An Anatomy of International Trade: Evidence from French Firms[1]

Jonathan Eaton,[2] Samuel Kortum,[3] and Francis Kramarz[4]

February 2010

[1]We thank Costas Arkolakis, Christopher Flinn, Jeremy Fox, Xavier d'Haultfoeuille, and Azeem Shaikh for extremely useful comments and Dan Lu for research assistance. The paper has benefitted enormously from the reports of three anonymous referees and from comments received at many venues, in particular the conference to celebrate the scientific contributions of Robert E. Lucas, Jr., April 2007. Eaton and Kortum gratefully acknowledge the support of the National Science Foundation under grant number SES-0339085.

[2]The Pennsylvania State University (jxe22@psu.edu).

[3]University of Chicago (kortum@uchicago.edu).

[4]CREST-INSEE

Abstract

We examine the sales of French manufacturing firms in 113 destinations, including France itself. Several regularities stand out: (1) the number of French firms selling to a market, relative to French market share, increases systematically with market size; (2) sales distributions are similar across markets of very different size and extent of French participation; (3) average sales in France rise 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 over 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 less variation in sales in any given market. We use our results to simulate the effects of a 10 percent counterfactual decline in bilateral trade barriers on French firms. While total French sales rise by around US$16 billion, sales by the top decile of firms rise by nearly US$23 billion. Every lower decile experiences a drop in sales, due to selling less at home or exiting altogether.

1 Introduction

We exploit detailed data on the exports of French firms to confront a new generation of trade theories. Those theories resurrect technological heterogeneity as the force driving international trade. In Eaton and Kortum (2002) differences in efficiencies across countries in making different goods determine aggregate bilateral trade flows. Since they focus only on aggregate data, underlying heterogeneity across individual producers remains hidden. Another literature, particularly Melitz (2003) and Bernard, Eaton, Jensen, and Kortum (henceforth BEJK, 2003), have developed models in which firm heterogeneity explicitly underlies comparative advantage. An implication is that data on individual firms can provide another window on the determinants of international trade.

On the purely empirical side a literature has established a number of regularities about firms in trade.[1] Another literature has modeled and estimated the export decision of individual firms in partial equilibrium.[2] However, the task of building a structure that can simultaneously embed behavior at the firm level into aggregate relationships and dissect aggregate shocks into their firm-level components remains incomplete. This paper seeks to further this mission.

To this end we examine the sales of French manufacturing firms in 113 destinations, including France itself. Combining these data with observations on aggregate trade and production reveals striking regularities in: (1) patterns of entry across markets, (2) the distribution of sales across markets, (3) how export participation connects with sales at home, and (4) how sales abroad relate to sales at home.

[1] For example, Bernard and Jensen (1995), for the United States, and Aw, Chung, and Roberts (1998), for Taiwan and Korea, document the size and productivity advantage of exporters.

[2] A pioneering paper here is Roberts and Tybout (1997).

We adopt Melitz (2003), as augmented by Helpman, Melitz, and Yeaple (2004) and Chaney (2008), as a basic framework for understanding these relationships. Core 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.

The basic model fails to come to terms with some features of the data, however: (1) Firms don't enter markets according to an exact hierarchy. (2) Their sales where they do enter deviate from the exact correlations the basic model would insist upon. (3) Firms that export sell too much in France. (4) In the typical destination there are too many firms selling small amounts.

To reconcile the basic model with the first two failures we introduce market and firm-specific heterogeneity in entry costs and demand. We deal with the last two by incorporating Arkolakis's (2008) formulation of market access costs. 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 main parameters using the method of simulated moments.

Our parameter estimates imply 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 fifty-seven 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 ignoring features of the firm that are universal across markets misses much.

We conclude by using our parameterized model to examine a world with lower trade barriers. To do so we embed our model into a general equilibrium framework, calibrating it to data on production and bilateral trade. A striking finding is the extent to which lower trade barriers, while raising welfare in every country, favor the largest firms at the expense of others. Total output of French firms rises by 3.8 percent, with all of the growth accounted for by firms in the top decile. Sales in every other decile fall. Import competition leads to the exit of 11.5 percent of firms, 43 percent of which are in the smallest decile.

Section 2 which follows explores four 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 model and considers some implications of the parameters. Section 5 explores the consequences of lower trade costs.

2 Empirical Regularities

Our data are the sales, translated into U.S. dollars, of 230,423 French manufacturing firms to 113 markets in 1986. (Table 3 lists the countries.) Among them only 34,558 sell outside France. The firm that exports most widely sells to 110 out of the 113 destinations. All but 523 of these firms indicate positive sales in France. Since much of our analysis focusses on the relationship between exporting and activity in France, we exclude these firms from much of our analysis, leaving 34,035 exporters also selling in France.[3]

[3] Appendix A describes the data. Biscourp and Kramarz (2007) and Eaton, Kortum, and Kramarz (EKK, 2004) use the same sources. EKK (2004) partition firms into 16 manufacturing sectors. While features vary across industries, enough similarity remains to lead us to ignore the industry dimension here. If a firm's

We cut the data in four different ways, each revealing sharp regularities.

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.[4] While the number of firms selling to a market tends to increase with the size of the market, the relationship is a cloudy one.

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 N_{nF}, the number of French firms selling to a market, with N_{nF} divided by French market share, π_{nF}, defined as:

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

where X_{nF} is total exports of our French firms to market 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.[5]

total exports declared to French customs exceed its total sales from mandatory reports to the French fiscal administration, we treat the firm as not selling in France. The 523 such firms represent 1.51 percent of all French exporters and account for 1.23 percent of the total French exports to our 112 export destinations.

[4] Manufacturing absorption is calculated as total production plus imports minus exports. See EKK (2004) for details.

[5] 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.

While the number of firms selling to a market rises with market size, so do sales per firm. Figure 1c shows the 95th, 75th, 50th, and 25th percentile sales in each market (on the y axis) against market size (on the x axis). The upward drift is apparent across the board.

We now turn to firm entry into different sets of markets. 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. Consider exporters to the top seven foreign destinations. Table 1 reports these destinations and the number of firms selling to each, as well as the total number of exporters. The last column of the table reports, for each top 7 destination, the marginal probability of a French exporter selling there.

Table 2 lists each of the strings of top 7 destinations that obey a hierarchical structure, together with the number of firms selling to each string (irrespective of their export activity outside the top 7). Overall, 27 percent of exporters (9260/34035) obey a hierarchy among these most popular destinations. The next column of Table 2 uses the marginal probabilities from Table 1 to predict the number selling to each hierarchical string, if selling in one market is independent of selling in any other of the top 7. Under independence, the number adhering to a hierarchy would be less than half of what we see in the data.

2.2 Sales Distributions

Our second exercise expands upon Figure 1c by looking at the entire distribution of sales within individual markets. We plot the sales of each firm in a particular market (relative to

We can then interpret Figure 1b as telling us how the number of sellers varies with market size.

mean sales there) against the fraction of firms selling in the market who sell at least that much.[6] Doing so for all our 113 destinations a remarkable similarity emerges. Figure 2 plots the results for Belgium, France, Ireland, and the United States, on common axes. The basic shape is common for size distributions.[7]

2.3 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.

[6]Following Gabaix and Ibragimov (2010) 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 $(r_n(j) - .5)/N_{nF}$.

[7]Sales distributions are often associated with the Pareto, at least in the upper tail (Simon and Bonini, 1958). To interpret Figure 2 as distributions, let x_n^q be the q'th percentile 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. If the distribution is Pareto with parameter $a > 1$ (so that the minimum sales relative to the mean is $(a-1)/a$), we have:

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

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 $-1/a$. 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). Note, however, that the distributions appear to deviate from the Pareto, especially at the lower end.

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, on a log scale, average sales in France of firms selling to k or more markets against the number of firms selling to k or more markets. The relationship is strikingly linear with a regression slope of -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 with 17,699 exporters. The least popular is Nepal, where only 43 French firms sell. 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. Firms selling to less popular markets and 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 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).[8]

[8] We were able to observe the relationship between market popularity and sales in France for the 1992 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 correlate as they should with changes in the

2.4 Export Intensity

Having looked separately at what exporters sell abroad and what they sell in France, we now examine the ratio of the two. We introduce the concept of a firm 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 the median and 95th percentile 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: (i) As a destination becomes more popular, normalized export intensity rises. The slope for the median is 0.38, but the relationship is a noisy one. (ii) Normalized "export" intensity for France itself is identically one, while median export intensity in Figure 4 is usually two orders of magnitude or more below one. Even among exporting firms, sales abroad are small compared to sales at home.

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, as the relationship would dictate.

3 Theory

In seeking to explain these relationships we turn to a parsimonious model which explains where firms sell and how much they sell there. 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.[9]

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:

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

[9] We go with Melitz's (2003) monopolistic competition approach rather than the Ricardian framework with a fixed range of commodities used in BEJK (2003), since it more readily delivers the feature that a larger market attracts more firms, as we see in our French data.

where $\theta > 0$ and $T_i \geq 0$ are parameters.[10] 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:

$$\mu_{ni}(c) = \mu_i^z \left(\frac{w_i d_{ni}}{c} \right) = \Phi_{ni} c^{\theta}, \tag{3}$$

where:

$$\Phi_{ni} = T_i (w_i d_{ni})^{-\theta}.$$

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 country i selling good j in country n must incur a fixed cost:

$$E_{ni}(j) = \varepsilon_n(j) E_{ni} M(f). \tag{4}$$

Here $\varepsilon_n(j)$ is a fixed-cost shock specific to good j in market n and E_{ni} is the component of the cost shock faced by all sellers from country i in destination n. The function $M(f)$, the same

[10] We follow Chaney (2008) in taking T_i as exogenous and Helpman, Melitz, and Yeaple (2004) and Chaney (2008) in treating the underlying heterogeneity in efficiency as Pareto. 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 (2007). 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 BEJK (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.)

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):

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

where the parameter $\lambda > 0$ reflects the increasing cost of reaching a larger fraction of potential buyers.[11] This function has the desirable properties that the cost of reaching zero buyers in a market is zero and that the total cost is increasing (and the marginal cost weakly increasing) in the fraction f of buyers reached.

Each potential buyer in market n has the same probability f of being reached by a particular seller which 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 (CES) aggregator with elasticity σ, where we require $\theta - 1 > \sigma > 1$. Hence we can write the aggregate demand for good j, if

[11] By L'Hôpital's Rule, at $\lambda = 1$ the function becomes $M(f) = -\ln(1 - f)$. Arkolakis (2008) provides an extensive discussion of this functional form, deriving it from a model of the microfoundations of marketing. He parameterizes the function in terms of $\beta = 1/\lambda$. The case $\beta > 0$ corresponds to an increasing marginal cost of reaching additional buyers, which always results in an outcome with $0 \leq f < 1$. The case $\beta \leq 0$ means a constant or decreasing marginal cost of reaching additional buyers, in which case the firm would go to a corner ($f = 0$ or $f = 1$) as in Melitz (2003). Our restriction $\lambda > 0$ thus covers all cases (including Melitz as $\lambda \rightarrow \infty$) that we could observe. A virtue of this formulation is that it can simultaneously explain why some firms sell a very small amount in a market while others stay out entirely.

it has price p and reaches a fraction f of the buyers in market n, as:

$$X_n(j) = \alpha_n(j) f X_n \left(\frac{p}{P_n}\right)^{-(\sigma-1)} \tag{5}$$

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:

$$\Pi_{ni}(p, f) = \left(1 - \frac{c_n(j)}{p}\right) \alpha_n(j) f \left(\frac{p}{P_n}\right)^{-(\sigma-1)} X_n - \varepsilon_n(j) E_{ni} M(f). \tag{6}$$

To maximize profit a producer will set the standard Dixit-Stiglitz (1977) markup over unit cost:

$$p_n(j) = \overline{m} c_n(j) \tag{7}$$

where:

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

It will seek a fraction:

$$f_{ni}(j) = \max\left\{1 - \left[\eta_n(j) \frac{X_n}{\sigma E_{ni}} \left(\frac{\overline{m} c_n(j)}{P_n}\right)^{-(\sigma-1)}\right]^{-\lambda}, 0\right\} \tag{8}$$

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:

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

We can now 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:

$$c \leq \overline{c}_{ni}(\eta) \tag{9}$$

where:

$$\overline{c}_{ni}(\eta) = \left(\eta \frac{X_n}{\sigma E_{ni}}\right)^{1/(\sigma-1)} \frac{P_n}{\overline{m}}. \tag{10}$$

We can use the expression for (10) to simplify the expression for the fraction of buyers a producer with unit cost $c \leq \overline{c}_{ni}(\eta)$ will reach:

$$f_{ni}(j) = 1 - \left(\frac{c}{\overline{c}_{ni}(\eta)}\right)^{\lambda(\sigma-1)}. \tag{11}$$

Substituting (7), (11) and (10) into (5) gives:

$$X_{ni}(j) = \varepsilon \left[1 - \left(\frac{c}{\overline{c}_{ni}(\eta)}\right)^{\lambda(\sigma-1)}\right] \left(\frac{c}{\overline{c}_{ni}(\eta)}\right)^{-(\sigma-1)} \sigma E_{ni}. \tag{12}$$

Conditioning on $\overline{c}_{ni}(\eta)$, ε replaces α in (5) as the shock to sales.

