical a seller would then

[8] At $\lambda = 1$ the function becomes $M(f) = \ln(1 - f)$.

choose to reach either every potential buyer in a market or none at all, an outcome equivalent to Melitz (2003). As shown by Arkolakis (2008), and as we replicate below, with λ finite a seller might still stay out of a market entirely or make the effort to reach only a small fraction of buyers there. A seller with a lower unit cost in a market undertakes greater effort to reach more buyers there.

Each potential buyer in market n has the same probability f of being reached by a particular seller that is independent across sellers. Hence each buyer can purchase the same measure of goods, although the particular goods in question vary across buyers. Buyers combine goods according to a constant elasticity of substitution aggregator with elasticity σ, where we require $\theta - 1 > \sigma > 1$. Hence we can write the aggregate demand for good j, if it has price p and reaches a fraction f of the buyers in market n, as:

$$X_n(j) = \alpha_n(j) f X_n \left(\frac{p}{P_n}\right)^{1-\sigma}$$

where X_n is total spending there. The term $\alpha_n(j)$ reflects an exogenous demand shock specific to good j in market n. The term P_n is the CES price index, which we derive below.

Conditional on selling in a market the producer of good j with unit cost $c_n(j)$ who charges a price p and reaches a fraction f of buyers earns a profit:

$$\Pi_n(p, f) = \left(1 - \frac{c_n(j)}{p}\right) \alpha_n(j) f \left(\frac{p}{P_n}\right)^{1-\sigma} X_n - \varepsilon_n(j) E_n \frac{1 - (1-f)^{1-1/\lambda}}{1 - 1/\lambda}. \tag{6}$$

Given its unit cost $c_n(j)$ and idiosyncratic sales and fixed-cost shifters $\alpha_n(j)$ and $\varepsilon_n(j)$ this expression is the same for any seller in market n regardless of its location. We now turn to the profit maximizing choices of p and f.

A producer will set the standard Dixit-Stiglitz (1977) markup over unit cost:

$$p_n(j) = \overline{m} c_n(j)$$

where:

$$\overline{m} = \frac{\sigma}{\sigma - 1}.$$

and seek a fraction:

$$f_n(j) = \max\left\{1 - \left[\eta_n(j)\frac{X_n}{\sigma E_n}\left(\frac{\overline{m}c_n(j)}{P_n}\right)^{1-\sigma}\right]^{-\lambda}, 0\right\} \tag{7}$$

of buyers in the market where:

$$\eta_n(j) = \frac{\alpha_n(j)}{\varepsilon_n(j)},$$

is the entry shock in market n given by the ratio of the demand shock to the fixed-cost shock. Note that it won't sell at all, hence avoiding any fixed cost there, if:

$$\eta_n(j)\left(\frac{\overline{m}c_n(j)}{P_n}\right)^{1-\sigma}\frac{X_n}{\sigma} \geq E_n.$$

Having now solved for the profit maximizing price $p_n(j)$ and entry effort $f_n(j)$ we can describe a seller's behavior in market n in terms of its unit cost $c_n(j) = c$, demand shock $\alpha_n(j) = \alpha$, and entry shock $\eta_n(j) = \eta$. From the condition above, a firm enters market n if and only if:

$$c \leq \overline{c}_n(\eta) \tag{8}$$

where:

$$\overline{c}_n(\eta) = \left(\eta\frac{X_n}{\sigma E_n}\right)^{1/(\sigma-1)}\frac{P_n}{\overline{m}}. \tag{9}$$

We can use the expression for (9) to simplify the expression for the fraction of buyers a producer with unit cost $c \leq \overline{c}_n(\eta)$ will reach:

$$f_n(\eta, c) = 1 - \left(\frac{c}{\overline{c}_n(\eta)}\right)^{\lambda(\sigma-1)}. \tag{10}$$

Its total sales are then:

$$X_n(j) = \alpha f_n(\eta, c) \left(\frac{\overline{m}c}{P_n}\right)^{1-\sigma} X_n. \tag{11}$$

Since it charges a markup $\overline{m}$ over unit cost its total gross profit is simply:

$$\Pi^G(j) = X_n(j)/\sigma \tag{12}$$

some of which is covering its fixed cost:

$$E_n(j) = \frac{\alpha}{\eta} E_n M(f_n(\eta, c)). \tag{13}$$

To summarize, the relevant characteristics of market n that apply across sellers are total purchases X_n, the price index P_n, and the common component of the fixed cost E_n. The particular situation of a potential seller of product j in market n is captured by three magnitudes: the unit cost $c_n(j)$ and the demand and entry shocks $\alpha_n(j)$ and $\eta_n(j)$. We treat $\alpha_n(j)$ and $\eta_n(j)$ as the realizations of producer-specific shocks drawn from a joint density $g(\alpha, \eta)$ that is the same across destinations n and independent of $c_n(j)$.

Equations (8) and (9), governing entry, and (11), governing sales conditional on entry, link our theory to the data on French firms' entry and sales in different markets of the world described in Section 2. Before returning to the data, however, we need to solve for the price index P_n in each market.

3.3 The Price Index

As described above, each buyer in market n has access to the same measure of goods (even though they are not necessarily the same goods). Every buyer faces the same probability $f_n(\eta, c)$ of purchasing a good with cost c and entry shock η for any value of α. Hence we can write the price index P_n faced by a representative buyer in market n as:

$$P_n = \overline{m}\left[\int\int\left(\int_0^{\overline{c}_n(\eta)} \alpha f_n(\eta, c)c^{1-\sigma}d\mu_n(c)\right)g(\alpha,\eta)d\alpha d\eta\right]^{-1/(\sigma-1)}.$$

To solve we use, respectively, (3), (10), and the laws of integration, to get:

$$\begin{aligned}
P_n &= \overline{m}\left[\Phi_n\int\int\alpha\left(\int_0^{\overline{c}_n(\eta)} f_n(\eta,c)\theta c^{\theta-\sigma}dc\right)g(\alpha,\eta)d\alpha d\eta\right]^{-1/(\sigma-1)} \\
&= \overline{m}\left[\Phi_n\int\int\alpha\left(\int_0^{\overline{c}_n(\eta)} \theta c^{\theta-\sigma}dc - \overline{c}_n(\eta)^{-\lambda(\sigma-1)}\int_0^{\overline{c}_n(\eta)}\theta c^{\theta-\sigma+\lambda(\sigma-1)}dc\right)g(\alpha,\eta)d\alpha d\eta\right]^{-1/(\sigma-1)} \\
&= \overline{m}\left[\Phi_n\left(\frac{\theta}{\theta-(\sigma-1)} - \frac{\theta}{\theta+(\sigma-1)(\lambda-1)}\right)\int\int\alpha\overline{c}_n(\eta)^{\theta-(\sigma-1)}g(\alpha,\eta)d\alpha d\eta\right]^{-1/(\sigma-1)}.
\end{aligned}$$

Substituting the expression for the entry hurdle (9) into this last expression and simplifying gives:

$$P_n = \overline{m}\left(\kappa_1\Phi_n\right)^{-1/\theta}\left(\frac{X_n}{\sigma E_n}\right)^{(1/\theta)-1/(\sigma-1)} \tag{14}$$

where:

$$\kappa_1 = \left[\frac{\theta}{\theta-(\sigma-1)} - \frac{\theta}{\theta+(\sigma-1)(\lambda-1)}\right]\int\int\alpha\eta^{[\theta-(\sigma-1)]/(\sigma-1)}g(\alpha,\eta)d\alpha d\eta. \tag{15}$$

Note that the price index relates to total expenditure relative to the entry cost with an elasticity of $1/\theta - 1/(\sigma-1)$. Our restriction that $\theta > \sigma - 1$ assures that the effect is negative: A larger market enjoys lower prices, a manifestation of Krugman's (1980) "home market effect" common across models of monopolistic competition. Our parameter estimates will give us a sense of its magnitude.

Having solved for the price index P_n we return to the sales and entry of an individual firm there.

3.4 Firm Entry and Sales

We can now restate the conditions for entry, (8) and (9), and the expression for sales conditional on entry, (11), in terms of the parameters underlying the price index. A firm j with unit cost c and sales and entry shocks α and η will enter market n if c and η satisfy:

$$c \leq \overline{c}_n(\eta).$$

where, substituting the price index (14) into (9):

$$\overline{c}_n(\eta) = \eta^{1/(\sigma-1)} \left(\frac{X_n}{\sigma E_n \kappa_1 \Phi_n} \right)^{1/\theta}. \tag{16}$$

Substituting (10) and (14) into (11), conditional on entry its sales there are:

$$\begin{aligned} X_n(j) &= \alpha \left[1 - \left(\frac{c}{\overline{c}_n(\eta)} \right)^{\lambda(\sigma-1)} \right] c^{-(\sigma-1)} \left(\frac{X_n}{\sigma E_n \kappa_1 \Phi_n} \right)^{(\sigma-1)/\theta} \sigma E_n \\ &= \varepsilon \left[1 - \left(\frac{c}{\overline{c}_n(\eta)} \right)^{\lambda(\sigma-1)} \right] \left(\frac{c}{\overline{c}_n(\eta)} \right)^{-(\sigma-1)} \sigma E_n \end{aligned} \tag{17}$$

Note that ε has replaced α as the shock to sales. Firms that have overcome a higher entry hurdle must sell more for entry to be worthwhile.

Knowing now what an individual firm does in market n, we turn to aggregate firm behavior in that market.

3.5 Aggregate Entry and Sales

For firms with a given η in market n a measure $\mu_n(\overline{c}_n(\eta))$ will pass the entry hurdle. Integrating across the marginal density $g_2(\eta)$, the measure of entrants into market n is:

$$J_n = \int \left[\mu_n(\overline{c}_n(\eta)) \right] g_2(\eta) d\eta = \Phi_n \int [\overline{c}_n(\eta_n)^\theta] g_2(\eta) d\eta = \frac{\kappa_2}{\kappa_1} \frac{X_n}{\sigma E_n} \tag{18}$$

where:

$$\kappa_2 = \int \eta^{\theta/(\sigma-1)} g_2(\eta) d\eta \tag{19}$$

Note that this measure rises in proportion to X_n.[9]

Suppliers to market n have heterogeneous costs. But, conditional on entry, suppliers from each source country i have the same distribution of unit costs in n. To see why, consider good j in market n with entry shock η. For any cost c less than the entry threshold, the fraction of suppliers from i with $c_{ni}(j) \leq c$ among those with $c_{ni}(j) \leq \overline{c}_n(\eta)$ is simply

$$\mu_{ni}(c)/\mu_{ni}(\overline{c}_n(\eta)) = [c/\overline{c}_n(\eta)]^\theta$$

for any $c \leq \overline{c}_n(\eta)$. Hence for any η this proportion does not depend on source i. Since we assume that the distribution of η is independent i, different sources will have different measures of suppliers selling in market n, but all who do sell will have the same distribution of unit costs.

Hence, given the constant markup over unit cost, suppliers from any source have the same distribution of prices in n and, hence, of sales. An implication is that the fraction of entrants into n coming from i, π_{ni}, is also the fraction of spending by country n on goods originating from country i:

$$\pi_{ni} = \frac{X_{ni}}{X_n}, \tag{20}$$

where X_{ni} is n's purchases on goods originating from i. This relationship gives us a connection between the cluster of parameters embedded in π_{ni} in (4) above and data on trade shares.

[9] In describing the data we used N to indicate a number of firms (an integer, of course). Since the theory implies a continuum of firms we use J to denote a measure of producers.

Combining (18) and (20), we get that the measure of firms from country i selling in country n is:

$$J_{ni} = \pi_{ni} J_n = \frac{\kappa_2}{\kappa_1} \frac{\pi_{ni} X_n}{\sigma E_n}. \tag{21}$$

Hence the measure of firms from source i in destination n is proportional to i's trade share π_{ni} there and to market size X_n relative to E_n. Average sales of these firms $\overline{X}_{ni}$ is:

$$\overline{X}_{ni} = \frac{\pi_{ni} X_n}{\pi_{ni} J_n} = \frac{\kappa_1}{\kappa_2} \sigma E_n. \tag{22}$$

The distribution of sales in a particular market, and hence mean sales there, is invariant to the location i of the supplier.

3.6 Fixed Costs and Profits

Since our model is one of monopolistic competition, producers charge a markup over unit cost. If total spending in a market is X_n then gross profits earned across firms in that market are X_n/σ. If firms were homogeneous then fixed costs would fully dissipate profits. But, with producer heterogeneity, firms with a unit cost below the entry cutoff in a market earn a positive profit there. Here we solve for the share of profits that are dissipated by fixed costs. While not used in our estimation below, this derivation delivers a useful implication of the model which we can quantify once we have estimated the model's parameters.

We return to the expression for a firm's fixed cost in destination n (13), substituting (7):

$$E_n(\alpha, \eta, c) = \frac{\alpha}{\eta} E_n \frac{1 - \left(\frac{c}{\overline{c}_n(\eta)}\right)^{(\lambda-1)(\sigma-1)}}{1 - 1/\lambda}.$$

Integrating across the range of unit costs consistent with entry into destination n (given η)

while using (3) and (9), gives us:

$$
\begin{aligned}
E_n(\alpha,\eta) &= \frac{\alpha}{\eta}E_n\frac{1-\Phi_n\theta\int_0^{\bar{c}_n(\eta)}\left(\frac{c}{\bar{c}_n(\eta)}\right)^{(\lambda-1)(\sigma-1)}c^{\theta-1}dc}{1-1/\lambda} \\
&= \alpha\eta^{(\theta-(\sigma-1))/(\sigma-1)}\left(\frac{X_n}{\sigma\kappa_1}\right)\left[\frac{\lambda}{\theta/(\sigma-1)+\lambda-1}\right].
\end{aligned}
$$

Integrating across the joint density of α and η, inserting (15), we get that total fixed costs in a market $\overline{E}_n$ are:

$$
\overline{E}_n = \int\int E_n(\alpha,\eta)g(\alpha,\eta)d\alpha d\eta = \frac{X_n}{\sigma\theta}[\theta-(\sigma-1)]. \tag{23}
$$

Thus total entry costs are a fraction $[\theta-(\sigma-1)]/\theta$ of the gross profits X_n/σ earned in any destination n. Net profits earned in market n are simply $X_n/(\overline{m}\theta)$.

3.7 A Streamlined Representation

We now employ a change of variables that simplifies the model in two respects. First, it allows us to characterize unit cost heterogeneity in terms of a uniform measure. Second, it allows us to consolidate parameters.

To isolate the heterogeneous component of unit costs we transform the efficiency draw of any potential producer in France as:

$$
u(j) = T_F z_F(j)^{-\theta}. \tag{24}
$$

We refer to $u(j)$ as firm j's standardized unit cost. From (2), the measure of firms with standardized unit cost below u equals the measure with efficiency above $(T_F/u)^{1/\theta}$ which is simply $\mu_F^z((T_F/u)^{1/\theta}) = u$. Hence standardized costs have a uniform measure that doesn't depend on any parameters.

Substituting (24) into (1) and using (4), we can write unit cost in market n in terms of $u(j)$ as:

$$c_n(j) = \frac{w_F d_{nF}}{z_F(j)} = \left(\frac{u(j)}{\pi_{nF}}\right)^{1/\theta} \Phi_n^{-1/\theta}. \tag{25}$$

Associated with the entry hurdle $\overline{c}_n(\eta)$ is an entry hurdle $\overline{u}_n(\eta)$ satisfying:

$$\overline{c}_n(\eta) = \left(\frac{\overline{u}_n(\eta)}{\pi_{nF}}\right)^{1/\theta} \Phi_n^{-1/\theta}. \tag{26}$$

Firm j will enter market n if its $u(j)$ and $\eta_n(j)$ satisfy:

$$u(j) \leq \overline{u}_n(\eta_n(j)) = \left(\frac{X_{nF}}{\kappa_1 \sigma E_n}\right) \eta_n(j)^{\widetilde{\theta}} \tag{27}$$

where

$$\widetilde{\theta} = \frac{\theta}{\sigma - 1} > 1. \tag{28}$$

Conditional on firm j's passing this hurdle we can use (25) and (26) to rewrite firm j's sales in market n, expression (17), in terms of $u(j)$ as:

$$X_{nF}(j) = \varepsilon_n(j) \left[1 - \left(\frac{u(j)}{\overline{u}_n(\eta_n(j))}\right)^{\lambda/\widetilde{\theta}}\right] \left(\frac{u(j)}{\overline{u}_n(\eta_n(j))}\right)^{-1/\widetilde{\theta}} \sigma E_n \tag{29}$$

Equations (27) and (29) reformulate the entry and sales equations (16) and (17) in terms of $u(j)$ rather than $c_n(j)$.

Since standardized unit cost $u(j)$ applies across all markets it gets to the core of a firm's underlying efficiency as it applies to its entry and sales in different markets. Notice that in reformulating the model as (27) and (29), the two parameters θ and σ enter only collectively through the parameter $\widetilde{\theta}$. It translates unobserved heterogeneity in $u(j)$ into observed heterogeneity in sales. A higher value of θ implies less heterogeneity in efficiency while a higher value of σ means that a given level of heterogeneity in efficiency translates into greater heterogeneity in sales. Since we observe sales and not underlying efficiency we are able to identify only $\widetilde{\theta}$.

3.8 Connecting the Model to the Empirical Regularities

We now show how the model can deliver the features of the data about entry and sales described in Section 2. In doing so we quantify total French sales X_{nF} in each of our 113 destinations with the actual data. We quantify the measure of French firms J_{nF} selling in each destination with the actual (integer) number N_{nF}.

Aggregate Entry. From (21) we get:

$$\frac{N_{nF}}{\pi_{nF}} = \frac{\kappa_2}{\kappa_1}\frac{X_n}{\sigma E_n}, \tag{30}$$

a relationship between the number of French firms selling to market n relative to French market share and the size of market n, just like the one plotted in Figure 1b. The fact that the relationship is tight with a slope that is positive but less than one suggests that entry cost σE_n rises systematically with market size, but not proportionately so.

We don't impose any such relationship, but rather employ (30) to calculate:

$$\sigma E_n = \frac{\kappa_2}{\kappa_1}\frac{X_{nF}}{N_{nF}} = \frac{\kappa_2}{\kappa_1}\overline{X}_{nF} \tag{31}$$

directly from the data.

We can use equation (31) to examine how fixed costs vary with country characteristics. Regressing $\overline{X}_{nF}$ against our market size measure (both in logarithms) yields a slope of 0.31 (with a standard error of 0.02). This relationship relates to the slope in Figure 1b, showing that the number of entrants rises with market size with an elasticity of 0.65. Larger markets attract more firms, but not in proportion, since the cost of entry rises as well. The firms that do enter sell more, generating an overall elasticity of total sales with respect to market size of 0.96 (in line with the gravity literature).[10]

Firm Entry. Using (31) we can write (27) in terms of observables as:

$$u(j) \leq \overline{u}_n(\eta_n(j)) = \frac{N_{nF}}{\kappa_2}\eta_n(j)^{\tilde{\theta}}. \tag{32}$$

Without variation in the firm and market specific entry shock $\eta_n(j)$, (32) would imply efficiency is all that would matter for entry, dictating a deterministic ranking of destinations with a less efficient firm (with a higher $u(j)$) selling to a subset of the destinations served by any more efficient firm. Hence deviations from market hierarchies identify variation in $\eta_n(j)$. As Table 2 illustrates, there is some tendency for firms to enter markets according to a hierarchy, but it is a loose one.

