

DISCUSSION PAPER SERIES

No. 7111

**AN ANATOMY OF
INTERNATIONAL TRADE:
EVIDENCE FROM FRENCH FIRMS**

Jonathan Eaton, Samuel S Kortum and
Francis Kramarz

***INTERNATIONAL TRADE AND
REGIONAL ECONOMICS***



Centre for Economic Policy Research

www.cepr.org

Available online at:

www.cepr.org/pubs/dps/DP7111.asp

AN ANATOMY OF INTERNATIONAL TRADE: EVIDENCE FROM FRENCH FIRMS

**Jonathan Eaton, New York University
Samuel S Kortum, University of Chicago
Francis Kramarz, Center for Research in Economics and Statistics (CREST)
and CEPR**

Discussion Paper No. 7111
January 2009

Centre for Economic Policy Research
53–56 Gt Sutton St, London EC1V 0DG, UK
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820
Email: cepr@cepr.org, Website: www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INTERNATIONAL TRADE AND REGIONAL ECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Jonathan Eaton, Samuel S Kortum and Francis Kramarz

ABSTRACT

An Anatomy of International Trade:

Evidence from French Firms

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 very similar across markets of very different size and extent of French participation; (3) average sales in France rise very systematically with selling to less popular markets and to more markets. We adopt a model of firm heterogeneity and export participation which we estimate to match moments of the French data using the method of simulated moments. The results imply that nearly half the variation across firms that we see in market entry can be attributed to a single dimension of underlying firm heterogeneity, efficiency. Conditional on entry, underlying efficiency accounts for a much smaller variation in sales in any given market. Parameter estimates imply that fixed costs eat up a little more than half of gross profits. We use our results to simulate the effects of a counterfactual decline in bilateral trade barriers on French firms. 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.

JEL Classification: F12 and L25

Keywords: export, firms, general equilibrium, productivity and trade

Jonathan Eaton
Department of Economics
New York University
269 Mercer Street
New York, NY 10003
USA

Email: jonathan.eaton@nyu.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=100602

Samuel S Kortum
Department of Economics
University of Chicago
1126 East 59th Street
Chicago, IL 60637
USA

Email: kortum@uchicago.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=121874

Francis Kramarz
CREST, Director
15 boulevard Gabriel Péri
92245 Malakoff
France

Email: kramarz@ensae.fr

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=125317

Submitted 18 December 2008

*We thank Costas Arkolakis, Christopher Flinn, Jeremy Fox, Xavier d'Haultfoeuille, and Azeem Shaikh for extremely useful comments. The authors have presented the paper 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.

Executive Summary

We exploit a detailed set of 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. Eaton and Kortum (2002) develop a model relating differences in efficiencies across countries in making different goods to aggregate bilateral trade flows. Since they focus only on aggregate data underlying heterogeneity across individual producers remains hidden. Subsequent papers, 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. Another literature has modeled and estimated the export decision of individual firms in partial equilibrium. 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 exploit detailed data on 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.

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 second 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.

Our main estimation procedure eschews productivity measures since connecting them to our model requires additional assumptions about various parameters of production that are irrelevant to the model's predictions for entry and sales. Looking at one simple measure of productivity, value added per worker, we find that French firms that export are 22 per cent more productive than average. We fit additional parameters to match this figure as well as the share of intermediates in gross production.

With the parameterization complete we can put a number on the "home market effect," the extent to which a larger market benefits from more variety. We find the benefit of bigness substantial: A doubling of demand for manufactures in a market lowers the manufacturing price index by around 20 percent.

We conclude by using our parameterized model to examine two alternative scenarios: a world with lower trade barriers and a world with lower entry costs. To do so we embed our model into a general equilibrium framework. Calibrating the framework to data on production and bilateral trade from our 113 countries and the rest of the world, we can examine the implications of changes in exogenous parameters for income, wages, and prices in each country and for bilateral trade. We can also use these counterfactual outcomes and our parameter estimates to simulate the implications for French firms. A striking finding is that lower trade barriers, while raising welfare in every country, result in substantially more inequality in the distribution of firm size. Even though total output of French firms rises by 3.8 percent, all of the growth is 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.

1 Introduction

We exploit a detailed set of 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. Eaton and Kortum (2002) develop a model relating differences in efficiencies across countries in making different goods to aggregate bilateral trade flows. Since they focus only on aggregate data underlying heterogeneity across individual producers remains hidden. Subsequent papers, 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 exploit detailed data on 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.

¹For example, Aw, Chung, and Roberts (1998), for Taiwan and Korea, and Bernard and Jensen (1999), for the United States, document the size and productivity advantage of exporters.

²A 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 second 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.

Our main estimation procedure eschews productivity measures since connecting them to our model requires additional assumptions about various parameters of production that are irrelevant to the model's predictions for entry and sales. Looking at one simple measure of productivity, value added per worker, we find that French firms that export are 22 per cent more productive than average. We fit additional parameters to match this figure as well as the share of intermediates in gross production.

With the parameterization complete we can put a number on the “home market effect,” the extent to which a larger market benefits from more variety. We find the benefit of bigness substantial: A doubling of demand for manufactures in a market lowers the manufacturing price index by around 20 percent.

We conclude by using our parameterized model to examine two alternative scenarios: a world with lower trade barriers and a world with lower entry costs. To do so we embed our model into a general equilibrium framework. Calibrating the framework to data on production and bilateral trade from our 113 countries and the rest of the world, we can examine the implications of changes in exogenous parameters for income, wages, and prices in each country and for bilateral trade. We can also use these counterfactual outcomes and our parameter estimates to simulate the implications for French firms. A striking finding is that lower trade barriers, while raising welfare in every country, result in substantially more inequality in the distribution of firm size. Even though total output of French firms rises by 3.8 percent, all of the growth is 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 five empirical regularities. With these in mind in Section 3 we turn to a model of exporting by heterogeneous firms. Section 4 explains how we estimate the parameters of the model while section 5 explores the implications of lowering entry and trade costs.

2 Empirical Regularities

Our data are the sales, translated into U.S. dollars, of 229,900 French manufacturing firms to 113 markets in 1986. (Table 3 lists the countries.) Among them only 34,035 sell elsewhere than in France. The firm that exports most widely sells to 110 out of the 113 destinations.³

We assemble our complex data in four different ways, each revealing sharp regularities:

³Appendix 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. We have dropped from our analysis 523 firms whose total exports declared to French customs exceed their total sales from mandatory reports to the French fiscal administration. These 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. In our estimation procedure, we will interpret these 523 French firms as exporters that did not enter the domestic market.

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 clearly to increase with the size of the market, the relationship is a cloudy one. Note in particular that more French firms sell to France than its market size would suggest.

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

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

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

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

Models of perfect and Bertrand competition and the standard model of monopolistic com-

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

petition without market-specific entry costs predict that the number of sellers in a market is invariant to market size. Figures 1a and 1b compel us to abandon these approaches.

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, indicating the tendency of sales per firm to increase with market size.

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 k th most popular destination as well. Not surprisingly firms are less orderly in their choice of destinations. A good metric of how far they depart from a hierarchy is elusive. We can get some sense, however, by looking simply at exporters to the top seven foreign destinations. Table 1 reports these destinations and the number of firms selling to each, 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. An upper bound on the fraction of firms that obey a hierarchy is the 52 percent of exporters who do sell in Belgium.

Table 2 lists each of the strings of 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). Note, for example, that 52 percent of the firms that sell in Belgium adhere to a hierarchy. 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. Independence implies that only 13 percent of exporters would obey

a hierarchy (e.g., selling to the UK only when also selling to both Belgium and Germany), so that 87 percent would deviate from it (e.g., by selling to the UK but not also selling to both Belgium and Germany). In fact more than twice that number, 27 percent, adhere to the hierarchy. Note that many more sell to the shortest string (only Belgium) and to the three longest strings than independence would imply. We conclude that a model needs to recognize both a tendency for firms to export according to a hierarchy while allowing them significant latitude to depart from it.

2.2 Sales Distributions

Our second exercise is to look at the distribution of sales within individual markets. We plot the sales of each firm in a particular market (relative to mean sales there) against the fraction of firms selling in the market who sell at least that much.⁵ 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. Since there are many fewer firms exporting than selling in France the upper percentiles in the foreign destinations are empty. Nonetheless, the shape is about the same.

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

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

where x_n is sales of a firm in market n relative to the mean. Suppose the sales distribution is

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

Pareto with parameter $a > 1$ (so that the minimum sales relative to the mean is $(a - 1)/a$).

We could then write:

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

or:

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

implying a straight line with slope $-1/a$. At the top percentiles the slope does appear nearly constant, and below -1 , but at the lower tails it is much steeper, reflecting the presence of suppliers selling very small amounts.⁶ This shape is well known in the industrial organization literature looking at various size measures.⁷ What we find here is that this shape is inherited across markets looking at the same set of potential sellers.