Since it charges a markup $\overline{m} = \sigma/(\sigma - 1)$ over unit cost its total gross profit is simply $X_n(j)/\sigma$. Some of this profit is eaten up by its fixed cost, which, from (4) and (11), is:

$$E_{ni}(j) = \varepsilon E_{ni} \frac{1 - (c/\overline{c}_{ni}(\eta))^{(\lambda-1)(\sigma-1)}}{1 - 1/\lambda}. \tag{13}$$

To summarize, the relevant characteristics of market n that apply across sellers are total purchases X_n, the price index P_n, and, for sellers from country i, the common component of the fixed cost E_{ni}. 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 $c_n(j)$.[12]

Equations (9) and (10), governing entry, and (12), 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 $f_{ni}(\eta, c)$ 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:

$$P_n = \overline{m} \left[\int \int \left(\sum_{i=1}^{N} \int_0^{\bar{c}_{ni}(\eta)} \alpha f_{ni}(\eta, c) c^{1-\sigma} d\mu_{ni}(c) \right) g(\alpha, \eta) d\alpha d\eta \right]^{-1/(\sigma-1)}.$$

To solve we use (3), (10), (11), and the laws of integration, to get:

$$P_n = \overline{m} \left[\kappa_0 \int \int \alpha \sum_{i=1}^{N} \Phi_{ni} \left[\left(\eta \frac{X_n}{\sigma E_{ni}} \right)^{1/(\sigma-1)} \frac{P_n}{\overline{m}} \right]^{\theta-(\sigma-1)} g(\alpha, \eta) d\alpha d\eta \right]^{-1/(\sigma-1)}$$

where:

$$\kappa_0 = \frac{\theta}{\theta - (\sigma - 1)} - \frac{\theta}{\theta + (\sigma - 1)(\lambda - 1)}.$$

Moving P_n to one side of the equation gives:

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

[12] We require $E[\eta^{\theta/(\sigma-1)}]$ and $E[\alpha\eta^{[\theta/(\sigma-1)]-1}]$ both to be finite. For any given $g(\alpha, \eta)$, these restrictions will imply an upper bound on the parameter θ.

where:

$$\kappa_1 = \kappa_0 \int \int \alpha \eta^{[\theta-(\sigma-1)]/(\sigma-1)} g(\alpha, \eta) d\alpha d\eta. \tag{15}$$

and:

$$\Psi_n = \sum_{i=1}^{N} \Phi_{ni} (\sigma E_{ni})^{-[\theta-(\sigma-1)]/(\sigma-1)}. \tag{16}$$

Note that the price index has an elasticity of $(1/\theta)-1/(\sigma-1)$ with respect to total expenditure (given the terms in Ψ_n). Our restriction that $\theta > \sigma - 1$ makes the effect negative: A larger market enjoys lower prices, for reasons similar to the price index effect in Krugman (1980) and common across models of monopolistic competition.[13]

3.4 Entry, Sales, and Fixed Costs

An individual firm enters market n if its cost c is below the threshold given by (10), which we can now rewrite using (14) as:

$$\overline{c}_{ni}(\eta) = \eta^{1/(\sigma-1)} \left(\frac{X_n}{\kappa_1 \Psi_n} \right)^{1/\theta} (\sigma E_{ni})^{-1/(\sigma-1)}. \tag{17}$$

For firms from country i with a given value of η, a measure $\mu_{ni}(\overline{c}_{ni}(\eta))$ will pass the entry hurdle.

To obtain the total measure of firms from i that sell in n, denoted J_{ni}, we integrate across the marginal density $g_2(\eta)$:

$$J_{ni} = \int \left[\mu_{ni}(\overline{c}_{ni}(\eta)) \right] g_2(\eta) d\eta = \frac{\kappa_2}{\kappa_1} \frac{\pi_{ni} X_n}{\sigma E_{ni}} \tag{18}$$

[13] In models of monopolistic competition with homogeneous firms, consumers anywhere consume all varieties regardless of where they are made. Since more varieties are produced in a larger market, local consumers benefit from their ability to buy these goods without incurring trade costs. With heterogeneous firms and a fixed cost of market entry, an additional benefit to consumers in a large market is greater variety.

where:

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

and:

$$\pi_{ni} = \frac{\Phi_{ni}(\sigma E_{ni})^{-[\theta-(\sigma-1)]/(\sigma-1)}}{\Psi_n}. \tag{20}$$

From (12), integrating over the measure of costs, given by (3), substituting (17), and replacing ε with α/η, firms from country i with a given value of α and η will sell a total amount:

$$X_{ni}(\alpha, \eta) = \alpha \eta^{[\theta-(\sigma-1)]/(\sigma-1)} \frac{\kappa_0}{\kappa_1} \pi_{ni} X_n.$$

in market n.

To obtain total sales by firms from country i in market n, X_{ni}, we integrate across the joint density $g(\alpha, \eta)$ to get:

$$X_{ni} = \pi_{ni} X_n.$$

Hence the trade share of source i in market n is just π_{ni}.[14]

To obtain an expression for the fixed costs incurred by firms from i in n, given their values of α and η, we integrate (13), replacing ε with α/η, over the measure with c below $\bar{c}_{ni}(\eta)$ to

[14] Of the total measure J_n of firms selling in n, the fraction from i is:

$$\frac{J_{ni}}{J_n} = \frac{\pi_{ni}/E_{ni}}{\sum_{k=1}^{N} \pi_{nk}/E_{nk}}.$$

If E_{ni} doesn't vary according to i then a country's share in the measure of firms selling in n and its share in total sales there are both π_{ni}, where:

$$\pi_{ni} = \frac{T_i(w_i d_{ni})^{-\theta}}{\sum_{k=1}^{N} T_k(w_k d_{nk})^{-\theta}}.$$

get:

$$E_{ni}(\alpha,\eta) = \alpha\eta^{[\theta-(\sigma-1)]/(\sigma-1)}\left(\frac{\pi_{ni}X_n}{\sigma\kappa_1}\right)\left[\frac{\lambda}{[\theta/(\sigma-1)]+\lambda-1}\right].$$

Integrating across the joint density $g(\alpha,\eta)$, total fixed costs incurred in market n by firms from i, $\overline{E}_{ni}$, are:

$$\overline{E}_{ni} = \frac{\pi_{ni}X_n}{\sigma\theta}[\theta-(\sigma-1)].$$

Summing across sources i, the total fixed costs $\overline{E}_n$ incurred in market n are simply:

$$\overline{E}_n = \frac{\theta-(\sigma-1)}{\theta}\frac{X_n}{\sigma}. \tag{21}$$

Note that $\overline{E}_n$ (spending on fixed costs) does not depend on the E_{ni}, the country-pair components of the fixed cost per firm, just as in standard models of monopolistic competition. A drop in E_{ni} leads to more entry and ultimately the same total spending on fixed costs.

If total spending in a market is X_n then gross profits earned by firms in that market are X_n/σ. If firms were homogeneous then fixed costs would fully dissipate gross profits. With producer heterogeneity, firms with a unit cost below the entry cutoff in a market retain a net profit there. Total entry costs are a fraction $[\theta-(\sigma-1)]/\theta$ of gross profits and net profits are $X_n/(\overline{m}\theta)$.

3.5 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 of any potential producer in France as:

$$u(j) = T_F z_F(j)^{-\theta}. \tag{22}$$

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 parameter.

Substituting (22) into (1) and using (20), we can write unit cost in market n in terms of $u(j)$ as:

$$c_{nF}(j) = \frac{w_F d_{nF}}{z_F(j)} = \left(\frac{u(j)}{\Phi_{nF}}\right)^{1/\theta}. \tag{23}$$

Associated with the entry hurdle $\overline{c}_{nF}(\eta)$ is a standardized entry hurdle $\overline{u}_{nF}(\eta)$ satisfying:

$$\overline{c}_{nF}(\eta) = \left(\frac{\overline{u}_{nF}(\eta)}{\Phi_{nF}}\right)^{1/\theta}. \tag{24}$$

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

$$u(j) \leq \overline{u}_{nF}(\eta_n(j)) = \left(\frac{\pi_{nF} X_n}{\kappa_1 \sigma E_{nF}}\right) \eta_n(j)^{\widetilde{\theta}} \tag{25}$$

where:

$$\widetilde{\theta} = \frac{\theta}{\sigma - 1} > 1. \tag{26}$$

Conditional on firm j's passing this hurdle we can use (23) and (24) to rewrite firm j's sales in market n, expression (12), in terms of $u(j)$ as:

$$X_{nF}(j) = \varepsilon_n(j) \left[1 - \left(\frac{u(j)}{\overline{u}_{nF}(\eta_n(j))}\right)^{\lambda/\widetilde{\theta}}\right] \left(\frac{u(j)}{\overline{u}_{nF}(\eta_n(j))}\right)^{-1/\widetilde{\theta}} \sigma E_{nF}. \tag{27}$$

Equations (25) and (27) reformulate the entry and sales equations (17) and (12) 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 (25) and (27), 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. Observing just entry and sales we are able to identify only $\widetilde{\theta}$.

3.6 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. We equate the measure of French firms J_{nF} selling in each destination with the actual (integer) number N_{nF}, and $\overline{X}_{nF}$ with their average sales there.

3.6.1 Entry

From (18) we get:

$$\frac{N_{nF}}{\pi_{nF}} = \frac{\kappa_2}{\kappa_1}\frac{X_n}{\sigma E_{nF}}, \tag{28}$$

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_{nF} rises systematically with market size, but not proportionately so. We don't impose any such relationship, but rather employ (28) to calculate:

$$\sigma E_{nF} = \frac{\kappa_2}{\kappa_1}\frac{\pi_{nF} X_{nF}}{N_{nF}} = \frac{\kappa_2}{\kappa_1}\overline{X}_{nF} \tag{29}$$

directly from the data.[15]

Using (29) we can write (25) as:

$$u(j) \leq \overline{u}_{nF}(\eta_n(j)) = \frac{N_{nF}}{\kappa_2}\eta_n(j)^{\widetilde{\theta}}. \tag{30}$$

Without variation in the firm and market specific entry shock $\eta_n(j)$, (30) 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, as we see in Table 2, identify variation in $\eta_n(j)$

3.6.2 Sales in a Market

Conditional on a firm's entry into market n, the term:

$$v_{nF}(j) = \frac{u(j)}{\overline{u}_{nF}(\eta_n(j))} \tag{31}$$

is distributed uniformly on $[0, 1]$. Replacing $u(j)$ with $v_{nF}(j)$ in expression (27) and exploiting (29) we can write sales as:

$$X_{nF}(j) = \varepsilon_n(j)\left[1 - v_{nF}(j)^{\lambda/\widetilde{\theta}}\right] v_{nF}(j)^{-1/\widetilde{\theta}}\frac{\kappa_2}{\kappa_1}\overline{X}_{nF}. \tag{32}$$

[15] We can use equation (29) to infer how fixed costs vary with country characteristics. Simply regressing $\overline{X}_{nF}$ against our market size measure X_n (both in logs) yields a coefficient of 0.35, one minus the slope in Figure 1b, in which N_{nF}/π_{nF} rises with market size with an elasticity of 0.65. (The connection between the two regressions is a result of the accounting identity: $\overline{X}_{nF} * N_{nF}/\pi_{nF} = X_n$.) If GDP per capita is added to the regression it has a negative effect on entry costs. French data from 1992 and data from Denmark (from Pedersen, 2009) and from Uruguay (compiled by Raul Sampognaro) show similar results for market size but not a robust effect of GDP per capita. We also find that mean sales are higher in the home country. Appendix B reports these results.

Not only does v_{nF} have the same distribution in each market n, so does ε_n.[16] 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_{nF}), consistent with the common shapes of sales distributions exhibited in Figure 2. If the only source of variation in sales were the term $v_{nF}(j)^{-1/\widetilde{\theta}}$ the sales distribution would be Pareto with parameter $\widetilde{\theta}$. The term in square brackets, however, implies a downward deviation from the Pareto that is more pronounced as $v_{nF}(j)$ gets small, consistent with the curvature in the lower end of the sales distributions that we observe in Figure 2. The sales distribution will also inherit properties of the distribution of $\varepsilon_n(j)$.

3.6.3 Sales in France Conditional on Entry 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 (32) as it applies to France and use (31) and (30) to replace $v_{FF}(j)$ with $v_{nF}(j)$:

$$X_{FF}(j)|_n = \frac{\alpha_F(j)}{\eta_n(j)} \left[1 - v_{nF}(j)^{\lambda/\widetilde{\theta}} \left(\frac{N_{nF}}{N_{FF}} \right)^{\lambda/\widetilde{\theta}} \left(\frac{\eta_n(j)}{\eta_F(j)} \right)^{\lambda} \right] v_{nF}(j)^{-1/\widetilde{\theta}} \left(\frac{N_{nF}}{N_{FF}} \right)^{-1/\widetilde{\theta}} \frac{\kappa_2}{\kappa_1} \overline{X}_{FF}. \tag{33}$$

Since $v_{nF}(j)$ and $\eta_n(j)$ have the same distributions across destinations n, the only systematic source of variation across n is N_{nF}/N_{FF}.

[16] 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:

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

which does not depend on n. The term $\eta^{\widetilde{\theta}}/\kappa_2$ captures the fact that entrants are a selected sample with typically better than average entry shocks.

Consider first its presence in the term in square brackets, representing the fraction of buyers reached in France. Since N_{nF}/N_{FF} is near zero everywhere but France, the term in square brackets is close to one for all $n \neq F$. Hence the relationship between N_{nF} and $X_{FF}(j)$ is dominated by the appearance of N_{nF} outside the square bracket, implying that sales in France of firms fall with N_{nF} with an elasticity of $-1/\widetilde{\theta}$. Interpreting Figure 3c in terms of Equation (32), the slope of -0.57 implies a $\widetilde{\theta}$ of 1.75.

Expression (33) also suggests how we can identify other parameters of the model. The gap between the percentiles in Figure 3d is governed by the variation in the demand shock $\alpha_F(j)$ in France together with variation in the entry shock $\eta_n(j)$ in country n.

Together (32) and (33) 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. An exporting firm may be close to the entry hurdle in the export market, and hence selling to a small fraction of buyers there, while reaching most consumers at home. Hence looking at the home sales of exporters isolates firms that reach most of the French market. These equations also explains why France itself is somewhat below the trend line in Figures 3a and 3b: The many nonexporting firms that reach just a small fraction of the French market appear only in the data point for France.

