

ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

*An International Society for the Advancement of Economic
Theory in its Relation to Statistics and Mathematics*

<http://www.econometricsociety.org/>

Econometrica, Vol. 79, No. 5 (September, 2011), 1453–1498

AN ANATOMY OF INTERNATIONAL TRADE: EVIDENCE FROM FRENCH FIRMS

JONATHAN EATON

Pennsylvania State University, University Park, PA 16802, U.S.A.

SAMUEL KORTUM

University of Chicago, Chicago, IL 60637, U.S.A.

FRANCIS KRAMARZ

CREST (ENSAE), 92245, Malakoff, France

The copyright to this Article is held by the Econometric Society. It may be downloaded, printed and reproduced only for educational or research purposes, including use in course packs. No downloading or copying may be done for any commercial purpose without the explicit permission of the Econometric Society. For such commercial purposes contact the Office of the Econometric Society (contact information may be found at the website <http://www.econometricsociety.org> or in the back cover of *Econometrica*). This statement must be included on all copies of this Article that are made available electronically or in any other format.

AN ANATOMY OF INTERNATIONAL TRADE: EVIDENCE FROM FRENCH FIRMS

BY JONATHAN EATON, SAMUEL KORTUM, AND FRANCIS KRAMARZ¹

We examine the sales of French manufacturing firms in 113 destinations, including France itself. Several regularities stand out: (i) the number of French firms selling to a market, relative to French market share, increases systematically with market size; (ii) sales distributions are similar across markets of very different size and extent of French participation; (iii) average sales in France rise systematically with selling to less popular markets and to more markets. We adopt a model of firm heterogeneity and export participation which we estimate to match moments of the French data using the method of simulated moments. The results imply that over half the variation across firms in market entry can be attributed to a single dimension of underlying firm heterogeneity: efficiency. Conditional on entry, underlying efficiency accounts for much less of the variation in sales in any given market. We use our results to simulate the effects of a 10 percent counterfactual decline in bilateral trade barriers on French firms. While total French sales rise by around \$16 billion (U.S.), sales by the top decile of firms rise by nearly \$23 billion (U.S.). Every lower decile experiences a drop in sales, due to selling less at home or exiting altogether.

KEYWORDS: Export destinations, firms, efficiency, sales distribution, simulated moments.

1. INTRODUCTION

WE EXPLOIT DETAILED DATA on the exports of French firms to confront a new generation of trade theories. These theories resurrect technological heterogeneity as the force that drives international trade. In Eaton and Kortum (2002), differences in efficiencies across countries in making different goods determine aggregate bilateral trade flows. Since they focused only on aggregate data, underlying heterogeneity across individual producers remains hidden. Subsequently, Melitz (2003) and Bernard, Eaton, Jensen, and Kortum (2003; henceforth BEJK) 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.² Another literature has modeled and estimated the

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

²For example, Bernard and Jensen (1995), for the United States, and Aw, Chung, and Roberts (2000), for Taiwan and Korea, documented the size and productivity advantage of exporters.

export decision of individual firms in partial equilibrium.³ However, the task of building a structure that can simultaneously embed behavior at the firm level into aggregate relationships and dissect aggregate shocks into their firm-level components remains incomplete. This paper seeks to further this mission.

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

We adopt the approach of 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, however, to come to terms with some features of the data: (i) Firms do not enter markets according to an exact hierarchy. (ii) Their sales where they do enter deviate from the exact correlations the basic model insists on. (iii) Firms that export sell too much in France. (iv) In the typical destination, there are too many firms selling small amounts.

To reconcile the basic model with the first two failures, we introduce market and firm-specific heterogeneity in entry costs and demand. We deal with the last two by incorporating Arkolakis's (2010) 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 57 percent. Conditional on a firm selling in a market, knowing its efficiency improves our ability to predict how much it sells there, but by much less. While these results leave much to be explained by the idiosyncratic interaction between individual firms and markets, they tell us that any theory that ignores features of the firm that are universal across markets misses much.

We conclude by using our parameterized model to examine a world with 10 percent lower trade barriers. To do so, we embed our model, with our parameter estimates, into a general equilibrium framework, calibrating it to aggregate data on production and bilateral trade. A striking finding is the extent to which lower trade barriers, while raising welfare in every country, favor the largest firms at the expense of others. Total output of French firms rises by

³A pioneering paper here is Roberts and Tybout (1997).

3.8 percent, with all of the growth accounted for by firms in the top decile. Sales in every other decile fall. Import competition leads to the exit of 11.5 percent of firms, 43 percent of which are in the smallest decile.

Section 2, which follows, explores four empirical regularities. With these in mind, we turn, in Section 3, to a model of exporting by heterogeneous firms. Section 4 explains how we estimate the model and considers some implications of the parameters. Section 5 explores the consequences of lower trade costs.

2. EMPIRICAL REGULARITIES

Our data are the sales, translated into U.S. dollars, of 230,423 French manufacturing firms to 113 markets in 1986. (Table III in Section 5.3 lists the countries.) Among them, only 34,558 sell outside France. The firm that exports most widely sells to 110 out of the 113 destinations. All but 523 of these firms indicate positive sales in France. Since much of our analysis focuses on the relationship between exporting and activity in France, we exclude these firms from much of our analysis, leaving 34,035 exporters also selling in France.⁴

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

2.1. *Market Entry*

Figure 1A plots the number of French manufacturing firms N_{nF} selling to a market against total manufacturing absorption X_n in that market across our 113 markets.⁵ While the number of firms selling to a market tends to increase with the size of the market, the relationship is cloudy.