Sales in a Market. To get further insight into what our specification implies for the distribution of sales within a given market n note that, conditional on a firm's entry, the term:

$$v_n(j) = \frac{u(j)}{\overline{u}_n(\eta_n(j))} \tag{33}$$

[10] If we add the logarithm of 1986 real GDP per capita (from the World Bank's *World Development Indicators*) as an additional right-hand side variable, the coefficient on the logarithm of market size rises to 0.41 (standard error 0.03) while the coefficient on the logarithm of real GDP per capita is -0.29 (standard error 0.06). Hence while larger markets have a higher fixed cost of entry, given size the cost is lower in richer countries.We were not able to obtain 1986 real GDP per capita for 10 of our 113 destinations (Albania, Angola, Bulgaria, Czechoslovakia, East Germany, Libya, USSR, Vietnam, Yugoslavia, and Zaire). Hence this second regression was performed on the remaining 103 countries. (The coefficient on the logarithm of market size in the univariate regression for this smaller group is .30, hardly different from that for the larger group.)

is distributed uniformly on $[0,1]$. Replacing $u(j)$ with $v_n(j)$ in expression (29) and exploiting (31) we can write sales as:

$$X_{nF}(j) = \varepsilon_n(j) \left[1 - v_n(j)^{\lambda/\widetilde{\theta}}\right] v_n(j)^{-1/\widetilde{\theta}} \frac{\kappa_2}{\kappa_1} \overline{X}_{nF}. \tag{34}$$

Not only does v_n have the same distribution in each market n, so does ε_n.[11] Hence the distribution of sales in any market n is identical up to a scaling factor equal to $\overline{X}_{nF}$ (reflecting variation in σE_n). Hence we can generate the common shapes of sales distributions exhibited in Figure 2. The variation introduced by ε_n explains why the sales distribution in a market might inherit the lognormal characteristics apparent in that Figure. A further source of curvature is the term in square brackets, representing the fraction of buyers reached. As $v_n(j)$ goes to one, with finite λ, the fraction approaches zero, capturing the curvature of sales distributions at the lower end, as observed in Figures 2. Finally the term $v_n(j)^{-1/\widetilde{\theta}}$ instills Pareto features into the distribution. These features will be more pronounced as $v_n(j)$ approaches zero since very efficient firms will be reaching almost all buyers. The component of $v_n(j)$ attributable to $u(j)$ carries across markets.

Sales in France conditional on Selling in a Foreign Market. We can also look at the sales in France of French firms selling to any market n. To condition on these firms' selling in market n we take (34) as it applies to France and use (33) and (32) to replace $v_F(j)$ with

[11] To see that the distribution of $\varepsilon_n(j)$ is the same in any n consider the joint density of α and η conditional on entry into market n:

$$\frac{\overline{u}_n(\eta)}{\int \overline{u}_n(\eta')g_2(\eta')d\eta'} g(\alpha,\eta) = \frac{\eta^{\widetilde{\theta}}}{\kappa_2} g(\alpha,\eta)$$

which does not depend on n.

$v_n(j)$:

$$X_{FF}(j)|_n = \frac{\alpha_F(j)}{\eta_n(j)} \left[1 - v_n(j)^{\lambda/\widetilde{\theta}} \left(\frac{N_{nF}}{N_{FF}} \right)^{\lambda/\widetilde{\theta}} \left(\frac{\eta_n(j)}{\eta_F(j)} \right)^{\lambda} \right] v_n(j)^{-1/\widetilde{\theta}} \left(\frac{N_{nF}}{N_{FF}} \right)^{-1/\widetilde{\theta}} \frac{\kappa_2}{\kappa_1} \overline{X}_{FF}. \quad (35)$$

By expressing (35) in terms of $v_n(j)$ we can exploit the fact that (given entry in n) $v_n(j)$ is uniformly distributed on the unit interval. This expression relates to Figures 3c and 3d, which depict features of the distribution of sales in France of firms selling to some other market n according to the number N_{nF} that sell there. Note first that since all the other terms on the right-hand of (35) have the same value or distribution across markets, N_{nF} is the only systematic source of variation across n. Interpreting Figure 3c in terms of Equation (34), the slope of -0.57 suggests a value of $\widetilde{\theta}$ of 1.75.

To relate the equation to Figures 3c and 3d we begin with the expression in square brackets, representing the fraction of buyers reached in France. As in (34), a firm whose $v_n(j)$ approaches one will sell to only a small fraction of buyers (with finite λ). But compared with (34), this effect is muted by the lower number of firms selling in n relative to those selling in France. With numbers like those in the data, French exporters are very far from the margin of entry into France so will reach nearly all buyers there. Since the term in square brackets is close to one for most exporting firms, the component $N_{nF}^{-1/\widetilde{\theta}}$ to its right dominates the relationship. This relationship implies that the entire distribution of sales in France should shift with N_{nF} on a log scale. Hence together (34) and (35) reconcile the near loglinearity of sales in France with N_{nF} and the extreme curvature at the lower end of the sales distribution in any given market. The gap between the percentiles in Figure 3d is governed by the variation in the demand shock α_F in France together with variation in the entry shock $\eta_n(j)$ in country n. The entry shock in n produces variation across firms in how much entry into n says about the resulting distribution of a firm's normalized cost.

Normalized Export Intensity. Finally, we can calculate firm j's normalized export intensity in market n:

$$\frac{X_{nF}(j)}{X_{FF}(j)} / \left(\frac{\overline{X}_{nF}}{\overline{X}_{FF}}\right) = \frac{\alpha_n(j)}{\alpha_F(j)} \left[\frac{1 - v_n(j)^{\lambda/\widetilde{\theta}}}{1 - v_n(j)^{\lambda/\widetilde{\theta}} \left(\frac{N_{nF}}{N_{FF}}\right)^{\lambda/\widetilde{\theta}} \left(\frac{\eta_n(j)}{\eta_F(j)}\right)^{\lambda}}\right] \left(\frac{N_{nF}}{N_{FF}}\right)^{1/\widetilde{\theta}}. \qquad (36)$$

Figure 4 plots the median of this statistic for French firms across export markets. Note first how the presence of the sales shock $\alpha_n(j)$ accommodates random variation in sales in different markets conditional upon entry.

As in (35), the only systematic source of cross-country variation on the right-hand side is in the number of French firms. In contrast to (34) and (35), however, the firm's overall efficiency $v_n(j)$ has no direct effect on normalized export intensity since it cancels (having the same effect in n as it has in France). For $n = F$ the relationship collapses to an identity $1 = 1$. For $n \neq F$ $N_{nF} \ll N_{FF}$ implying that the term in square brackets is much less than one. The reason is that an exporter selling in France is likely to be very far from the entry cutoff so reaches most buyers while it can be quite marginal in its export destination. Hence our model explains the low numbers on the y axis of Figure 4.

Aside from this general collapse of sales to any export market n relative to sales in France, the last term in Equation (36) predicts that normalized export intensity will increase with the number of French firms selling there, the relationship portrayed in Figure 4. The reason is that the harder the market is to enter (i.e., the lower N_{nF}/N_{FF}), the lower is unit cost in France, but the distribution in foreign destination n is the same. According to (36) the elasticity of normalized export intensity with respect to N_{nF}/N_{FF} is $1/\widetilde{\theta}$. The slope coefficient of 0.38 reported in Section 2.5 suggests a value of $\widetilde{\theta}$ of 2.63.[12]

To say more about the connection between the model and the data we need to go to the computer.

4 Estimation

We estimate the parameters of the model by the method of simulated moments. We begin by parameterizing the joint density of the market-specific shocks α and η. We then explain how we simulate a set of artificial French exporters given a particular set of parameter values, with each firm assigned a cost draw u and an α and η in each market. From this artificial data set we calculate a set of moments and compare them with moments from the actual data on French exporters. We search for parameter values that bring the artificial moments close to the actual ones. We report our results and examine the model's fit.

4.1 Parameterization

To complete the specification, we assume that $g(\alpha, \eta)$ is joint lognormal. Specifically, $\ln \alpha$ and $\ln \eta$ are normally distributed with zero means and variances σ_a^2, σ_h^2, and correlation $\rho_.$ Under

[12]In relating equation (35) to Figures 3c and 3d and equation (36) to Figure 4 a slight discrepancy arises. All the firms in our data sell in France. The theory admits the possibility that a firm might not sell in France but could still sell elsewhere. Hence $X_{FF}(j)$ in these equations actually applies to what firm j's sales in France would be if it did enter. But because the French share in France is so much larger than the French share elsewhere, a French firm selling in another market but not in France is very unlikely. Hence the discrepancy is minor.

these assumptions we may write (15) and (19) as:

$$\kappa_1 = \left[\frac{\widetilde{\theta}}{\widetilde{\theta}-1} - \frac{\widetilde{\theta}}{\widetilde{\theta}+\lambda-1}\right] \exp\left\{\frac{\sigma_a^2 + 2\rho\sigma_a\sigma_h(\widetilde{\theta}-1) + \sigma_h^2(\widetilde{\theta}-1)^2)}{2}\right\} \tag{37}$$

and:

$$\kappa_2 = \exp\left\{\frac{\left(\widetilde{\theta}\sigma_h\right)^2}{2}\right\}. \tag{38}$$

Since the entry cost shock is given by $\ln\varepsilon = \ln\alpha - \ln\eta$, the implied variance of the fixed-cost shock is

$$\sigma_e^2 = \sigma_a^2 + \sigma_h^2 - 2\rho\sigma_a\sigma_h,$$

which is decreasing in ρ.

Our estimation conditions on the actual data for: (i) French sales in each of our 113 destinations, X_{nF}, and (ii) the number of French firms selling there, N_{nF}. With this conditioning our model has only five parameters

$$\Theta = \{\widetilde{\theta}, \lambda, \sigma_a, \sigma_h, \rho\}.$$

We use (31) to back out the cluster of parameters σE_n using our data on $\overline{X}_{nF} = X_{nF}/N_{nF}$ and the κ_1 and κ_2 implied by (37) and (38). Similarly, we use (32) to back out a firm's entry hurdle in each market $\overline{u}_n(\eta_n)$ given its η_n and the κ_2 implied by (38).

4.2 Simulation

For given values of Θ we create an artificial set of French exporters that operate as the model tells them. We truncate the distribution of the cost draw u to ensure that each artificial firm sells in France and exports to at least one foreign destination (weighting the firm appropriately, as described below).

We refer to an artificial French exporter by s and the number of such exporters by S. The number S does not bear any relationship to the number of actual French exporters. A larger S implies less sampling variation in our simulations.

Our simulation proceeds in 3 stages:

1. As we search for parameters we want to hold fixed the realizations of the stochastic components of the model. Hence Stage 1 does not require any parameter values. It involves two steps:

   (a) We draw realizations of $v(s)$'s independently from the uniform distribution $U[0, 1]$, for $s = 1, \ldots, S$, putting them aside to construct standardized unit cost $u(s)$ in Stage 3.

   (b) We draw $S \times 113$ realizations of $a_n(s)$ and $h_n(s)$ independently from:

$$\left[\begin{array}{c} a_n(s) \\ h_n(s) \end{array} \right] \sim N \left[\left(\begin{array}{c} 0 \\ 0 \end{array} \right), \left(\begin{array}{cc} 1 & 0 \\ 0 & 1 \end{array} \right) \right]$$

   putting them aside to construct the $\alpha_n(s)$ and $\eta_n(s)$ in Stage 3.

2. Stage 2 requires a set of parameters Θ and data for each destination n on total sales X_{nF} by French exporters and the number N_{nF} of French firms selling there. It involves two steps:

   (a) Using (37) and (38) we calculate κ_1 and κ_2.

   (b) Using (31) we calculate σE_n for each destination n.

3. Stage 3 combines the simulation draws from Stage 1 and the parameter values and destination variables from Stage 2. It involves seven steps:

(a) We use the draws from 1b and the parameter values from 2a to construct $S \times 113$ realizations for each of $\ln \alpha_n(s)$ and $\ln \eta_n(s)$ as:

$$\begin{bmatrix} \ln \alpha_n(s) \\ \ln \eta_n(s) \end{bmatrix} = \begin{bmatrix} \sigma_a\sqrt{1-\rho^2} & \sigma_a\rho \\ 0 & \sigma_h \end{bmatrix} \begin{bmatrix} a_n(s) \\ h_n(s) \end{bmatrix}$$

(b) We construct the $S \times 113$ entry hurdles:

$$\overline{u}_n(s) = \frac{N_{nF}}{\kappa_2}\eta_n(s)^{\tilde{\theta}}. \tag{39}$$

where $\overline{u}_n(s)$ stands for $\overline{u}_n(\eta_n(s))$.

(c) The simulation algorithm has the flexibility to simulate either exporting firms that sell in France (as we do in estimating parameters) or all firms selling in France (as we do when we explore broader implications of the model) Thus we calculate:

$$\overline{u}^X(s) = \max_{n \neq F}\{\overline{u}_n(s)\},$$

the maximum u consistent with exporting somewhere, and:

$$\overline{u}(s) = \begin{cases} \min\{\overline{u}_F(s), \overline{u}^X(s)\} & \text{to simulate exporters selling in France} \\ \overline{u}_F(s) & \text{to simulate all firms selling in France} \end{cases}$$

In either case, we want $u(s) \leq \overline{u}(s)$ for each artificial exporter s. In other words, $u(s)$ should be a realization from the uniform distribution over the interval $[0, \overline{u}(s)]$. Therefore we construct:

$$u(s) = v(s)\overline{u}(s).$$

using the $v(s)$ from Stage 1.

(d) In the model a measure $\overline{u}$ of firms have standardized unit cost below $\overline{u}$. Our artificial French exporter s therefore gets an importance weight $\overline{u}(s)$. This importance weight will be used in constructing statistics on artificial French exporters that relate to statistics on actual French exporters.[13]

(e) We calculate $\delta_{nF}(s)$, which indicates whether artificial exporter s enters market n, as determined by the entry hurdles:

$$\delta_{nF}(s) = \begin{cases} 1 \text{ if } u(s) \leq \overline{u}_n(s) \\ 0 \text{ otherwise.} \end{cases}$$

where of course $\delta_{FF}(s)$ and a least one other $\delta_{nF}(s)$ necessarily equal 1.

(f) Wherever $\delta_{nF}(s) = 1$ we calculate sales as:

$$X_{nF}(s) = \frac{\alpha_n(s)}{\eta_n(s)} \left[1 - \left(\frac{u(s)}{\overline{u}_n(s)} \right)^{\lambda/\widetilde{\theta}} \right] \left(\frac{u(s)}{\overline{u}_n(s)} \right)^{-1/\widetilde{\theta}} \sigma E_n.$$

This procedure gives us the behavior of S artificial French exporters. We know three things about each one: where it sells, $\delta_{nF}(s)$, how much it sells there, $X_{nF}(s)$, and its importance weight, $\overline{u}(s)$. From these we can compute any moment that could have been constructed from the actual French data.

Our moments are all constructed from the number of firms $\widetilde{N}^k$ predicted to achieve some outcome k. Let $\delta^k(s)$ be an indicator for when artificial firm s achieves outcome k. We simulate the number of firms achieving that outcome as:

$$\widetilde{N}^k = \frac{1}{S} \sum_{s=1}^{S} \overline{u}(s) \delta^k(s). \tag{40}$$

We now explain the moments that we actually seek to match.

[13]See Gouriéroux and Monfort (1995, Chapter 5) for a discussion of the use of importance weights in simulation.

4.3 Moments

In our estimation, we simulate firms that make it into at least one foreign market and into France as well. The reason for the first requirement is that firms that sell only in France are very numerous, and hence capturing them would consume a large portion of simulation draws. But since their activity is so limited they add little to parameter identification.[14] The reason for the second requirement is that key moments in our estimation procedure are based on sales in France by exporters, which we can compute only for firms that sell in the home market.[15] Given parameter estimates, we later explore the implications of the model for nonexporters as well.

For a candidate value Θ we use the algorithm above to simulate the sales of $500,000$ artificial French exporting firms in 113 markets. From these artificial data we compute a vector of moments $\widehat{m}(\Theta)$ analogous to particular moments m in the actual data.

Our moments are all proportions of firms that fall into sets of exhaustive and mutually exclusive bins, where the number of firms in each bin is counted in the data and is simulated from the model using (40). Our bins capture four features of French firms' behavior: (1) their sales in export destination n, (2) their sales in France conditional on selling to n, (3) their sales in n relative to their sales in France conditional on selling to n, and (4) their entry into particular subsets of export markets:

1. The first set of moments relates to the sales distributions presented in Section 2.2. For

[14] Hence we set $\overline{u}(s) = \min\{\overline{u}_F(s), \overline{u}^X(s)\}$ in Stage 3, step c of the simulation algorithm. We also estimated the model matching moments of nonexporting firms as well. Coefficient estimates were similar to those we report below but the estimation algorithm, given estimation time, was much less precise.

[15] There are very few firms (500) apparently not selling in France.

firms selling in each of the 112 export destinations n we compute the qth percentile of sales $s_n^q(1)$ *in that market* (i.e., the level of sales such that a fraction q of firms selling in n sells less than $s_n^q(1)$) for $q = 50, 75, 95$. Using these $s_n^q(1)$ we assign firms that sell in n into four mutually exclusive and exhaustive bins determined by these three sales levels. We compute the proportions $\widehat{m}_n(1;\Theta)$ of artificial firms falling into each bin analogous to the actual proportion $m_n(1) = (0.5, 0.25, 0.2, 0.05)'$. Stacking across the 112 countries gives us $\widehat{m}(1;\Theta)$ and $m(1)$, each with 448 elements (subject to 112 adding-up constraints).

2. The second set of moments relates to the sales in France of exporting firms discussed in Section 2.4. For firms selling in each of the 112 export destinations n we compute the qth percentile of sales $s_n^q(2)$ *in France* for $q = 50, 75, 95$. Proceeding as above we get $\widehat{m}(2;\Theta)$ and $m(2)$, each with 448 elements (subject to 112 adding-up constraints).

3. The third set of moments relates to normalized export intensity by market discussed in Section 2.5. For firms selling in each of the 112 export destinations n we compute the qth percentile of the ratio $s_n^q(3)$ of sales in n to sales in France for $q = 50, 75$. Proceeding as above we get $\widehat{m}(3;\Theta)$ and $m(3)$, each with 336 elements (subject to 112 adding-up constraints).

4. The fourth set of moments relates to the entry strings discussed in Section 2.3. We compute the proportion $\widehat{m}^k(4;\Theta)$ of simulated exporters selling to each possible combination k of the seven most popular export destinations (listed in Table 1). One possibility is exporting yet selling to none of the top seven, giving us 2^7 possible combinations (so that $k = 1, ..., 128$). The corresponding moments from the actual data are simply the

proportion $m^k(4)$ of exporters selling to combination k. Stacking these proportions gives us $\widehat{m}(4;\Theta)$ and $m(4)$, each with 128 elements (now subject to 1 adding up constraint).

Stacking the four sets of moments gives us a 1360-element vector of deviations between the moments of the actual and artificial data:

$$y(\Theta) = m - \widehat{m}(\Theta) = \begin{bmatrix} m(1) - \widehat{m}(1,\Theta) \\ m(2) - \widehat{m}(2,\Theta) \\ m(3) - \widehat{m}(3,\Theta) \\ m(4) - \widehat{m}(4,\Theta) \end{bmatrix}.$$

We base our estimation procedure on the moment condition:

$$\mathbf{E}[y(\Theta_0)] = 0$$

where Θ_0 is the true value of Θ.

4.4 Estimation Procedure

We seek a value of Θ that achieves:

$$\widehat{\Theta} = \arg\min_{\Theta} \left\{ y(\Theta)'\mathbf{W}y(\Theta) \right\},$$

where $\mathbf{W}$ is a 1360×1360 weighting matrix. We search for Θ using the simulated annealing algorithm described in Appendix B. The weighting matrix is the generalized inverse of the estimated variance-covariance matrix $\mathbf{\Omega}$ of the 1360 moments calculated from the data m. We calculate $\mathbf{\Omega}$ using the following bootstrap procedure:

1. We resample, with replacement, 229,900 firms from our initial dataset 2000 times.

2. For each resampling b we calculate m^b, the proportion of firms that fall into each of the 1360 bins, holding the $s_n^q(\tau)$ fixed to calculate $m^b(\tau)$ for $\tau = 1, 2, 3$ and holding the destination strings k fixed for $m^b(4)$.

3. We calculate:

$$\mathbf{\Omega} = \frac{1}{2000} \sum_{b=1}^{2000} \left(m^b - m\right)\left(m^b - m\right)'.$$

Because of the adding up constraints this matrix has rank 1023, forcing us to take its generalized inverse to compute $\mathbf{W}$.

We calculate standard errors using a bootstrap technique, taking into account both sampling error and simulation error. To account for sampling error each bootstrap b replaces m with a different m^b. To account for simulation error each bootstrap b samples a new set of 500,000 v^b's, a_n^b's, and h_n^b's from stage 1 of our simulation algorithm, thus generating a new $\widehat{m}^b(\Theta)$.[16] Defining $y^b(\Theta) = m^b - \widehat{m}^b(\Theta)$ for each b we estimate:

$$\widehat{\Theta}_b = \arg\min_{\Theta} \left\{ y^b(\Theta)'\mathbf{W}y^b(\Theta) \right\}$$

using the same simulated annealing procedure. Doing this exercise 25 times we calculate:

$$V(\Theta) = \frac{1}{25} \sum_{b=1}^{25} \left(\widehat{\Theta}_b - \widehat{\Theta}\right)\left(\widehat{\Theta}_b - \widehat{\Theta}\right)'$$

and take the square roots of the diagonal elements as the standard errors.