3.6.4 Normalized Export Intensity

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

$$\frac{X_{nF}(j)}{X_{FF}(j)} / \left(\frac{\overline{X}_{nF}}{\overline{X}_{FF}} \right) = \frac{\alpha_n(j)}{\alpha_F(j)} \left[\frac{1 - v_{nF}(j)^{\lambda/\widetilde{\theta}}}{1 - v_{nF}(j)^{\lambda/\widetilde{\theta}} \left(\frac{N_{nF}}{N_{FF}} \right)^{\lambda/\widetilde{\theta}} \left(\frac{\eta_n(j)}{\eta_F(j)} \right)^{\lambda}} \right] \left(\frac{N_{nF}}{N_{FF}} \right)^{1/\widetilde{\theta}}. \tag{34}$$

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 (33), the only systematic source of cross-country variation on the right-hand side of (34) is N_{nF}/N_{FF}. The relationship is consistent with three features of Figure 4. First, trivially, the observation for France is identically 1. Second, normalized export intensity is substantially below 1 for destinations served by only a small fraction of French firms, as is the case for all foreign markets. Third, normalized export intensity increases with the number of French firms selling there. According to (34) the elasticity of normalized export intensity with respect to N_{nF}/N_{FF} is $1/\widetilde{\theta}$ (ignoring N_{nF}/N_{FF}'s role in the denominator of the term in the square bracket, which is tiny since N_{nF}/N_{FF} is close to zero for $n \neq F$). The slope coefficient of 0.38 reported in Section 2.4 suggests a value of $\widetilde{\theta}$ of 2.63.[17]

4 Estimation

We estimate the parameters of the model by the method of simulated moments. We simulate firms that make it into at least one foreign market and into France as well.[18] Given parameter

[17] Equations (33) and (34) apply to the latent sales in France of firms that sell in n but don't enter France. In Figures 3c and 3d we can only look at the firms that sell in both places, of course. Since the French share in France is so much larger than the French share elsewhere, our theory predicts that a French firm selling in another market but not in France is very unlikely. Indeed, the number of such firms is small.

[18] The reason for the selling-in-France requirement is that key moments in our estimation procedure involve sales in France by exporters, which we can compute only for firms that enter the home market. The reason for the foreign-market requirement is that firms selling only in France are very numerous, so that capturing them would consume a large portion of simulation draws. But since their activity is so limited they add little to parameter identification. We also estimated the model matching moments of nonexporting firms as well.

estimates, we later explore the implications of the model for nonexporters as well.

We first complete our parameterization of the model. Second, we 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. Third, we describe how we calculate a set of moments from these artificial data to compare with moments from the actual data. Finally, we explain our estimation procedure, 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_α, σ_η, and correlation ρ.[19] Under these assumptions we may write (15) and (19) as:

$$\kappa_1 = \left[\frac{\widetilde{\theta}}{\widetilde{\theta} - 1} - \frac{\widetilde{\theta}}{\widetilde{\theta} + \lambda - 1} \right] \exp \left\{ \frac{\sigma_\alpha + 2\rho\sigma_\alpha\sigma_\eta(\widetilde{\theta} - 1) + \sigma_\eta(\widetilde{\theta} - 1)^2)}{2} \right\} \tag{35}$$

and:

$$\kappa_2 = \exp \left\{ \frac{\left(\widetilde{\theta}\sigma_\eta\right)^2}{2} \right\}. \tag{36}$$

As above, in estimating the model we equate the measure of French firms J_{nF} selling in each destination with the actual (integer) number N_{nF} and equate $\overline{X}_{nF}$ with their average

Coefficient estimates were similar to those we report below but the estimation algorithm, given estimation time, was much less precise.

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

$$\sigma_\varepsilon^2 = \sigma_\alpha^2 + \sigma_\eta^2 - 2\rho\sigma_\alpha\sigma_\eta,$$

which is decreasing in ρ.

sales there.[20] We are left with only five parameters to estimate:

$$\Theta = \{\widetilde{\theta}, \lambda, \sigma_\alpha, \sigma_\eta, \rho\}.$$

For a given Θ we use (29) to back out the cluster of parameters σE_{nF} using our data on $\overline{X}_{nF} = X_{nF}/N_{nF}$ and the κ_1 and κ_2 implied by (35) and (36). Similarly, we use (30) to back out a firm's entry hurdle in each market $\overline{u}_{nF}(\eta_n)$ given its η_n and the κ_2 implied by (36).

4.2 Simulation Algorithm

For estimating parameters, for assessing the implications of those estimates, and for performing counterfactual experiments, we will need to construct sets of artificial French firms that operate as the model tells them, given some Θ. 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.

As we search over different parameters Θ we want to hold fixed the realizations of the stochastic components of the model. Hence, prior to running any simulations: (i) We draw S realizations of $v(s)$ independently from the uniform distribution $U[0,1]$, putting them aside to construct standardized unit cost $u(s)$ below. (ii) We draw $S \times 113$ realizations of $a_n(s)$ and $h_n(s)$ independently from the standard normal distribution $N(0,1)$, putting them aside to construct the $\alpha_n(s)$ and $\eta_n(s)$ below.

[20] The model predicts that some French firms will export while not selling domestically. Consequently, the data for N_{nF} and X_{nF} that we condition on in our estimation include the 523 French exporters, mentioned in Footnote 3) who don't enter the domestic market.

A given simulation of the model 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 eight steps:

1. Using (35) and (36) we calculate κ_1 and κ_2.

2. Using (29) we calculate σE_{nF} for each destination n.

3. We use the $a_n(s)$'s and $h_n(s)$'s to construct $S \times 113$ realizations for each of $\ln \alpha_n(s)$ and $\ln \eta_n(s)$ as

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

4. We construct the $S \times 113$ entry hurdles:

$$\overline{u}_n(s) = \frac{N_{nF}}{\kappa_2} \eta_n(s)^{\widetilde{\theta}}, \tag{37}$$

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

5. We calculate

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

   the maximum u consistent with exporting somewhere, and

$$\overline{u}(s) = \min\{\overline{u}_F(s), \overline{u}^X(s)\}, \tag{38}$$

   the maximum u consistent with selling in France and exporting somewhere.

6. To simulate exporters that sell in France, $u(s)$ 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 $v(s)$'s that were drawn prior to the simulation.

7. 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.[21]

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

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

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

$$X_{nF}(s) = \frac{\alpha_n(s)}{\eta_n(s)} \left[1 - \left(\frac{u(s)}{\overline{u}_n(s)}\right)^{\lambda/\widetilde{\theta}}\right] \left(\frac{u(s)}{\overline{u}_n(s)}\right)^{-1/\widetilde{\theta}} \sigma E_{nF}. \tag{39}$$

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.[22]

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

[22] In principle, in any finite sample, the number of simulated firms overcoming the entry hurdle for a destination n could be zero, even though the distribution of efficiencies is unbounded from above. Helpman, Melitz, and Rubinstein (2008) account for zeros by truncating the upper tail of the Pareto distribution. Zero exports to a destination then arise simply because not even the most efficient possible firm could surmount the entry barrier there.

4.3 Moments

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 the number 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. Let N^k be the number of firms achieving some outcome k in the actual data and $\widehat{N}^k$ the corresponding number in the simulated data. Using $\delta^k(s)$ as an indicator for when artificial firm s achieves outcome k, we calculate $\widehat{N}^k$ as:

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

We now describe the moments that we seek to match.[23]

We have chosen our moments to capture the four features of French firms' behavior described in Section 2:

1. The first set of moments relate to the entry strings discussed in Section 2.1. We compute the proportion $\widehat{m}^k(1;\Theta)$ 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 $m^k(1)$ of exporters selling to combination k. Stacking these proportions gives us $\widehat{m}(1;\Theta)$ and $m(1)$, each with 128 elements (subject to 1 adding up constraint).

[23] Notice, from (37), (38), and (19), that the average of the importance weights $\overline{u}(s)$ is the simulated number of French firms that export and sell in France.

2. The second set of moments relate to the sales distributions presented in Section 2.2. For firms selling in each of the 112 export destinations n we compute the qth percentile sales $s_n^q(2)$ *in that market* (i.e., the level of sales such that a fraction q of firms selling in n sells less than $s_n^q(2)$) for q = 50, 75, 95. Using these $s_n^q(2)$ we assign firms that sell in n into four mutually exclusive and exhaustive bins determined by these three sales levels. We compute the proportions $\widehat{m}_n(2;\Theta)$ of artificial firms falling into each bin analogous to the actual proportion $m_n(2) = (0.5, 0.25, 0.2, 0.05)'$. Stacking across the 112 countries gives us $\widehat{m}(2;\Theta)$ and $m(2)$, each with 448 elements (subject to 112 adding-up constraints).

3. The third set of moments relate to the sales in France of exporting firms discussed in Section 2.3. For firms selling in each of the 112 export destinations n we compute the qth percentile sales $s_n^q(3)$ *in France* for q = 50, 75, 95. Proceeding as above we get $\widehat{m}(3;\Theta)$ and $m(3)$, each with 448 elements (subject to 112 adding-up constraints).

4. The fourth set of moments relate to normalized export intensity by market discussed in Section 2.4. For firms selling in each of the 112 export destinations n we compute the qth percentile ratio $s_n^q(4)$ of sales in n to sales in France for $q = 50$, 75. Proceeding as above we get $\widehat{m}(4;\Theta)$ and $m(4)$, each with 336 elements (subject to 112 adding-up constraints).

For the last three sets we emphasize higher percentiles because they (i) appear less noisy in the data and (ii) account for much more of total sales.

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

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

By inserting the actual data on $\overline{X}_{nF}$ (to get σE_{nF}) and N_{nF} (to get $\overline{u}_n(s)$) in our simulation routine we are ignoring sampling error in these measures. The first has no effect on our estimate of Θ. The reason is that a change in σE_{nF} would shift the sales in n of each artificial firm, mean sales $\overline{X}_{nF}$, and the percentiles of sales in that market, all by the same proportion, leaving our moments unchanged. We only need $\overline{X}_{nF}$ to get estimates of σE_{nF} given Θ. Our estimate of Θ does, however, depend on the data for N_{nF}. Appendix C reports on Monte Carlo simulations examining the sensitivity of our parameter estimates to this form of sampling error.[24]

4.4 Estimation Procedure

We base our estimation procedure on the moment condition:

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

where Θ_0 is the true value of Θ. We thus seek a $\widehat{\Theta}$ that achieves:

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

[24] The deviation between the theoretical value and the realization has a coefficient of deviation approximated by $1/\sqrt{N_{nF}}$, which is highest for Nepal, with 43 sellers, at 0.152. The median number of sellers is 686 (Malaysia) implying a coefficient of variation of 0.038.

where $\mathbf{W}$ is a 1360×1360 weighting matrix.[25] We search for $\widehat{\Theta}$ using the simulated annealing algorithm.[26] At each function evaluation involving a new value of Θ we compute a set of 500,000 artificial firms and construct the moments for them as described above. The simulated annealing algorithm converges in 1 to 3 days on a standard PC.

We calculate standard errors using a bootstrap technique, taking into account both sam-

[25]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 destination strings fixed to calculate $m^b(1)$ and the $s_n^q(\tau)$ fixed to calculate $m^b(\tau)$ for $\tau = 2, 3, 4$. (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 $\mathbf{W}$.

[26]Goffe, Ferrier, and Rogers (1994) describe the algorithm. We use a version developed specifically for Gauss and available on the web from William Goffe (Simann).

pling error and simulation error.[27]

4.5 Results

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

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

As a check on our procedure, given our sample size, we conduct a Monte Carlo analysis, described in Appendix C. A basic finding is that the standard errors above are good indicators of the ability of our procedure to recover parameters. We also analyze the sensitivity of our

[27] 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 as described in Section 4.2, thus generating a new $\widehat{m}^b(\Theta)$. (Just like m^b, $\widehat{m}^b$ is calculated according to the bins defined from the actual data.) Defining $y^b(\Theta) = m^b - \widehat{m}^b(\Theta)$, for each b we search for:

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

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. Since we pursue our bootstrapping procedure only to calculate standard errors rather than to perform tests we do not recenter the moments to account for the initial misfit of our model. Recentering would involve setting:

$$y^b(\Theta) = m^b - \widehat{m}^b(\Theta) - (m - \widehat{m}(\widehat{\theta})).$$

above. In fact, our experiments with recentered moments yielded similar estimates of the standard errors. See Horowitz (2001) for an authoritative explanation.

results, as described in Appendix D, to different moments. A basic finding is that the results are largely insensitive to the alternatives we explore.

We turn to some implications of our parameter estimates.

Our discussion in Section 3.6 foreshadowed our estimate of $\widetilde{\theta}$, which lies between the values implied by the slopes in Figures 3c and Figure 4. From equations (32), (30), and (31), the characteristic of a firm determining both entry and sales conditional on entry, is $v^{-1/\widetilde{\theta}}$, where $v \sim U[0,1]$. 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 (21), our estimate of 2.46 implies that fixed costs dissipate 59 percent of gross profit in any destination.

Our estimate of σ_α implies enormous idiosyncratic variation in a firm's 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_η means much less idiosyncratic variation in the entry shock η, with a ratio of the 75th to 25th percentile equal to 1.58. Given σ_α and σ_η, the variance of sales within a market decreases in ρ, as can be seen from equation (39). Hence the negative estimate of ρ reflects high variation of sales in a market.

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.[28] 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

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

to the luck of the draw in individual markets.

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 $\widehat{\Theta}$ to simulate a set of artificial firms including nonexporters.[29] We then compare four features of these artificial firms with corresponding features of the actual ones.