Figure 1c provides a different picture of how the distribution of sales is very similar across markets. The lines for the four percentiles are roughly parallel on a log scale (although the 25th percentile is decidedly flatter) indicating the common shape of sales distribution across destinations.

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.

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.

⁶Considering 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).

⁷See Simon and Bonini (1958) and Luttmer (2007), among many, for a discussion and explanations.

Figure 3a depicts average sales in France on the y axis for the group of firms that sell to at least k markets with k on the x axis. Note the near monotonicity with which sales in France rise with the number of foreign markets served.

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

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

We conclude that firms 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). A challenge for modeling is reconciling the stark linear (in logs) relationships in Figures 3b, 3c, and 3d with the more nuanced size distributions in Figure 2.⁸

⁸We were able to observe the relationship between market popularity and sales in France for the 1992

2.4 Export Intensity

Having looked separately at what exporters sell abroad and what they sell in the French market, we now examine the ratio of the two. We introduce the concept of 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.

First, as a destination becomes more popular, normalized export intensity rises. The slope for the median is 0.38. Hence, if the number of sellers to a market rises by 10 percent, normalized export intensity rises by around 4 percent, but the relationship is a noisy one.

Second, while we have excluded France from the figure, its y coordinate would be 1. Note that the y coordinates in the figure are at least an order of magnitude below one. An additional challenge for the model is to account for the tendency of exporters to sell so much at home.

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

3 Theory

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

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

3.1 Producer Heterogeneity

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

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

The measure of potential producers in country i who can produce their good with efficiency

at least z is:

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

where $\theta > 0$ is a parameter.⁹ Using (1), the measure of goods that can be delivered from country i to country n at unit cost below c is $\mu_{ni}(c)$ defined as:

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

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

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

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

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

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

⁹We follow Helpman, Melitz, and Yeaple (2004) and Chaney (2007) in treating the underlying heterogeneity in efficiency as Pareto. Our observations above on patterns of sales by French firms in different markets are very suggestive of an underlying Pareto distribution. A Pareto distribution of efficiencies can arise naturally from a dynamic process that is a history of independent shocks, as shown by Simon (1956), Gabaix (1999), and Luttmer (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: $\mu_i^Z(z; J) = J \{1 - \exp[-(T/J)z^{-\theta}]\}$ (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 $J \rightarrow \infty$ gives (2). (To take the limit rewrite the expression as $\{1 - \exp[-(T/J)z^{-\theta}]\} / J^{-1}$ and apply L'Hôpital's rule.)

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

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

3.2 Demand, Market Structure, and Entry

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

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

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

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

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

where the parameter $\lambda \geq 0$ reflects the increasing cost of reaching a larger fraction of potential buyers.¹⁰ This function has the desirable properties that the cost of reaching 0 buyers in a market is 0 and that the total cost is increasing (and the marginal cost weakly increasing) in the fraction f of buyers reached. Taking the limit $\lambda \rightarrow \infty$ implies a constant marginal cost of reaching an added buyer. Since buyers in a market turn out to be identical a seller would then

¹⁰By L'Hôpital's Rule, at $\lambda = 1$ the function becomes $M(f) = -\ln(1 - f)$.

choose to reach either every potential buyer in a market or none at all, an outcome equivalent to Melitz (2003). With λ finite a seller might still stay out of a market entirely or make the effort to reach only a small fraction of buyers there. A seller with a lower unit cost in a market will optimally undertake greater effort to reach more buyers there.

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

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

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

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

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

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

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

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

where:

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

It will seek a fraction:

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

of buyers in the market where:

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

is the entry shock in market n given by the ratio of the demand shock to the fixed-cost shock.

Note that it won't sell at all, hence avoiding any fixed cost there, if:

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

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:

$$c \leq \bar{c}_n(\eta) \quad (8)$$

where:

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

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

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

Its total sales are then:

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

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

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

some of which is covering its fixed cost:

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

To summarize, the relevant characteristics of market n that apply across sellers are total purchases X_n , the price index P_n , and the common component of the fixed cost E_n . The particular situation of a potential seller of product j in market n is captured by three magnitudes: the unit cost $c_n(j)$ and the demand and entry shocks $\alpha_n(j)$ and $\eta_n(j)$. We treat $\alpha_n(j)$ and $\eta_n(j)$ as the realizations of producer-specific shocks drawn from a joint density $g(\alpha, \eta)$ that is the same across destinations n and independent of $c_n(j)$.¹¹

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

3.3 The Price Index

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

¹¹We require that $E[\eta^{\theta/(\sigma-1)}]$ and $E[\alpha\eta^{\theta/(\sigma-1)-1}]$ are both finite. For any given $g(\alpha, \eta)$, these restrictions will imply an upper bound on the parameter θ .

write the price index P_n faced by a representative buyer in market n as:

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

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

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

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

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

where:

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

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

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

3.4 Firm Entry and Sales

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

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

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

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

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

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

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

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

3.5 Aggregate Entry and Sales

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

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

where:

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

Note that this measure rises in proportion to X_n .¹²

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

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

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

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

$$\pi_{ni} = \frac{X_{ni}}{X_n}, \quad (20)$$

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

¹²In describing the data we used N to indicate a number of firms (an integer, of course). Since the theory implies a continuum of firms we use J to denote a measure of producers.

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

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

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

$$\bar{X}_{ni} = \frac{\pi_{ni} X_n}{\pi_{ni} J_n} = \frac{\kappa_1}{\kappa_2} \sigma E_n. \quad (22)$$

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

3.6 Fixed Costs and Profits

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

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

$$E_n(\alpha, \eta, c) = \frac{\alpha}{\eta} E_n \frac{1 - (c/\bar{c}_n(\eta))^{(\lambda-1)(\sigma-1)}}{1 - 1/\lambda}. \quad (23)$$

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

using (3) and (16), gives us:

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

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

$$\begin{aligned} \bar{E}_n &= \int \int E_n(\alpha, \eta) g(\alpha, \eta) d\alpha d\eta \\ &= \frac{X_n}{\sigma \theta} [\theta - (\sigma - 1)]. \end{aligned} \tag{24}$$

Thus total entry costs are a fraction $[\theta - (\sigma - 1)]/\theta$ of the gross profits X_n/σ earned in any destination n . Net profits earned in market n are simply $X_n/(\bar{m}\theta)$. Note that \bar{E}_n (spending on fixed costs) does not depend on E_n , the country component of the fixed cost per firm, just as in standard models of monopolistic competition. A drop in E_n leads to more entry and hence the same total spending on fixed costs.

3.7 A Streamlined Representation

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

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

$$u(j) = T_F z_F(j)^{-\theta}. \tag{25}$$

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

Substituting (25) into (1) and using (4), we can write unit cost in market n in terms of $u(j)$ as:

$$c_n(j) = \frac{w_F d_{nF}}{z_F(j)} = \left(\frac{u(j)}{\pi_{nF}} \right)^{1/\theta} \Phi_n^{-1/\theta}. \quad (26)$$

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

$$\bar{c}_n(\eta) = \left(\frac{\bar{u}_n(\eta)}{\pi_{nF}} \right)^{1/\theta} \Phi_n^{-1/\theta}. \quad (27)$$

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

$$u(j) \leq \bar{u}_n(\eta_n(j)) = \left(\frac{X_{nF}}{\kappa_1 \sigma E_n} \right) \eta_n(j)^{\tilde{\theta}} \quad (28)$$

where

$$\tilde{\theta} = \frac{\theta}{\sigma - 1} > 1. \quad (29)$$

Conditional on firm j 's passing this hurdle we can use (26) and (27) to rewrite firm j 's sales in market n , expression (17), in terms of $u(j)$ as:

$$X_{nF}(j) = \varepsilon_n(j) \left[1 - \left(\frac{u(j)}{\bar{u}_n(\eta_n(j))} \right)^{\lambda/\tilde{\theta}} \right] \left(\frac{u(j)}{\bar{u}_n(\eta_n(j))} \right)^{-1/\tilde{\theta}} \sigma E_n \quad (30)$$

Equations (28) and (30) reformulate the entry and sales equations (16) and (17) in terms of $u(j)$ rather than $c_n(j)$.

Since standardized unit cost $u(j)$ applies across all markets, it gets to the core of a firm's underlying efficiency as it applies to its entry and sales in different markets. Notice that in

reformulating the model as (28) and (30), the two parameters θ and σ enter only collectively through the parameter $\tilde{\theta}$. 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 $\tilde{\theta}$.

3.8 Connecting the Model to the Empirical Regularities

We now show how the model can deliver the features of the data about entry and sales described in Section 2. We quantify the measure of French firms J_{nF} selling in each destination with the actual (integer) number N_{nF} and their average sales their with $\bar{X}_{nF} = X_{nF}/N_{nF}$.