4.5 Results

The best fit is achieved at the following parameter values (with bootstrapped standard errors in parentheses):

$\widetilde{\theta}$	λ	σ_a	σ_h	ρ
2.46	0.91	1.69	0.34	−0.65
(0.10)	(0.12)	(0.03)	(0.01)	(0.03)

Our discussion in Section 3.7 foreshadowed our estimate of $\widetilde{\theta}$, which lies between the slopes in Figures 3c and Figure 4. From equations (34), (32), and (33), the characteristic of a firm

[16] Just like m^b, $\widehat{m}^b$ is calculated according to the bins defined from the actual data.

determining both entry and sales conditional on entry, is $v^{-1/\widetilde{\theta}}$, where $v\tilde{~}U[0,1]$. Our estimate of $\widetilde{\theta}$ implies that the ratio of the 75th to the 25th percentile of this term is 1.56. Another way to assess the magnitude of $\widetilde{\theta}$ is by its implication for aggregate fixed costs of entry. Using expression (23), our estimate of 2.463 implies that fixed costs dissipate about 59 percent of gross profit in any destination.

Our estimate of σ_a implies enormous idiosyncratic variation in sales across destinations. In particular, the ratio of the 75th to the 25th percentile of the sales shock α is 9.78. In contrast, our estimate of σ_h means much less idiosyncratic variation in the entry shock η, with a ratio of the 75th to 25th percentile equal to 1.58. As we show more systematically below, the feature of a firm that is common across countries explains relatively little of the variation in sales conditional on entry, but about half of the variation in entry.

A feature of the data is the entry of firms into markets where they sell very little, as seen in Figure 1c. Two features of our estimates reconcile these small sales with a fixed cost of entry. First, our estimate of λ, which is close to one, means that a firm that is close to the entry cutoff incurs a very small entry cost.[17] Second, the negative covariance between the sales and entry shocks explains why a firm with a given u might enter a market and sell relatively little. The first applies to firms that differ systematically in their efficiency while the second applies to the luck of the draw in individual markets.[18]

[17] Arkolakis (2008) finds a value around one consistent with various observations from several countries.

[18] We perform a Monte Carlo test of the ability of our estimation procedure to recover parameter values. We simulate 230,000 artificial French firms with the estimated parameter values reported above. We then apply the estimation procedure, exactly as described, to the simulated data to estimate Θ (using the same weighting matrix $\mathbf{W}$ as in the original estimation). The table below reports the values of used to create the simulated

4.6 Model Fit

We can evaluate the model by seeing how well it replicates features of the data described in Section 2. To glean a set of predictions of our model we use our parameter estimates to simulate $230,000$ artificial firms including nonexporters.[19] We then compare four features of these simulated firms with corresponding features of the actual ones. For the first three we plot the simulated data (x's) and actual data (circles) against market-level characteristics.

Equation (34) in Section 3.7 motivates Figure 5a, which plots the simulated and actual median and 95th percentile of sales to each market against actual mean French sales in that market. The model captures very well both the distance between the two percentiles in any given market and how each percentile varies across markets. The model also nearly matches the amount of noise in these percentiles, especially in markets where mean sales are small.

Equation (35) in Section 3.7 motivates Figure 5b, which plots the median and 95th percentile of sales in France of firms selling to each market against the actual number of firms selling there. Again, the model picks up the spread in the distribution as well as the slope. It also captures the fact that the data point for France is below the line, reflecting the marketing

data (the "truth") and the parameter estimates the estimation procedure delivers:

	$\widetilde{\theta}$	λ	σ_a	σ_h	ρ
"truth"	2.46	0.91	1.69	0.34	−0.65
estimates	2.54	0.67	1.69	0.32	−0.56

Our estimates land in the same ballpark as the true parameters, with deviations in line with the standard errors reported above..

[19] To reduce sampling error our estimation procedure simulates 500,000 firms, restricting attention to exporters. Here we simulate the behavior of 230,000 firms, both non-exporters and exporters, to mimic more closely features of the raw data behind our analysis. In step c of stage 3 in the simulation algorithm we set $\overline{u}(s) = \overline{u}_F(s)$.

technology parameterized by λ. The model understates noise in these percentiles in markets served by a small number of French firms.

Equation (36) in Section 3.7 motivates Figure 5c, which plots the median of normalized export intensity in each market against the actual number of French firms selling there. The model picks up the low magnitude of normalized export intensity and how it varies with the number of firms selling in a market. Despite our high estimate of σ_a, however, the model understates the noisiness of the relationship.

Aside from these Figures we report the number of firms selling to the 7 most popular export destinations in order of their popularity (that is, obeying a hierarchy), both in the actual and simulated data. That is, we report the number of firms selling to Belgium and no other top 7, to Belgium and Germany and no other top 7, etc. with the following results:

Country	Actual Number	Simulated Number
B	3988	4417
B-DE	863	912
B-DE-CH	579	402
B-DE-CH-I	330	275
B-DE-CH-I-UK	313	297
B-DE-CH-I-UK-NL	781	505
B-DE-CH-I-UK-NL-USA	2406	2840
Total	9260	9648

.

In the actual data 27.2 percent of exporters adhere to hierarchies compared with 30.3 percent in the simulated data, 13.5 percent implied by simply predicting on the basis of the marginal probabilities, and 100 percent had there been no entry shock ($\sigma_h = 0$).

4.7 Sources of Variation

In our model variation across firms in entry and sales reflects both differences in their underlying efficiency, which applies across all markets, and idiosyncratic entry and sales shocks in individual markets. We ask how much of the variation in entry and in sales can be explained by the universal rather than the idiosyncratic components.

4.7.1 Variation in Entry

We first calculate the fraction of the variance of entry in each market that can be explained by the cost draw u alone. By the law of large numbers, the fraction of French firms selling in n is a close approximation to the probability that a French firm will sell in n. Thus we write this probability as:

$$q_n = \frac{N_{nF}}{N_{FF}}.$$

The unconditional variance of entry for a randomly chosen French firm is therefore:

$$V_n^U = q_n \left(1 - q_n\right). \tag{41}$$

Conditional on its standardized unit cost u a firm enters market n if its entry shock η_n satisfies:

$$\eta_n \geq \left(u\kappa_2 / N_{nF}\right)^{1/\widetilde{\theta}}.$$

Since η_n is lognormally distributed with mean 0 and variance σ_h the probability that this condition is satisfied is:

$$q_n(u) = 1 - \Phi\left(\frac{\ln(u\kappa_2/N_{nF})}{\widetilde{\theta}\sigma_h}\right)$$

where Φ is the standard normal cumulative density. The variance conditional on u is therefore:

$$V_n^C(u) = q_n(u)[1 - q_n(u)].$$

A natural measure (similar to R^2 in a regression) of the explanatory power of the firm's cost draw for market entry is

$$R_n^E = 1 - \frac{E\left[V_n^C(u)\right]}{V_n^U}.$$

We simulated the term $E\left[V_n^C(u)\right]$ using the same techniques employed in our estimation routine, with 230,000 simulated firms, obtaining a value of R_n^E for each of our 112 export markets. The average value across markets is 0.57 (with a standard deviation across markets of only 0.01). Hence we can attribute 57 percent of the variation in entry in a market to the core efficiency of the firm rather than its draw of η in that market.[20]

4.7.2 Variation in Sales

Looking at the firms that enter a particular market, how much does the variation in u explain the variation in their sales there. Consider firm j selling in market n. Inserting (32) into (29), the log of its sales there is:

$$\ln X_{nF}(j) = \underbrace{\ln \alpha_n(j)}_{1} + \underbrace{\ln\left[1 - \left(\frac{u(j)\kappa_2}{N_{nF}\left[\eta_n(j)\right]^{\tilde{\theta}}}\right)^{\lambda/\tilde{\theta}}\right]}_{2} - \underbrace{\frac{1}{\tilde{\theta}}\ln u(j)}_{3} + \underbrace{\ln\left((N_{nF}/\kappa_2)^{1/\tilde{\theta}}\sigma E_n\right)}_{4}.$$

where we have divided sales into four components. Component 4 is common to all firms selling in market n so does not contribute to variation in sales there. The first component involves firm j's idiosyncratic sales shock in market n while component 3 involves its efficiency shock that applies across all markets. Complicating matters is component 2, which involves both firm j's idiosyncratic entry shock in market j, $\eta_n(j)$, and its overall efficiency shock, $u(j)$. We

[20]Not conditioning on u the probability q_n that a firm sells in any market n other than France is small. It is straightforward to show that taking the limit as $q_n \to 0$ the term R_n^E is independent of n. Hence the systematic variation in R_n^E across markets is small.

deal with this issue by first asking how much of the variation in $\ln X_n(j)$ is due to variation in component 3 and then in the variation in components 2 and 3 together.

We simulate sales of 230,000 firms across our 113 markets, and divide the contribution of each component to its sales in each market where it sells. We find that component 3 contributes a 0.048 share to the variation in $\ln X_{nF}(j)$, averaging across markets (with a standard error of 0.0003). Again averaging across markets, the share of components 2 and 3 together in the variation of $\ln X_{nF}(j)$ is 0.39 (with a standard deviation of 0.0025).[21]

Together these results indicate that the general efficiency of a firm is very important in explaining its entry into different markets, but makes a much smaller contribution to the variation in the sales of firms actually selling in a market. An explanation is that in order to enter a market a firm has already to have a low value of u. Hence the various sellers present in a market already have low values of u, so that differences among them in their sales are dominated by their market-specific sales shock $\alpha_n(j)$.

This finding does, of course, depend on our parameter estimates. A lower value of $\widetilde{\theta}$ (implying more sales heterogeneity attributable to efficiency) or lower values of σ_a or σ_h would lead us to attribute more to the firm's underlying efficiency rather that destination-specific shocks.

4.8 Productivity

Our methodology so far has allowed us to estimate $\widetilde{\theta}$, which incorporates both underlying heterogeneity in efficiency, as reflected in θ, and how this heterogeneity in efficiency gets

[21] The presence of N_{nF} in component 2 means that the contribution of each component varies across markets, but our simulation indicates that the differences are very small.

translated into sales, through σ. In order to break down $\widetilde{\theta}$ into these components we turn to data on firm productivity, as measured by value added per worker, and how it differs among firms selling to different numbers of markets.[22]

A common observation is that exporters are more productive (according to various measures) than the average firm.[23] The same is true of our exporters here: The average value added per worker of exporters is 1.22 times the average for all firms. Moreover, value added per worker, like sales in France, tends to rise with the number of markets served, but not with nearly as much regularity.

A reason for this relationship in our model is that a more efficient firm, with a lower normalized unit cost $u(j)$, will typically both enter more markets and sell more widely in any given market. As its fixed costs are not proportionately higher, larger sales get translated into higher value added relative to inputs used, including those used in fixed costs. An offsetting factor is that iceberg transport costs make serving foreign markets a less productive endeavor than supplying the home market. Determining the net effect requires a quantitative assessment.

How do we calculate productivity among our simulated firms? Because it provides a simple analytic expression we first look at the productivity of a firm's operations in selling to a particular market n. (Because it connects better with our model we calculate value added per unit of spending on factors rather than per worker.) We define its value added $V_n(j)$ in

[22] Because value added per worker is a crude productivity measure, we didn't incorporate these numbers into our method of simulated moments estimation above, as we didn't need them to estimate the other parameters of the model.

[23] See, for example, Bernard and Jensen (1999), Lach, Roberts, and Tybout (1997), and BEJK (2003).

market n as:

$$V_n(j) = X_n(j) - I_n(j)$$

where $I_n(j)$ is firm j's spending on intermediates to supply that market. We calculate this intermediate spending as:

$$I_n(j) = (1-\beta)\overline{m}^{-1}X_n(j) + (1-\beta^F)E_n(j),$$

where β is the share of factor costs in variable costs and β^F is the share of factor costs in fixed costs.

Value added per unit of factor cost $q_n(j)$ is then:

$$\begin{aligned} q_n(j) &= \frac{V_n(j)}{\beta\overline{m}^{-1}X_n(j) + \beta^F E_n(j)} \\ &= \frac{\left[1-(1-\beta)\overline{m}^{-1}\right]X_n(j) - (1-\beta^F)E_n(j)}{\beta\overline{m}^{-1}X_n(j) + \beta^F E_n(j)} \\ &= \frac{\left[\overline{m}-(1-\beta)\right] - \overline{m}(1-\beta^F)\left[E_n(j)/X_n(j)\right]}{\beta + \overline{m}\beta^F\left[E_n(j)/X_n(j)\right]}. \end{aligned} \tag{42}$$

The only source of cross-firm heterogeneity in productivity arises through the ratio $E_n(j)/X_n(j)$: Firms having more sales $X_n(j)$ relative to entry costs $E_n(j)$ are more productive.

Using (5) and (10) for the numerator and (29) for the denominator, exploiting (33), we can write this ratio in terms of $v_n(j)$ as:

$$\frac{E_n(j)}{X_n(j)} = \frac{\lambda}{\sigma(\lambda-1)}\frac{v_n(j)^{(1-\lambda)/\widetilde{\theta}}-1}{v_n(j)^{-\lambda/\widetilde{\theta}}-1}.$$

Since $v_n(j)$ is distributed uniformly on $[0,1]$,in any market n the distribution of the ratio of fixed costs to sales revenue, and hence the distribution of productivity, is invariant to any market-specific feature such as size or location. In particular, the distribution of productivity is not affected by trade openness.

What we have said so far applies to the productivity of units selling in a market, which are not the same thing as the firms producing there. To measure the overall productivity of a firm we need to sum its sales, value added, and factor costs across its activities in different markets. Defining total sales $X(j) = \sum_n X_n(j)$ and total entry costs $E(j) = \sum_n E_n(j)$, firm j's productivity is:

$$q(j) = \frac{[\overline{m} - (1-\beta)] - \overline{m}(1-\beta^F)\,[E(j)/X(j)]}{\beta + \overline{m}\beta^F\,[E(j)/X(j)]}. \tag{43}$$

To simplify we assume $\beta^F = 0$ so that all fixed costs are purchased services. We then calibrate β from the average share of manufacturing value added in gross production across UNIDO, 0.36.[24]

Note that the expression for firm productivity (43) depends on the elasticity of substitution σ (through $\overline{m} = \sigma/(\sigma-1)$) but not on θ. Following BEJK (2003) we find an $\overline{m}$ that makes the productivity advantage of exporters in our simulated data match their productivity advantage in the actual data (1.22). This exercise delivers $\overline{m} = 1.51$ or $\sigma = 2.98$, implying $\beta = 0.34$. We use these values in the counterfactual exercises described in the next section.

Figure 6 reports average value added per worker against the logarithm of the minimum number of markets where the firms sell, with nonexporters included in the first datapoint. Circles represent the actual data and asterisks our simulation, based on (43) with $\overline{m} = 1.51$. Note that, among the actual firms, value added per worker more than doubles as we move from all of our firms (selling in at least one market) to those selling to at least 74 markets,

[24] Because of profits and fixed costs this measured value added share, denoted V, differs from β^V. The two are connected by the relationship:

$$\beta^V = \overline{m}V - 1/\theta,$$

so that β^V is determined from V simultaneously with our estimates of $\overline{m}$ and θ.

and then plummets as we look at firms that sell more widely (although we are looking at fewer than one hundred firms in this upper tail). The model picks up the rise in measured productivity corresponding to wider entry at the low end (representing the vast majority of firms) but fails to reproduce the spike at the high end.

5 General Equilibrium and Counterfactuals

We now consider how changes in policy and the environment would affect individual firms. To do so we need to consider how such changes would affect wages and prices. So far we have conditioned on a given equilibrium outcome. We now have to ask how the world reequilibrates.

5.1 Embedding the Model in a General Equilibrium Framework

Embedding our analysis in general equilibrium requires additional assumptions:

1. We introduce factors as in Ricardo (1821). Each country is endowed with an amount L_i of labor (or a composite factor), which is freely mobile across activities within a country but does not migrate. Its wage in country i is W_i.

2. We introduce intermediates as in Eaton and Kortum (2002). Manufacturing inputs are a Cobb-Douglas combination of intermediates, which are a representative bundle of manufactures with price index P_i given in (14), and labor, with labor having a share β. Hence we can write the cost of an input bundle:

$$w_i = \kappa_3 W_i^{\beta} P_i^{1-\beta},$$

where $\kappa_3 = \beta^{-\beta}(1-\beta)^{-(1-\beta)}$.

3. We introduce nonmanufacturing as in Alvarez and Lucas (2007). Final output, which is nontraded, is a Cobb-Douglas combination of manufactures and labor, with manufactures having a share γ. Labor is the only input into nonmanufactures. Hence the price of final output in country i is proportional to $P_i^{\gamma} W_i^{1-\gamma}$.

4. Fixed costs go to labor in the local market. Hence the demand for labor in manufacturing overhead is $\overline{E}_n / W_n$, with $\overline{E}_n$ as derived in (23).

Equilibrium in the world market for manufactures requires that the sum across countries of absorption of manufactures from country i equal its gross output Y_i, or:

$$Y_i = \sum_{n=1}^{N} \pi_{ni} X_n \tag{44}$$

for each country i. To determine equilibrium wages around the world requires that we turn these expressions into conditions for equilibrium in world labor markets.

Country i's total absorption of manufactures is the sum of final demand and use as intermediates:

$$X_i = \gamma (Y_i^A + D_i^A) + [(1-\beta)(\sigma - 1)/\sigma] \, Y_i \tag{45}$$

where Y_i^A is GDP and D_i^A the trade deficit.

To relate Y_i^A to labor income we write:

$$Y_i^A = Y_i^L + \Pi_i, \tag{46}$$

where $Y_i^L = W_i L_i$ is labor income and Π_i are total net profits earned by country i's manufacturing producers from their sales at home and abroad. Since our model is static we treat deficits as exogenous.

Net profits earned in destination n both by domestic firms and by exporters selling there, which we denote Π_n^D are gross profits X_n/σ less fixed costs incurred there, $\overline{E}_n$. Using (23) for $\overline{E}_n$:

$$\Pi_n^D = \frac{(\sigma-1)}{\sigma\theta} X_n.$$

Producers from country i earn a share π_{ni} of these profits. Hence:

$$\Pi_i = \sum_{n=1}^{N} \pi_{ni} \Pi_n^D = \frac{(\sigma-1)}{\sigma\theta} Y_i, \qquad (47)$$

where the second equality comes from applying the conditions (44) for equilibrium in the market for manufactures.

Substituting (46) into (45) and using the fact that gross manufacturing production Y_i is gross manufacturing absorption X_i less the manufacturing trade deficit D_i:[25]

$$Y_i + D_i = \gamma \left[Y_i^L + \frac{(\sigma-1)}{\sigma\theta} Y_i + D_i^A \right] + \frac{(1-\beta)(\sigma-1)}{\sigma} Y_i.$$

Solving for Y_i:

$$Y_i = \frac{\gamma\sigma\left(Y_i^L + D_i^A\right) - \sigma D_i}{1 + (\sigma-1)(\beta - \gamma/\theta)}. \qquad (48)$$

Our conditions for equilibrium in the world market for manufactures (given deficits) thus translate into the following conditions for labor market equilibrium:

$$\frac{\gamma\sigma\left(Y_i^L + D_i^A\right) - \sigma D_i}{1 + (\sigma-1)(\beta - \gamma/\theta)} = \sum_{n=1}^{N} \pi_{ni} \frac{\gamma\sigma\left(Y_n^L + D_n^A\right) - (\sigma-1)(1 - \beta + \gamma/\theta) D_n}{1 + (\sigma-1)(\beta - \gamma/\theta)}. \qquad (49)$$

[25] For simplicity we reconcile the differences between manufacturing and overall trade deficits by thinking of them as transfers of the final good, which is otherwise not traded. For large economies, the manufacturing deficit is the largest component of the overall trade deficit. See Dekle, Eaton, and Kortum (2008) for a fuller treatment of deficits in a similar model of bilateral trade.