Entry. Since our estimation routine conditions entry hurdles on the actual number of French firms selling in each market, our simulation would hit these numbers were it not for simulation error. The total number of exporters is a different matter since the model determines the extent to which the same firms are selling to multiple countries. We simulate 31,852 exporters, somewhat below the actual number of 34,035. Table 2 displays all the export strings that obey a hierarchy out of the 128 subsets of the 7 most popular export destinations. The first column is the actual number of French firms selling to that string of countries while the last column display the simulated number. In the actual data 27.2 percent of exporters adhere to hierarchies compared with 30.3 percent in the model simulation. In addition the model captures very closely the number selling to each of the seven different strings that obey a hierarchy.

Sales in a Market. Equation (32) motivates Figure 5a, which plots the simulated (x's) and actual (circles) values of the median and 95th percentile sales to each market against

[29] Here we simulate the behavior of $S = 230,000$ artificial firms, both nonexporters and exporters that sell in France, to mimic more closely features of the raw data behind our analysis. Thus in step 5 in the simulation algorithm we reset $\overline{u}(s) = \overline{u}_F(s)$.

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.

Sales in France Conditional on Entry in a Foreign Market. Equation (33) motivates Figure 5b, which plots the median and 95th percentile 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 technology parameterized by λ. The model understates noise in these percentiles in markets served by a small number of French firms.

Export Intensity. Equation (34) motivates Figure 5c, which plots median 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_α, however, the model understates the noisiness of the relationship.

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

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

where V_n^U is the variance of entry and $E\left[V_n^C(u)\right]$ is the expected variance of entry given the firm's standardized unit cost u.[30]

We simulated the term $E\left[V_n^C(u)\right]$ using the 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. Looking across markets, the results indicate that we can attribute on average 57 percent (with a standard deviation of .01) of the variation in entry in a market to the core efficiency of the firm rather than to its draw of η in that market.

[30] By the law of large numbers, the fraction of French firms selling in n approximates the probability that a French firm will sell in n. Writing this probability as $q_n = N_{nF}/N_{FF}$, the unconditional variance of entry for a randomly chosen French firm is $V_n^U = q_n\left(1 - q_n\right)$. Conditional on its standardized unit cost u a firm enters market n if its entry shock η_n satisfies $\eta_n \geq (u\kappa_2/N_{nF})^{1/\widetilde{\theta}}$. Since $\ln\eta_n$ is distributed $N(0, \sigma_\eta)$, the probability that this condition is satisfied is:

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

where Φ is the standard normal cumulative density. The variance conditional on u is therefore $V_n^C(u) = q_n(u)[1 - q_n(u)]$.

4.7.2 Variation in Sales

Looking at the firms that enter a particular market, how much does variation in u explain sales variation? Consider firm j selling in market n. Inserting (30) into (27), the log of sales is:

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

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 n, $\eta_n(j)$, and its overall efficiency shock, $u(j)$. We 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. Looking across markets, we find that component 3 itself contributes only 4.8 percent of the variation in $\ln X_{nF}(j)$ and components 2 and 3 together around 39 percent (with very small standard deviations).[31]

The dominant role of α is consistent with the shapes of the sales distributions in Figure 2. Even though the Pareto term u must eventually dominate the lognormal α at the very upper tail, such large observations are unlikely in a reasonably sized sample. With our parameter values the lognormal term is evident both in the curvature and the slope even among the

[31] In comparison, Munch and Nguyen (2009) find that firm effects drive around 40 percent of the sales variation of Danish sellers across markets.

biggest sellers.[32]

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.[33]

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 translated into sales, through σ. In order to separate $\widetilde{\theta}$ into these components we turn to data on firm productivity.

A common observation is that exporters are more productive (according to various measures) than the average firm.[34] The same is true of our French exporters: 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 will typically both

[32] Our estimate of $\widetilde{\theta}$ implies that the sales distribution asymptotes to a slope of around 0.41, much lower than the values reported in footnote 6 among the top 1 percent of firms selling in a market. Sornette (2000, p.80-82) discusses how a lognormal with a large variance can mimic a Pareto over a very wide range of the upper tail. Hence α continues to play a role even among the largest sellers in our sample.

[33] Our finding that u plays a larger role in entry than in sales conditional on entry is consistent with our higher estimate of σ_α relative to σ_η. A lower value of $\widetilde{\theta}$ (implying more sales heterogeneity attributable to efficiency), for given σ_α and σ_η, would lead us to attribute more variation in both entry and in sales to the firm's underlying efficiency.

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

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 more efficient firms enter tougher markets where sales are lower relative to fixed costs. Determining the net effect of efficiency on productivity requires a quantitative assessment.

We calculate productivity among our simulated firms as value added per unit of factor cost, proceeding as follows.[35] The value added $V_i(j)$ of firm j from country i is its gross production $Y_i(j) = \sum_n X_{ni}(j)$, less spending on intermediates $I_i(j)$:

$$V_i(j) = Y_i(j) - I_i(j).$$

We calculate intermediate spending as:

$$I_i(j) = (1 - \beta)\overline{m}^{-1} Y_i(j) + E_i(j)$$

where β is the share of factor costs in variable costs and $E_i(j) = \sum_n E_{ni}(j)$.[36] Value added per unit of factor cost is:

$$q_i(j) = \frac{V_i(j)}{\beta \overline{m}^{-1} Y_i(j)} = \frac{[\overline{m} - (1 - \beta)] - \overline{m}\,[E_i(j)/Y_i(j)]}{\beta} \tag{41}$$

The only potential source of heterogeneity is in the ratio $E_i(j)/Y_i(j)$.[37]

[35] In our model value added per worker and value added per unit of factor cost are proportional since all producers in a country face the same wage and input costs, with labor having the same share.

[36] We treat all fixed costs as purchased services so exclude factor costs. See EKK (2008) for a more general treatment.

[37] An implication of (41) is that the distribution of productivity of sellers in supplying any given market is invariant to any market-specific feature such as size, location, or trade openness. To see this result replace $E_i(j)/Y_i(j)$ in (41) with $E_{ni}(j)/X_{ni}(j)$. Using (13) for the numerator and (27) for the denominator, exploiting

The expression for firm productivity (41) depends on the elasticity of substitution σ (through $\overline{m} = \sigma/(\sigma - 1)$) independently of $\widetilde{\theta}$. We can thus follow BEJK (2003) and find the σ that makes the productivity advantage of exporters in our simulated data match their productivity advantage in the actual data (1.22). To find the productivity of a firm in the simulated data we take its sales $X_{ni}(j)$ and fixed cost $E_{ni}(j)$, using (13), in each market where it sells and then sum to get $Y_i(j)$ and $E_i(j)$.[38] We then calculate the average $q_i(j)$ for exporters relative to that for all firms. The σ that delivers the same advantage in the simulated data as the actual is $\sigma = 2.98$, implying $\beta = 0.34$. Using our estimate of $\widetilde{\theta} = \theta/(\sigma - 1) = 2.46$ the implied value of θ is 4.87.[39]

(31), we can write this ratio in terms of $v_{ni}(j)$ as:

$$\frac{E_{ni}(j)}{X_{ni}(j)} = \begin{cases} \frac{\lambda}{\sigma(\lambda-1)} \frac{v_{ni}(j)^{(1-\lambda)/\widetilde{\theta}}-1}{v_{ni}(j)^{-\lambda/\widetilde{\theta}}-1} & \lambda \neq 1 \\ \frac{-\ln v_{ni}(j)}{\sigma\widetilde{\theta}\left(v_{ni}(j)^{-1/\widetilde{\theta}}-1\right)} & \lambda = 1 \end{cases}$$

Since $v_{ni}(j)$ is distributed uniformly on $[0,1]$, in any market n, the distribution of the ratio of fixed costs to sales revenue depends only on λ, σ, and $\widetilde{\theta}$, and nothing specific to destination n.

[38] We calibrate β from data on the share of manufacturing value added in gross production. Denoting the value-added share as s^V, averaging across UNIDO gives us $s^V = 0.36$. Taking into account profits and fixed costs we calculate:

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

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

[39] While the model here is different, footnote 9 suggests that the parameter θ plays a similar role here to its role in Eaton and Kortum (2002) and in BEJK (2003). Our estimate of $\theta = 4.87$ falls between the estimates of 8.28 and 3.60, respectively, from those papers. Our estimate of $\sigma = 2.98$ is not far below the estimate of 3.79 from BEJK (2003). Using Chaney (2008) as a basic framework, Crozet and Koenig (2010) estimate θ and σ (along with a parameter δ reflecting the effect of distance D on trade costs) using French firm level data on exports to countries contiguous to France. They employ a three-step procedure, estimating first a probit equation for firm entry, retrieving $\delta\theta$ from the coefficient $\ln D$, second an equation estimating firm sales,

4.9 The Price Index Effect

With values for the individual parameters θ and σ we can return to equation (14) and ask how much better off are buyers in a larger market. Taking E_{ni}'s as given, the elasticity of P_n with respect to X_n is $1/\theta - 1/(\sigma - 1) = -0.30$: Doubling market size leads to a 30 percent lower manufacturing price index.[40]

5 General Equilibrium and Counterfactuals

We now consider how our framework can be used to examine how counterfactual aggregate shocks 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. Factors are 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.

retrieving $\delta(\sigma - 1)$ from the coefficient $\ln D$, third a relationship between a productivity and sales, retrieving $\theta - (\sigma - 1)$]. Performing this procedure for 34 industries they obtain estimates of θ ranging from 1.34 to 4.43 (mean 2.59) and of σ ranging from 1.11 to 3.63 (mean 1.72).

[40] Recall that our analysis in Footnote 15 suggested that entry costs for French firms rise with market size with an elasticity of 0.35, attenuating this effect. Assuming that this elasticity applies to the entry costs E_{ni} for all sources i, a calculation using (14) and (16) still leaves us with an elasticity of $(1 - .35) * (-0.30) = -0.20$.

2. Intermediates are as in Eaton and Kortum (2002). Consistent with Section 4.8, manufacturing inputs are a Cobb-Douglas combination of labor and intermediates, where intermediates have the price index P_i given in (14). Hence an input bundle costs:

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

where β is the labor share.

3. Nonmanufacturing is 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 pay for labor in the destination. We thus decompose the country-specific component of the entry cost $E_{ni} = W_n F_{ni}$, where F_{ni} reflects the labor required for entry for a seller from i in n.[41]

5. Each country i's manufacturing deficit D_i and overall trade deficit D_i^A are held at their 1986 values (relative to world GDP).

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

$$Y_i = \sum_{n=1}^{N} \pi_{ni} X_n. \tag{42}$$

We can turn these equilibrium conditions for manufacturing output into conditions for labor market equilibrium determining wages W_i in each country as follows:

[41] Combined with our treatment of fixed costs as intermediates in our analysis of firm-level productivity, assumption 4 implies that these workers are outsourced manufacturing labor.

As shown in Appendix E, we can solve for Y_i in terms of wages W_i, and exogenous terms as:

$$Y_i = \frac{\gamma\sigma\left(W_i L_i + D_i^A\right) - \sigma D_i}{1 + (\sigma - 1)(\beta - \gamma/\theta)}. \tag{43}$$

Since $X_n = Y_n + D_n$:

$$X_n = \frac{\gamma\sigma\left(W_n L_n + D_n^A\right) - (\sigma - 1)(1 - \beta + \gamma/\theta) D_n}{1 + (\sigma - 1)(\beta - \gamma/\theta)}. \tag{44}$$

From expression (20) we can write:

$$\pi_{ni} = \frac{T_i\left(W_i^{\beta} P_i^{1-\beta} d_{ni}\right)^{-\theta} (\sigma F_{ni})^{-[\theta-(\sigma-1)]/(\sigma-1)}}{\sum_{k=1}^{N} T_k\left(W_k^{\beta} P_k^{1-\beta} d_{nk}\right)^{-\theta} (\sigma F_{nk})^{-[\theta-(\sigma-1)]/(\sigma-1)}}. \tag{45}$$

Substituting (43), (44), and (45) into (42) gives us a set of equations determining wages W_i given prices P_i. From expression (14) we have:

$$P_n = \overline{m}\kappa_1^{-1/\theta}\left[\sum_{i=1}^{N} T_i (W_i^{\beta} P_i^{1-\beta} d_{ni})^{-\theta} (\sigma F_{ni})^{-[\theta-(\sigma-1)]/(\sigma-1)}\right]^{-1/\theta}\left(\frac{X_n}{W_n}\right)^{(1/\theta)-1/(\sigma-1)}. \tag{46}$$

giving us prices P_i given wages W_i. Together the two sets of equations deliver W_i and P_i around the world in terms of exogenous variables.