The relationship comes into focus, however, when the number of firms is normalized by the share of France in a market. Figure 1B continues to report market size across the 113 destinations along the x axis. The y axis replaces N_{nF} , the number of French firms selling to a market, with N_{nF} divided by French market share, π_{nF} , defined as

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

where X_{nF} is total exports of our French firms to market n .

⁴ Appendix A describes the data. Appendices, data, and programs are available in the Supplemental Material (Eaton, Kortum, and Kramarz (2011)). Biscourp and Kramarz (2007) and Eaton, Kortum, and Kramarz (2004; EKK) used the same sources. EKK (2004) partitioned firms into 16 manufacturing sectors. While features vary across industries, enough similarity remains to lead us to ignore the industry dimension here. If a firm's total exports declared to French customs exceed its total sales from mandatory reports to the French fiscal administration, we treat the firm as not selling in France. The 523 such firms represent 1.51 percent of all French exporters and account for 1.23 percent of the total French exports to our 112 export destinations.

⁵ Manufacturing absorption is calculated as total production plus imports minus exports. See EKK (2004) for details.

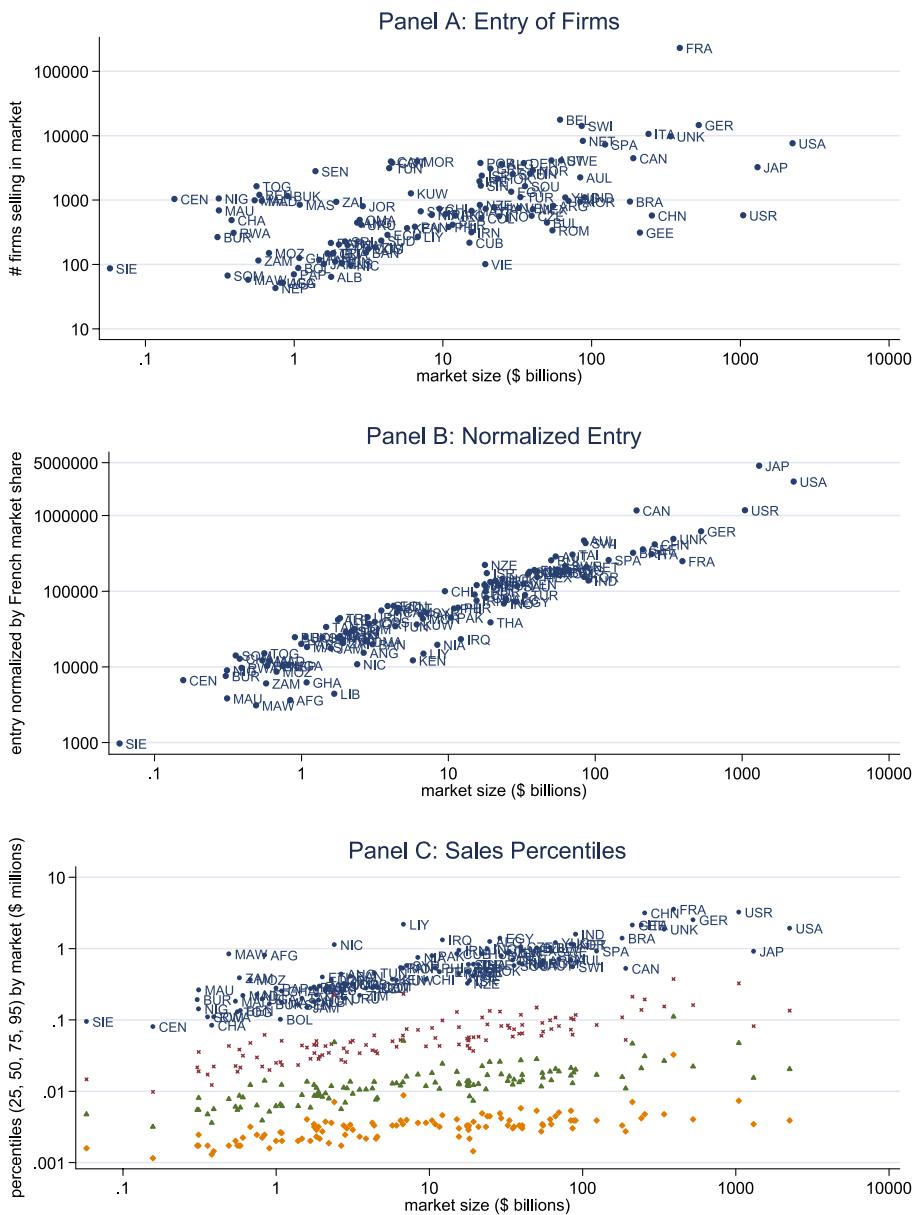


FIGURE 1.—Entry and sales by market size.

TABLE I
FRENCH FIRMS EXPORTING TO THE SEVEN MOST POPULAR DESTINATIONS

Export Destination	Number of Exporters	Fraction of Exporters
Belgium ^a (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)	9752	0.287
Netherlands (NL)	8294	0.244
United States (US)	7608	0.224
Any destination (all French exporters)	34,035	

^aBelgium includes Luxembourg.

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 0.65.⁶

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

We now turn to firm entry into different sets of markets. As a starting point for this examination, suppose firms obey a hierarchy in the sense that any firm selling to the $k + 1$ st most popular destination necessarily sells to the k th most popular destination as well. Not surprisingly, firms are less orderly in their choice of destinations. Consider exporters to the top seven foreign destinations. Table I 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-seven destination, the unconditional probability of a French exporter selling there.