From expression (4), we can write:

$$\pi_{ni} = \frac{T_i\left(W_i^{\beta}P_i^{1-\beta}d_{ni}\right)^{-\theta}}{\sum_{k=1}^{N} T_k\left(W_k^{\beta}P_k^{1-\beta}d_{nk}\right)^{-\theta}}. \tag{50}$$

From expression (14):

$$P_n = \overline{m}\kappa_1^{-1/\theta}\kappa_3\left[\sum_{i=1}^{N} T_i(W_i^{\beta}P_i^{1-\beta}d_{ni})^{-\theta}\right]^{-1/\theta}\left(\frac{X_n}{\sigma E_n}\right)^{(1/\theta)-1/(\sigma-1)}. \tag{51}$$

where:

$$X_n = \frac{\gamma\sigma\left(Y_n^L + D_n^A\right) - (\sigma-1)\left(1-\beta+\gamma/\theta\right)D_n}{1+(\sigma-1)\left(\beta-\gamma/\theta\right)} \tag{52}$$

and:

$$E_n = W_n F_n, \tag{53}$$

where F_n is a parameter reflecting the efficiency of labor in country n in providing overhead services.

Incorporating (50), (52), and (53) equations (49) and (51) determine wages W_i and manufacturing price indices P_i around the world as functions of each country's T_i, L_i, F_i, D_i, and D_i^A, each country pair's d_{ni},and the parameters $\gamma, \beta, \sigma, \theta$, and κ_1.

5.2 Perturbing the Equilibrium

We apply the method explained in Dekle, Eaton, and Kortum (2008) to calculate counterfactuals. Denote the counterfactual value of any variable x as x' and define $\widehat{x} = x'/x$. Equilibrium in world manufactures in the counterfactual requires:

$$Y_i' = \sum_{n=1}^{N} \pi_{ni}' X_n'. \tag{54}$$

We can write each of the components in terms of each country's baseline labor income, Y_i^L, baseline trade shares π_{ni}, and the parameters β, γ, θ, and σ and the change in wages $\widehat{W}_i$ and $\widehat{P}_i$ using (48), (52), (51), and (50) as follows:

$$\begin{aligned} Y_i' &= \frac{\gamma\sigma\left(Y_i^L \widehat{L}_i \widehat{W}_i + D_i^{A\prime}\right) - \sigma D_i'}{1 + (\sigma - 1)(\beta - \gamma/\theta)} \\ X_n' &= \frac{\gamma\sigma\left(Y_n^L \widehat{L}_n \widehat{W}_n + D_n^{A\prime}\right) - (\sigma - 1)(1 - \beta + \gamma/\theta) D_n'}{1 + (\sigma - 1)(\beta - \gamma/\theta)} \\ \pi_{ni}' &= \frac{\pi_{ni} \widehat{W}_i^{-\beta\theta} \widehat{P}_i^{-(1-\beta)\theta} \widehat{T}_i \widehat{d}_{ni}^{-\theta}}{\sum_{k=1}^N \pi_{nk} \widehat{W}_k^{-\beta\theta} \widehat{P}_k^{-(1-\beta)\theta} \widehat{T}_k \widehat{d}_{nk}^{-\theta}} \end{aligned}$$

where sticking these three equations into (54) yields a set of equations involving $\widehat{W}_i$ for given $\widehat{P}_i$'s. From (51) we can get an involving $\widehat{P}_i$ for given $\widehat{W}_i$'s:

$$\widehat{P}_n = \left[\sum_{i=1}^N \pi_{ni} \widehat{W}_i^{-\beta\theta} \widehat{P}_i^{-(1-\beta)\theta} \widehat{T}_i \widehat{d}_{ni}^{-\theta}\right]^{-1/\theta} \left(\frac{\widehat{X}_n}{\widehat{W}_n \widehat{F}_n}\right)^{(1/\theta)-1/(\sigma-1)} \tag{55}$$

where, from above:

$$\widehat{X}_n = \frac{\gamma\sigma\left(Y_n^L \widehat{L}_n \widehat{W}_n + D_n^{A\prime}\right) - (\sigma - 1)(1 - \beta + \gamma/\theta) D_n'}{\gamma\sigma\left(Y_n^L + D_n^A\right) - (\sigma - 1)(1 - \beta + \gamma/\theta) D_n}.$$

We can use the system of equations ((54)) and (55) to solve for the changes in wages and prices that would result from exogenous changes in trade barriers $\widehat{d}_{ni}$, entry costs $\widehat{F}_i$, labor forces $\widehat{L}_i$, technology parameters $\widehat{T}_i$, or deficits D_i' or $D_i^{A\prime}$.

5.3 A Counterfactual Calculation

We implement counterfactual simulations for our 113 countries in 1986, aggregating the rest of the world into a 114th country (ROW). We calibrate the π_{ni} with data on trade shares. We exploit (46), (47), and (48) to express labor income in terms of data on GDP and deficits:

$$Y_i^L = \frac{[1 + (\sigma - 1)(\beta - \gamma/\theta)]\, Y_i^A - [(\sigma - 1)\gamma/\theta]\, D_i^A - [(\sigma - 1)/\theta] D_i]}{1 + (\sigma - 1)\beta}.$$

We calibrate β (common across countries) and γ (country specific) from data on manufacturing production and trade. Appendix C describes our sources of data and our procedures for assembling them to execute the counterfactual.

We consider a ten percent drop in trade barriers, i.e., $\widehat{d}_{ni} = 1/1.1$ for $n \neq i$, with $\widehat{d}_{ii} = 1$. This change roughly replicates the increase in French import and export shares over the decade following 1986. Lower trade barriers raise the real wage in every country, ranging from a nearly imperceptible rise in remote nonmanufacturing countries to increases of 8 percent in Belgium and nearly 25 percent in Singapore.

We then calculate the implications of this change for individual firms, holding fixed all of the firm-specific shocks that underlie firm heterogeneity. The idea is to hold all else fixed while considering the microeconomic changes brought about by general-equilibrium forces.

Using our results for $\widehat{W}_n, \widehat{\pi}_{ni}$, and $\widehat{X}_n$ from our general equilibrium analysis, we produce a dataset, recording both baseline and counterfactual firm-level behavior, as follows:

1. We calculate the implied percentage change in French exports in each market n as $\widehat{X}_{nF} = \widehat{\pi}_{ni}\widehat{X}_n$ and, using (30), the number of French firms selling there as $\widehat{N}_{nF} = \widehat{\pi}_{ni}\widehat{X}_n/W_n$ (since we assume that fixed costs involve local labor).

2. We apply these percentage changes to our original data set to get counterfactual values of total French sales in each market X^C_{nF} and the number of French sellers there N^C_{nF}.[26]

[26] Recall that, for reasons such as manufacturing exports by nonmanufacturing firms, the aggregate exports described in Appendix C exceed total exports by our French firms. Having used the aggregate data to calculate the counterfactual equilibrium as described in the previous section, we applied the percentage changes from that exercise to the X_{nF} and N_{nF} in our firm dataset.

3. We run the simulation described in Section 4 Part B with $S = 500,000$.

   (a) Stage 1 is carried out just once, so that the same stochastic draws apply both to the baseline and the counterfactual.

   (b) In Stage 2 we set Θ to the parameter estimates reported in Section 4, Part E.

   (c) We use our baseline values X_{nF} and N_{nF} to calculate baseline σE_n's for each destination and baseline $\overline{u}_n(s)$'s for each destination and firm, using (31) and (39).

   (d) We use our counterfactual values X_{nF}^C and N_{nF}^C to calculate counterfactual σE_n^C's for each destination and counterfactual $\overline{u}_n^C(s)$'s for each destination and firm, again using (31) and (39).

   (e) To avoid wasting time simulating a firm that doesn't sell anywhere in either the baseline or the counterfactual we define:

   $$\overline{\overline{u}}(s) = \max_n \{\overline{u}_n(s), \overline{u}_n^C(s)\}$$

   and set:

   $$u(s) = v(s)\overline{\overline{u}}(s).$$

   where $v(s)$ was drawn from $U[0,1]$ in Stage 1. We consequently assign firm s an importance weight $\overline{\overline{u}}(s)$. A firm for which $u(s) \leq \overline{u}_n(s)$ sells in market n in the baseline while a firm for which $u(s) \leq \overline{u}_n^C(s)$ sells there in the counterfactual. Hence our simulation allows for entry, exit, and survival.

   (f) We calculate entry and sales in each of the 113 markets in the baseline and counterfactual, deflating counterfactual values to baseline prices.

5.4 Firm Level Implications of Globalization

Table 3 summarizes the results, which are dramatic. Total sales by French firms rise by \$14,528 million, the net effect of a \$32,738 million increase in exports and a \$14,528 million decline in domestic sales. Despite this rise in total sales, competition from imports drives 26,212 firms out of business, although 10,702 firms start exporting.

Tables 4 and 5 decompose these changes into the contributions of firms of different baseline size, with Table 4 considering the counts of firms. Most of the losses are at the bottom end: Nearly half the firms in the bottom decile are wiped out while only the top percentile avoids any attrition. Because so many firms in the top decile already export, the greatest number of new exporters emerge from the second highest decile. The biggest percentage increase in number of exporters is for firms in the third from the bottom decile.

Table 5 decomposes sales revenues. All of the increase is in the top decile, and most of that in the top percentile. For every other decile sales decline. Almost two-thirds of the increase in export revenue is from the top percentile, although lower deciles experience much higher percentage increases in their export revenues.

Comparing the numbers in Tables 4 and 5 reveals that, even among survivors, revenue per firm falls in every decile except the top. In summary, the decline in trade barriers improves the performance of the very top firms at the expense of the rest.[27]

Table 6 looks at an alternative decomposition of firms according to the number of markets where they initially sold. Most of the increase in export revenues is among the firms that were already exporting most widely. But the percentage increase falls with the initial number of

[27] The first row of the tables pertains to firms that entered only to export. There are only 1108 of them selling a total of \$4 million.

markets served. For firms that initially export to few markets, a substantial share of export growth comes from entering new markets.

Finally, Table 7 looks at what happens in each foreign market. The percentage change in exports and in number of exporters is what is delivered by our counterfactual equilibrium. Applying these changes to our baseline gives the absolute change in exports and in number of exporters. How much of the increase in exports to each market is due to entry (the extensive margin)? This fraction tends to be small (although, because of sampling error in small markets, it can vary a lot). In the last two columns we look at growth in sales by incumbent firms. As Arkolakis (2008) would predict, sales by firms with an initially smaller presence grow substantially more than those at the top.

6 Conclusion

We examine some key features of the sales of French firms across 113 different markets, including France itself. Much of what we see can be interpreted in terms of a standard model of heterogeneous producers. We think that the model provides a useful tool for linking what goes on at the aggregate level with the situation of individual firms.

References

Alvarez, Fernando and Robert E. Lucas (2007) "A General Equilibrium Analysis of the Eaton-Kortum Model of International Trade," forthcoming, *Journal of Monetary Economics.*

Arkolakis, Konstantinos (2008) "Market Penetration Costs and the New Consumers Margin in International Trade," mimeo, Yale University.

Aw, Bee Yan, Chung, Sukkyun, and Roberts, Mark J. (1998) "Productivity and the Decision to Export: Micro Evidence from Taiwan and South Korea." *NBER Working Paper No. 6558.*

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

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

Bernard, Andrew B., Redding, Stephen, and Schott, Peter (2006) "Multi-Product Firms and Trade Liberalization." *NBER Working Paper* #12782.

Biscourp, Pierre and Kramarz, Francis (2007) "Employment, Skill Structure, and International Trade: Firm-Level Evidence for France," *Journal of International Economics,* 72: 22-51.

Chaney, Thomas (2008) "Distorted Gravity: Heterogeneous Firms, Market Structure, and the Geography of International Trade," forthcoming *American Economic Review.*

Clerides, Sofronis, Lach, Saul, and Tybout, James R. (1998) "Is Learning by Exporting Important? Micro-Dynamic Evidence from Colombia, Mexico, and Morocco." *Quarterly Journal of Economics* 113: 903-947.

Dekle, Robert, Jonathan Eaton, and Samuel Kortum (2008) "Global Rebalancing with Gravity," *IMF Staff Papers,* forthcoming.

Eaton, Jonathan and Kortum, Samuel (2002) "Technology, Geography, and Trade." *Econometrica 70*: 1741-1780.

Eaton, Jonathan, Kortum, Samuel, and Francis Kramarz (2004) "Dissecting Trade: Firms, Industries, and Export Destinations." *American Economic Review, Papers and Proceedings*, 94: 150-154. (Data sources and tables available in, *NBER Working Paper #10344.*)

Feenstra, Robert C., Lipsey, Robert E. and Bowen, Henry P. (1997) "World Trade Flows, 1970-1992, with Production and Tariff Data." *NBER Working Paper No. 5910.*

Gabaix, Xavier (1991) "Zipf's Law for Cities: An Explanation," *Quarterly Journal of Economics,* 94: 739-767.

Gabaix, Xavier and Ibragimov, Rustam (2008) "Rank-1/2: A Simple Way to Improve the OLS Estimation of Tail Exponents," New York University.

Goffe William L., Gary D. Ferrier, and John Rogers (1994) "Global Optimization of Statistical Functions with Simulated Annealing," *Journal of Econometrics,* 60: 65-99.

Gouriéroux, Christian and Alain Monfort (1994) *Simulation Based Econometric Methods*, CORE Lecture Series, Université catholique de Louvain.

Helpman, Elhanan, Melitz, Marc J., and Yeaple, Stephen R. (2004) "Exports vs. FDI with Heterogeneous Firms," *American Economic Review,* 94: 300-316.

Krugman, Paul R. (1980) "Scale Economies, Product Differentiation, and the Patterns of Trade," *American Economic Review,* 70: 950-959.

Horowitz, Joel (2001) "The Bootstrap," Chapter 52, *Handbook of Econometrics, Volume 5,* J.J. Heckman and E.E. Leamer, editors. Amsterdam: Elsevier Science B.V. pp. 3159-3228.

Liu, Jun S. (2001) *Monte Carlo Strategies in Scientific Computing*, New York: Springer.

Luttmer, Erzo (2006) "Selection, Growth, and the Size Distribution of Firms," forthcoming, *Quarterly Journal of Economics.*

Lucas, Robert E. Jr. (1987) *Models of Business Cycles.* Oxford: Basil Blackwell.

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

OECD (1995) *The OECD STAN Database.*

Ricardo, David (1821) *On The Principles of Political Economy and Taxation.* London: John Murray, Albemarle-Street.

Roberts, Mark and Tybout, Robert (1997) "The Decision to Export in Colombia: An Empirical Model of Entry with Sunk Costs." *American Economic Review* 87: 545-564.

Simon, Herbert A., and Charles P. Bonini (1958) "The Size Distribution of Business Firms," *American Economic Review,* 98: 607-617.

Simon, Herbert A. (1955) "On a Class of Skew Distribution Functions," *Biometrica,* 42: 425-440.

Sutton, John (1991) *Sunk Costs and Market Structure.* Cambridge, MA: The MIT Press.

UNIDO (1999) *Industrial Statistics Database.*

United Nations Statistics Division (2007), National Accounts Available on the Internet: http://unstats.un.org/unsd/snaama.

A Appendix A: Constructing the Firm Level Data

Up to 1992 all shipments of goods entering or leaving France were declared to French customs either by their owners or by authorized customs commissioners. These declarations constitute the basis of all French trade statistics. Each shipment generates a record. Each record contains the firm identifier, the SIREN, the country of origin (for imports) or destination (for exports), a product identifier (a 6-digit classification), and a date. All records are aggregated first at the monthly level. In the analysis files accessible to researchers, these records are further aggregated by year and by 3-digit product (NAP 100 classification, the equivalent of the 3-digit SIC code). Therefore, each observation is identified by a SIREN, a NAP code, a country code, an import or export code, and a year. In our analysis, we restrict attention to exporting firms in the manufacturing sector in year 1986 and in year 1992. Hence, we aggregate across manufacturing products exported. We can thus measure each firm's amount of total exports in years 1986 and 1992 by country of destination. Transactions are recorded in French Francs and reflect the amount received by the firm (i.e., including discounts, rebates, etc.). Even though our file is exhaustive, i.e., all exported goods are present, direct aggregation of all movements may differ from published trade statistics, the second being based on list prices and thus exclude rebates.

We match this file with the Base d'Analyse Longitudinale, Système Unifié de Statistiques d'Entreprises (BAL-SUSE) database, which provides firm-level information. The BAL-SUSE database is constructed from the mandatory reports of French firms to the fiscal administration. These reports are then transmitted to INSEE where the data are validated. It includes all firms subject to the "Bénéfices Industriels et Commerciaux" regime, a fiscal regime mandatory for all manufacturing firms with a turnover above 3,000,000FF in 1990 (1,000,000FF in the service sector). In 1990, these firms comprised more than 60% of the total number of firms in France while their turnover comprised more than 94% of total turnover of firms in France. Hence, the BAL-SUSE is representative of French enterprises in all sectors except the public sector.

From this source, we gather balance sheet information (total sales, total labor costs, total wage-bill, sales, value-added, total employment). Matching the Customs database and the BAL-SUSE database leaves us 229,900 firms in manufacturing (excluding construction, mining and oil industries) in 1986 with valid information on sales and exports. In 1992, the equivalent number is 217,346. In 1986, 34,035 firms export to at least one country Among them 17,699 export to Belgium, the most popular destination. To match our data with aggregate trade and production data, we restrict attention to 113 countries (including France).

A Appendix B: Estimation Algorithm

The algorithm we use to fit theoretical moments to their empirical counterparts is simulated annealing. We rely on a version specifically developed for Gauss and available on the web from William Goffe (Simann). Goffe, Ferrier, and Rogers (1994) describe the algorithm. Simulated annealing, in contrast with other optimization algorithms (Newton-Raphson, for instance), explores the entire surface and moves both uphill and downhill to optimize the function. It is therefore largely independent of starting values. Because it goes both downhill and uphill, it escapes local maxima. Finally, the function to optimize does not need to have stringent properties; differentiability and continuity, for instance, are not needed. The version developed by these authors, and implemented in Simann, possesses some features that make it more efficient (in particular, less time-consuming) than previous implementations of simulated annealing. The program includes a precise explanation of the various parameters that must be set in advance. It also suggests reasonable starting values. The program, as well as the starting values, is available from the authors. Optimizing our admittedly complex function, with 5 parameters, on a standard personal computer takes around one week.

We use bootstrap to compute standard errors for our parameters. We follow the bootstrap procedure suggested by Horowitz (2001) closely. More precisely, we use the bootstrap with recentering as suggested when using a method of moments estimation strategy (Horowitz, 2001, subsection 3.7, pages 3186-3187). Because each bootstrap repetition requires one week for estimation, we used only 10 repetitions. The small number is unlikely to have any effect given the concentration of the bootstrap estimates around the estimated values.

A Appendix C: The Data for Counterfactuals

Our quantitative methodology for performing general-equilibrium counterfactuals largely follows Dekle, Eaton, and Kortum (2008). While our analysis in that paper was based on 44 countries, here we consider 113 plus a rest-of-world aggregate, making the total 114. All data are for 1986, translated into millions of U.S. dollars at the 1986 exchange rate.

For each country n, data on GDP Y_n^A and the trade deficit in goods and services D_n^A are from the United Nations Statistics Division (2007).[28] We took total absorption of manufactures X_n from our earlier work, EKK (2004). Bilateral trade in manufactures is from Feenstra, Lipsey and Bowen (1997). Starting with the file WBEA86.ASC, we aggregate across all manufacturing industries. Given these trade flows $\pi_{ni}X_n$ we calculate the share of exporter i in n's purchases π_{ni} and manufacturing trade deficits D_n. The home shares π_{ii} are residuals.

The shares of manufactures in final output γ_n are calibrated to achieve consistency between our observations for the aggregate economy and the manufacturing sector. In particular:

$$\gamma_n = \frac{X_n - (1-\beta)(1-1/\sigma)(X_n - D_n)}{Y_n^A + D_n^A}.$$

[28] A couple of observations were missing from the data available on line. To separate GDP between East and West Germany, we went to the 1992 hardcopy. For the USSR and Czechoslovakia, we set the trade deficit in goods and services equal to the trade deficit in manufactures.

Table 1: Exports to the Seven Most Popular Destinations

Country	Number of French Exporters	Marginal Probability
Belgium-Luxembourg	17,699	0.59
Germany	14,579	0.49
Switzerland	14,173	0.47
Italy	10,643	0.36
United Kingdom	9,752	0.33
Netherlands	8,294	0.28
United States	7,608	0.25
Total firms exporting to at least one of the top 7	29,929	1.00
Total firms exporting	34,035	

Source: Customs data, year 1986. Marginal probability is the ratio of the number exporting to the destination relative to the number exporting to at least one of these seven countries

Table 2: Do Firms Obey Country Hierarchies ?