5.2 Calculating Counterfactual Outcomes

We apply the method used in Dekle, Eaton, and Kortum (henceforth DEK, 2008) to equations (42) and (46) to calculate counterfactuals.[42] Denote the counterfactual value of any variable x as x' and define $\widehat{x} = x'/x$ as its change. Equilibrium in world manufactures in the

[42] DEK (2008) calculated counterfactual equilibria in a perfectly competitive 44-country world. Here we adopt their procedure to accommodate the complications posed by monopolistic competition and firm heterogeneity, expanding coverage to 114 countries (our 113 plus rest of world).

counterfactual requires:

$$Y_i' = \sum_{n=1}^{N} \pi_{ni}' X_n'. \tag{47}$$

We can write each of the components in terms of each country's baseline labor income, $Y_i^L = W_i L_i$, baseline trade shares π_{ni}, baseline deficits, and the change in wages $\widehat{W}_i$ and prices $\widehat{P}_i$ using (43), (44), and (45) as follows:

$$\begin{aligned}
Y_i' &= \frac{\gamma\sigma\left(Y_i^L \widehat{W}_i + D_i^A\right) - \sigma D_i}{1 + (\sigma-1)(\beta - \gamma/\theta)} \\
X_n' &= \frac{\gamma\sigma\left(Y_n^L \widehat{W}_n + D_n^A\right) - (\sigma-1)(1-\beta+\gamma/\theta) D_n}{1 + (\sigma-1)(\beta - \gamma/\theta)} \\
\pi_{ni}' &= \frac{\pi_{ni} \widehat{W}_i^{-\beta\theta} \widehat{P}_i^{-(1-\beta)\theta} \widehat{d}_{ni}^{-\theta} \widehat{F}_{ni}^{-[\theta-(\sigma-1)]/(\sigma-1)}}{\sum_{k=1}^{N} \pi_{nk} \widehat{W}_k^{-\beta\theta} \widehat{P}_k^{-(1-\beta)\theta} \widehat{d}_{nk}^{-\theta} \widehat{F}_{nk}^{-[\theta-(\sigma-1)]/(\sigma-1)}}
\end{aligned}$$

where sticking these three equations into (47) yields a set of equations involving $\widehat{W}_i$'s for given $\widehat{P}_i$'s. From (46) we can get a set of equations involving $\widehat{P}_i$'s for given $\widehat{W}_i$'s:

$$\widehat{P}_n = \left[\sum_{i=1}^{N} \pi_{ni} \widehat{W}_i^{-\beta\theta} \widehat{P}_i^{-(1-\beta)\theta} \widehat{d}_{ni}^{-\theta} \widehat{F}_{ni}^{-[\theta-(\sigma-1)]/(\sigma-1)}\right]^{-1/\theta} \left(\frac{\widehat{X}_n}{\widehat{W}_n}\right)^{(1/\theta)-1/(\sigma-1)}. \tag{48}$$

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 calibrate $\beta = 0.34$ (common across countries) as described in Section 4.8 and Y_i^L and country-specific γ_i's from data on GDP, manufacturing production, and trade. Appendix F describes our sources of data and our procedures for assembling them to execute the counterfactual.

A simple iterative algorithm solves jointly for changes in wages and prices, giving us $\widehat{W}_n$, $\widehat{P}_n$, $\widehat{\pi}_{ni}$, and $\widehat{X}_n$. From these values we calculate: (i) the implied change in French exports in each market n, using the French price index for manufactures as numéraire, as $\widehat{X}_{nF} = \widehat{\pi}_{nF} \widehat{X}_n / \widehat{P}_F$ and (ii) the change in the number of French firms selling there, using (28), as $\widehat{N}_{nF} = \left(\widehat{\pi}_{ni} \widehat{X}_n\right) / \left(\widehat{W}_n \widehat{F}_{nF}\right)$.

We then calculate the implications of this change for individual firms. We hold fixed all of the firm-specific shocks that underlie firm heterogeneity in order to isolate the microeconomic changes brought about by general-equilibrium forces. We produce a dataset, recording both baseline and counterfactual firm-level behavior, as follows:

1. We apply the changes $\widehat{X}_{nF}$ and $\widehat{N}_{nF}$ derived above to our original data set to get counterfactual values of total French sales in each market X_{nF}^C and the number of French sellers there N_{nF}^C.

2. We run the simulation algorithm described in Section 4.2 with $S = 500,000$ and the same stochastic draws applying both to the baseline and to the counterfactual. We set Θ to the parameter estimates reported in Section 4.5. To accommodate counterfactuals, we tweaked our algorithm in four places:

   (a) In step 2, we use our baseline values X_{nF} and N_{nF} to calculate baseline σE_{nF}'s and our counterfactual values X_{nF}^C and N_{nF}^C to calculate counterfactual σE_{nF}^C's in each destination.

   (b) In step 4 we use our baseline values X_{nF} and N_{nF} to calculate baseline $\overline{u}_n(s)$'s and our counterfactual values X_{nF}^C and N_{nF}^C to calculate counterfactual $\overline{u}_n^C(s)$'s for each destination and firm, using (29) and (37).

   (c) In step 5, we set

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

   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.

(d) In step 8 we calculate entry and sales in each of the 113 markets in the baseline and in the counterfactual.

5.3 Counterfactual: Implications of Globalization

We consider a ten percent drop in trade barriers, i.e., $\widehat{d}_{ni} = 1/(1.1)$ for $i \neq n$ and $\widehat{d}_{nn} = \widehat{F}_{ni} = 1$. This change roughly replicates the increase in French import share over the decade following 1986.[43] Table 3 shows the aggregate general equilibrium consequences of this change: (i) the change in the real wage $\widehat{\omega}_n = \left(\widehat{W}_n/\widehat{P}_n\right)^{\gamma_n}$, (ii) the change in the relative wage $\widehat{W}_n/\widehat{W}_F$, (iii) the change in the sales of French firms to each market $\widehat{X}_{nF}$, and (iv) the change in the number of French firms selling to each market $\widehat{N}_{nF}$. Lower trade barriers raise the real wage in every country, typically by less than 5 percent.[44] Relative wages move quite a bit more, capturing terms of trade effects from globalization.[45] The results that matter at the firm level are French sales and the number of French firms active in each market. While French sales

[43] Using time-series data from the OECD's STAN database, we calculated the ratio of manufacturing imports to manufacturing absorption (gross production + imports - exports) for the 16 OECD countries with uninterrupted annual data from 1986-2000. By 1997 this share had risen for all 16 countries, with a minimum increase of 2.4 for Norway and a maximum of 21.1 percentage points for Belgium. France, with a 10.0 and Greece with an 11.0 percentage point increase straddled the median.

[44] There are a several outliers on the upper end, with Belgium experiencing a 9 percent gain, Singapore a 24 percent gain, and Liberia a 49 percent gain. These results are associated with anomalies in the trade data due to entrepot trade or (for Liberia) ships. These anomalies have little consequence for our overall results.

[45] A good predictor of the change in country n's relative wage is its baseline share of exports in manufacturing production, as the terms of trade favor small open economies as trade barriers decline. A regression in logs across the 112 export destinations yields an R^2 of 0.72.

declines by 5 percent in the home market, exports increase substantially, with a maximum 80 percent increase in Japan.[46] The number of French exporters increases roughly in parallel with French exports.[47]

Table 4 summarizes the results, which are dramatic. Total sales by French firms rise by \$16.4 million, the net effect of a \$34.5 million increase in exports and a \$18.1 million decline in domestic sales. Despite this rise in total sales, competition from imports drives almost 27 thousand firms out of business, although almost 11 thousand firms start exporting.

Tables 5 and 6 decompose these changes into the contributions of firms of different baseline size, with Table 5 considering the counts of firms. Nearly half the firms in the bottom decile are wiped out while only the top percentile avoid 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 6 decomposes sales revenues. All of the increase is in the top decile, and most of that

[46] A good predictor of the change in French exports to n comes from log-linearizing π_{nF}, noting that $\ln X_{nF} = \ln \pi_{nF} + \ln X_n$. The variable capturing the change in French cost advantage relative to domestic producers in n is $x_1 = \pi_{nn}\left[\beta \ln(\widehat{W}_n/\widehat{W}_F) - \ln \widehat{d}_{nF}\right]$ with a predicted elasticity of θ (we ignore changes in the relative value of the manufacturing price index, since they are small). The variable capturing the percentage change in n's absorption is $x_2 = \ln \widehat{W}_n$ with a predicted elasticity of 1. A regression across the 112 export destinations of $\ln X_{nF}$ on x_1 and x_2 yields an R^2 of 0.88 with a coefficient on x_1 of 5.66 and on x_2 of 1.30.

[47] We can compare these counterfactual outcomes with the actual change in the number of French sellers in each market between 1986 and 1992. We tend to overstate the increase and to understate heterogeneity across markets. The correlation between the counterfactual change and the change in the data is nevertheless 0.40. Our counterfactual, treating all parameters as given except for the iceberg trade costs, obviously misses much of what happened in those 6 years.

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 5 and 6 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.[48]

In results not shown we decompose the findings according to the number of markets where firms 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 markets served. For firms that initially export to few markets, a substantial share of export growth comes from entering new markets.

Finally, we can also look at what happens in each foreign market. Very little of the increase in exports to each market is due to entry (the extensive margin). We can also 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.[49]

[48] 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.

[49] We can examine another counterfactual, a uniform change $\widehat{F}$ in labor required for entry, analytically. Proportional changes in $\widehat{F}$ cancel out in the expression for π'_{ni}, so that the only action is in equation (48) for the manufacturing price index:

$$\widehat{P} = \widehat{F}^{[\theta-(\sigma-1)]/[\beta\theta(\sigma-1)]}.$$

With $\widehat{F} = 1/1.1$ the number of French firms selling in every market rises by 10 percent and total sales increase by the factor $1/0.919 = 1.088$ in each market.

6 Conclusion

A literature documenting the superior performance of exporters, notably Bernard and Jensen (1995), inspired a new generation of trade theories incorporating producer heterogeneity. These theories, in turn, delivered predictions beyond the facts that motivated them. Accounting for the data presented here, breaking down firms' foreign sales into individual export destinations, confronts these theories with a formidable challenge.

We find that the Melitz (2003) model, with parsimonious modification, succeeds in explaining the basic features of such data along a large number of dimensions: entry into markets, sales distributions conditional on entry, and the link between entry and sales in individual foreign destinations and sales at home. Not only does the framework explain the facts, it provides a link between firm-level and aggregate observations that allows for a general equilibrium examination of the effect of aggregate shocks on individual firms.

Our framework does not, however, constitute a reductionist explanation of what firms do in different markets around the world. In particular, it leaves the vastly different performance of the same firm in different markets as a residual. Our analysis points to the need for further research into accounting for this variation.

References

Alvarez, Fernando and Robert E. Lucas (2007) "A General Equilibrium Analysis of the Eaton-Kortum Model of International Trade," *Journal of Monetary Economics,* 54: 1726-1768.

Arkolakis, Konstantinos (2008) "Market Penetration Costs and the New Consumers Margin in International Trade," unpublished, 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 (1995) "Exporters, Jobs, and Wages in U.S. Manufacturing: 1976-1987," *Brookings Papers on Economic Activity: Microeconomics,* 67-119.

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," *American Economic Review,* 98: 1707-1721.

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.

Crozet, Matthieu and Koenig, Pamina (2010) "Structural Gravity Equations with Intensive and Extensive Margins," *Canadian Journal of Economics,* 43: 41-62.

Dekle, Robert, Jonathan Eaton, and Samuel Kortum (2008) "Global Rebalancing with Gravity," *IMF Staff Papers,* 55: 511-540.

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 No. 10344.*)

Eaton, Jonathan, Kortum, Samuel, and Francis Kramarz (2008) "An Anatomy of International Trade: Evidence from French Firms," *NBER Working Paper No. 14610.*

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 (1999) "Zipf's Law for Cities: An Explanation," *Quarterly Journal of Economics,* 94: 739-767.

Gabaix, Xavier and Ibragimov, Rustam (2010) "Rank-1/2: A Simple Way to Improve the

OLS Estimation of Tail Exponents," forthcoming, *Journal of Business and Economic Statistics.*

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 Rubinstein, Yona (2008) "Estimating Trade Flows: Trading Partners and Trading Volumes," *Quarterly Journal of Economics,* 123: 441-487.

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.

Luttmer, Erzo (2007) "Selection, Growth, and the Size Distribution of Firms," *Quarterly Journal of Economics,* 122: 1103-1144.

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

Munch, Jakob R. and Nguyen, Daniel X. (2009) "Decomposing Firm-Level Sales Variation," *EPRU Working Paper No. 2009-05.*

OECD (1995) *The OECD STAN Database.*

Pedersen, Niels K. (2009), *Essays in International Trade.* Ph.D. dissertation, Northwestern University, Evanston, Illinois.

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

Roberts, Mark and Tybout, James R. (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.

Sornette, Didier (2000) *Critical Phenomena in Natural Sciences.* Berlin: Springer.

UNIDO (1999) *Industrial Statistics Database.*

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

A Constructing the French 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 excluding 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. All values are translated into U.S. dollars at the 1986 exchange rate.

B Comparison with Danish and Uruguayan Firms

Here we present a number of regression of the form

$$\ln \overline{X}_{ni} = \beta_0 + \beta_1 \ln X_n + \beta_2 \ln y_n + \varepsilon_{ni}$$

where $\overline{X}_{ni}$ is the mean sales in n among exporters from i, X_n is n's market size, and y_n is GDP per capita in n (taken from the World Bank's *World Development Indicators*) We have data for i = France in 1986 and 1992, for i = Denmark (from Pedersen, 2009) in 1993, and for i = Uruguay (compiled by Raul Sampognaro) averaged across 1992 and 1993. We weren't able to get y_n for every country so the regressions with that variable included have fewer observations.

	France				Denmark		Uruguay	
	1986		1992		1993		1992/93	
Intercept, $\widehat{\beta}_0$	-5.1	-4.3	-4.6	-4.0	-4.7	-5.3	-5.3	-4.6
	(0.23)	(0.34)	(0.22)	(0.38)	(0.66)	(0.94)	(0.66)	(1.09)
Market size, $\widehat{\beta}_1$	0.35	0.42	0.34	0.41	0.32	0.30	0.32	0.36
	(0.02)	(0.04)	(0.02)	(0.04)	(0.05)	(0.07)	(0.06)	(0.10)
GDP per capita, $\widehat{\beta}_2$		-0.17		-0.14		0.09		-0.16
		(0.05)		(0.07)		(0.12)		(0.19)
Observations	113	100	98	94	39	37	63	61
Adjusted R^2	0.68	0.71	0.72	0.74	0.33	0.33	0.24	0.23

The estimated elasticity with respect to market size, $\widehat{\beta}_1$, is similar across the 3 sources, while the effect of y_n is not robust (significantly different from zero only for France).