3.8.1 Entry

From (21) we get:

$$\frac{N_{nF}}{\pi_{nF}} = \frac{\kappa_2}{\kappa_1} \frac{X_n}{\sigma E_n}, \quad (31)$$

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

We don't impose any such relationship, but rather employ (31) to calculate:

$$\sigma E_n = \frac{\kappa_2}{\kappa_1} \frac{X_{nF}}{N_{nF}} = \frac{\kappa_2}{\kappa_1} \bar{X}_{nF} \quad (32)$$

directly from the data.

We can use equation (32) to examine how fixed costs vary with country characteristics. Regressing \bar{X}_{nF} against our market size measure (both in logs) yields a slope of 0.31 (with a standard error of 0.02). This relationship relates to the slope in Figure 1b, showing that the number of entrants rises with market size with an elasticity of 0.65. Larger markets attract more firms, but not in proportion, since the cost of entry rises as well. The firms that do enter sell more, generating an elasticity of total sales with respect to market size of 0.96 (close to the gravity prediction that imports vary with market size with a unit elasticity).¹³

Using (32) we can write (28) in terms of observables as:

$$u(j) \leq \bar{u}_n(\eta_n(j)) = \frac{N_{nF}}{\kappa_2} \eta_n(j)^{\tilde{\theta}}. \quad (33)$$

Without variation in the firm and market specific entry shock $\eta_n(j)$, (33) would imply efficiency is all that would matter for entry, dictating a deterministic ranking of destinations with a less efficient firm (with a higher $u(j)$) selling to a subset of the destinations served by any more efficient firm. Hence deviations from market hierarchies identify variation in $\eta_n(j)$. As Table 2 illustrates, there is some tendency for firms to enter markets according to a hierarchy, but it is a loose one.

¹³If we include the log of 1986 real GDP per capita (from the World Bank's *World Development Indicators*) as an additional right-hand side variable, the coefficient on the log of market size rises to 0.41 (standard error 0.03) while the coefficient on the log of real GDP per capita is -0.29 (standard error 0.06). Hence while larger markets have a higher fixed cost of entry, given size the cost is lower in richer countries. We were not able to obtain 1986 real GDP per capita for 10 of our 113 destinations (Albania, Angola, Bulgaria, Czechoslovakia, East Germany, Libya, USSR, Vietnam, Yugoslavia, and Zaire). Hence this second regression was performed on the remaining 103 countries. (The coefficient on the log of market size in the univariate regression for this smaller group is 0.30, hardly different from that for the larger group.)

3.8.2 Sales in a Market

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

$$v_n(j) = \frac{u(j)}{\bar{u}_n(\eta_n(j))} \quad (34)$$

is distributed uniformly on $[0, 1]$. Replacing $u(j)$ with $v_n(j)$ in expression (30) and exploiting (32) we can write sales as:

$$X_{nF}(j) = \varepsilon_n(j) \left[1 - v_n(j)^{\lambda/\tilde{\theta}} \right] v_n(j)^{-1/\tilde{\theta}} \frac{\kappa_2}{\kappa_1} \bar{X}_{nF}. \quad (35)$$

Not only does v_n have the same distribution in each market n , so does ε_n .¹⁴ Hence the distribution of sales in any market n is identical up to a scaling factor equal to \bar{X}_{nF} (reflecting variation in σE_n). Hence we can generate the common shapes of sales distributions exhibited in Figure 2. The variation introduced by ε_n explains why the sales distribution in a market might inherit the lognormal characteristics apparent in that Figure. A further source of curvature is the term in square brackets, representing the fraction of buyers reached. As $v_n(j)$ goes to one, with finite λ , the fraction approaches zero, capturing the curvature of sales distributions at the lower end, as observed in Figure 2. Finally the term $v_n(j)^{-1/\tilde{\theta}}$ instills Pareto features into the distribution. These features will be more pronounced as $v_n(j)$ approaches zero since

¹⁴To 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 :

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

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

very efficient firms will be reaching almost all buyers.

3.8.3 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 (35) as it applies to France and use (34) and (33) to replace $v_F(j)$ with $v_n(j)$:

$$X_{FF}(j)|_n = \frac{\alpha_F(j)}{\eta_n(j)} \left[1 - v_n(j)^{\lambda/\tilde{\theta}} \left(\frac{N_{nF}}{N_{FF}} \right)^{\lambda/\tilde{\theta}} \left(\frac{\eta_n(j)}{\eta_F(j)} \right)^\lambda \right] v_n(j)^{-1/\tilde{\theta}} \left(\frac{N_{nF}}{N_{FF}} \right)^{-1/\tilde{\theta}} \frac{\kappa_2}{\kappa_1} \bar{X}_{FF}. \quad (36)$$

By expressing (36) in terms of $v_n(j)$ we can exploit the fact that (given entry in n) $v_n(j)$ is uniformly distributed on the unit interval. Note that since all the other terms on the right-hand of (36) have the same value or distribution across markets, N_{nF} is the only systematic source of variation across n , entering in two places both times as a ratio to N_{FF} .

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 $-1/\tilde{\theta}$. Interpreting Figure 3c in terms of Equation (35), the slope of -0.57 implies a $\tilde{\theta}$ of 1.75.

Expression (36) 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 α_F in France together with variation in the entry shock $\eta_n(j)$ in country n .

Together (35) and (36) 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 cutoff unit cost 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 market. These equations also explains why France itself is somewhat below the trend line in Figures 3a and 3b. Once we include nonexporters we suddenly have many firms reaching a small fraction of the market.

3.8.4 Normalized Export Intensity

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

$$\frac{X_{nF}(j)}{X_{FF}(j)} / \left(\frac{\bar{X}_{nF}}{\bar{X}_{FF}} \right) = \frac{\alpha_n(j)}{\alpha_F(j)} \left[\frac{1 - v_n(j)^{\lambda/\tilde{\theta}}}{1 - v_n(j)^{\lambda/\tilde{\theta}} \left(\frac{N_{nF}}{N_{FF}} \right)^{\lambda/\tilde{\theta}} \left(\frac{\eta_n(j)}{\eta_F(j)} \right)^\lambda} \right] \left(\frac{N_{nF}}{N_{FF}} \right)^{1/\tilde{\theta}}. \quad (37)$$

Figure 4 plots the median and 95th percentile of this statistic for French firms across export markets. Note first how the presence of the sales shock $\alpha_n(j)$ accommodates random variation in sales in different markets conditional upon entry.

As in (36), the only systematic source of cross-country variation on the right-hand side is in the number of French firms. In contrast to (35) and (36), however, the firm's overall efficiency $v_n(j)$ has no direct effect on normalized export intensity since it cancels (having the same effect in n as it has in France). For $n = F$ the relationship collapses to an identity $1 = 1$. As discussed above, for $n \neq F$ $N_{nF} \ll N_{FF}$, implying that the term in square brackets is much less than one. Hence our model explains the low numbers on the y axis of Figure 4.

Aside from this general collapse of sales to any export market n relative to sales in France, the last term in Equation (37) predicts that normalized export intensity will increase with the number of French firms selling there, the relationship portrayed in Figure 4. The reason is that the harder the market is to enter (i.e., the lower N_{nF}/N_{FF}), the lower is unit cost

in France, but the distribution in foreign destination n is the same. According to (37) the elasticity of normalized export intensity with respect to N_{nF}/N_{FF} is $1/\tilde{\theta}$. The slope coefficient of 0.38 reported in Section 2.5 suggests a value of $\tilde{\theta}$ of 2.63.¹⁵

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

4 Estimation

We estimate the parameters of the model by the method of simulated moments. We begin by completing our parameterization of the model. We then explain how we simulate a set of artificial French exporters given a particular set of parameter values, with each firm assigned a cost draw u and an α and η in each market. We then describe how we calculate a set of moments from these artificial data to compare with moments from the actual data. We then 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_a^2 , σ_h^2 , and correlation ρ . Under

¹⁵In relating equation (36) to Figures 3c and 3d and equation (37) to Figure 4 a slight discrepancy arises. The theory admits the possibility that a firm might not sell in France but could still sell elsewhere. Hence $X_{FF}(j)$ in these equations actually applies to what firm j ’s sales in France would be if it did enter. Obviously, we don’t have data on latent sales in France for firms that don’t sell there. But because the French share in France is so much larger than the French share elsewhere, a French firm selling in another market but not in France is very unlikely. Hence the discrepancy is minor.

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

$$\kappa_1 = \left[\frac{\tilde{\theta}}{\tilde{\theta} - 1} - \frac{\tilde{\theta}}{\tilde{\theta} + \lambda - 1} \right] \exp \left\{ \frac{\sigma_a^2 + 2\rho\sigma_a\sigma_h(\tilde{\theta} - 1) + \sigma_h^2(\tilde{\theta} - 1)^2}{2} \right\} \quad (38)$$

and:

$$\kappa_2 = \exp \left\{ \frac{(\tilde{\theta}\sigma_h)^2}{2} \right\}. \quad (39)$$

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

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

which is decreasing in ρ .