Table II lists each of the strings of top-seven destinations that obey a hierarchical structure, together with the number of firms selling to each string (irrespective of their export activity outside the top 7). Overall, 27 percent of exporters (9260/34,035) obey a hierarchy among these most popular destinations. The next column of Table II uses the probabilities from Table I to predict the number of firms selling to each hierarchical string, if selling in one market is independent of selling in any other of the top seven. Under independence,

⁶If we make the assumption that French firms do not 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.

TABLE II
FRENCH FIRMS SELLING TO STRINGS OF TOP-SEVEN COUNTRIES

Export String ^a	Number of French Exporters		
	Data	Under Independence	Model
BE ^a	3988	1700	4417
BE-DE	863	1274	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	2406	15	2840
Total	9260	4532	9648

^aThe string BE means selling to Belgium but no other among the top 7; BE-DE means selling to Belgium and Germany but no other, and so forth.

the number of firms adhering to a hierarchy would be less than half of what we see in the data.

2.2. Sales Distributions

Our second exercise expands on Figure 1C by looking at the entire distribution of sales within individual markets. We plot the sales of each firm in a particular market (relative to mean sales there) against the fraction of firms selling in the market who sell at least that much.⁷ By doing so for all our 113 destinations, a remarkable similarity emerges. Figure 2 plots the results for Belgium, France, Ireland, and the United States on common axes. The basic shape is common for size distributions.⁸

⁷Following Gabaix and Ibragimov (2011), 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 , we calculate $(r_n(j) - 0.5)/N_{nF}$. To preserve confidentiality, our plots report the geometric mean of a group of adjacent sales. For the top 1000 sales, a group contains four observations, with larger groups used for lower ranks.

⁸Sales distributions are often associated with a Pareto distribution, at least in the upper tail (Simon and Bonini (1958)). To interpret Figure 2 as distributions, let x_n^q be the q th percentile French sales in market n normalized by mean sales in that market. We can write

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

where x_n is sales of a firm in market n relative to the mean. If the distribution is Pareto with parameter $a > 1$ (so that the minimum sales relative to the mean is $(a - 1)/a$), we have

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

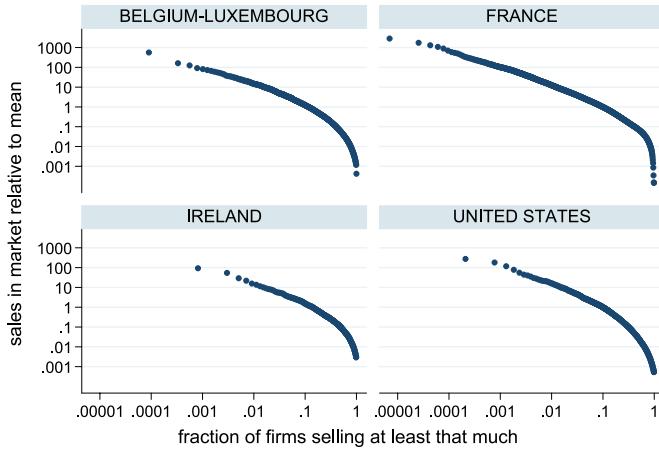


FIGURE 2.—Sales distributions of French firm: Graphs by country.

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 sells to all 113 destinations. Figure 3A depicts average sales in France on the y axis for the group of firms that sell to at least k markets with k on the x axis. Note the near monotonicity with which sales in France rise with the number of foreign markets served.

Figure 3B reports, on a log scale, average sales in France of firms selling to k or more markets against the number of firms selling to k or more markets. The relationship is strikingly linear with a regression slope of -0.66 .

Second, we rank countries according to their popularity as destinations for exports. The most popular destination is France itself, where all of our firms sell, followed by Belgium with 17,699 exporters. The least popular is Nepal, where only 43 French firms sell. Figure 3C depicts average sales in France on the y axis plotted against the number of firms selling to the k th most popular

or

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

implying a straight line with slope $-1/a$. Considering only sales by the top 1 percent of French firms selling in the four destinations depicted in Figure 2, regressions yield slopes of -0.74 (Belgium), -0.87 (France), -0.69 (Ireland), and -0.82 (United States). Note, however, that the distributions appear to deviate from a Pareto distribution, especially at the lower end.

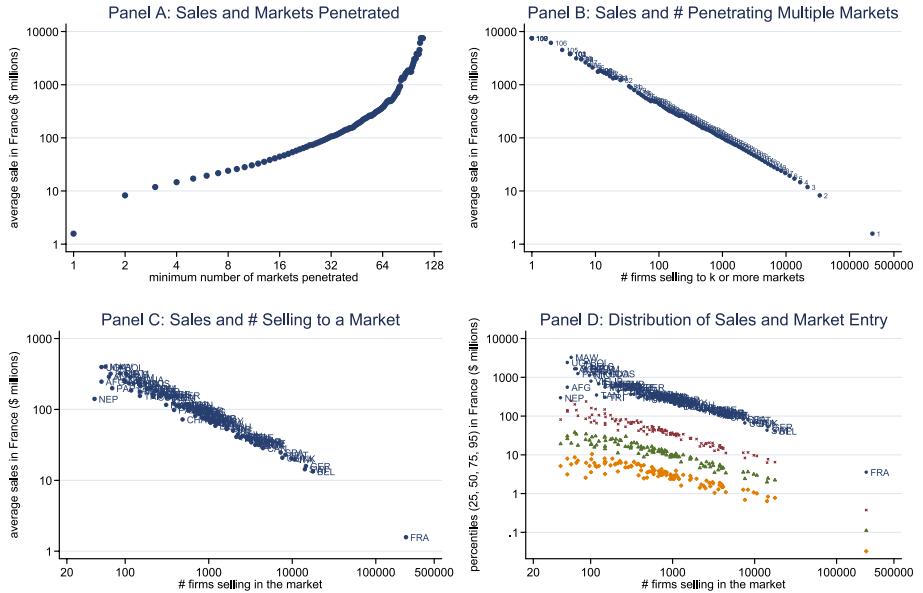


FIGURE 3.—Sales in France and market entry.

market on the x axis. The relationship is tight and linear in logs as in Figure 3B, although slightly flatter, with a slope of -0.57 . Firms selling to less popular markets and to more markets systematically sell more in France.