By pooling the four samples we can test for a common value of β_1 and we can explore whether there is a significant home-country effect (kicking in when the source and destination are the same, which we observe for France and Denmark, but not for Uruguay). We cannot reject the hypothesis of a common slope coefficient on market size if we allow for source-specific intercepts and a home-country effect. Imposing that restriction on the pooled regression yields:

	Estimate	Standard Error
Intercept, $\widehat{\beta}_0$	-4.30	(0.27)
Source-country effects:		
France (1986)	0	dropped
France (1992)	0.35	(0.09)
Denmark (1992)	0.09	(0.13)
Uruguay (1992-1993)	-0.40	(0.11)
Home-country effect	1.60	(0.38)
Market size, $\widehat{\beta}_1$	0.38	(0.03)
GDP per capita, $\widehat{\beta}_2$	-0.13	(0.04)
Observations	292	
Adjusted R^2	0.59	

Notice that the home effect is significantly positive and with a magnitude greatly exceeding other differences by source.

C Monte Carlo Analysis

We perform a Monte Carlo evaluation of our estimation procedure, examining its performance in general and, more particularly, the sensitivity of our parameter estimates to sampling error in N_{nF}. We use the algorithm described in Section 4.2, with the estimated parameter values reported above along with the actual values of N_{nF}, to simulate 230,423 artificial French firms (the same as the total number in our actual sample).[50] We then proceed from scratch with the artificial dataset as we did with the actual French data, calculating the moments described in Section 4.3 and the number of firms selling in each market. We apply our estimation procedure, exactly as described, to these simulated data, using the new moments and counts of firms, to estimate Θ (using the same weighting matrix $\mathbf{W}$ as in the original estimation). We performed this exercise 25 times.

The table below reports the values used to create the simulated data (the "truth"), the same as our estimate above, the mean estimate across the 25 experiments, and the standard deviation:

	$\widetilde{\theta}$	λ	σ_α	σ_η	ρ
"truth"	2.46	0.91	1.69	0.34	-0.65
mean estimate	2.47	0.94	1.71	0.34	-0.65
standard deviation	0.04	0.19	0.02	0.00	0.04

The apparent bias in the estimates is very small although sampling error in N_{nF} appears to create uncertainty about the value of λ.

[50] Since we want to simulate the universe of firms rather than those exporting and selling in France, step 5 of the prodecure sets $\overline{u}(s) = \max_n \{\overline{u}_n(s)\}$.

D Robustness Checks

The table below shows how our parameter estimates differ when we employ different sets of moments in the estimation procedure. The first row repeats the baseline estimates described in the paper. The second row shows the results of dropping all moments that are based on the 95th percentiles (in the second and third set of moments listed in Section 4.3). The third row drops the 95th percentiles but also adds the 25th percentiles (to the second, third, and fourth set of moments). The fourth row employs the moments constructed as in the baseline but dropping the 56 countries below the median in terms of their popularity to French exporters (for the second, third, and fourth set of moments).

	$\widetilde{\theta}$	λ	σ_α	σ_η	ρ
Baseline moments	2.46	0.91	1.69	0.34	−0.65
Moments based on 50th and 75th percentile	2.57	0.75	1.69	0.34	−0.65
Moments based on 25th, 50th, and 75th percentile	2.46	0.56	1.47	0.36	−0.51
Only most popular destinations	2.57	0.63	1.68	0.34	−0.64

Note that the parameters are generally quite robust to these alternative moments, although λ (the parameter with the largest standard error in our baseline estimates) does vary considerably.

Our analysis pools all manufacturing firms. To what extent do the patterns on which our moments are based vary across different types of firms? To start to answer this question we have classified our firms into consumption or investment goods producers, and replicated Figures 1, 3, and 4 for each group, available on request.[51] The patterns are very similar although some intriguing differences emerge. For producers of consumer goods, for example,

[51] We classify firms with NAP (nomenclature d'activités et de produits) codes 22-34 as investment good producers.

sales in France tend to be smaller, conditioning on number of destinations and popularity of destination. An interesting question is the extent to which our framework can account for these by treating the two types of firms as representing different regions of the distribution of the u's, α's and η's disproportionately, or whether accommodating these differences requires a richer parameterization. We leave this question for future research.

E Deriving the General Equilibrium Equations

To derive (43) we write country i's total absorption of manufactures is the sum of final demand and use as intermediates as:

$$X_i = \gamma(Y_i^A + D_i^A) + [(1-\beta)(\sigma-1)/\sigma]\, Y_i, \tag{49}$$

where Y_i^A is GDP and D_i^A the trade deficit. To relate Y_i^A to wages we write:

$$Y_i^A = W_i L_i + \Pi_i, \tag{50}$$

where Π_i are total net profits earned by country i's manufacturing producers from their sales at home and abroad.

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 total entry costs incurred there, $\overline{E}_n$. Using (21) 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:

$$\Pi_i = \sum_{n=1}^{N} \pi_{ni} \Pi_n^D = \frac{(\sigma-1)}{\sigma\theta} Y_i, \tag{51}$$

where the second equality comes from applying the conditions (42) for equilibrium in the market for manufactures.

Substituting (50) into (49) and using the fact that gross manufacturing production Y_i is gross manufacturing absorption X_i less the manufacturing trade deficit D_i:

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

Solving for Y_i yields expression (43).

F The Data for Counterfactuals

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).[52] 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.[53] In particular:

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

[52] 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.

[53] The value of γ_n lies in the interval $(0.15, 0.55)$ for 100 or our 113 countries. The average is 0.36.

We exploit (50), (51), and (43) in order to derive baseline labor Y_i^L from data on GDP and deficits:

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

All data are for 1986, translated into U.S. dollars at the 1986 exchange rate. See DEK (2008) for further details.

Table 1 - French Firms Exporting to the Seven Most Popular Destinations

Country	Number of Exporters	Fraction of Exporters
Belgium* (BE)	17,699	0.520
Germany (DE)	14,579	0.428
Switzerland (CH	14,173	0.416
Italy (IT)	10,643	0.313
United Kingdom (UK)	9,752	0.287
Netherlands (NL)	8,294	0.244
United States (US)	7,608	0.224
Total Exporters	34,035	

* Belgium includes Luxembourg

Table 2 - French Firms Selling to Strings of Top Seven Countries

	Number of French Exporters		
Export String	Data	Under Independence	Model
BE*	3,988	1,700	4,417
BE-DE	863	1,274	912
BE-DE-CH	579	909	402
BE-DE-CH-IT	330	414	275
BE-DE-CH-IT-UK	313	166	297
BE-DE-CH-IT-UK-NL	781	54	505
BE-DE-CH-IT-UK-NL-US	2,406	15	2,840
Total	9,260	4,532	9,648

* The string "BE" means selling to Belgium but no other among the top 7, "BE-DE" means selling to Belgium and Germany but no other, etc.

Table 3 - Aggregate Outcomes of Counterfactual Experiment (first of two panels)

		Counterfactual Changes (ratio of counterfactual to baseline)			
Country	Code	Real Wage	Relative Wage	Sales of French Firms	Number of French Firms
AFGHANISTAN	AFG	1.01	0.92	1.22	1.23
ALBANIA	ALB	1.01	0.94	1.35	1.34
ALGERIA	ALG	1.00	0.90	1.09	1.12
ANGOLA	ANG	1.00	0.90	1.08	1.10
ARGENTINA	ARG	1.01	0.96	1.57	1.52
AUSTRALIA	AUL	1.02	0.96	1.35	1.29
AUSTRIA	AUT	1.04	1.04	1.49	1.32
BANGLADESH	BAN	1.01	0.95	1.37	1.33
BELGIUM*	BEL	1.09	1.11	1.44	1.19
BENIN	BEN	1.02	0.94	1.12	1.09
BOLIVIA	BOL	1.02	0.94	1.21	1.18
BRAZIL	BRA	1.01	0.96	1.64	1.57
BULGARIA	BUL	1.02	0.95	1.38	1.34
BURKINA FASO	BUK	1.01	0.93	1.17	1.17
BURUNDI	BUR	1.01	0.92	1.21	1.21
CAMEROON	CAM	1.01	0.92	1.19	1.19
CANADA	CAN	1.04	1.05	1.43	1.26
CENTRAL AFRICAN REPUBLIC	CEN	1.02	1.03	1.33	1.20
CHAD	CHA	1.01	0.90	1.07	1.10
CHILE	CHI	1.03	1.02	1.53	1.38
CHINA	CHN	1.01	0.94	1.38	1.36
COLOMBIA	COL	1.01	0.92	1.22	1.23
COSTA RICA	COS	1.02	0.94	1.22	1.20
COTE D'IVOIRE	COT	1.03	0.98	1.36	1.28
CUBA	CUB	1.01	0.93	1.26	1.24
CZECHOSLOVAKIA	CZE	1.03	1.01	1.52	1.38
DENMARK	DEN	1.04	1.06	1.46	1.27
DOMINICAN REPUBLIC	DOM	1.04	0.99	1.36	1.28
ECUADOR	ECU	1.02	0.96	1.33	1.28
EGYPT	EGY	1.02	0.92	1.12	1.12
EL SALVADOR	ELS	1.02	0.93	1.10	1.10
ETHIOPIA	ETH	1.01	0.92	1.08	1.09
FINLAND	FIN	1.03	1.02	1.53	1.38
FRANCE	FRA	1.02	1.00	0.95	0.88
GERMANY, EAST	GEE	1.01	0.96	1.58	1.52
GERMANY, WEST	GER	1.03	1.02	1.60	1.45
GHANA	GHA	1.02	0.99	1.38	1.28
GREECE	GRE	1.02	0.97	1.37	1.30
GUATEMALA	GUA	1.01	0.92	1.16	1.17
HONDURAS	HON	1.02	0.95	1.19	1.16
HONG KONG	HOK	1.14	1.20	1.33	1.02
HUNGARY	HUN	1.05	1.04	1.41	1.25
INDIA	IND	1.01	0.95	1.40	1.37
INDONESIA	INO	1.02	0.96	1.44	1.38
IRAN	IRN	1.01	0.93	1.16	1.16
IRAQ	IRQ	1.04	0.94	1.08	1.06
IRELAND	IRE	1.07	1.09	1.43	1.21
ISRAEL	ISR	1.04	1.01	1.44	1.31
ITALY	ITA	1.02	0.99	1.57	1.46
JAMAICA	JAM	1.05	1.02	1.35	1.22
JAPAN	JAP	1.01	0.98	1.80	1.69
JORDAN	JOR	1.03	0.95	1.16	1.13
KENYA	KEN	1.01	0.93	1.18	1.18
KOREA, SOUTH	KOR	1.04	1.04	1.58	1.40
KUWAIT	KUW	1.02	0.94	1.14	1.12

* Belgium includes Luxembourg

Country	Code	Counterfactual Changes (ratio of counterfactual to baseline) Real Wage	Relative Wage	Sales of French Firms	Number of French Firms
LIBERIA	LIB	1.49	1.03	1.27	1.14
LIBYA	LIY	1.02	0.95	1.10	1.07
MADAGASCAR	MAD	1.01	0.94	1.22	1.20
MALAWI	MAW	1.01	0.92	1.17	1.17
MALAYSIA	MAY	1.07	1.08	1.45	1.23
MALI	MAL	1.02	0.95	1.15	1.12
MAURITANIA	MAU	1.08	1.19	1.36	1.05
MAURITIUS	MAS	1.07	1.05	1.47	1.29
MEXICO	MEX	1.01	0.94	1.32	1.30
MOROCCO	MOR	1.02	0.98	1.37	1.30
MOZAMBIQUE	MOZ	1.01	0.94	1.29	1.26
NEPAL	NEP	1.01	1.01	1.35	1.24
NETHERLANDS	NET	1.06	1.14	1.41	1.14
NEW ZEALAND	NZE	1.03	1.00	1.45	1.33
NICARAGUA	NIC	1.01	0.89	1.06	1.09
NIGER	NIG	1.02	1.09	1.47	1.25
NIGERIA	NIA	1.00	0.89	1.07	1.12
NORWAY	NOR	1.04	1.03	1.38	1.23
OMAN	OMA	1.04	0.99	1.11	1.03
PAKISTAN	PAK	1.02	0.97	1.41	1.34
PANAMA	PAN	1.09	0.96	1.15	1.10
PAPUA NEW GUINEA	PAP	1.07	1.09	1.33	1.13
PARAGUAY	PAR	1.01	0.93	1.21	1.20
PERU	PER	1.02	0.97	1.39	1.32
PHILIPPIPINES	PHI	1.02	0.97	1.50	1.43
PORTUGAL	POR	1.03	1.03	1.47	1.32
ROMANIA	ROM	1.01	0.97	1.68	1.61
RWANDA	RWA	1.00	0.90	1.14	1.17
SAUDI ARABIA	SAU	1.02	0.95	1.15	1.11
SENEGAL	SEN	1.03	1.01	1.36	1.24
SIERRA LEONE	SIE	1.03	1.17	1.36	1.08
SINGAPORE	SIN	1.24	1.15	1.37	1.10
SOMALIA	SOM	1.03	0.96	1.09	1.05
SOUTH AFRICA	SOU	1.03	1.01	1.56	1.43
SPAIN	SPA	1.02	0.97	1.49	1.42
SRI LANKA	SRI	1.03	0.99	1.34	1.24
SUDAN	SUD	1.00	0.91	1.13	1.15
SWEDEN	SWE	1.04	1.05	1.51	1.33
SWITZERLAND	SWI	1.05	1.05	1.48	1.31
SYRIA	SYR	1.02	0.96	1.20	1.15
TAIWAN	TAI	1.04	1.05	1.64	1.44
TANZANIA	TAN	1.01	0.94	1.15	1.13
THAILAND	THA	1.03	0.99	1.50	1.40
TOGO	TOG	1.03	0.96	1.11	1.07
TRINIDAD AND TOBOGO	TRI	1.04	1.01	1.22	1.12
TUNISIA	TUN	1.04	1.00	1.36	1.26
TURKEY	TUR	1.01	0.95	1.37	1.33
UGANDA	UGA	1.00	0.90	1.06	1.08
UNITED KINGDOM	UNK	1.03	1.00	1.46	1.35
UNITED STATES	USA	1.01	0.96	1.45	1.40
URUGUAY	URU	1.02	1.00	1.65	1.53
USSR	USR	1.00	0.92	1.32	1.33
VENEZUELA	VEN	1.01	0.91	1.18	1.20
VIETNAM	VIE	1.01	0.95	1.37	1.33
YUGOSLAVIA	YUG	1.02	0.97	1.48	1.41
ZAIRE	ZAI	1.06	1.21	1.37	1.04
ZAMBIA	ZAM	1.03	1.12	1.49	1.22
ZIMBABWE	ZIM	1.02	0.97	1.43	1.36

Table 4 - Counterfactuals: Firm Totals

		Counterfactual	
	Baseline	Change from Baseline	Percentage Change
Number:			
All Firms	231,402	-26,589	-11.5
Exporting	32,969	10,716	32.5
Values ($ millions):			
Total Sales	436,144	16,442	3.8
Domestic Sales	362,386	-18,093	-5.0
Exports	73,758	34,534	46.8

Counterfactual simulation of a 10% decline in trade costs.