Our estimation conditions on the actual data for: (i) French sales in each of our 113 destinations, X_{nF} , and (ii) the number of French firms selling there, N_{nF} .¹⁶ With this conditioning our model has only five parameters

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

For a given Θ we use (32) to back out the cluster of parameters σE_n using our data on $\bar{X}_{nF} = X_{nF}/N_{nF}$ and the κ_1 and κ_2 implied by (38) and (39). Similarly, we use (33) to back out a firm's entry hurdle in each market $\bar{u}_n(\eta_n)$ given its η_n and the κ_2 implied by (39).

4.2 Simulation

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

¹⁶The 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 who don't enter the domestic market. See footnote 3 above.

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 realizations of $v(s)$'s independently from the uniform distribution $U[0, 1]$, for $s = 1, \dots, S$, 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:

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

putting them aside to construct the $\alpha_n(s)$ and $\eta_n(s)$ below.

Here we describe how to simulate a set of exporting firms that sell in France, as these are the firms whose moments we seek to match in our estimation. As we explain below, it is easy to modify the routine to include nonexporters and firms not selling in France, as we do in drawing out implications of the model and in performing counterfactual experiments.

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 seven steps:

1. Using (38) and (39) we calculate κ_1 and κ_2 .
2. Using (32) we calculate σE_n 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

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

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

$$\bar{u}_n(s) = \frac{N_{nF}}{\kappa_2} \eta_n(s)^{\tilde{\theta}}, \quad (40)$$

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

5. We calculate

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

the maximum u consistent with exporting somewhere, and

$$\bar{u}(s) = \min\{\bar{u}_F(s), \bar{u}^X(s)\}, \quad (41)$$

the maximum u consistent with selling in France and exporting somewhere. To simulate exporters that sell in France we want $u(s) \leq \bar{u}(s)$ for each artificial exporter s , hence $u(s)$ should be a realization from the uniform distribution over the interval $[0, \bar{u}(s)]$.

Therefore we construct:

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

using the $v(s)$'s that were drawn prior to the simulation.

6. In the model a measure \bar{u} of firms have standardized unit cost below \bar{u} . Our artificial French exporter s therefore gets an importance weight $\bar{u}(s)$. This importance weight will be used in constructing statistics on artificial French exporters that relate to statistics on actual French exporters.¹⁷

¹⁷See Gouriéroux and Monfort (1995, Chapter 5) for a discussion of the use of importance weights in simulation.

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

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

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

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

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

4.3 Moments

In our estimation, we simulate firms that make it into at least one foreign market and into France as well. The reason for the first requirement is that firms that sell only in France are very numerous, and hence capturing them would consume a large portion of simulation draws. But since their activity is so limited they add little to parameter identification.¹⁸ The reason for the second requirement is that key moments in our estimation procedure are based on sales in France by exporters, which we can compute only for firms that sell in the home market.¹⁹ Given parameter estimates, we later explore the implications of the model for nonexporters as well.

¹⁸We also estimated the model matching moments of nonexporting firms as well. Coefficient estimates were similar to those we report below but the estimation algorithm, given estimation time, was much less precise.

¹⁹As mentioned in footnote 3, there are only 523 firms that apparently do not sell in France.

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 \widetilde{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 \widetilde{N}^k as:

$$\widetilde{N}^k = \frac{1}{S} \sum_{s=1}^S \bar{u}(s) \delta^k(s). \quad (42)$$

We now describe the moments that we seek to match.²⁰

We have chosen our moments to capture the four features of French firms' behavior described in Section 2: (1) their entry into particular subsets of export markets, (2) their sales in export destination n , (3) their sales in France conditional on selling to n , and (4) their sales in n relative to their sales in France conditional on selling to n :

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 $k = 1, \dots, 128$). The corresponding moments from the actual data are simply the proportion $m^k(1)$ of exporters selling to combination k . Stacking these proportions gives

²⁰Notice, from (40), (41), and (19), that the average of the importance weights $\bar{u}(s)$ is the simulated number of French firms that export and sell in France.

us $\widehat{m}(1; \Theta)$ and $m(1)$, each with 128 elements (subject to 1 adding up constraint).

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 q th 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 q th 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 q th 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:

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

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:

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

where \mathbf{W} is a 1360×1360 weighting matrix. We search for Θ using the simulated annealing algorithm.²¹ 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.

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.

²¹Goffe, Ferrier, and Rogers (1994) describe the algorithm. We use a version developed specifically for Gauss and available on the web from William Goffe (Simann).

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:

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

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

We calculate standard errors using a bootstrap technique, taking into account both sampling error and simulation error. To account for sampling error each bootstrap b replaces m with a different m^b . To account for simulation error each bootstrap b samples a new set of 500,000 v^b 's, a_n^b 's, and h_n^b 's as described in Section 4.2, thus generating a new $\hat{m}^b(\Theta)$.²² Defining $y^b(\Theta) = m^b - \hat{m}^b(\Theta)$ for each b we search for:

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

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

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

and take the square roots of the diagonal elements as the standard errors.²³

²²Just like m^b , \hat{m}^b is calculated according to the bins defined from the actual data.

²³Since 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:

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

4.5 Results

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

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

Our discussion in Section 3.7 foreshadowed our estimate of $\tilde{\theta}$, which lies between the slopes in Figures 3c and Figure 4. From equations (35), (33), and (34), the characteristic of a firm determining both entry and sales conditional on entry, is $v^{-1/\tilde{\theta}}$, where $v \sim U[0, 1]$. Our estimate of $\tilde{\theta}$ implies that the ratio of the 75th to the 25th percentile of this term is 1.56. Another way to assess the magnitude of $\tilde{\theta}$ is by its implication for aggregate fixed costs of entry. Using expression (24), our estimate of 2.46 implies that fixed costs dissipate about 59 percent of gross profit in any destination.

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

A feature of the data is the entry of firms into markets where they sell very little, as seen in Figure 1c. Two features of our estimates reconcile these small sales with a fixed cost of entry. First, our estimate of λ , which is close to one, means that a firm that is close to the entry

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

cutoff incurs a very small entry cost.²⁴ Second, the negative covariance between the sales and entry shocks explains why a firm with a given u might enter a market and sell relatively little. The first applies to firms that differ systematically in their efficiency while the second applies to the luck of the draw in individual markets.²⁵

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 $\hat{\Theta}$ to simulate a set of artificial firms including nonexporters.²⁶ 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

²⁴ Arkolakis (2008) finds a value around one consistent with various observations from several countries.

²⁵ We perform a Monte Carlo test of the ability of our estimation procedure to recover parameter values. We simulate 230,000 artificial French firms with the estimated parameter values reported above. To do so step 5 of the simulation algorithm was altered to $\bar{u}(s) = \bar{u}_F(s)$. We then apply the estimation procedure, exactly as described, to these simulated data to estimate Θ (using the same weighting matrix \mathbf{W} as in the original estimation). The table below reports the values used to create the simulated data (the “truth”) and the parameter estimates the estimation procedure delivers:

	$\tilde{\theta}$	λ	σ_a	σ_h	ρ
“truth”	2.46	0.91	1.69	0.34	-0.65
estimates	2.54	0.67	1.69	0.32	-0.56

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

²⁶ 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 $\bar{u}(s) = \bar{u}_F(s)$.

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 data 27.2 percent of exporters adhere to hierarchies compared with 30.3 percent in the model simulation, 13.3 percent implied by simply predicting on the basis of the marginal probabilities, and 100 percent had there been no entry shock ($\sigma_h = 0$). 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 (35) in Section 3.7 motivates Figure 5a, which plots the simulated (x's) and actual (circles) values of the median and 95th percentile sales to each market against actual mean French sales in that market. The model captures very well both the distance between the two percentiles in any given market and how each percentile varies across markets. The model also nearly matches the amount of noise in these percentiles, especially in markets where mean sales are small.