Exporting to:	Observed Number of Firms	Under Independence	
		Predicted Probability	Number of Firms
1000000	3988	0.037	1119
1100000	863	0.036	1063
1110000	579	0.032	956
1111000	330	0.018	528
1111100	313	0.009	255
1111110	781	0.003	98
1111111	2406	0.001	33
Total	9260	0.135	4052

Source: Customs data, year 1986. All firms export at least to one of these 7 countries: Belgium, Germany, Switzerland, Italy, United Kingdom, Netherlands, United States, in this order. The string of 0s and 1s in column labelled "Exporting to" refer to this order: 1000000 means that these firms only export to Belgium; 1111111 means that these firms export to all seven countries. We ignore in this table export behavior outside these 7 countries

Table 3 - Counterfactuals: Aggregate Outcomes

		Counterfactual	
	Baseline	Change from Baseline	Percentage Change
Number:			
All Firms	230,260	-26,212	-11.4
Selling in France	229,357	-27,840	-12.1
Exporting	32,796	10,702	32.6
Values ($ millions):			
Total Sales	433,768	14,528	3.3
Domestic Sales	364,892	-18,210	-5.0
Exports	68,877	32,738	47.5

Counterfactual simulation of a 10% decline in trade costs.

Table 4 - Counterfactuals: Firm Entry and Exit by Initial Size

Initial Size Interval (percentile)	All Firms Baseline # of Firms	All Firms Counterfactual Change from Baseline	All Firms Counterfactual Change in %	Exporters Baseline # of Firms	Exporters Counterfactual Change from Baseline	Exporters Counterfactual Change in %
not active	0	1,108	---	0	1,108	---
0 to 10	22,974	-11,323	-49.3	746	17	2.3
10 to 20	23,036	-5,743	-24.9	137	93	67.8
20 to 30	23,038	-3,626	-15.7	182	192	105.6
30 to 40	23,030	-2,413	-10.5	412	393	95.5
40 to 50	23,034	-1,749	-7.6	749	591	78.8
50 to 60	23,024	-1,098	-4.8	1,433	891	62.2
60 to 70	23,031	-711	-3.1	2,530	1,313	51.9
70 to 80	23,027	-438	-1.9	4,242	1,781	42.0
80 to 90	23,028	-178	-0.8	7,504	2,322	30.9
90 to 99	20,729	-41	-0.2	12,700	1,936	15.2
99 to 100	2,310	0	0.0	2,161	64	3.0
Totals	230,260	-26,212		32,796	10,702	

Table 5 - Counterfactuals: Firm Growth by Initial Size

Initial Size Interval (percentile)	Total Sales: Baseline in $millions	Total Sales: Counterfactual Change from Baseline	Total Sales: Counterfactual Change in %	Exports: Baseline in $millions	Exports: Counterfactual Change from Baseline	Exports: Counterfactual Change in %
not active	0	4	---	0	4	---
0 to 10	41	-24	-58.0	1	2	367.7
10 to 20	193	-92	-47.8	1	2	251.8
20 to 30	475	-184	-38.8	1	3	306.8
30 to 40	967	-306	-31.6	2	9	377.9
40 to 50	1,822	-487	-26.7	5	20	368.2
50 to 60	3,347	-711	-21.2	18	51	276.2
60 to 70	6,267	-1,050	-16.8	60	139	230.5
70 to 80	12,661	-1,544	-12.2	202	382	189.3
80 to 90	31,322	-2,028	-6.5	1,023	1,409	137.8
90 to 99	148,502	3,685	2.5	15,663	11,587	74.0
99 to 100	228,170	17,265	7.6	51,901	19,131	36.9
Totals	433,768	14,528		68,877	32,738	

Table 6 - Counterfactuals: Firms Growth by Initial Export Penetration

	Baseline		Counterfactual		
Initial # of Export Destinations	Number of Firms	Exports in $millions	Change in Exports	Percentage Change in Exports	% Contribution of New Destinations
0	197,464	0	114	---	100.0
1	13,739	133	317	238.8	41.8
2	5,010	204	400	196.1	34.1
3	2,863	259	430	165.9	28.8
4	1,834	291	472	162.6	27.1
5 to 10	4,343	1,996	2,461	123.3	19.6
10 to 25	3,262	8,237	6,744	81.9	11.0
Over 25	1,745	57,757	21,799	37.7	2.6
Totals	230,260	68,877	32,738		

Table 7 - Firm Export Behavior by Country and Firm Size in Country

Country	Baseline Exports in $millions	Counterfactual: Change in Exports	Counterfactual: Percentage Change in Exports	Counterfactual: Percentage Change in # Exporters	Counterfactual: Contribution of New Exporters, %	Counterfactual: Growth of Incumbant Exporters, by Initial Size in Market: Below Median	Counterfactual: Growth of Incumbant Exporters, by Initial Size in Market: Above 95th Percentile
AFGHAN	9	2	22.2	18.2	16.8	74.6	13.8
ALBANIA	3	1	37.0	31.8	6.8	97.8	31.3
ALGERIA	1,411	95	6.7	11.3	4.1	42.4	4.1
ANGOLA	63	5	8.2	10.3	2.8	52.4	6.4
ARGENTIN	230	132	57.4	53.6	12.4	212.7	41.7
AUSTRALI	342	128	37.4	30.7	7.0	109.0	28.9
AUSTRIA	688	357	51.9	33.3	7.0	163.6	41.1
BANGLADE	19	8	40.3	35.3	7.8	116.9	29.6
BELGIUML	6,039	2,661	44.1	19.3	2.5	129.0	38.4
BENIN	57	7	11.9	9.1	2.1	37.9	9.8
BOLIVIA	6	1	21.2	19.0	6.9	61.7	17.8
BRAZIL	505	317	62.8	54.7	8.8	297.4	44.8
BULGARIA	117	37	31.7	34.9	4.9	130.2	24.0
BURKINAF	39	6	16.3	17.8	4.0	55.0	11.8
BURUNDI	15	3	17.1	21.4	4.1	59.5	12.7
CAMEROON	400	68	17.0	20.1	4.0	71.1	12.9
CANADA	685	297	43.3	25.3	5.1	131.6	34.4
CENAFREP	23	8	33.9	19.6	3.2	93.4	26.8
CHAD	14	1	7.0	11.8	3.1	30.8	5.0

CHILE	58	31	53.8	39.9	7.0	153.2	39.1
CHINA	406	132	32.4	38.6	6.1	167.2	23.4
COLOMBIA	89	20	22.6	24.9	6.6	99.9	16.7
COSTAR	13	3	21.6	14.7	1.4	63.3	17.7
COTEDI	253	96	38.0	30.1	5.9	110.9	29.1
CUBA	30	7	22.9	24.4	7.5	81.7	16.5
CZECH	94	50	53.4	42.3	8.9	132.1	38.6
DENMARK	641	308	48.0	28.3	4.6	137.8	38.5
DOMREP	16	5	34.0	29.5	10.0	88.3	25.6
ECUADOR	13	7	54.4	30.0	37.7	93.0	27.5
EGYPT	422	53	12.6	12.4	4.6	52.7	8.5
ELSALV	9	1	9.2	15.3	2.2	28.5	7.2
ETHIOPIA	10	1	7.7	8.6	2.7	27.4	5.7
FINLAND	400	232	58.0	39.4	8.1	179.2	43.8
GERMANYE	137	83	60.7	49.6	14.8	201.1	38.9
GERMANYW	12,330	7,370	59.8	45.6	8.2	302.2	45.6
GHANA	22	8	36.1	30.0	4.9	120.7	28.8
GREECE	440	163	37.1	29.7	8.8	120.9	27.7
GUATEM	17	3	20.4	15.3	2.2	52.0	19.5
HONDURAS	17	3	18.5	16.3	0.9	57.9	17.2
HONGKONG	291	96	33.1	2.3	0.0	43.1	32.4
HUNGARY	130	52	39.9	23.5	5.4	111.1	30.1
INDIA	681	279	41.0	36.6	7.5	221.4	30.2
INDONES	157	65	41.5	40.3	9.4	203.0	28.3
IRAN	84	14	16.5	17.6	3.5	77.7	13.6
IRAQ	224	20	8.8	5.9	2.4	29.4	6.4
IRELAND	301	128	42.4	21.6	3.0	111.9	35.4
ISRAEL	201	94	47.0	33.6	7.7	142.9	35.9
ITALY	8,204	4,748	57.9	46.6	10.1	296.9	43.1

JAMAICA	6	2	37.4	20.9	7.7	88.8	29.4
JAPAN	947	761	80.3	67.8	13.3	307.8	56.4
JORDAN	52	9	16.6	14.0	1.6	42.0	13.6
KENYA	166	30	18.3	19.4	4.6	89.2	15.9
KOREAS	644	335	52.1	40.3	6.9	245.7	42.2
KUWAIT	175	24	13.8	12.5	3.0	48.4	10.6
LIBERIA	51	16	30.8	14.3	1.2	190.6	29.0
LIBYA	89	9	10.5	8.0	1.9	25.3	9.0
MADAGASC	50	10	20.3	20.8	3.9	62.6	14.9
MALAWI	12	3	23.0	18.5	1.6	66.1	22.5
MALAYSIA	71	34	47.1	24.9	3.6	116.4	39.2
MALI	33	5	15.7	11.1	2.7	43.0	11.9
MAURITAN	48	17	36.1	5.1	0.9	60.6	34.6
MAURITIU	59	25	42.5	31.2	4.3	117.4	35.6
MEXICO	216	64	29.6	31.5	4.9	135.5	22.8
MOROCCO	571	222	38.8	30.9	7.0	124.3	29.6
MOZAMBIQ	11	4	33.3	26.5	17.0	77.5	18.8
NEPAL	2	1	37.8	16.5	5.1	85.0	32.0
NETHERL	3,255	1,350	41.5	14.1	1.5	110.4	37.5
NEWZEAL	65	28	43.3	34.9	6.2	122.9	32.6
NICARAG	16	1	5.9	8.8	3.7	19.7	2.3
NIGER	60	25	41.9	28.0	4.1	105.2	36.3
NIGERIA	256	19	7.2	12.8	4.8	58.6	5.0
NORWAY	665	251	37.7	23.0	3.4	115.2	32.1
OMAN	60	6	10.3	2.6	0.6	20.4	9.5
PAKISTAN	166	61	36.7	31.9	5.2	135.6	28.0
PANAMA	61	8	13.7	10.4	2.0	45.1	11.9
PAPNEWGU	3	1	32.3	11.2	0.6	63.8	29.4
PARAGUAY	20	4	19.2	20.9	2.9	67.0	15.6

PERU	77	30	39.1	34.3	7.7	134.2	31.1
PHILIPP	62	31	50.5	43.4	6.5	176.2	39.3
PORTUGAL	607	306	50.4	32.4	6.1	151.1	40.6
ROMANIA	139	80	57.8	66.3	7.8	236.0	46.1
RWANDA	15	2	13.2	16.9	2.5	55.6	11.2
SAUDI	572	88	15.4	11.4	2.0	52.1	12.4
SENEGAL	146	53	36.5	24.9	4.4	100.1	29.0
SIERRAL	3	1	37.2	7.3	2.5	64.2	34.4
SINGAPO	276	101	36.6	11.5	0.9	77.2	33.7
SOMALIA	4	0	7.3	4.0	0.1	17.6	6.6
SOUTHAFR	265	160	60.1	45.3	10.9	180.5	44.3
SPAIN	3,211	1,711	53.3	41.7	9.0	221.4	41.1
SRILANKA	12	4	35.5	24.9	9.1	95.8	25.7
SUDAN	18	2	10.8	18.5	3.0	48.0	8.6
SWEDEN	1,082	570	52.7	33.4	7.4	202.5	40.8
SWITZER	2,637	1,296	49.1	31.1	5.7	153.6	39.4
SYRIA	79	16	20.8	17.8	3.1	62.5	15.3
TAIWAN	197	127	64.1	46.1	8.8	193.3	48.5
TANZANIA	5	1	14.8	11.4	4.7	35.0	12.0
THAILAND	386	178	46.1	36.7	5.7	296.2	38.0
TOGO	54	6	11.2	7.3	0.9	28.1	9.4
TRINTOB	5	1	22.2	11.9	0.7	44.3	18.6
TUNISIA	438	145	33.1	25.2	4.3	101.0	26.2
TURKEY	461	171	37.1	36.3	6.8	150.3	29.2
UGANDA	3	0	7.6	3.3	0.7	33.8	6.5
UK	6,404	3,007	47.0	34.6	7.0	213.1	36.0
US	6,379	2,716	42.6	41.4	9.0	247.4	31.6
URUGUAY	27	17	64.2	54.5	8.5	182.0	51.0
USSR	523	149	28.5	29.1	7.0	142.0	19.8

VENEZUEL	173	34	19.6	18.9	5.2	79.7	14.4
VIETNAM	14	6	41.2	32.3	4.3	117.7	37.2
YUGOSLAV	298	153	51.3	42.3	9.8	189.1	33.9
ZAIRE	94	35	36.7	3.3	0.1	52.4	35.7
ZAMBIA	9	5	52.1	17.5	2.6	122.8	42.8
ZIMBABWE	10	5	48.5	35.5	7.0	120.3	38.5
Totals	68,877	32,738					