Delving further into the French sales of exporters to markets of varying popularity, Figure 3D reports the 95th, 75th, 50th, and 25th percentile sales in France (on the y axis) against the number of firms selling to each market. Note the tendency of sales in France to rise with the unpopularity of a destination across all percentiles (less systematically so for the 25th percentile).⁹

2.4. Export Intensity

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

⁹We were able to observe the relationship between market popularity and sales in France for the 1992 cross section as well. The analog (not shown) of Figure 3C is nearly identical. Furthermore, the changes between 1986 and 1992 in the number of French firms selling in a market correlate negatively 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, as the relationship would dictate.

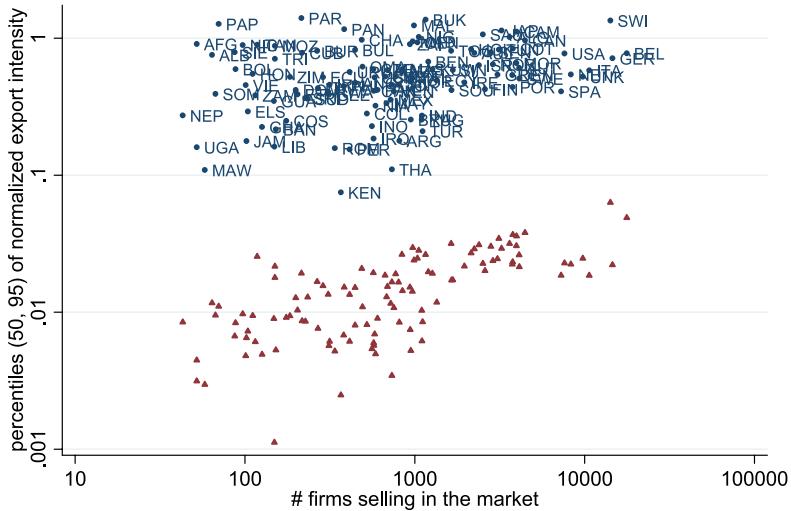


FIGURE 4.—Distribution of export intensity, by market.

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 \bar{X}_{nF} are average sales by French firms in market n ($X_{FF}(j)$ and \bar{X}_{FF} are the corresponding magnitudes in France). Scaling by \bar{X}_{nF} normalizes by the effect of market n on the size of all French firms there. Scaling by $X_{FF}(j)$ normalizes by firm j 's size in France.

Figure 4 plots the median and 95th percentile normalized export intensity for each foreign market n (on the y axis) against the number of firms selling to that market (on the x axis) on log scales. Two aspects stand out: (i) As a destination becomes more popular, normalized export intensity rises. The slope for the median is 0.39, but the relationship is noisy. (ii) Normalized “export” intensity for France itself is identically 1, while median export intensity in Figure 4 is usually 2 orders of magnitude or more below 1. Even among exporting firms, sales abroad are small compared to sales at home.

3. THEORY

In seeking to interpret these relationships, we turn to a parsimonious model which explains where firms sell and how much they sell there. The basic structure is monopolistic competition: Goods are differentiated, with each one corresponding to a firm. Selling in a market requires a fixed cost, while moving

goods from country to country incurs iceberg transport costs. Firms are heterogeneous in efficiency as well as in other characteristics, while countries vary in size, location, and fixed cost of entry.¹⁰

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 1 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 1 unit of good j to country n from country i is

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

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

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

where $\theta > 0$ and $T_i \geq 0$ are parameters.¹¹ Using (1), the measure of goods that can be delivered from country i to country n at unit cost below c is defined as

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

where

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

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

¹⁰We go with Melitz's (2003) monopolistic competition approach rather than the Ricardian framework with a fixed range of commodities used in BEJK (2003), since it more readily delivers the feature that a larger market attracts more firms, as we see in our French data.

¹¹We follow Chaney (2008) in taking T_i as exogenous and Helpman, Melitz, and Yeaple (2004) and Chaney (2008) in treating the underlying heterogeneity in efficiency as Pareto. A Pareto distribution of efficiencies can arise naturally from a dynamic process that is a history of independent shocks, as shown by Simon (1955), 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, 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.)

3.2. Demand, Market Structure, and Entry

A market n contains a measure of potential buyers. To sell to a fraction f of them, a producer in country i selling good j in country n must incur a fixed cost:

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

Here $\varepsilon_n(j)$ is a fixed-cost shock specific to good j in market n and E_{ni} is the component of the cost shock faced by all sellers from country i in destination n . The function $M(f)$, the same 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 (2010),

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

where the parameter $\lambda > 0$ reflects the increasing cost of reaching a larger fraction of potential buyers.¹² This function has the desirable properties that the cost of reaching zero buyers in a market is zero and that the total cost increases continuously (and the marginal cost weakly increases) in the fraction f of buyers reached. This formulation can explain why some firms sell a very small amount in a market while others stay out entirely.