Sales in France Conditional on Entry in a Foreign Market. Equation (36) in Section 3.7 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 (37) in Section 3.7 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_a , 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. By the law of large numbers, the fraction of French firms selling in n is a close approximation to the probability that a French firm will sell in n . Thus we write this probability as:

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

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

$$V_n^U = q_n (1 - q_n). \tag{43}$$

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

satisfies:

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

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

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

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

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

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

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

We simulated the term $E[V_n^C(u)]$ 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. The results indicate that we can attribute around 57 percent 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.²⁷

4.7.2 Variation in Sales

Looking at the firms that enter a particular market, how much does the variation in u explain the variation in their sales there. Consider firm j selling in market n . Inserting (33) into (30),

²⁷The average value across markets is 0.57, with a standard deviation of 0.01. The reason for so little systematic variation in R_n^E across markets is that, conditioning on u , the probability q_n that a firm sells in any market n other than France is small. Taking the limit as $q_n \rightarrow 0$, R_n^E is independent of n .

the log of its sales there is:

$$\ln X_{nF}(j) = \underbrace{\ln \alpha_n(j)}_1 + \underbrace{\ln \left[1 - \left(\frac{u(j)\kappa_2}{N_{nF} [\eta_n(j)]^{\tilde{\theta}}} \right)^{\lambda/\tilde{\theta}} \right]}_2 - \underbrace{\frac{1}{\tilde{\theta}} \ln u(j)}_3 + \underbrace{\ln \left((N_{nF}/\kappa_2)^{1/\tilde{\theta}} \sigma E_n \right)}_4.$$

where we have divided sales into four components. Component 4 is common to all firms selling in market n so does not contribute to variation in sales there. The first component involves firm j 's idiosyncratic sales shock in market n while component 3 involves its efficiency shock that applies across all markets. Complicating matters is component 2, which involves both firm j 's idiosyncratic entry shock in market 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. We find that component 3 itself contributes only around 4.8 percent of the variation in $\ln X_{nF}(j)$, and components 2 and 3 together around 39 percent.²⁸

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. This finding follows from our much higher estimate of σ_a relative to σ_h .²⁹

²⁸Again, 0.048 and 0.39 are averages across markets, with respective standard deviations 0.0003 and 0.0025.

The presence of N_{nF} in component 2 makes the contribution of each component vary across markets, but the differences are very small.

²⁹A lower value of $\tilde{\theta}$ (implying more sales heterogeneity attributable to efficiency), for given σ_a and σ_h , would lead us to attribute more to the firm's underlying efficiency rather than to destination-specific shocks.

4.8 Productivity

Our methodology so far has allowed us to estimate $\tilde{\theta}$, 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 break down $\tilde{\theta}$ into these components we turn to data on firm productivity, as measured by value added per worker, and how it differs among firms selling to different numbers of markets.³⁰

A common observation is that exporters are more productive (according to various measures) than the average firm.³¹ The same is true of our exporters here: The average value added per worker of exporters is 1.22 times the average for all firms. Moreover, value added per worker, like sales in France, tends to rise with the number of markets served, but not with nearly as much regularity.

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

How do we calculate productivity among our simulated firms? Because it provides a

³⁰Because value added per worker is a crude productivity measure, we didn't incorporate these numbers into our method of simulated moments estimation above, as we didn't need them to estimate the other parameters of the model.

³¹See, for example, Bernard and Jensen (1999), Lach, Roberts, and Tybout (1997), and BEJK (2003).

simple analytic expression we first look at the productivity of a firm's operations in selling to a particular market n . We define its value added $V_n(j)$ in market n as:

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

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

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

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

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

$$\begin{aligned} q_n(j) &= \frac{V_n(j)}{\beta\bar{m}^{-1}X_n(j) + \beta^F E_n(j)} \\ &= \frac{[1 - (1 - \beta)\bar{m}^{-1}] X_n(j) - (1 - \beta^F)E_n(j)}{\beta\bar{m}^{-1}X_n(j) + \beta^F E_n(j)} \\ &= \frac{[\bar{m} - (1 - \beta)] - \bar{m}(1 - \beta^F) [E_n(j)/X_n(j)]}{\beta + \bar{m}\beta^F [E_n(j)/X_n(j)]}. \end{aligned} \tag{44}$$

The only source of cross-firm heterogeneity in productivity arises through the ratio $E_n(j)/X_n(j)$:

Firms having more sales $X_n(j)$ relative to entry costs $E_n(j)$ are more productive.³²

Using (23) for the numerator and (30) for the denominator, exploiting (34), we can write this ratio in terms of $v_n(j)$ as:

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

³²While in the actual data we look at value added per worker, it is more direct to calculate our model's prediction for value added per unit of factor cost. In our model the two are proportional since we assume that all producers in a country face the same wage and input costs, with labor having the same share.

More efficient firms will typically have a low ratio of entry costs to sales and hence, according to (44), relatively high productivity.³³

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

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

$$q(j) = \frac{[\bar{m} - (1 - \beta)] - \bar{m}(1 - \beta^F) [E(j)/X(j)]}{\beta + \bar{m}\beta^F [E(j)/X(j)]}. \quad (45)$$

To simplify we assume $\beta^F = 0$ so that all fixed costs are purchased services. We then calibrate β from data on the share of manufacturing value added in gross production.³⁴

³³At $\lambda = 1$ the equation becomes:

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

For any λ the ratio of entry costs to sales is bounded between 0 for the most efficient firms ($v_n(j)$ near zero) and $1/\sigma$ for the least efficient ($v_n(j)$ near one) entering country n . Their productivity in market n is thus bounded between 1 and $[\bar{m} - (1 - \beta)]/\beta$. A higher \bar{m} or lower β will generate more variation in productivity.

³⁴Denoting the value-added share as V , averaging across UNIDO gives us $V = 0.36$. Taking into account profits and fixed costs we calculate:

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

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

Note that the expression for firm productivity (45) depends on the elasticity of substitution σ (through $\bar{m} = \sigma/(\sigma - 1)$) but not on θ . Following BEJK (2003) we find an \bar{m} that makes the productivity advantage of exporters in our simulated data match their productivity advantage in the actual data (1.22). This exercise delivers $\bar{m} = 1.51$ or $\sigma = 2.98$, implying $\beta = 0.34$. Using our estimate of $\tilde{\theta} = \theta/(\sigma - 1) = 2.46$ the implied value of θ is 4.87.³⁵

Figure 6 reports average value added per worker against the minimum number of markets where the firms sell. We normalize the average of all firms (including nonexporters) to one. Circles represent the actual data and x's our simulation, based on (45) with $\bar{m} = 1.51$. Note that, among the actual firms, value added per worker more than doubles as we move from all of our firms (selling in at least one market) to those selling to at least 74 markets, and then plummets as we look at firms that sell more widely (although we are looking at fewer than one hundred firms in this upper tail). The model picks up the rise in measured productivity corresponding to wider entry at the low end (representing the vast majority of firms) but fails to reproduce the spike at the high end.

4.9 The Home Market 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_n as given, the elasticity of P_n with respect to X_n is -0.30 : Doubling market size leads to a 30 percent lower manufacturing

³⁵While 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).

price index. Recall, however, that our analysis inspired by the relationship (32) suggested that entry costs rise with market size, with an elasticity of 0.31. Putting these two calculations together, however, still leaves us with an elasticity of -0.21 : Our results imply a substantially lower price index in larger markets.

5 General Equilibrium and Counterfactuals

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

5.1 Embedding the Model in a General Equilibrium Framework

Embedding our analysis in general equilibrium requires additional assumptions:

1. 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 .
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:

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

where β is the labor share and $\kappa_3 = \beta^{-\beta}(1 - \beta)^{-(1-\beta)}$.

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 labor in the destination. We thus decompose the country-specific component of the entry cost $E_n = W_n F_n$, where F_n reflects the efficiency of workers in country n in producing the entry cost component of production.³⁶

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

$$Y_i = \sum_{n=1}^N \pi_{ni} X_n \quad (46)$$

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

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

$$X_i = \gamma(Y_i^A + D_i^A) + [(1 - \beta)(\sigma - 1)/\sigma] Y_i \quad (47)$$

where Y_i^A is GDP and D_i^A the trade deficit. Since our model is static we treat deficits as exogenous.

To relate Y_i^A to wages we write:

$$Y_i^A = W_i L_i + \Pi_i, \quad (48)$$

³⁶Combined 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.

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, \bar{E}_n . Using (24) for \bar{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:

$$\Pi_i = \sum_{n=1}^N \pi_{ni} \Pi_n^D = \frac{(\sigma - 1)}{\sigma\theta} Y_i, \quad (49)$$

where the second equality comes from applying the conditions (46) for equilibrium in the market for manufactures.

Substituting (48) into (47) and using the fact that gross manufacturing production Y_i is gross manufacturing absorption X_i less the manufacturing trade deficit D_i :³⁷

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

Solving for Y_i :

$$Y_i = \frac{\gamma\sigma (W_i L_i + D_i^A) - \sigma D_i}{1 + (\sigma - 1)(\beta - \gamma/\theta)}. \quad (50)$$

Since $X_i = Y_i + D_i$:

$$X_n = \frac{\gamma\sigma (W_n L_n + D_n^A) - (\sigma - 1)(1 - \beta + \gamma/\theta) D_n}{1 + (\sigma - 1)(\beta - \gamma/\theta)}. \quad (51)$$

³⁷For simplicity we reconcile the differences between manufacturing and overall trade deficits by thinking of them as transfers of the final good, which is otherwise not traded. For large economies, the manufacturing deficit is the largest component of the overall trade deficit. See Dekle, Eaton, and Kortum (2008) for a fuller treatment of deficits in a similar model of bilateral trade.