Table 5 - Counterfactuals: Firm Entry and Exit by Initial Size

	All Firms			Exporters		
		Counterfactual			Counterfactual	
Initial Size Interval (percentile)	Baseline # of Firms	Change from Baseline	Change in %	Baseline # of Firms	Change from Baseline	Change in %
not active	0	1,118	---	0	1,118	---
0 to 10	23,140	-11,551	-49.9	767	15	2.0
10 to 20	23,140	-5,702	-24.6	141	78	55.1
20 to 30	23,140	-3,759	-16.2	181	192	106.1
30 to 40	23,140	-2,486	-10.7	357	357	100.0
40 to 50	23,140	-1,704	-7.4	742	614	82.8
50 to 60	23,138	-1,141	-4.9	1,392	904	65.0
60 to 70	23,142	-726	-3.1	2,450	1,343	54.8
70 to 80	23,140	-405	-1.8	4,286	1,829	42.7
80 to 90	23,140	-195	-0.8	7,677	2,290	29.8
90 to 99	20,826	-38	-0.2	12,807	1,915	15.0
99 to 100	2,314	0	0.0	2,169	62	2.8
Totals	231,402	-26,589		32,969	10,716	

Table 6 - Counterfactuals: Firm Growth by Initial Size

	Total Sales			Exports		
		Counterfactual			Counterfactual	
Initial Size Interval (percentile)	Baseline in $millions	Change from Baseline	Change in %	Baseline in $millions	Change from Baseline	Change in %
not active	0	3	---	0	3	---
0 to 10	41	-24	-58.0	1	2	345.4
10 to 20	190	-91	-47.7	1	2	260.3
20 to 30	469	-183	-39.0	1	3	266.7
30 to 40	953	-308	-32.3	2	7	391.9
40 to 50	1,793	-476	-26.6	6	18	307.8
50 to 60	3,299	-712	-21.6	18	48	269.7
60 to 70	6,188	-1,043	-16.9	58	130	223.0
70 to 80	12,548	-1,506	-12.0	206	391	189.5
80 to 90	31,268	-1,951	-6.2	1,085	1,501	138.4
90 to 99	148,676	4,029	2.7	16,080	11,943	74.3
99 to 100	230,718	18,703	8.1	56,301	20,486	36.4
Totals	436,144	16,442		73,758	34,534	