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

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

¹²By l'Hôpital's rule, at $\lambda = 1$, the function becomes $M(f) = -\ln(1 - f)$. Arkolakis (2010) provided an extensive discussion of this functional form, deriving it from a model of the micro-foundations of marketing. He parameterized the function in terms of $\beta = 1/\lambda$. The case $\beta > 0$ corresponds to an increasing marginal cost of reaching additional buyers, which always results in an outcome with $0 \leq f < 1$. The case $\beta \leq 0$ means a constant or decreasing marginal cost of reaching additional buyers, in which case the firm would go to a corner ($f = 0$ or $f = 1$) as in Melitz (2003). Our restriction $\lambda > 0$ thus covers all cases (including Melitz as $\lambda \rightarrow \infty$) that we could observe.

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.

The producer of good j from country i selling in market n with a unit cost of $c_n(j)$, charging a price p and reaching a fraction f of buyers, earns a profit

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

To maximize (6), the producer sets the standard Dixit–Stiglitz (1977) markup over unit cost,

$$(7) \quad p_n(j) = \bar{m} c_n(j),$$

where

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

It seeks a fraction

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

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 will not sell at all, hence avoiding any fixed cost there, if

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

We can now describe a seller's behavior in market n in terms of its unit cost $c_n(j) = c$, demand shock $\alpha_n(j) = \alpha$, and entry shock $\eta_n(j) = \eta$. From the condition above, a firm enters market n if and only if

$$(9) \quad c \leq \bar{c}_{ni}(\eta),$$

where

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

We can use the expression for (10) to simplify expression (8) for the fraction of buyers reached by a producer with unit cost $c \leq \bar{c}_{ni}(\eta)$

$$(11) \quad f_{ni}(\eta, c) = 1 - \left(\frac{c}{\bar{c}_{ni}(\eta)} \right)^{\lambda(\sigma-1)}.$$

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

$$(12) \quad X_{ni}(\alpha, \eta, c) = \frac{\alpha}{\eta} \left[1 - \left(\frac{c}{\bar{c}_{ni}(\eta)} \right)^{\lambda(\sigma-1)} \right] \left(\frac{c}{\bar{c}_{ni}(\eta)} \right)^{-(\sigma-1)} \sigma E_{ni}.$$

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

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

$$(13) \quad E_{ni}(\alpha, \eta, c) = \frac{\alpha}{\eta} E_{ni} \frac{1 - (c/\bar{c}_{ni}(\eta))^{\lambda(\sigma-1)}}{1 - 1/\lambda}.$$

To summarize, the relevant characteristics of market n that apply across sellers are total purchases X_n , the price index P_n , and, for sellers from country i , the common component of the fixed cost E_{ni} . The particular situation of a potential seller of product j in market n is captured by three magnitudes: the unit cost $c_n(j)$, the demand shocks $\alpha_n(j)$, and the entry shocks $\eta_n(j)$. 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 $c_n(j)$.¹³

Equations (9) and (10), which govern entry, and (12), which governs sales conditional on entry, link our theory to the data on French firms' entry and sales in different markets of the world described in Section 2. Before returning to the data, however, we need to solve for the price index P_n in each market.

3.3. The Price Index

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

¹³We require $E[\eta^{\theta/(\sigma-1)}]$ and $E[\alpha\eta^{[\theta/(\sigma-1)]-1}]$ both to be finite. For any given $g(\alpha, \eta)$, these restrictions imply an upper bound on the parameter θ .

representative buyer in market n as

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

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

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

where

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

Moving P_n to one side of the equation gives

$$(14) \quad P_n = \bar{m} (\kappa_1 \Psi_n)^{-1/\theta} X_n^{(1/\theta)-1/(\sigma-1)},$$

where

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

and

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

Note that the price index has an elasticity of $(1/\theta) - 1/(\sigma - 1)$ with respect to total expenditure (given the terms in Ψ_n). Our restriction that $\theta > \sigma - 1$ makes the effect negative: A larger market enjoys lower prices, for reasons similar to the price index effect in Krugman (1980) and common across models of monopolistic competition.¹⁴

¹⁴In models of monopolistic competition with homogeneous firms, consumers anywhere consume all varieties regardless of where they are made. Since more varieties are produced in a larger market, local consumers benefit from their ability to buy these goods without incurring trade costs. With heterogeneous firms and a fixed cost of market entry, an additional benefit to consumers in a large market is greater variety.

3.4. Entry, Sales, and Fixed Costs

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

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

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

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

$$(18) \quad J_{ni} = \int [\mu_{ni}(\bar{c}_{ni}(\eta))] g_2(\eta) d\eta = \frac{\kappa_2}{\kappa_1} \frac{\pi_{ni} X_n}{\sigma E_{ni}},$$

where

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

and

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

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

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

in market n .

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

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

Hence the trade share of source i in market n is just π_{ni} .¹⁵

¹⁵Of the total measure J_n of firms selling in n , the fraction from i is

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

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

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

Integrating across the joint density $g(\alpha, \eta)$, total fixed costs incurred in market n by firms from i are

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

Summing across sources i , the total fixed costs incurred in market n are simply

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

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

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

3.5. A Streamlined Representation

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

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

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

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

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

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 does not depend on any parameter.

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

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

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

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

Firm j enters market n if its $u(j)$ and $\eta_n(j)$ satisfy

$$(25) \quad u(j) \leq \bar{u}_{nF}(\eta_n(j)) = \left(\frac{\pi_{nF} X_n}{\kappa_1 \sigma E_{nF}} \right) \eta_n(j)^{\tilde{\theta}},$$

where

$$\tilde{\theta} = \frac{\theta}{\sigma - 1} > 1.$$

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

$$(26) \quad X_{nF}(j) = \varepsilon_n(j) \left[1 - \left(\frac{u(j)}{\bar{u}_{nF}(\eta_n(j))} \right)^{\lambda/\tilde{\theta}} \right] \left(\frac{u(j)}{\bar{u}_{nF}(\eta_n(j))} \right)^{-1/\tilde{\theta}} \sigma E_{nF}.$$

Equations (25) and (26) reformulate the entry and sales equations (17) and (12) in terms of $u(j)$ rather than $c_n(j)$.