From expression (4) we can write:

$$\pi_{ni} = \frac{T_i \left(W_i^\beta P_i^{1-\beta} d_{ni} \right)^{-\theta}}{\sum_{k=1}^N T_k \left(W_k^\beta P_k^{1-\beta} d_{nk} \right)^{-\theta}}. \quad (52)$$

Substituting (50), (51), and (52) into (46) gives us a set of equations determining wages W_i around the world given prices P_i . From expression (14) we have:

$$P_n = \bar{m} \kappa_1^{-1/\theta} \kappa_3 \left[\sum_{i=1}^N T_i (W_i^\beta P_i^{1-\beta} d_{ni})^{-\theta} \right]^{-1/\theta} \left(\frac{X_n}{\sigma W_n F_n} \right)^{(1/\theta)-1/(\sigma-1)}. \quad (53)$$

giving us prices P_i around the world given wages W_i .

5.2 Calculating Counterfactual Outcomes

We could use (46) and (53) together to see how changes in T_i , L_i , F_i , D_i , D_i^A , or d_{ni} anywhere would affect wages W_i and manufacturing price indices P_i around the world (for given values of γ , β , σ , θ , and κ_1). Here we limit ourselves to considering the effect of changes in trade barriers d_{ni} and in entry costs F_i .³⁸

We apply the method explained in Dekle, Eaton, and Kortum (henceforth DEK, 2008) to calculate counterfactuals.³⁹ Denote the counterfactual value of any variable x as x' and define $\hat{x} = x'/x$ as its change. Equilibrium in world manufactures in the counterfactual requires:

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

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

³⁸We thus treat deficits as unchanged (as a share of world GDP).

³⁹DEK (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).

using (50), (51), and (52) as follows:

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

where sticking these three equations into (54) yields a set of equations involving \widehat{W}_i 's for given \widehat{P}_i 's. From (53) we can get a set of equations involving \widehat{P}_i 's for given \widehat{W}_i 's:

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

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 manufacturing production and trade. Appendix B 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 numeraire, 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 (31), as $\widehat{N}_{nF} = \left(\widehat{\pi}_{ni} \widehat{X}_n \right) / \left(\widehat{W}_n \widehat{F}_n \right)$.

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} 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 .⁴⁰
2. We run the simulation described in Section 4.2 with $S = 500,000$ and the same stochastic draws applying both to the baseline and 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_n 's and our counterfactual values X_{nF}^C and N_{nF}^C to calculate counterfactual σE_n^C 's in each destination.
 - (b) In step 4 we use our baseline values X_{nF} and N_{nF} to calculate baseline $\bar{u}_n(s)$'s and our counterfactual values X_{nF}^C and N_{nF}^C to calculate counterfactual $\bar{u}_n^C(s)$'s for each destination and firm, using (32) and (40).
 - (c) In step 5, we set

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

A firm for which $u(s) \leq \bar{u}_n(s)$ sells in market n in the baseline while a firm for which $u(s) \leq \bar{u}_n^C(s)$ sells there in the counterfactual. Hence our simulation allows for entry, exit, and survival.

⁴⁰Recall that, for reasons such as manufacturing exports by nonmanufacturing firms, the aggregate exports described in Appendix B exceed total exports by our French firms. Having used the aggregate data to calculate the counterfactual equilibrium as described in the previous section, we applied the percentage changes from that exercise to the X_{nF} and N_{nF} in our firm dataset.

- (d) In step 7 we calculate entry and sales in each of the 113 markets in the baseline and in the counterfactual.

5.3 Counterfactual Results

We consider two counterfactuals. The first is a ten percent drop in trade barriers, i.e., $\hat{d}_{ni} = 1/(1.1)$ for $i \neq n$ and $\hat{d}_{nn} = \hat{F}_n = 1$. This change roughly replicates the increase in French import share over the decade following 1986.⁴¹ The second is a ten percent drop in entry costs around the world, i.e. $\hat{F}_n = 1/(1.1)$ and $\hat{d}_{ni} = 1$. We describe the results in turn.

5.3.1 Implications of Globalization

Table 3 shows the aggregate general equilibrium consequences of ten percent lower trade costs: (i) the change in the real wage $\hat{\omega}_n = \left(\widehat{W}_n/\widehat{P}_n\right)^{\gamma_n}$, (ii) the change in the relative wage $\widehat{W}_n/\widehat{W}_F$, the change in the sales of French firms to each market \widehat{X}_{nF} , and 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.⁴² Relative wages move quite a bit more, capturing terms of trade effects from globalization. The results that matter at the firm level are French

⁴¹Using 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.

⁴²There 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.

sales and the number of French firms active in each market. While French sales declines by 5 percent in the home market, exports increase substantially, with a maximum 80 percent increase in Japan.⁴³ The number of French exporters increases roughly in parallel with French exports.

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 avoids any attrition. Because so many firms in the top decile already export, the greatest number of new exporters emerge from the second highest decile. The biggest percentage increase in number of exporters is for firms in the third from the bottom decile.

Table 6 decomposes sales revenues. All of the increase is in the top decile, and most of that

⁴³A 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.

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.

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.⁴⁴

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.

5.3.2 Implications of Lower Entry Costs

Our counterfactual of 10 percent lower entry costs yields very stark results. Changes in entry costs enter directly only in equation (55), leading to lower prices. But, since we assume a common drop in entry costs, manufacturing price indices drop in parallel, with trade shares and relative wages unchanged. We can thus solve analytically for the common change in

⁴⁴The 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.

manufacturing prices:

$$\hat{P} = \hat{F}^{[\theta - (\sigma - 1)] / [\beta \theta (\sigma - 1)]} = 0.919.$$

The number of French firms selling in every market rises by 10 percent and sales (in terms of the prices of manufactures) change by the factor $1/0.919 = 1.088$ in each market. Of the 10 percent increase in French firms, just over 1 percent enter as exporters. The increase in the number of French exporters comes primarily from the upper end of the size distribution of incumbent firms, although the percentage increase in the number of exporters is greatest near the middle of the distribution. While incumbent firms expand into more markets, entry of new, small firms leaves the size distribution essentially unchanged.

6 Conclusion

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

References

- Alvarez, Fernando and Robert E. Lucas (2007) "A General Equilibrium Analysis of the Eaton-Kortum Model of International Trade," *Journal of Monetary Economics*, 54: 1726-1768.
- Arkolakis, Konstantinos (2008) "Market Penetration Costs and the New Consumers Margin in International Trade," mimeo, Yale University.
- Aw, Bee Yan, Chung, Sukkyun, and Roberts, Mark J. (1998) "Productivity and the Decision to Export: Micro Evidence from Taiwan and South Korea." *NBER Working Paper No. 6558*.
- Bernard, Andrew B., Eaton, Jonathan, Jensen, J. Bradford, and Kortum, Samuel (2003) "Plants and Productivity in International Trade." *American Economic Review*, 93: 1268-1290.
- Bernard, Andrew B. and Jensen, J. Bradford (1999) "Exceptional Exporter Performance, Cause, Effect, or Both?" *Journal of International Economics*, 47: 1-25.
- 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.
- Dekle, Robert, Jonathan Eaton, and Samuel Kortum (2008) "Global Rebalancing with Gravity," *IMF Staff Papers*, forthcoming.
- Eaton, Jonathan and Kortum, Samuel (2002) "Technology, Geography, and Trade." *Econometrica* 70: 1741-1780.
- Eaton, Jonathan, Kortum, Samuel, and Francis Kramarz (2004) "Dissecting Trade: Firms, Industries, and Export Destinations." *American Economic Review, Papers and Proceedings*, 94: 150-154. (Data sources and tables available in, *NBER Working Paper #10344*.)
- Feenstra, Robert C., Lipsey, Robert E. and Bowen, Henry P. (1997) "World Trade Flows, 1970-1992, with Production and Tariff Data." *NBER Working Paper No. 5910*.
- Gabaix, Xavier (1991) "Zipf's Law for Cities: An Explanation," *Quarterly Journal of Economics*, 94: 739-767.
- Gabaix, Xavier and Ibragimov, Rustam (2008) "Rank-1/2: A Simple Way to Improve the OLS Estimation of Tail Exponents," New York University.
- Goffe William L., Gary D. Ferrier, and John Rogers (1994) "Global Optimization of Statistical Functions with Simulated Annealing," *Journal of Econometrics*, 60: 65-99.

- Gouriéroux, Christian and Alain Monfort (1994) *Simulation Based Econometric Methods*, CORE Lecture Series, Université catholique de Louvain.
- Helpman, Elhanan, Melitz, Marc J., and Yeaple, Stephen R. (2004) “Exports vs. FDI with Heterogeneous Firms,” *American Economic Review*, 94: 300-316.
- Krugman, Paul R. (1980) “Scale Economies, Product Differentiation, and the Patterns of Trade,” *American Economic Review*, 70: 950-959.
- Horowitz, Joel (2001) “The Bootstrap,” Chapter 52, *Handbook of Econometrics, Volume 5*, J.J. Heckman and E.E. Leamer, editors. Amsterdam: Elsevier Science B.V. pp. 3159-3228.
- 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.
- OECD (1995) *The OECD STAN Database*.
- Ricardo, David (1821) *On The Principles of Political Economy and Taxation*. London: John Murray, Albemarle-Street.
- Roberts, Mark and Tybout, 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,” *Biometrika*, 42: 425-440.