Figure 1a:
Entry of French Firms and Market Size

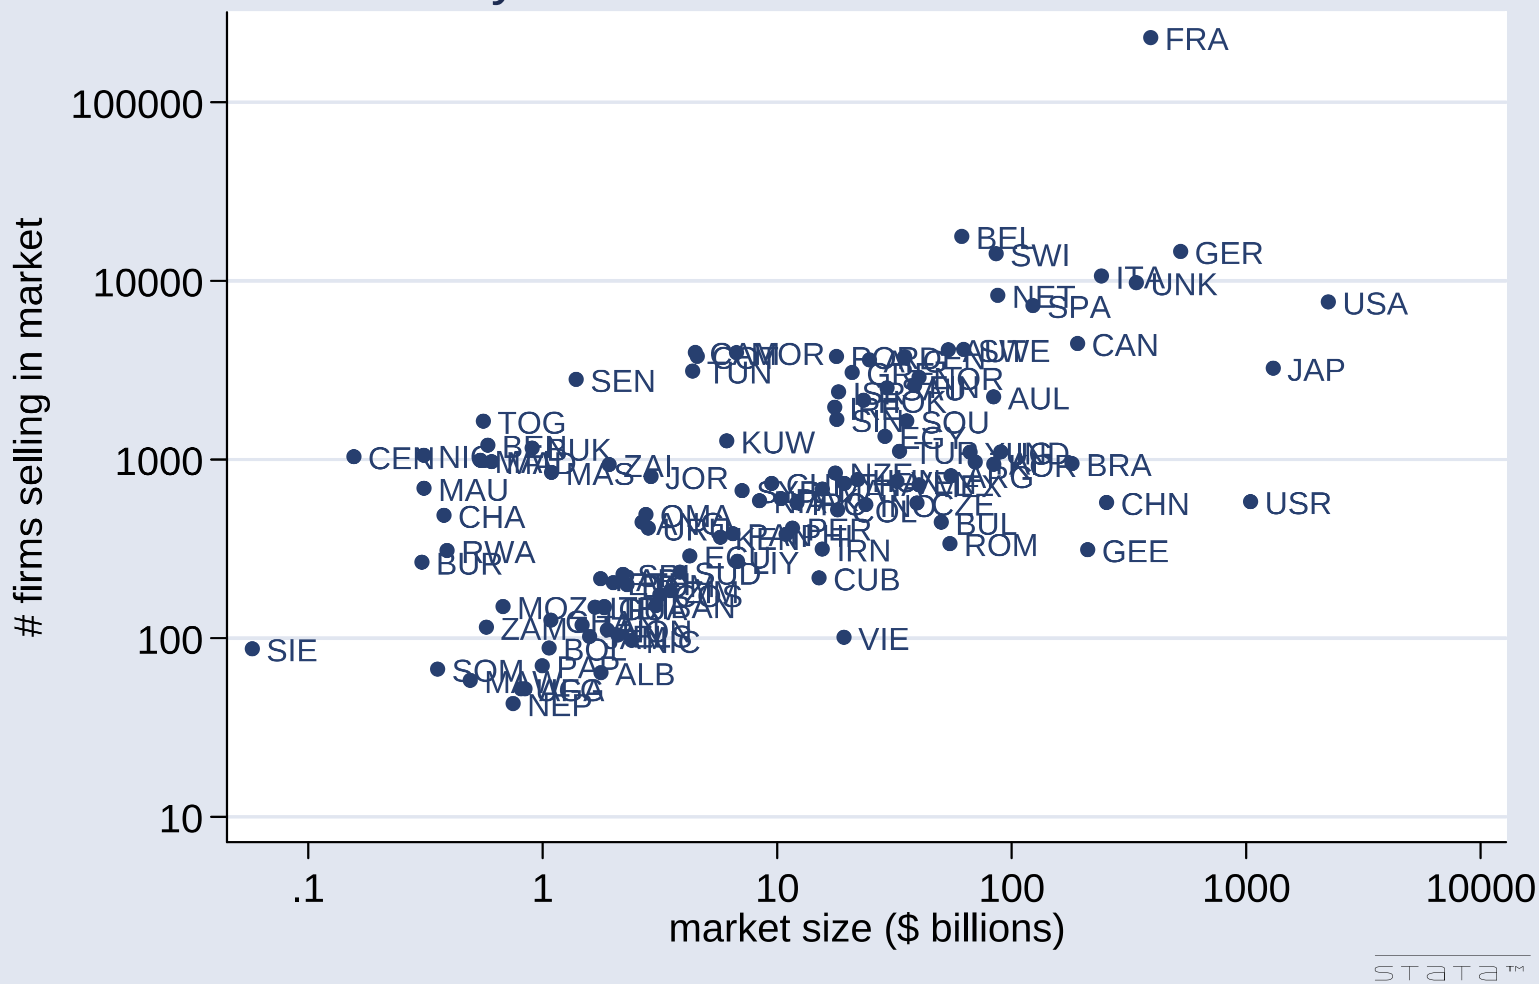

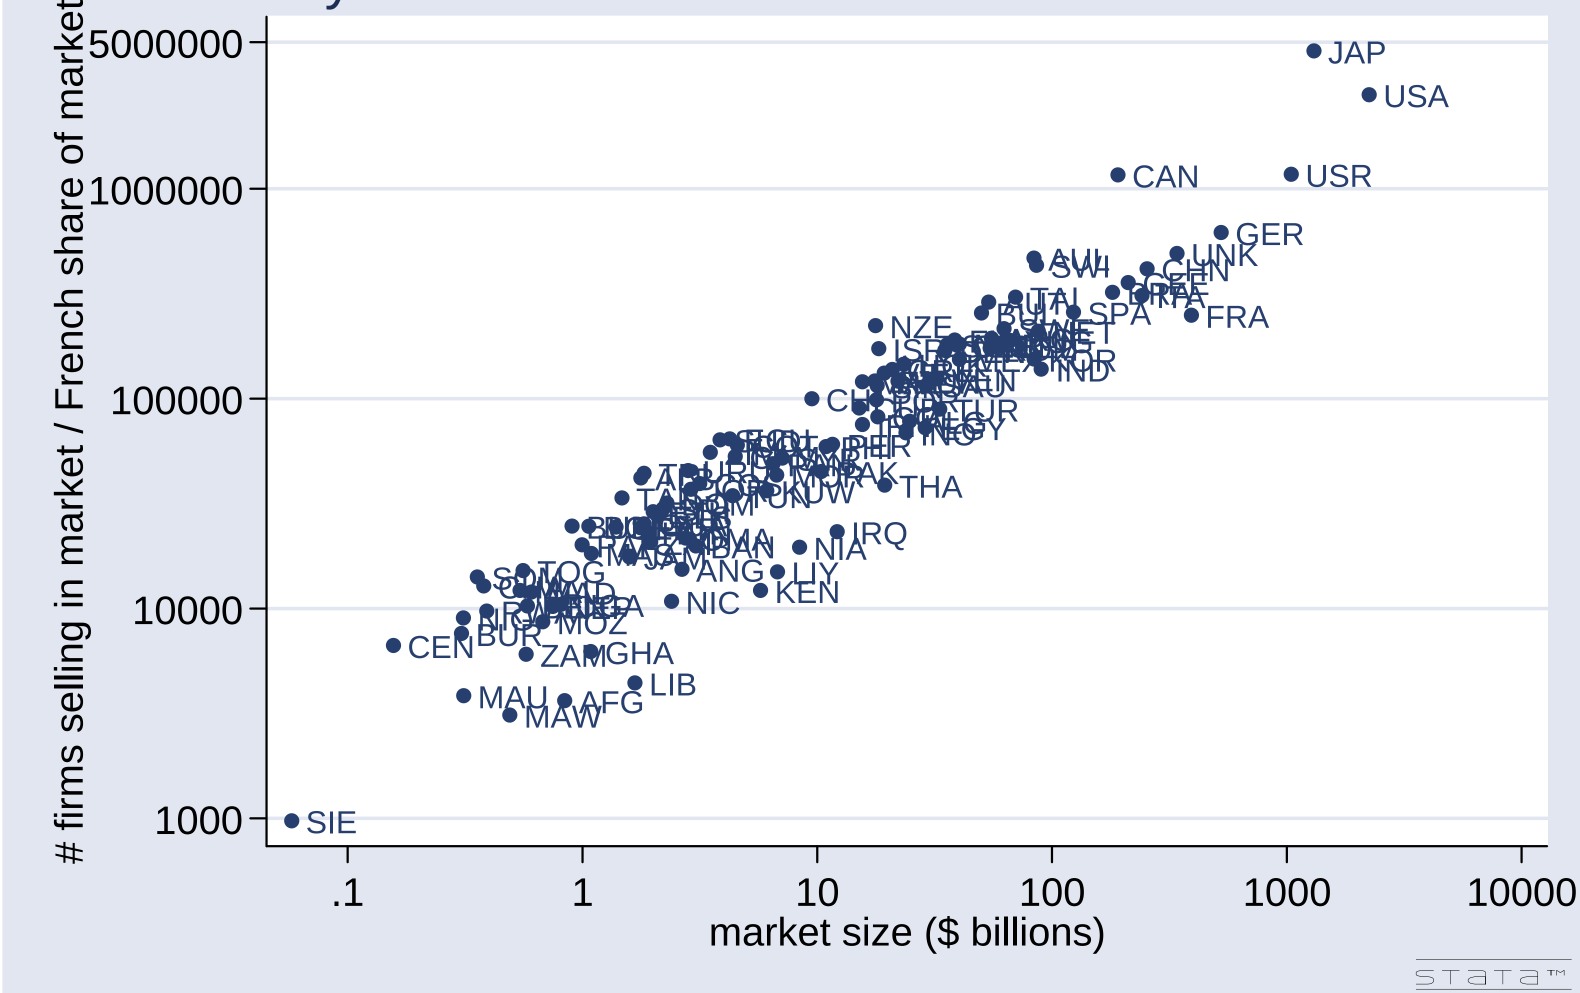

Figure 1b:
Entry Relative to French Market Share and Market Size

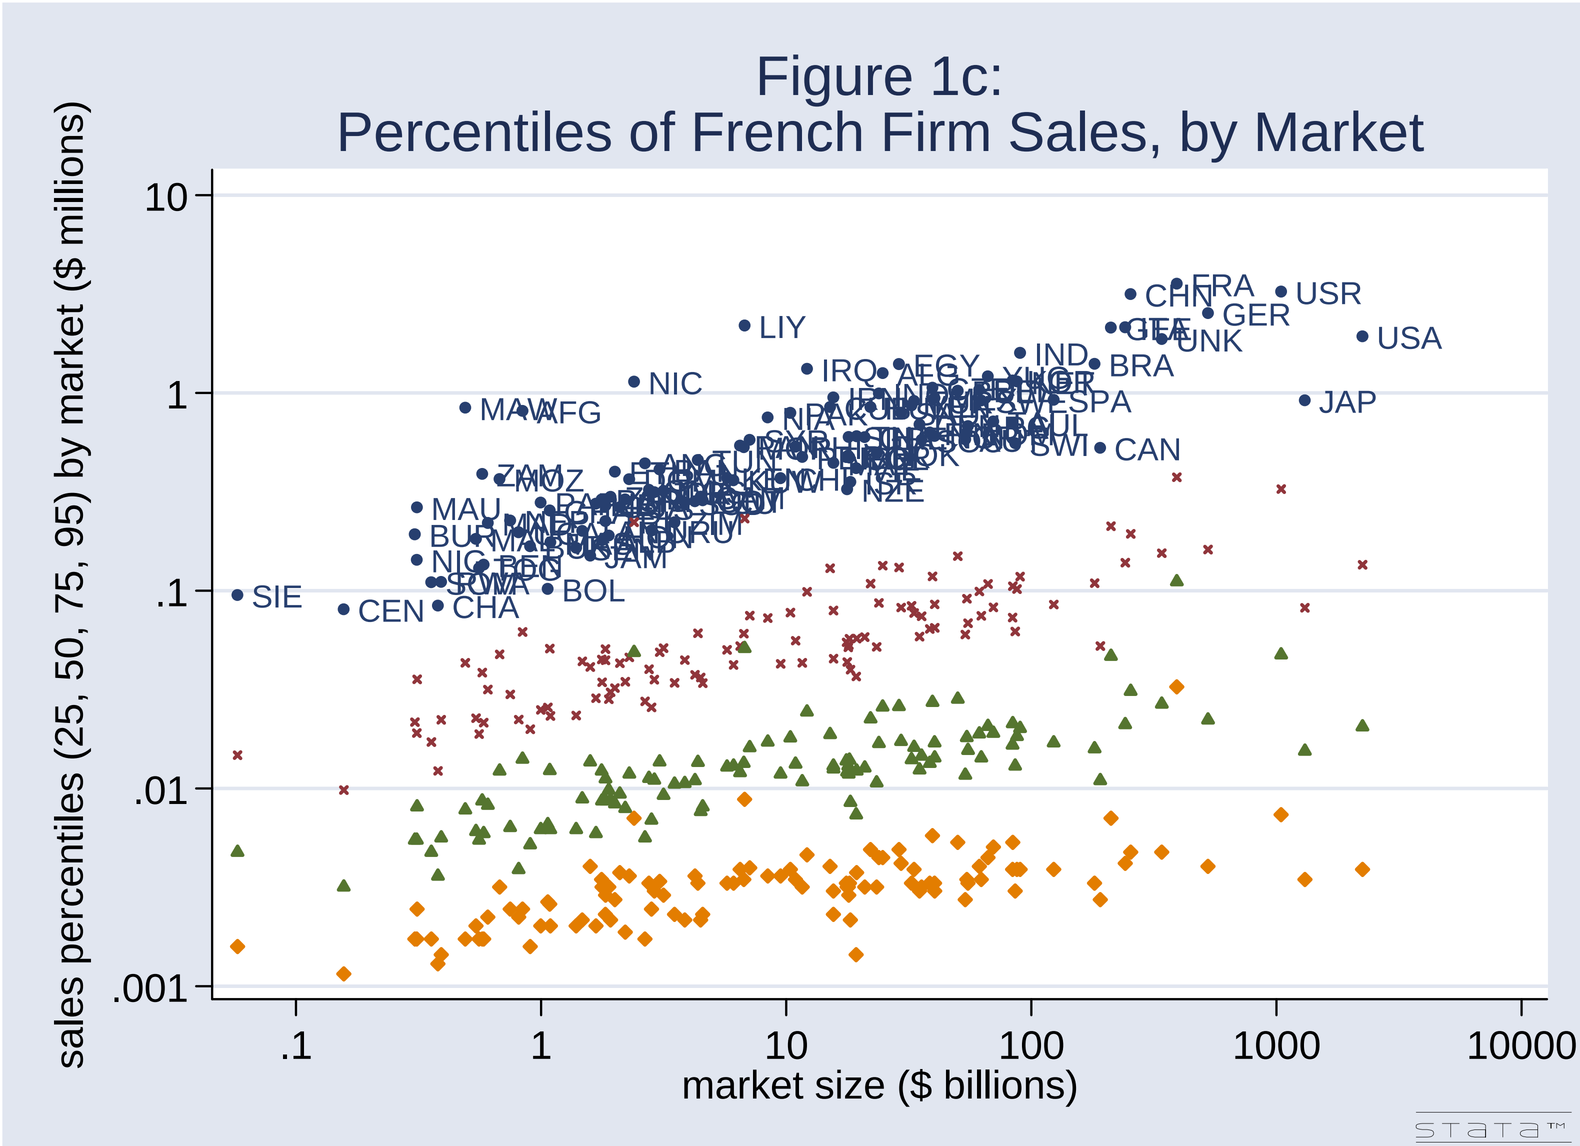
Figure 1c:
Percentiles of French Firm Sales, by Market
sales percentiles (25, 50, 75, 95) by market ($ millions)
10
1
.1
.01
.001
.1
1
10
100
1000
10000
market size ($ billions)
SIE
CEN
CHA
BOL
LIY
NIC
MAU
IRQ
EGY
IND
BRA
CHN
FRA
USR
GER
UNK
USA
JAP
CAN
SPA
SWI
NZE
JAM
stata™