Since standardized unit cost $u(j)$ applies across all markets, it gets to the core of a firm's underlying efficiency as it applies to its entry and sales in different markets. Notice that in reformulating the model as (25) and (26), the parameters θ , σ , and E_{nF} enter only collectively through $\tilde{\theta}$ and σE_{nF} . The parameter $\tilde{\theta}$ 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. By observing just entry and sales, we are able to identify only $\tilde{\theta}$.

3.6. Connecting the Model to the Empirical Regularities

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

3.6.1. Entry

From (18), we get

$$(27) \quad \frac{N_{nF}}{\pi_{nF}} = \frac{\kappa_2}{\kappa_1} \frac{X_n}{\sigma E_{nF}},$$

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 1 suggests that entry cost σE_{nF} rises systematically with market size, but not proportionately so. We do not impose any such relationship, but rather employ (27) to calculate

$$(28) \quad \sigma E_{nF} = \frac{\kappa_2}{\kappa_1} \frac{\pi_{nF} X_n}{N_{nF}} = \frac{\kappa_2}{\kappa_1} \bar{X}_{nF}$$

directly from the data.¹⁶

Using (28), we can write (25) as

$$(29) \quad u(j) \leq \bar{u}_{nF}(\eta_n(j)) = \frac{N_{nF}}{\kappa_2} \eta_n(j)^{\tilde{\theta}}.$$

Without variation in the firm-specific and market specific entry shock $\eta_n(j)$, (29) implies efficiency is all that matters for entry, dictating a deterministic ranking of destinations with a less efficient firm (with a higher $u(j)$) selling to a subset of the destinations served by any more efficient firm. Hence deviations from market hierarchies, as we see in Table II, identify variation in $\eta_n(j)$.

¹⁶We can use equation (28) to infer how fixed costs vary with country characteristics. Simply regressing \bar{X}_{nF} against our market size measure X_n (both in logs) yields a coefficient of 0.35, (1 minus the slope in Figure 1B), in which N_{nF}/π_{nF} rises with market size with an elasticity of 0.65. (The connection between the two regressions is a result of the accounting identity: $\bar{X}_{nF} * N_{nF}/\pi_{nF} = X_n$.) If gross domestic product (GDP) per capita is added to the regression, it has a negative effect on entry costs. French data from 1992, and data from Denmark (from Pedersen (2009)) and from Uruguay (compiled by Raul Sampognaro) show similar results for market size but not a robust effect of GDP per capita. We also find that mean sales are higher in the home country. Appendix B reports these results.

3.6.2. Sales in a Market

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

$$(30) \quad v_{nF}(j) = \frac{u(j)}{\bar{u}_{nF}(\eta_n(j))}$$

is distributed uniformly on $[0, 1]$. Replacing $u(j)$ with $v_{nF}(j)$ in expression (26), and exploiting (28), we can write sales as

$$(31) \quad X_{nF}(j) = \varepsilon_n(j) [1 - v_{nF}(j)^{\lambda/\tilde{\theta}}] v_{nF}(j)^{-1/\tilde{\theta}} \frac{\kappa_2}{\kappa_1} \bar{X}_{nF}.$$

Not only does v_{nF} have the same distribution in each market n , so does ε_n .¹⁷ 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_{nF}), consistent with the common shapes of sales distributions exhibited in Figure 2. If the only source of variation in sales were the term $v_{nF}(j)^{-1/\tilde{\theta}}$, the sales distribution would be Pareto with parameter θ . The term in square brackets, however, implies a downward deviation from a Pareto distribution that is more pronounced as $v_{nF}(j)$ approaches 1, consistent with the curvature in the lower end of the sales distributions that we observe in Figure 2. The sales distribution also inherits properties of the distribution of $\varepsilon_n(j)$.

3.6.3. Sales in France Conditional on Entry in a Foreign Market

We can also look at the sales in France of French firms selling to any market n . To condition on these firms, we take (31) as it applies to France and use (30) and (29) to replace $v_{FF}(j)$ with $v_{nF}(j)$:

$$(32) \quad X_{FF}(j)|_n = \frac{\alpha_F(j)}{\eta_n(j)} \left[1 - v_{nF}(j)^{\lambda/\tilde{\theta}} \left(\frac{N_{nF}}{N_{FF}} \right)^{\lambda/\tilde{\theta}} \left(\frac{\eta_n(j)}{\eta_F(j)} \right)^\lambda \right] \\ \times v_{nF}(j)^{-1/\tilde{\theta}} \left(\frac{N_{nF}}{N_{FF}} \right)^{-1/\tilde{\theta}} \frac{\kappa_2}{\kappa_1} \bar{X}_{FF}.$$

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

¹⁷To 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

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

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

Consider first the presence of N_{nF} 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 1 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 firms' sales in France fall with N_{nF} with an elasticity of $-1/\tilde{\theta}$. Interpreting Figure 3C in terms of Equation (32), the slope of -0.57 implies a $\tilde{\theta} = 1.75$.

Expression (32) also suggests how we can identify other parameters of the model. The gap between the percentiles in Figure 3D is governed by the variation in the demand shock $\alpha_F(j)$ in France together with the variation in the entry shock $\eta_n(j)$ in country n .

Together (31) and (32) reconcile the near log linearity of sales in France with N_{nF} and the extreme curvature at the lower end of the sales distribution in any given market. An exporting firm may be close to the entry hurdle in the export market and hence sells to a small fraction of buyers there, while reaching most consumers at home. Hence looking at the home sales of exporters isolates firms that reach most of the French market. These equations also explain why France itself is somewhat below the trend line in Figure 3A and B: The many nonexporting firms that reach just a small fraction of the French market appear only in the data point for France.