Figure 1: Entry and Sales by Market Size

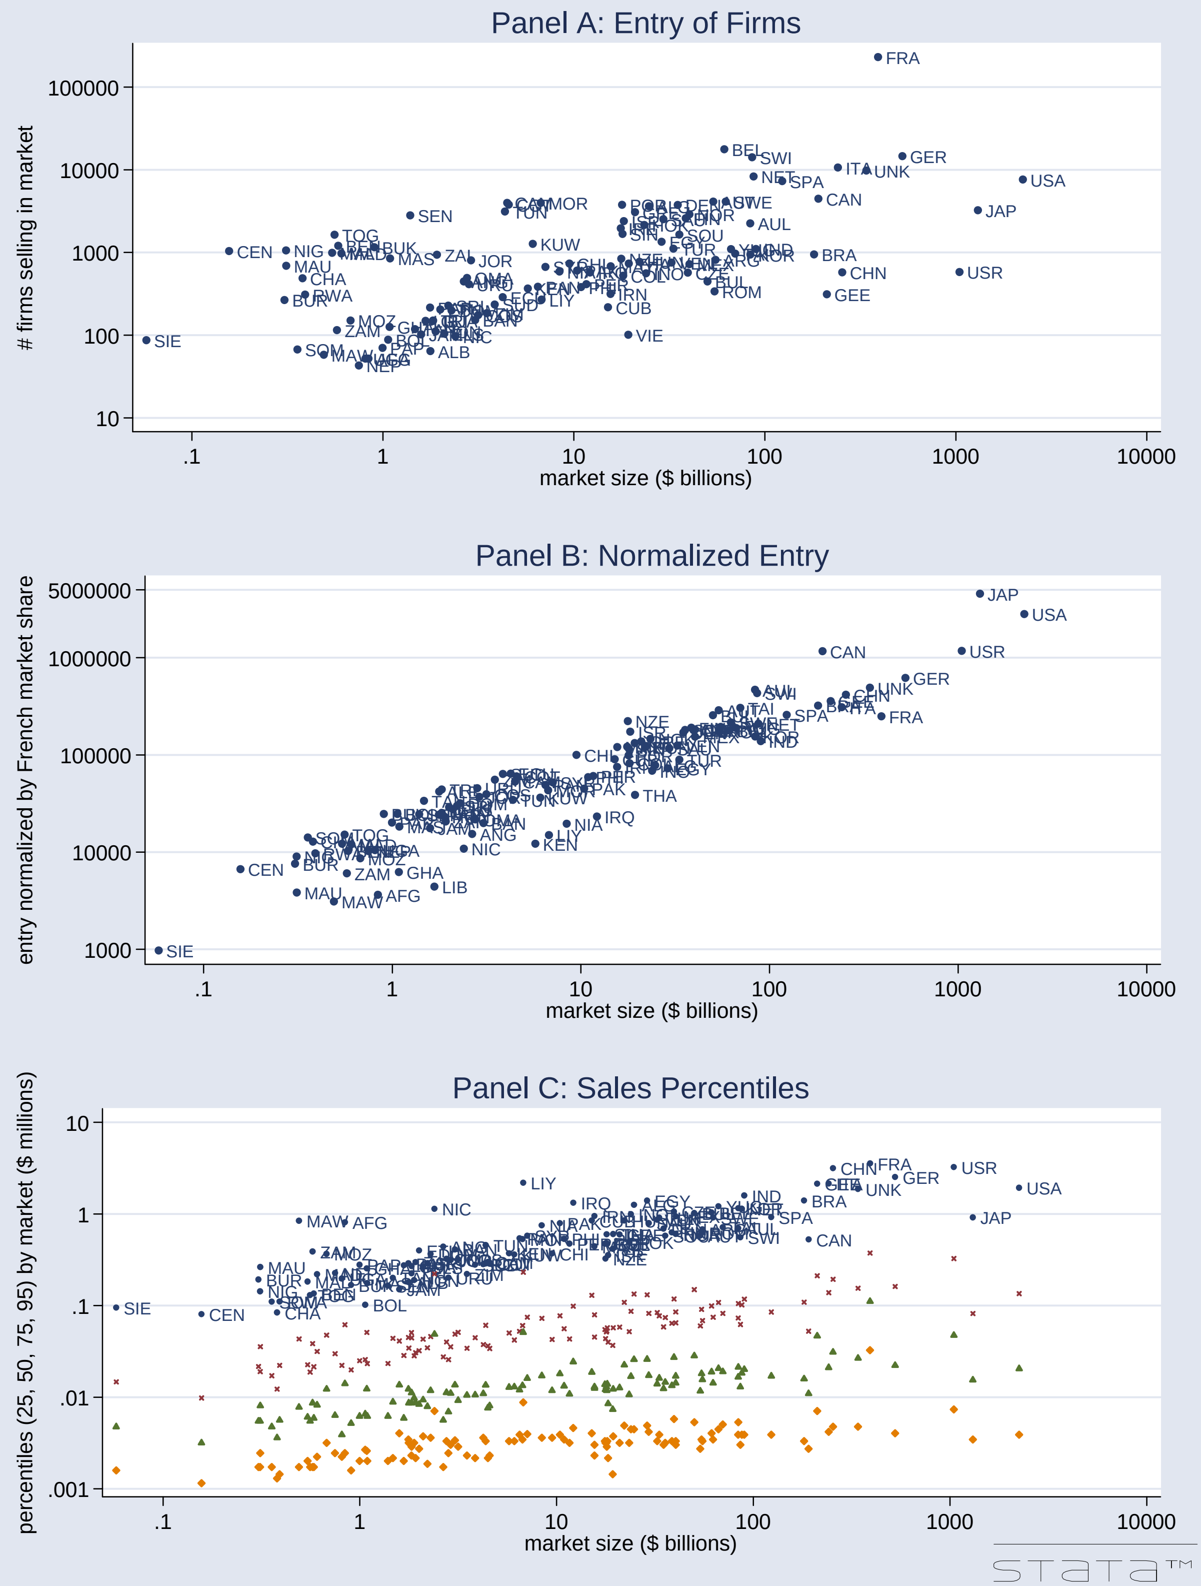

Figure 2
Sales Distributions of French Firms

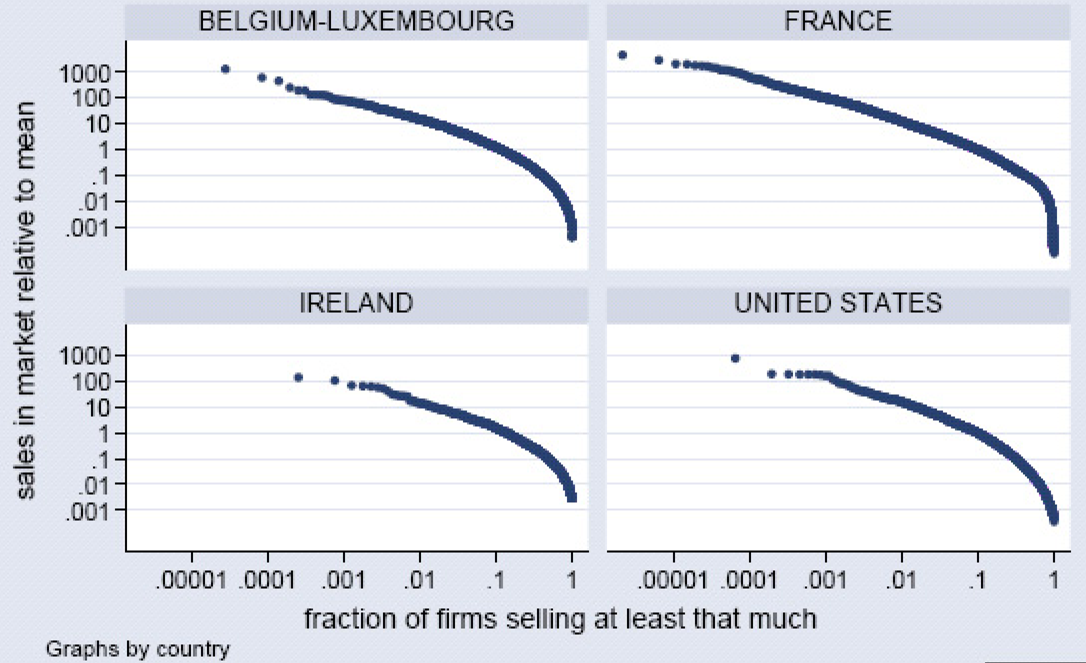

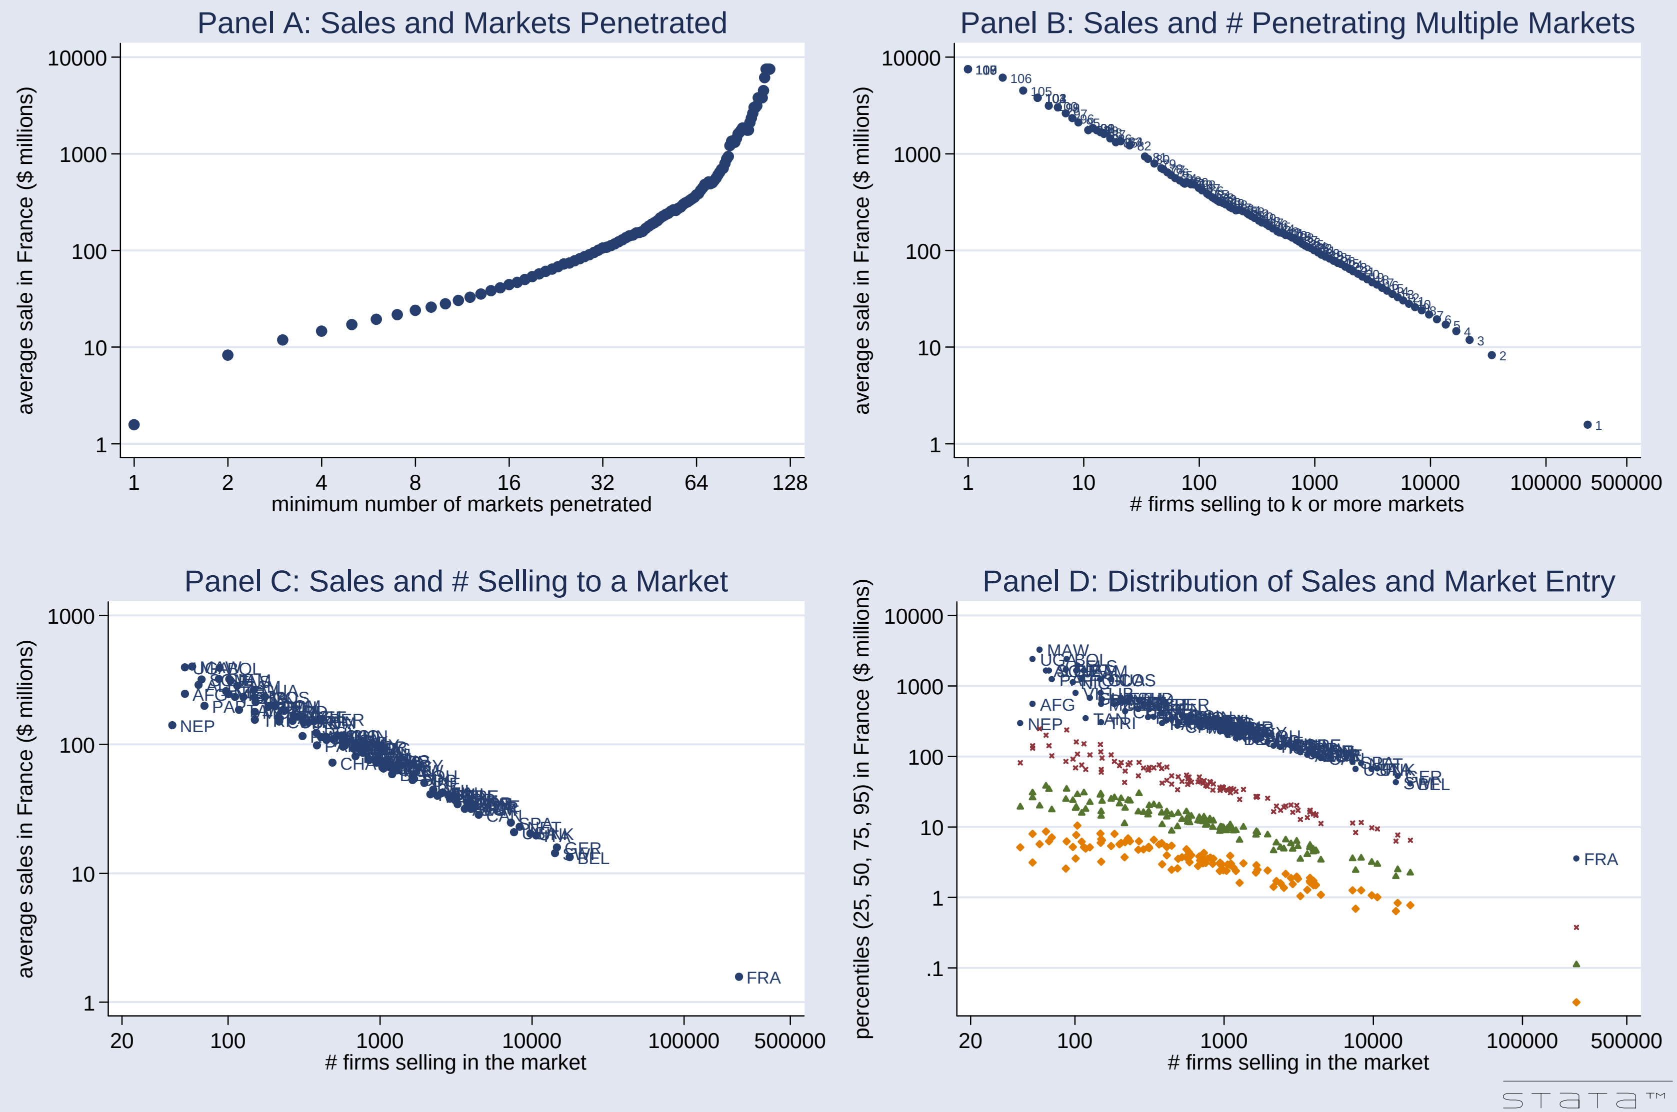

Figure 3: Sales in France and Market Entry

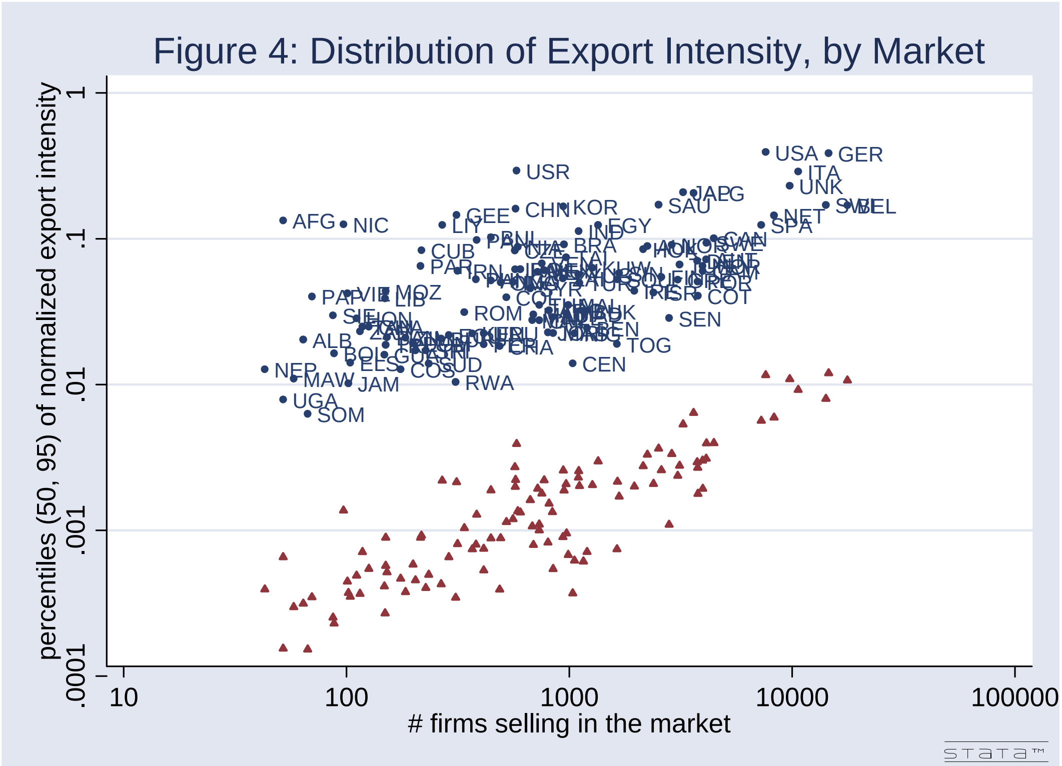

Figure 4: Distribution of Export Intensity, by Market

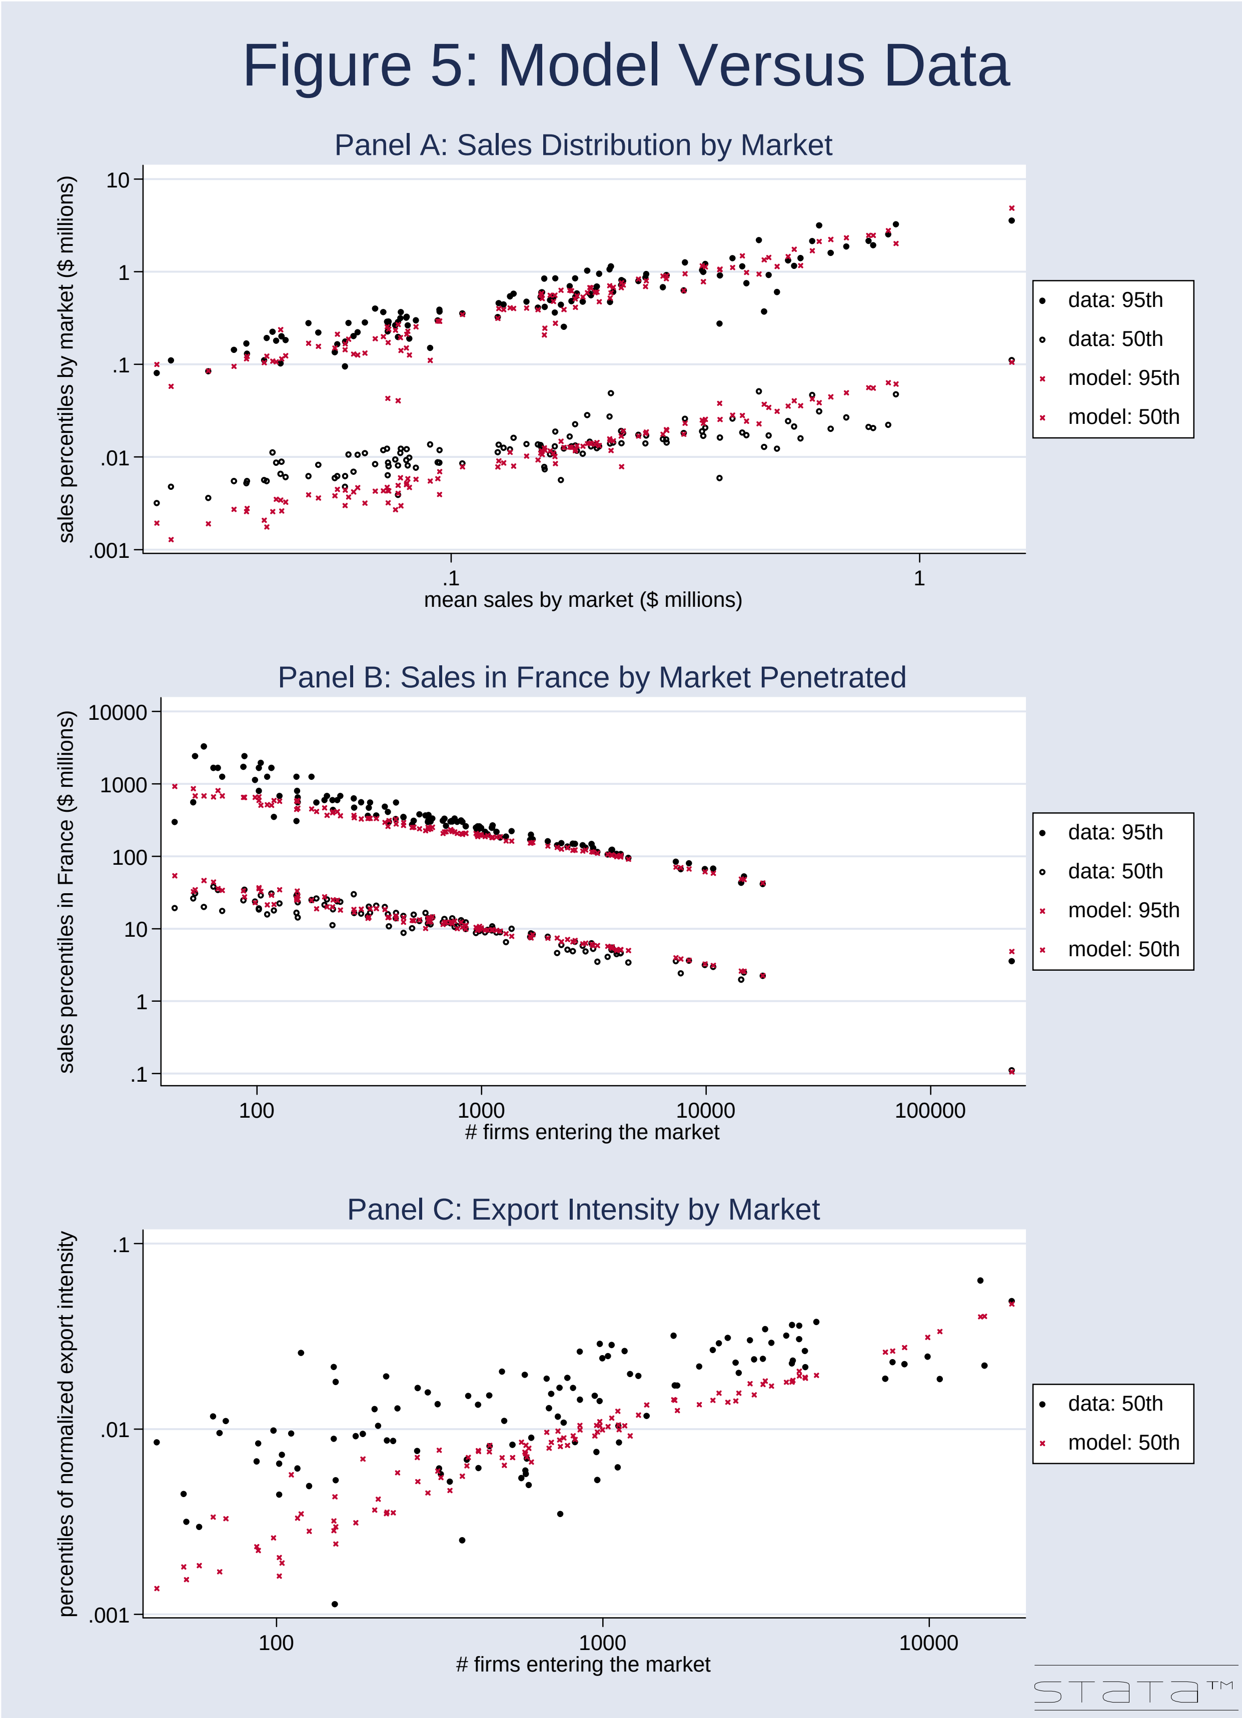
Figure 5: Model Versus Data
Panel A: Sales Distribution by Market
sales percentiles by market ($ millions)
10
1
.1
.01
.001
.1
1
mean sales by market ($ millions)
data: 95th
data: 50th
model: 95th
model: 50th
Panel B: Sales in France by Market Penetrated
sales percentiles in France ($ millions)
10000
1000
100
10
1
.1
100
1000
10000
100000
firms entering the market
data: 95th
data: 50th
model: 95th
model: 50th
Panel C: Export Intensity by Market
percentiles of normalized export intensity
.1
.01
.001
100
1000
10000
firms entering the market
data: 50th
model: 50th
STATA™