Figure 2
Sales Distributions of French Firms

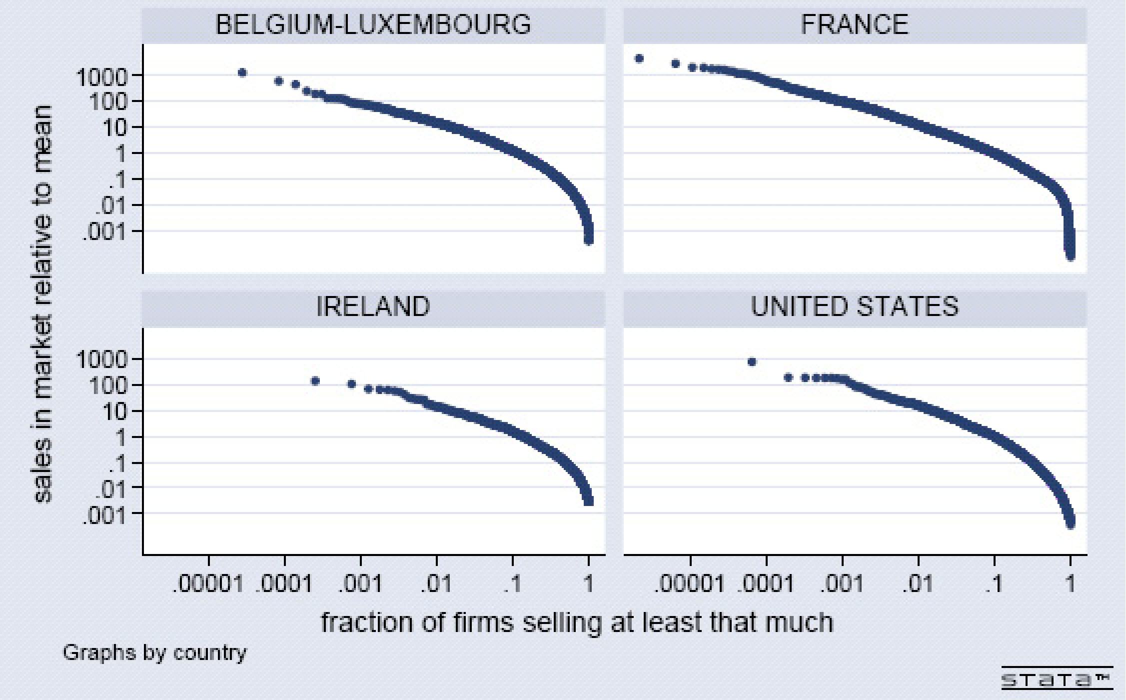

Figure 3a:
Size in France and Number of Markets Penetrated

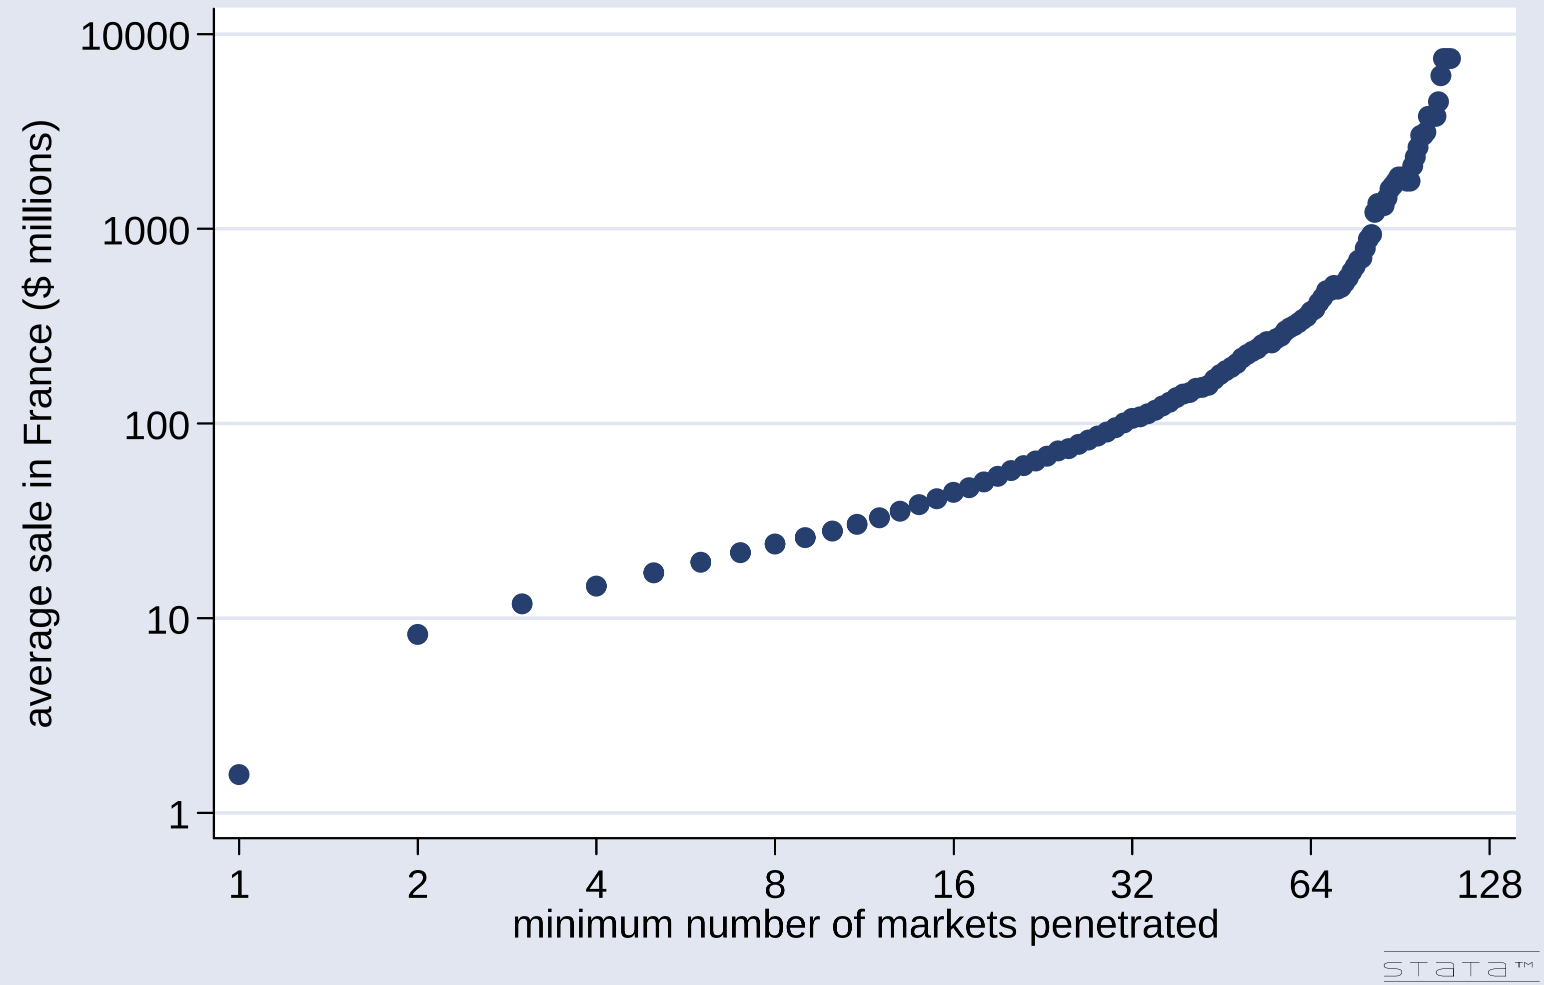

Figure 3b:
Size in France and Number Entering Multiple Markets

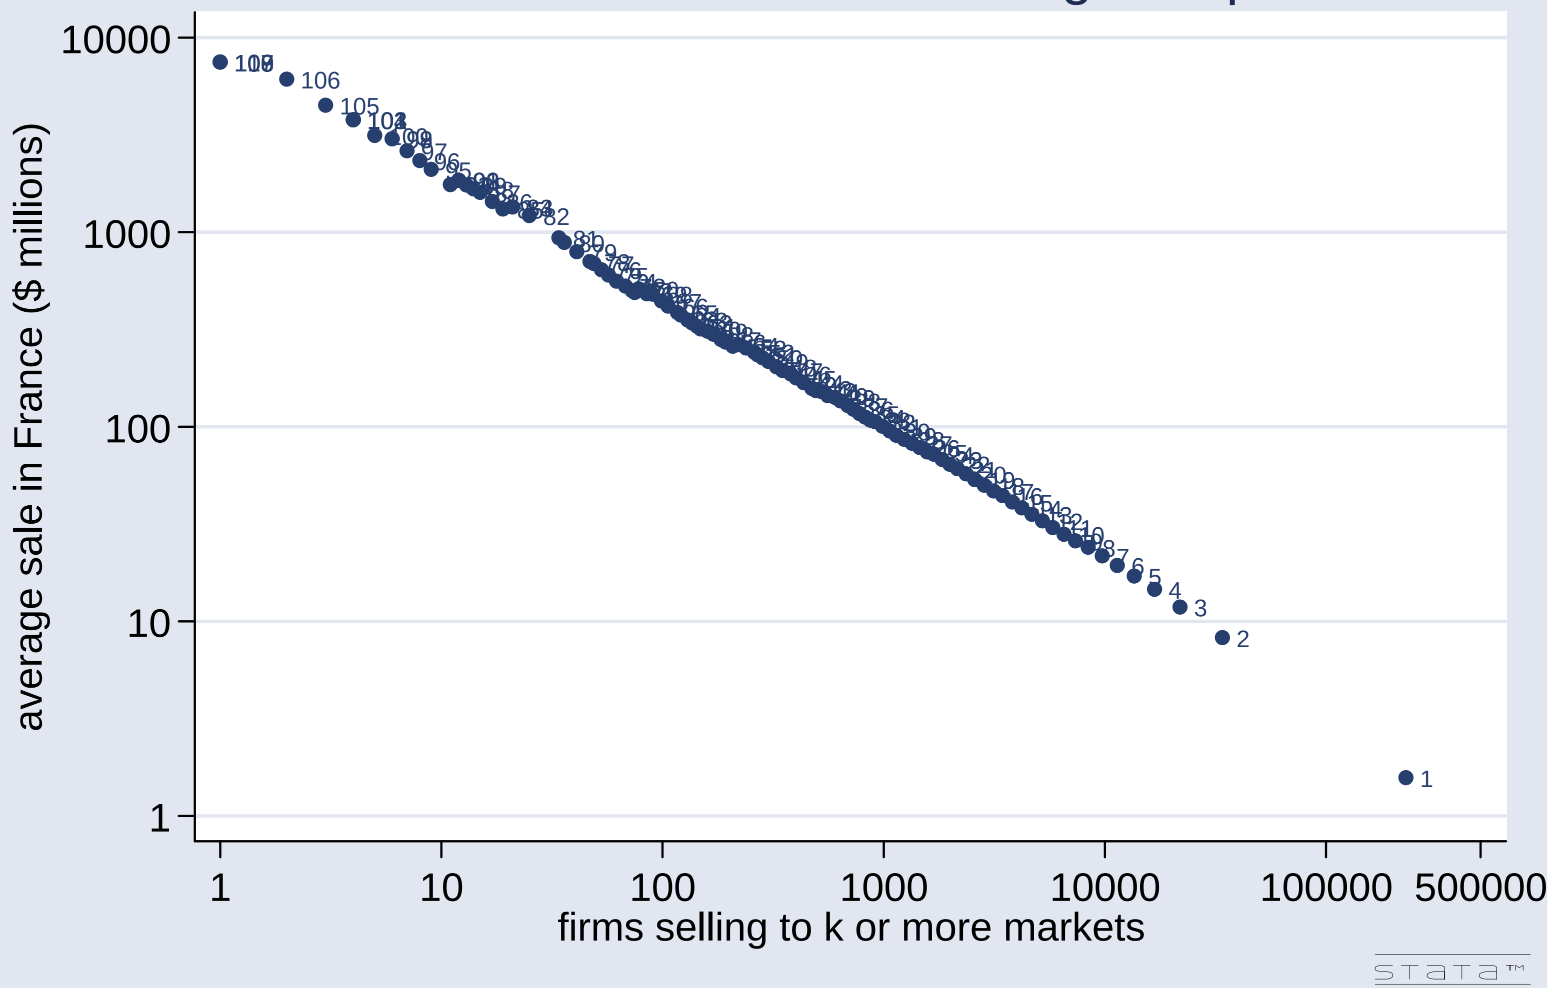

Figure 3c:
Size in France and Number Entering the Market

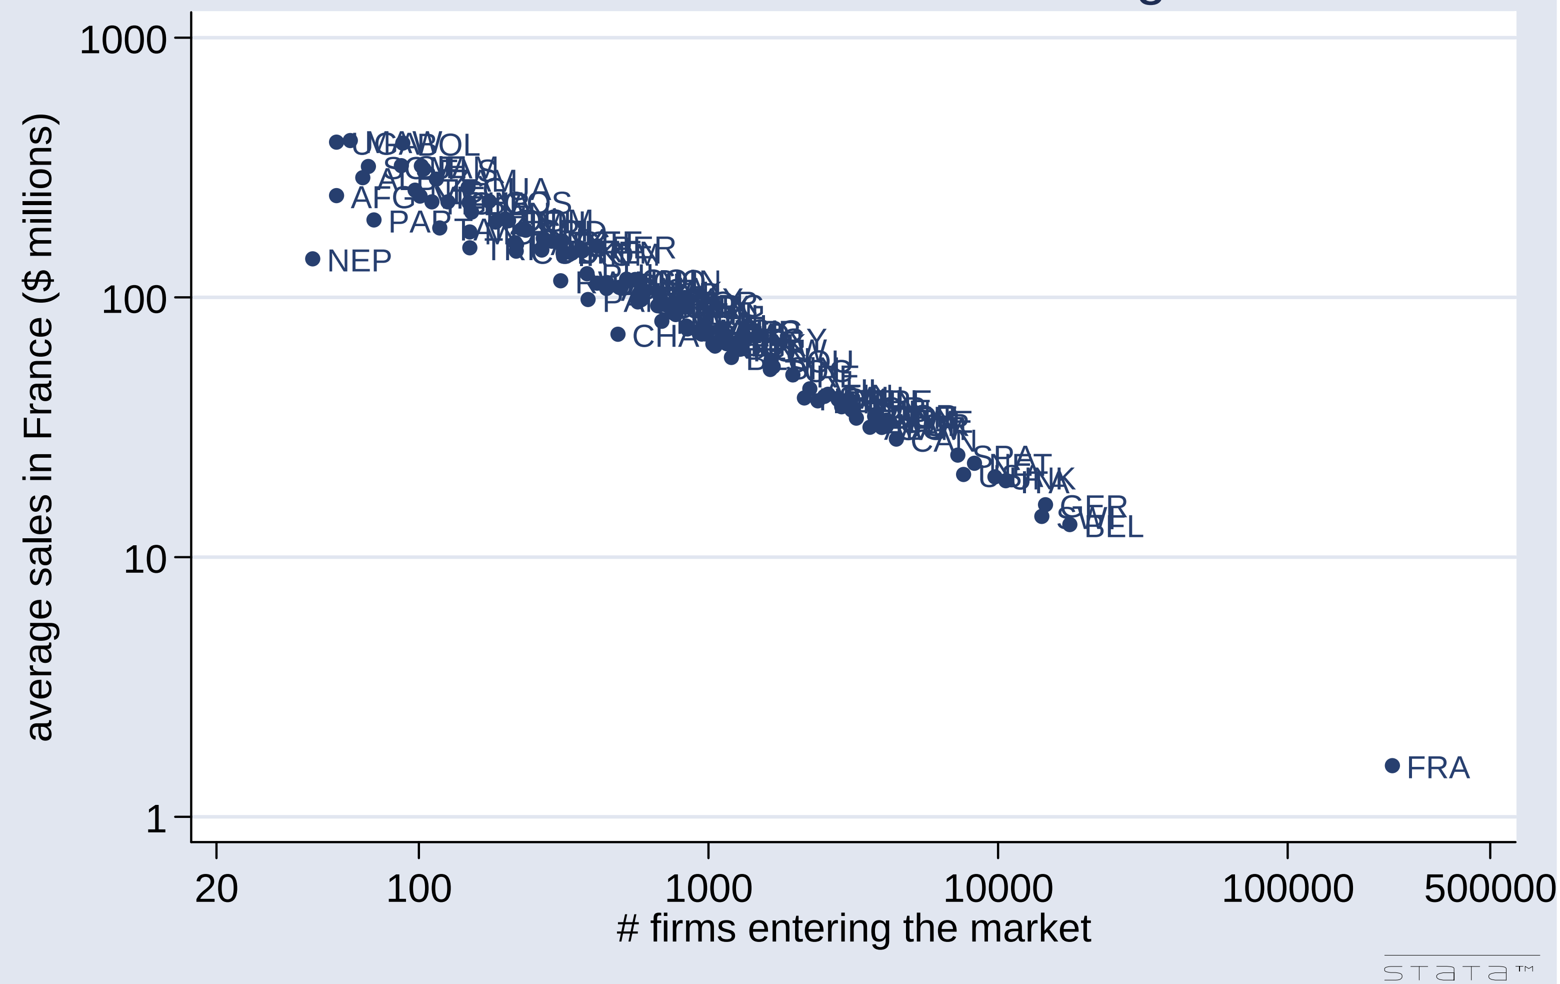

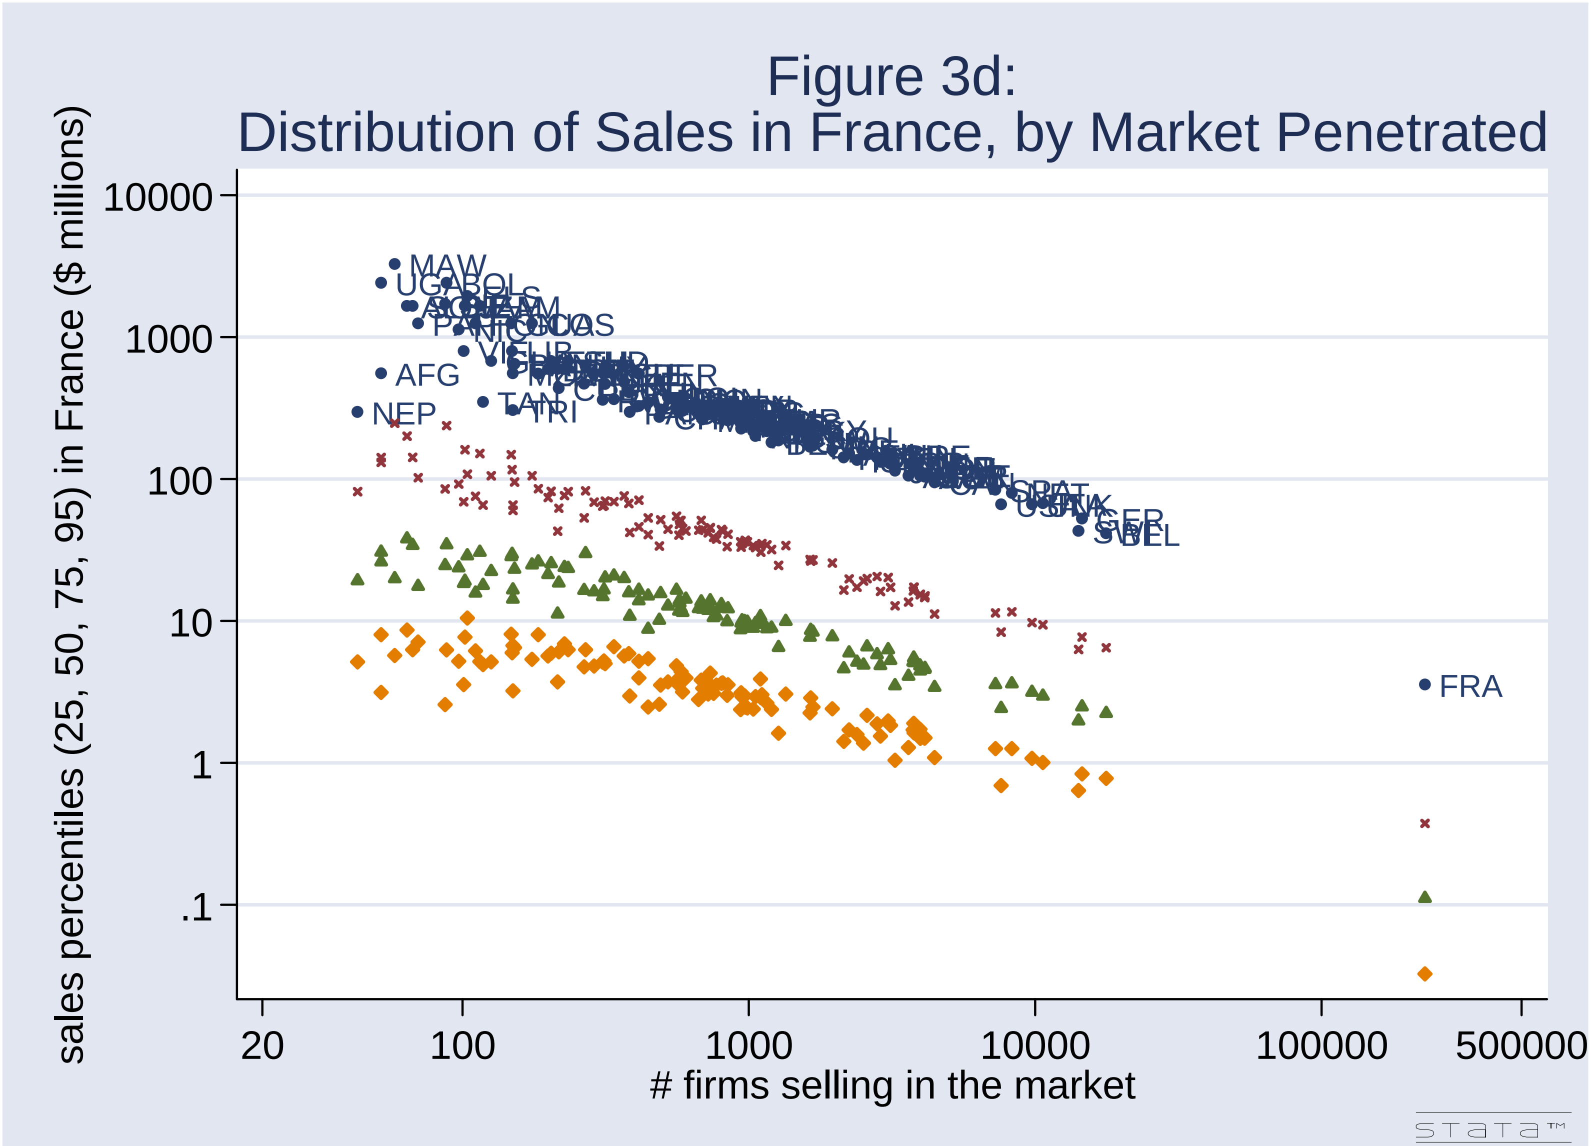
Figure 3d:
Distribution of Sales in France, by Market Penetrated
sales percentiles (25, 50, 75, 95) in France ($ millions)
10000
1000
100
10
1
.1
20
100
1000
10000
100000
500000
firms selling in the market
MAW
AFG
NEP
TAN
TRI
FRA
stata™

Figure 4:
Distribution of Export Intensity, by Market

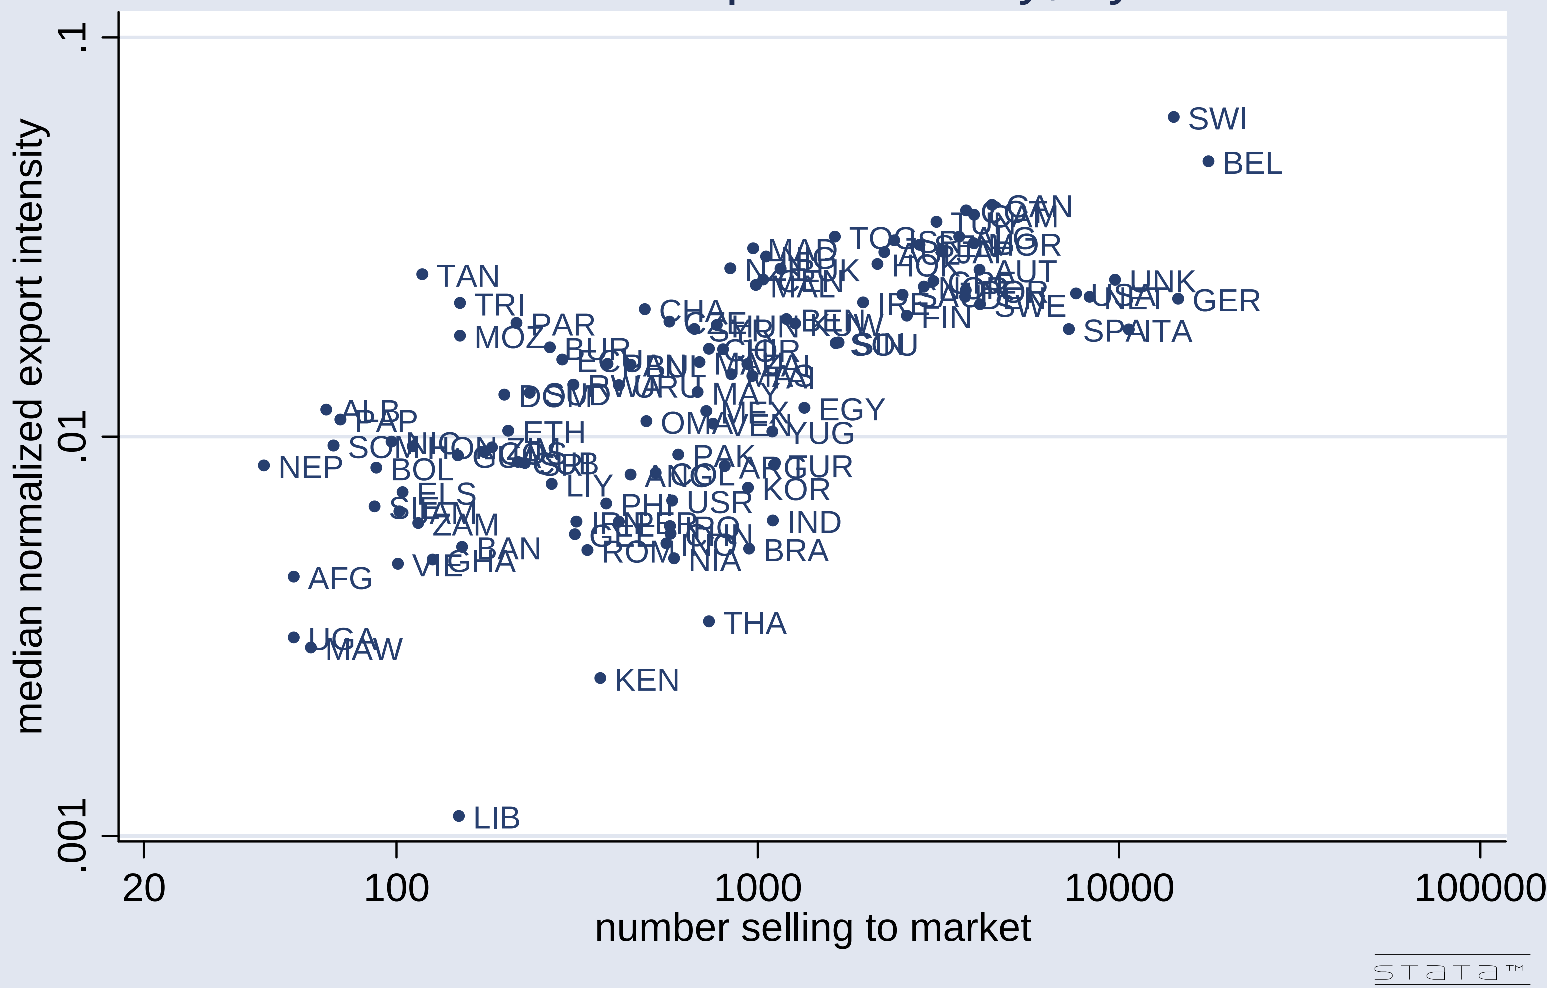

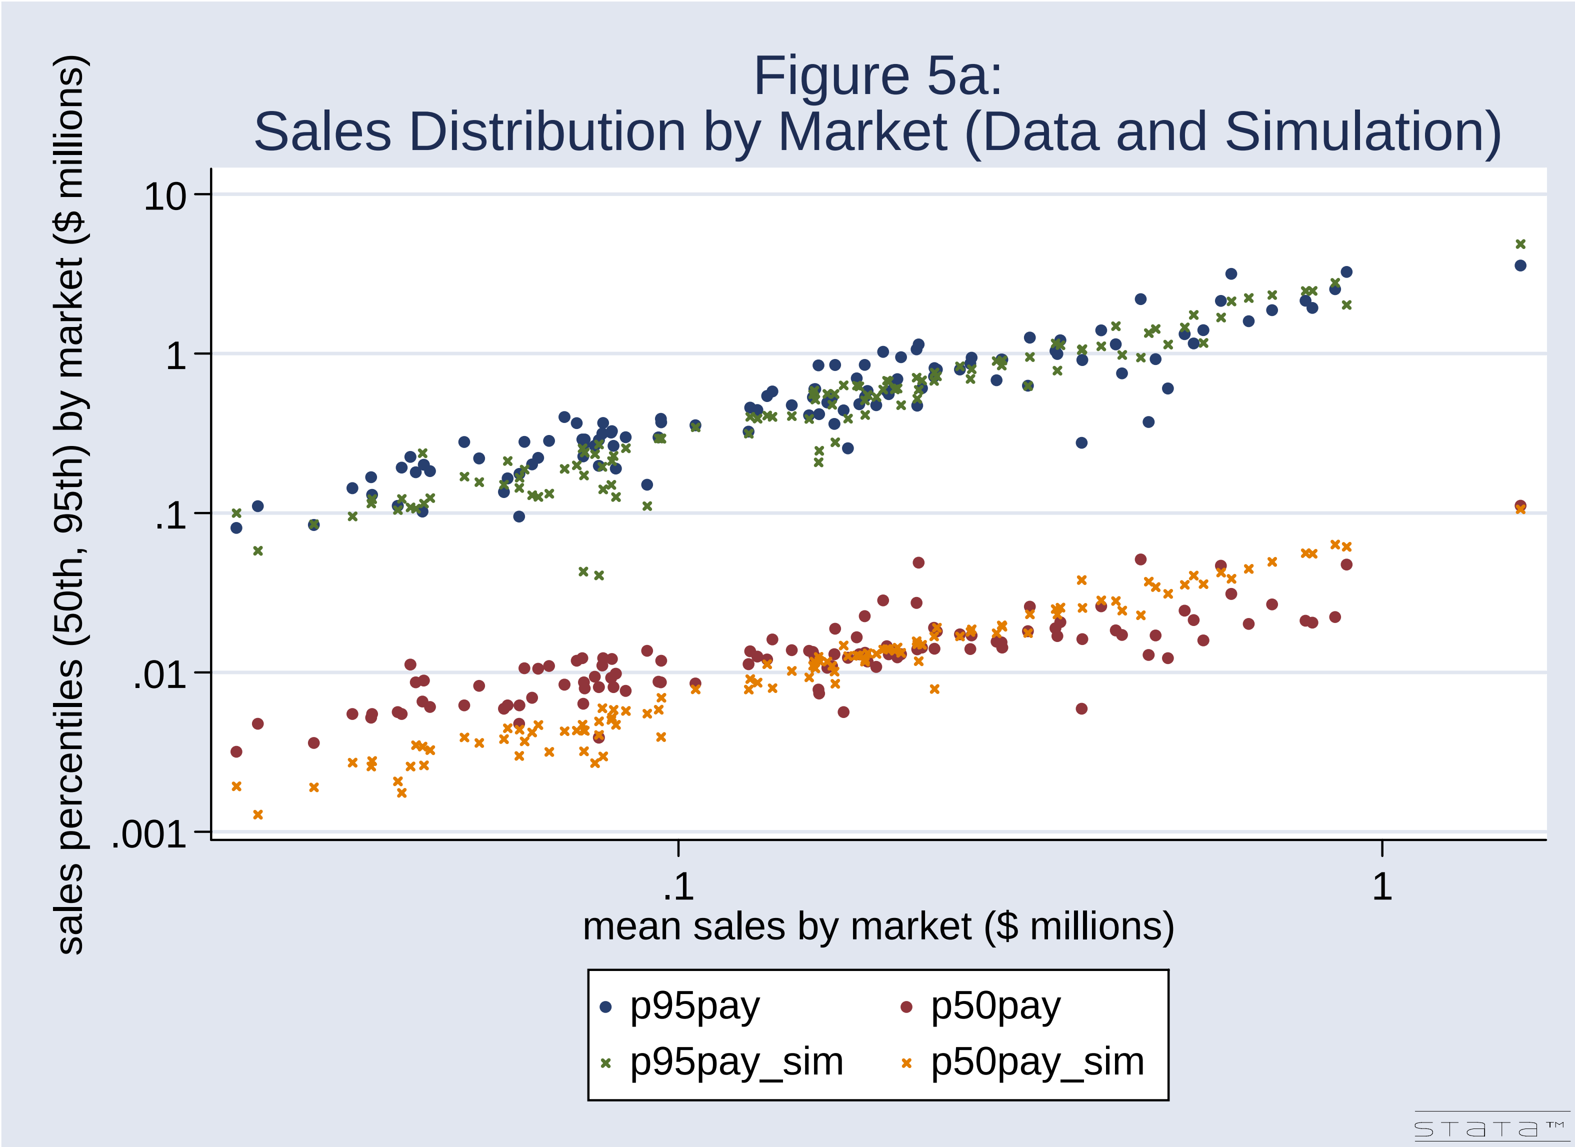

Figure 5a:
Sales Distribution by Market (Data and Simulation)
sales percentiles (50th, 95th) by market ($ millions)
10
1
.1
.01
.001
.1
1
mean sales by market ($ millions)
p95pay
p50pay
p95pay_sim
p50pay_sim
stata™