3.6.4. Normalized Export Intensity

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

$$(33) \quad \frac{X_{nF}(j)}{X_{FF}(j)} / \left(\frac{\bar{X}_{nF}}{\bar{X}_{FF}} \right) \\ = \frac{\alpha_n(j)}{\alpha_F(j)} \left[\frac{1 - v_{nF}(j)^{\lambda/\tilde{\theta}}}{1 - v_{nF}(j)^{\lambda/\tilde{\theta}} \left(\frac{N_{nF}}{N_{FF}} \right)^{\lambda/\tilde{\theta}} \left(\frac{\eta_n(j)}{\eta_F(j)} \right)^\lambda} \right] \left(\frac{N_{nF}}{N_{FF}} \right)^{1/\tilde{\theta}}.$$

Note first how the presence of the sales shock $\alpha_n(j)$ accommodates random variation in sales in different markets conditional on entry.

As in (32), the only systematic source of cross-country variation on the right-hand side of (33) is N_{nF} . The relationship is consistent with three features of Figure 4. First, trivially, the observation for France is identically 1. Second, normalized export intensity is substantially below 1 for destinations served by only a small fraction of French firms, as is the case for all foreign markets. Third, normalized export intensity increases with the number of French firms selling there. According to (33) the elasticity of normalized export intensity with respect to N_{nF}/N_{FF} is $1/\tilde{\theta}$ (ignoring N_{nF}/N_{FF} 's role in the denominator of the term in the square bracket, which is tiny since N_{nF}/N_{FF} is close to zero for

$n \neq F$). The slope coefficient of 0.38 reported in Section 2.4 suggests a value of $\tilde{\theta} = 2.63$.¹⁸

4. ESTIMATION

We estimate the parameters of the model by the method of simulated moments. We simulate firms that make it into at least one foreign market and into France as well.¹⁹ Given parameter estimates, we later explore the implications of the model for nonexporters as well.

We first complete our parameterization of the model. Second, we explain how we simulate a set of artificial French exporters given a particular set of parameter values, with each firm assigned a cost draw u , and an α and η in each market. Third, we describe how we calculate a set of moments from these artificial data to compare with moments from the actual data. Finally, we explain our estimation procedure, report our results, and examine the model's fit.

4.1. Parameterization

To complete the specification, we assume that $g(\alpha, \eta)$ is joint log normal. Specifically, $\ln \alpha$ and $\ln \eta$ are normally distributed with zero means and variances σ_α^2 and σ_η^2 , and correlation ρ .²⁰ Under these assumptions, we can write (15) and (19) as

$$(34) \quad \kappa_1 = \left[\frac{\tilde{\theta}}{\tilde{\theta} - 1} - \frac{\tilde{\theta}}{\tilde{\theta} + \lambda - 1} \right] \exp \left\{ \frac{\sigma_\alpha + 2\rho\sigma_\alpha\sigma_\eta(\tilde{\theta} - 1) + \sigma_\eta(\tilde{\theta} - 1)^2}{2} \right\}$$

¹⁸Equations (32) and (33) apply to the latent sales in France of firms that sell in n but do not enter France. In Figure 3C and D we can only look at the firms that sell in both places, of course. Since the French share in France is so much larger than the French share elsewhere, our theory predicts that a French firm selling in another market but not in France is very unlikely. Indeed, the number of such firms is small.

¹⁹The reason for the selling-in-France requirement is that key moments in our estimation procedure involve sales in France by exporters, which we can compute only for firms that enter the home market. The reason for the foreign-market requirement is that firms selling only in France are very numerous, so that capturing them would consume a large portion of simulation draws. However, since their activity is so limited, they add little to parameter identification. We also estimated the model matching moments of nonexporting firms. Coefficient estimates were similar to those we report below, but the estimation algorithm, given estimation time, was much less precise.

²⁰Since the entry-cost shock is given by $\ln \varepsilon = \ln \alpha - \ln \eta$, the implied variance of the entry-cost shock is

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

which is decreasing in ρ .

and

$$(35) \quad \kappa_2 = \exp\left\{\frac{(\tilde{\theta}\sigma_\eta)^2}{2}\right\}.$$

As above, in estimating the model, we equate the measure of French firms J_{nF} selling in each destination with the actual (integer) number N_{nF} and equate \bar{X}_{nF} with their average sales there.²¹ We are left with only five parameters to estimate:

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

For a given Θ , we use (28) to back out the cluster of parameters σE_{nF} using our data on $\bar{X}_{nF} = X_{nF}/N_{nF}$, and the κ_1 and κ_2 implied by (34) and (35). Similarly, we use (29) to back out a firm's entry hurdle in each market $\bar{u}_{nF}(\eta_n)$ given its η_n and the κ_2 implied by (35).

4.2. Simulation Algorithm

To estimate parameters, to assess the implications of those estimates, and to perform counterfactual experiments, we need to construct sets of artificial French firms that operate as the model tells them, given some Θ . We denote an artificial French exporter by s and the number of such exporters by S . The number S does not bear any relationship to the number of actual French exporters. A larger S implies less sampling variation in our simulations.

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

A given simulation of the model requires a set of parameters Θ , 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 the following eight steps:

- Step 1. Using (34) and (35), we calculate κ_1 and κ_2 .
- Step 2. Using (28), we calculate σE_{nF} for each destination n .
- Step 3. We use the $a_n(s)$'s and $h_n(s)$'s to construct $S \times 113$ realizations for each of $\ln \alpha_n(s)$ and $\ln \eta_n(s)$ as

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

²¹The model predicts that some French firms export but do not sell domestically. Consequently, the data for N_{nF} and X_{nF} that we condition on in our estimation include the 523 French exporters (mentioned in footnote 4) who do not enter the domestic market.