UNIDO (1999) *Industrial Statistics Database*.

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

A Appendix A: Constructing the Firm Level Data

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

movements may differ from published trade statistics, the second being based on list prices and thus 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.

A Appendix B: 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).⁴⁵ 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.⁴⁶ In particular:

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

We exploit (48), (49), and (50) in order to derive baseline labor Y_i^L from data on GDP and deficits:

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

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

⁴⁵A 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.

⁴⁶The value of γ_n lies in the interval (0.15, 0.55) for 100 or our 113 countries. The average is 0.36.

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

Export String	Number of French Exporters		
	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)

Country	Code	Counterfactual Changes (ratio of counterfactual to baseline)			
		Real Wage	Relative Wage	Sales of	Number of
				French Firms	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

		Counterfactual Changes (ratio of counterfactual to baseline)			
Country	Code	Real Wage	Relative Wage	Sales of	Number of
				French Firms	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
PHILIPPINES	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 TOBAGO	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

Initial Size Interval (percentile)	All Firms			Exporters		
	Baseline # of Firms	Counterfactual Change		Baseline # of Firms	Counterfactual Change	
		from Baseline	Change in %		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

Initial Size Interval (percentile)	Total Sales			Exports		
	Baseline in \$millions	Counterfactual Change		Baseline in \$millions	Counterfactual Change	
		from Baseline	Change in %		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

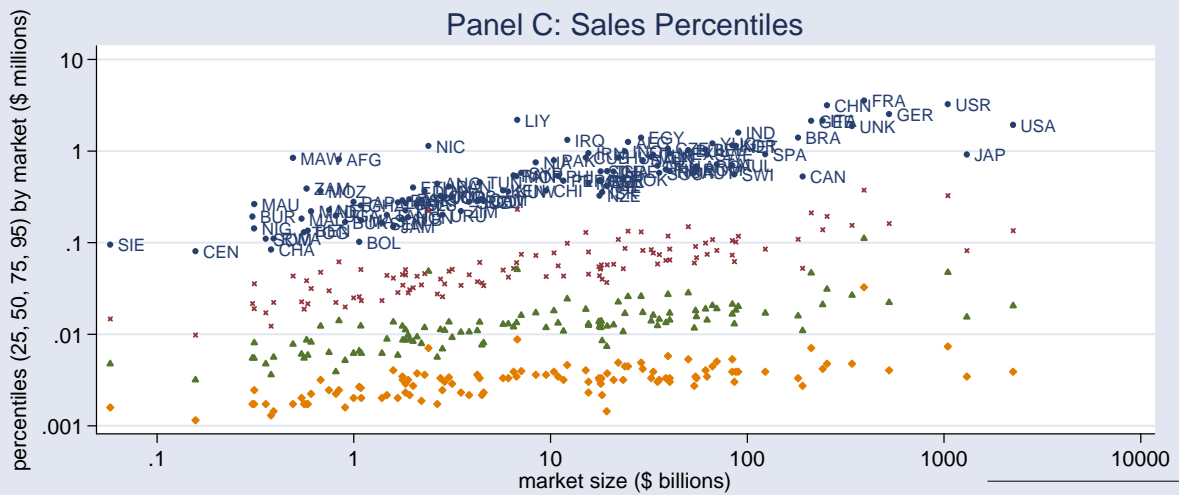
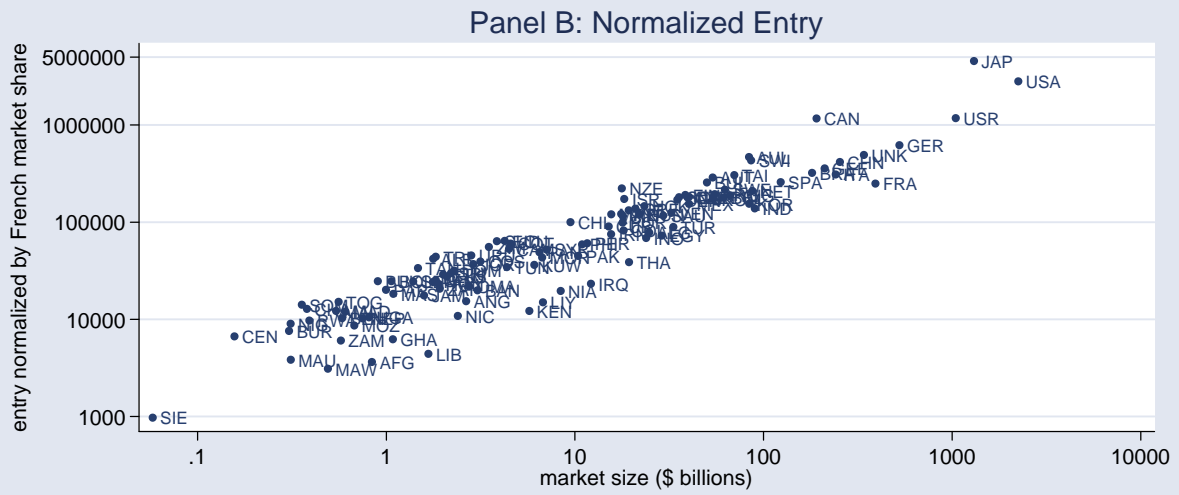
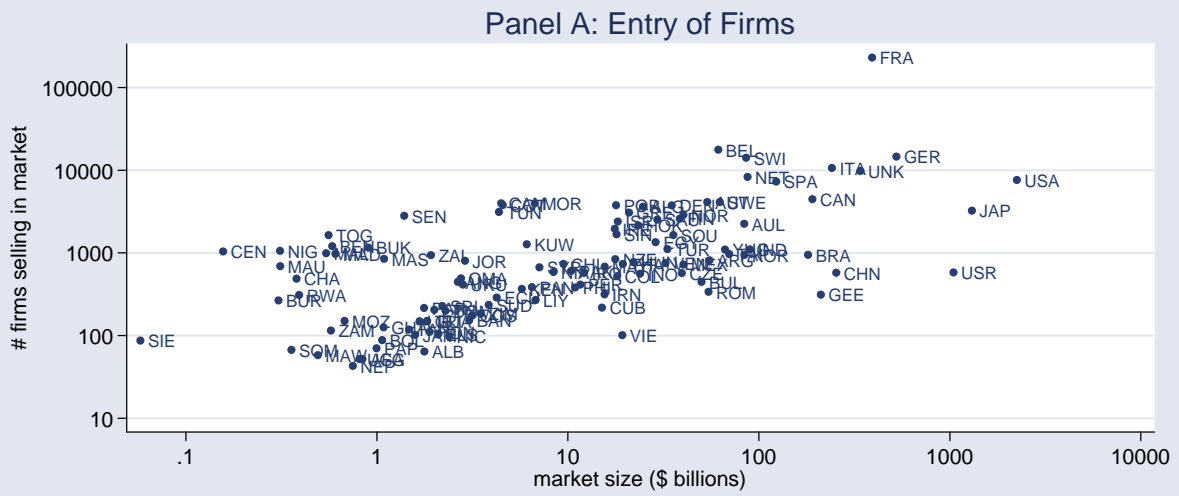
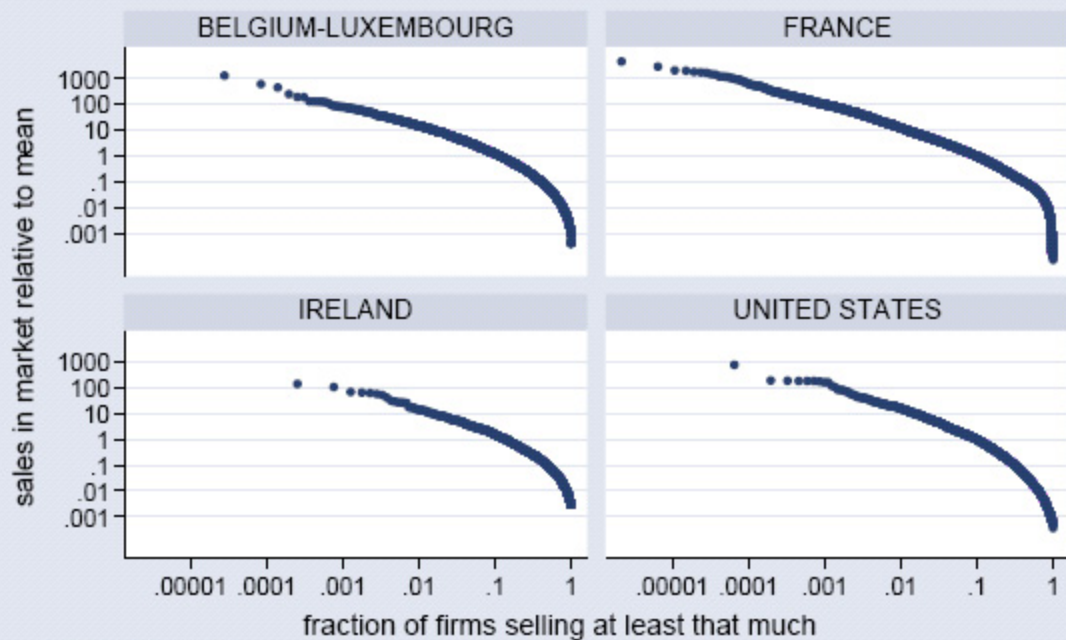