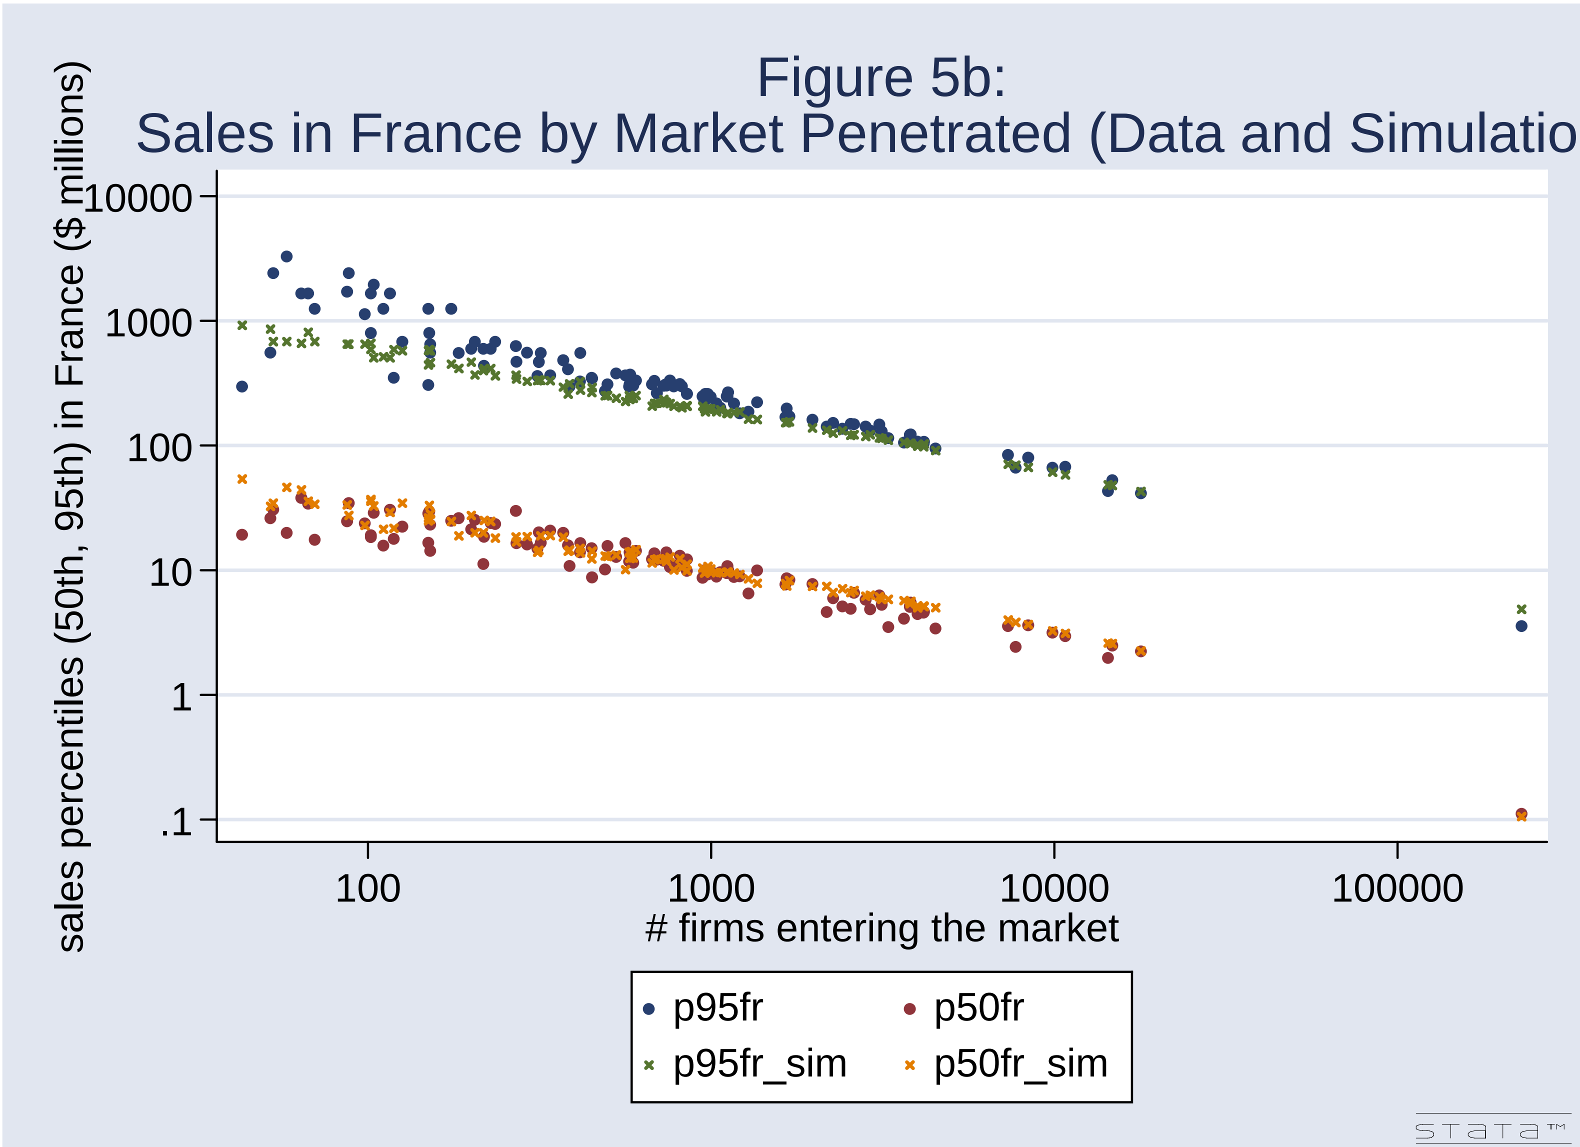
Figure 5b:
Sales in France by Market Penetrated (Data and Simulation)
sales percentiles (50th, 95th) in France ($ millions)
10000
1000
100
10
1
.1
100
1000
10000
100000
firms entering the market
p95fr
p50fr
p95fr_sim
p50fr_sim
stata™

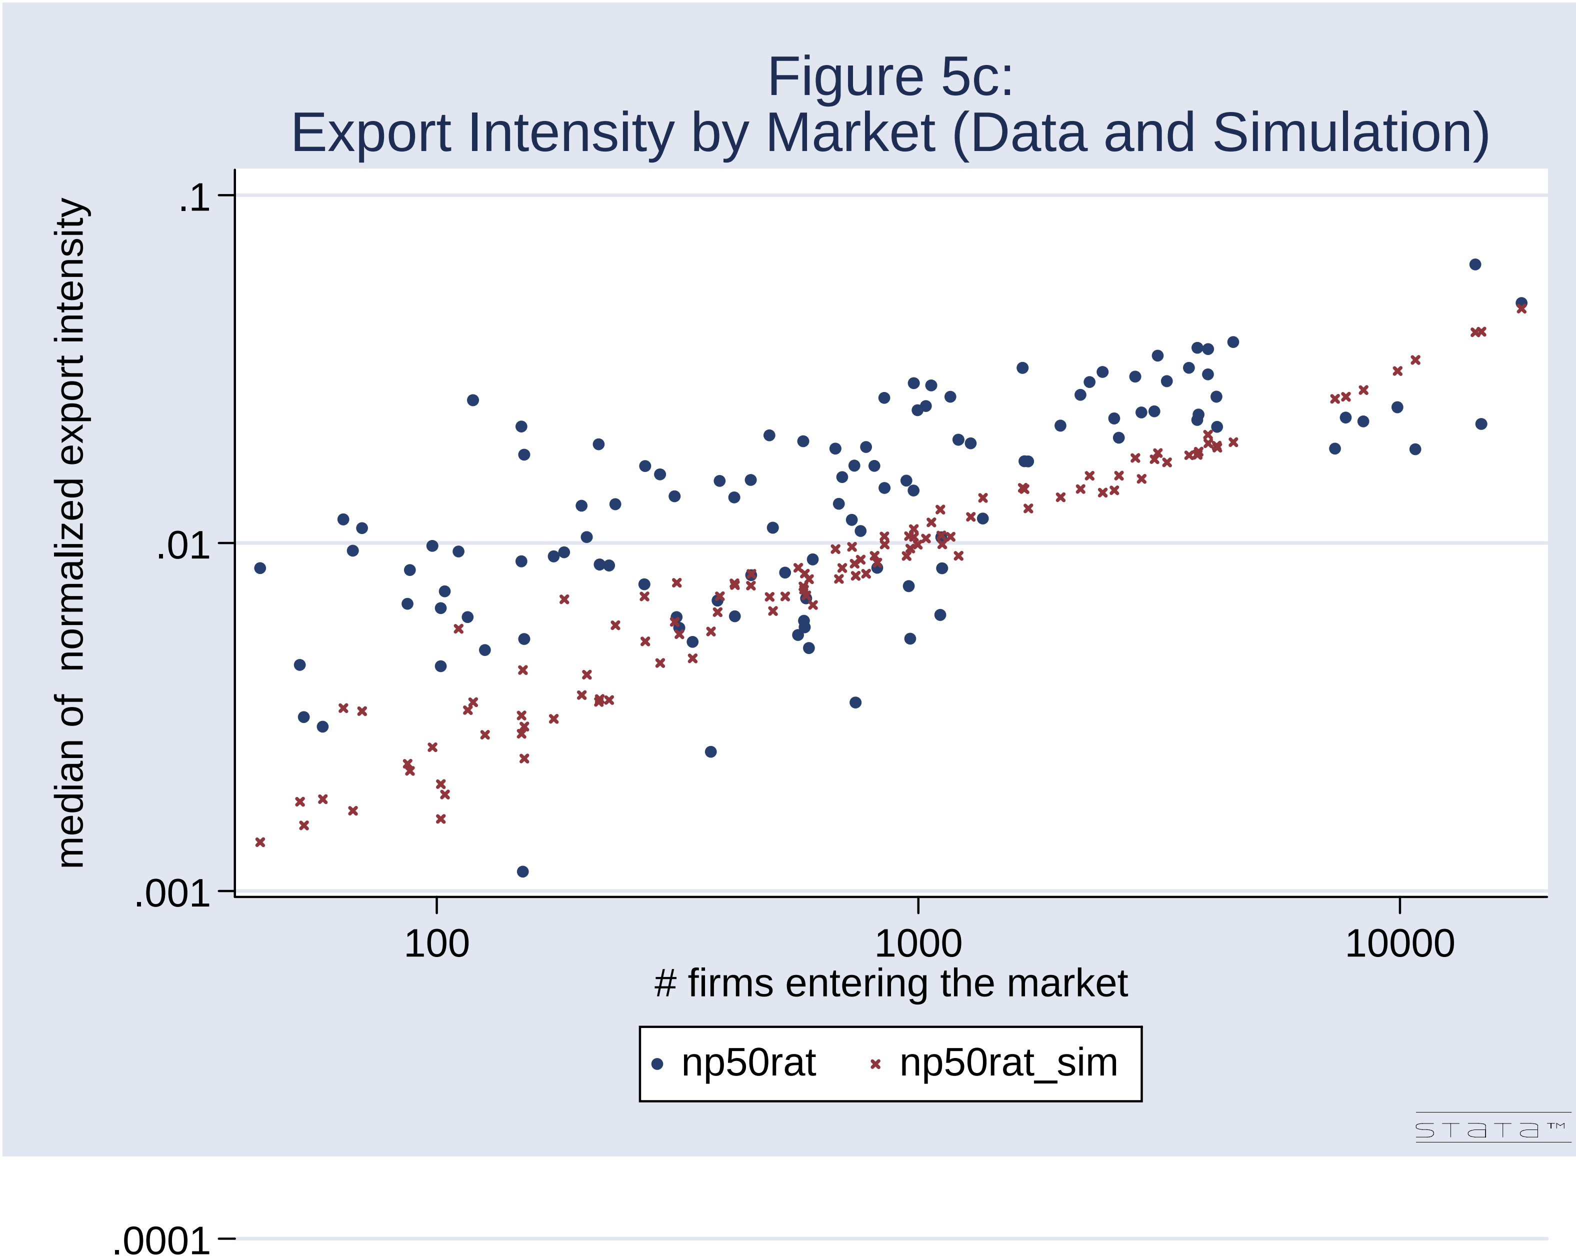
Figure 5c:
Export Intensity by Market (Data and Simulation)
median of normalized export intensity
.1
.01
.001
100
1000
10000
firms entering the market
np50rat
np50rat_sim
.0001

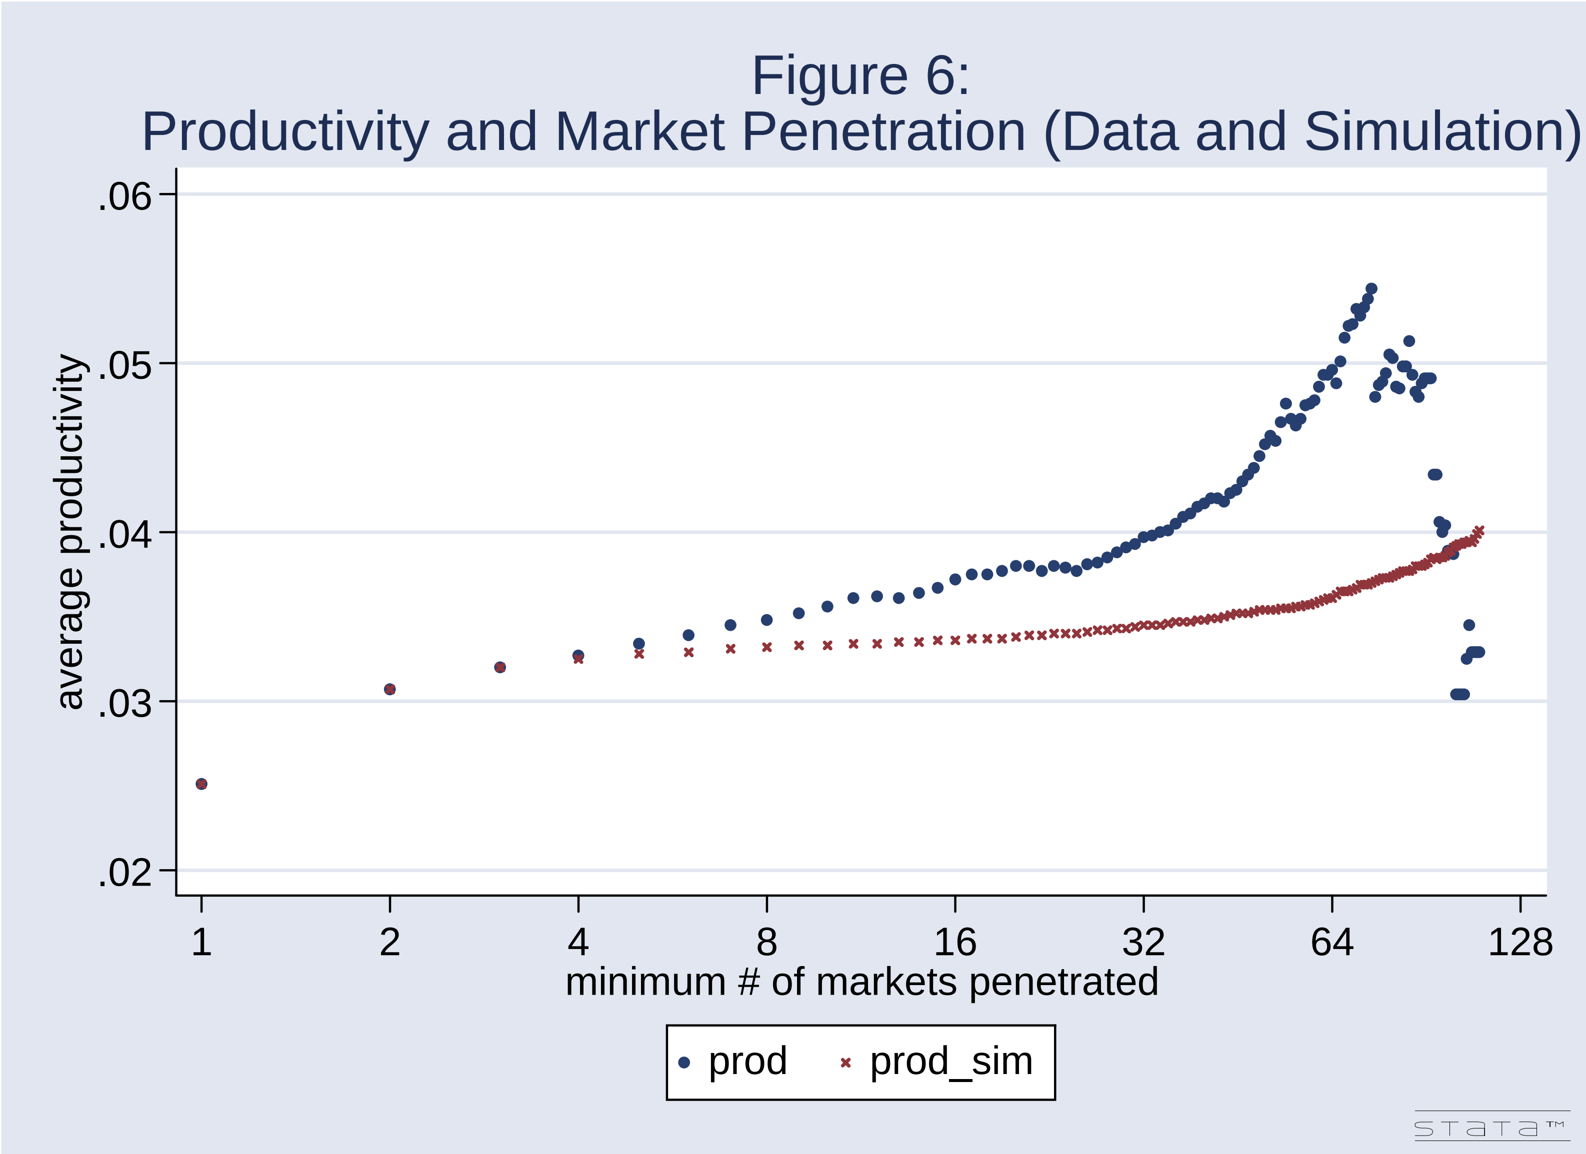

Figure 6:
Productivity and Market Penetration (Data and Simulation)
average productivity
.02
.03
.04
.05
.06
1
2
4
8
16
32
64
128
minimum # of markets penetrated
prod
prod_sim
stata™