Figure 2
Sales Distributions of French Firms

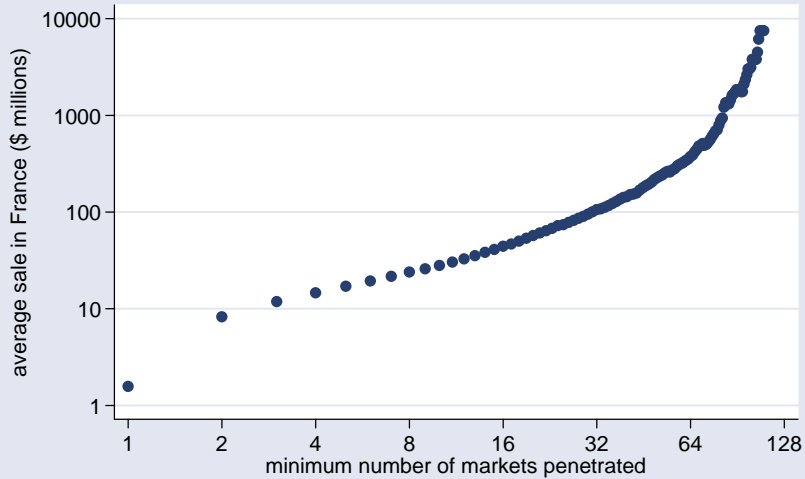


Graphs by country


STATA™

Figure 3: Sales in France and Market Entry

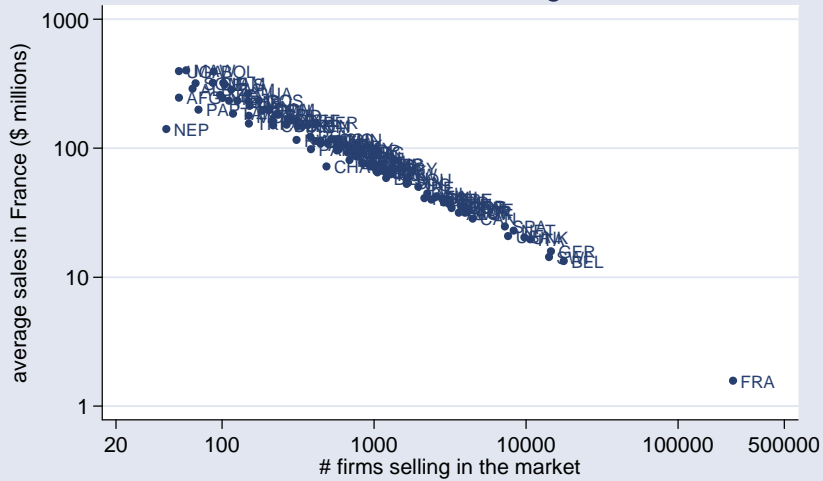
Panel A: Sales and Markets Penetrated



Panel B: Sales and # Penetrating Multiple Markets



Panel C: Sales and # Selling to a Market



Panel D: Distribution of Sales and Market Entry

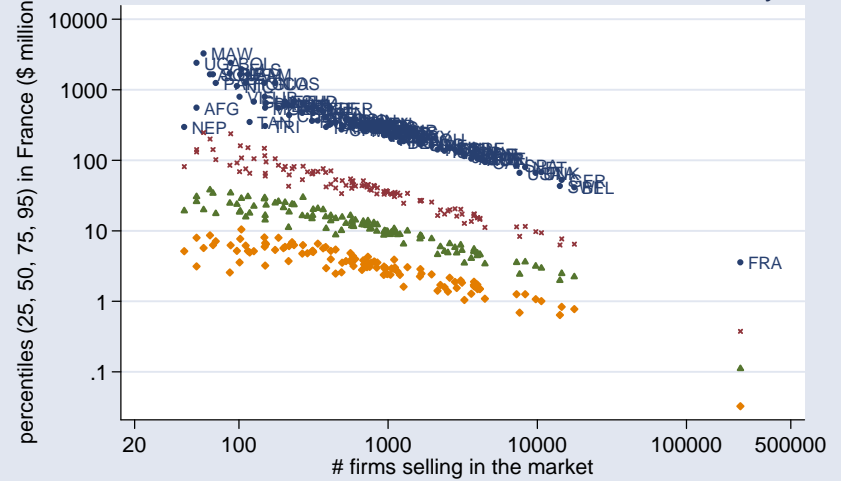
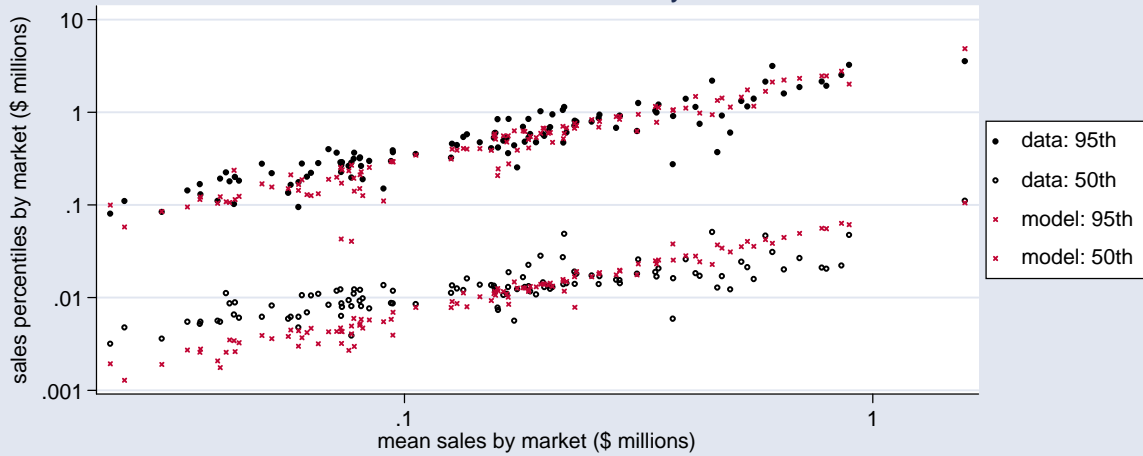


Figure 4: Distribution of Export Intensity, by Market




Figure 5: Model Versus Data

Panel A: Sales Distribution by Market



Panel B: Sales in France by Market Penetrated



Panel C: Export Intensity by Market

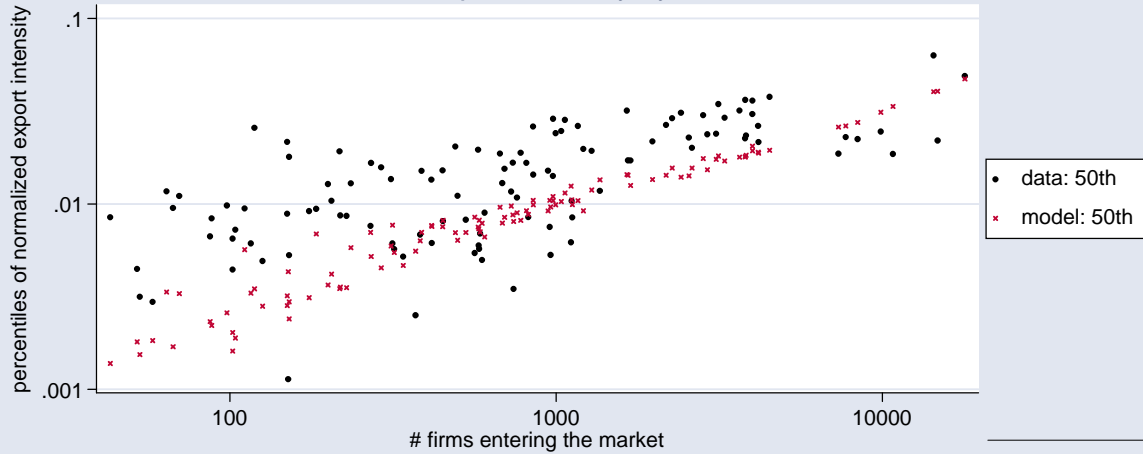


Figure 6: Productivity and Markets Penetrated
Model Versus Data