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

$$(36) \quad \bar{u}_n(s) = \frac{N_{nF}}{\kappa_2} \eta_n(s)^{\tilde{\theta}},$$

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

Step 5. We calculate

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

the maximum u consistent with exporting somewhere, and

$$(37) \quad \bar{u}(s) = \min\{\bar{u}_F(s), \bar{u}^X(s)\},$$

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

Step 6. To simulate exporters that sell in France, $u(s)$ 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.

Step 7. 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 is used to construct statistics on artificial French exporters that relate to statistics on actual French exporters.²²

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

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

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

$$(38) \quad X_{nF}(s) = \frac{\alpha_n(s)}{\eta_n(s)} \left[1 - \left(\frac{u(s)}{\bar{u}_n(s)} \right)^{\lambda/\tilde{\theta}} \right] \left(\frac{u(s)}{\bar{u}_n(s)} \right)^{-1/\tilde{\theta}} \sigma E_{nF}.$$

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 terms, we can compute any moment that could have been constructed from the actual French data.²³

²²See Gouriéroux and Monfort (1994, Chap. 5) for a discussion of the use of importance weights in simulation.

²³In principle, in any finite sample, the number of simulated firms that overcome the entry hurdle for a destination n could be zero, even though the distribution of efficiencies is unbounded

4.3. Moments

For a candidate value Θ , we use the algorithm above to simulate the sales of 500,000 artificial French exporting firms in 113 markets. From these artificial data, we compute a vector of moments $\hat{m}(\Theta)$ 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 let \hat{N}^k be the corresponding number in the simulated data. Using $\delta^k(s)$ as an indicator for when artificial firm s achieves outcome k , we calculate \hat{N}^k as

$$(39) \quad \hat{N}^k = \frac{1}{S} \sum_{s=1}^S \bar{u}(s) \delta^k(s).$$

We now describe the moments that we seek to match.²⁴

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

- The first set of moments relate to the entry strings discussed in Section 2.1. We compute the proportion $\hat{m}^k(1; \Theta)$ of simulated exporters selling to each possible combination k of the seven most popular export destinations (listed in Table I). One possibility is exporting yet selling to none of the top seven, giving us 2⁷ 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 us $\hat{m}(1; \Theta)$ and $m(1)$, each with 128 elements (subject to 1 adding up constraint).

- 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 $\hat{m}_n(2; \Theta)$ 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 $\hat{m}(2; \Theta)$ and $m(2)$, each with 448 elements (subject to 112 adding-up constraints).

- The third set of moments relate to the sales in France of exporting firms as discussed in Section 2.3. For firms selling in each of the 112 export destinations

from above. Helpman, Melitz, and Rubinstein (2008) accounted for zeros by truncating the upper tail of the Pareto distribution. Zero exports to a destination then arise simply because not even the most efficient firm could surmount the entry barrier there.

²⁴Notice, from (36), (37), 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.

n , we compute the q th percentile sales $s_n^q(3)$ in France for $q = 50, 75, 95$. Proceeding as above, we get $\hat{m}(3; \Theta)$ and $m(3)$, each with 448 elements (subject to 112 adding-up constraints).

- The fourth set of moments relate to normalized export intensity by market as 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 $\hat{m}(4; \Theta)$ 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:

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

By inserting the actual data on \bar{X}_{nF} (to get σE_{nF}) and N_{nF} (to get $\bar{u}_n(s)$) in our simulation routine, we are ignoring sampling error in these measures. Inserting \bar{X}_{nF} has no effect on our estimate of Θ . The reason is that a change in σE_{nF} shifts the sales in n of each artificial firm, the mean sales \bar{X}_{nF} , and the percentiles of sales in that market, all by the same proportion, leaving our moments unchanged. We only need \bar{X}_{nF} to get estimates of σE_{nF} given Θ . Our estimate of Θ does, however, depend on the data for N_{nF} . Appendix C reports on Monte Carlo simulations examining the sensitivity of our parameter estimates to this form of sampling error.²⁵

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

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

²⁵The coefficient of variation of N_{nF} is approximated by $1/\sqrt{N_{nF}}$, which is highest for Nepal, with 43 sellers, at 0.152. The median number of sellers is 686 (Malaysia), implying a coefficient of variation of 0.038.

where \mathbf{W} is a 1360×1360 weighting matrix.²⁶ We search for $\widehat{\boldsymbol{\Theta}}$ using the simulated annealing algorithm.²⁷ At each function evaluation involving a new value of $\boldsymbol{\Theta}$, we compute a set of 500,000 artificial firms and construct the moments for them as described above. The simulated annealing algorithm converges in 1–3 days on a standard PC.

We calculate standard errors using a bootstrap technique, taking into account both sampling error and simulation error.²⁸

²⁶The weighting matrix is the generalized inverse of the estimated variance–covariance matrix $\boldsymbol{\Omega}$ of the 1360 moments calculated from the data m . We calculate $\boldsymbol{\Omega}$ using the following bootstrap procedure: (i) We resample, with replacement, 229,900 firms from our initial data set 2000 times. (ii) 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 holding the $s_n^q(\tau)$ fixed to calculate $m^b(\tau)$ for $\tau = 2, 3, 4$. (iii) We calculate

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

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

²⁷Goffe, Ferrier, and Rogers (1994) described the algorithm. We use a version developed specifically for GAUSS by Goffe (1996).

²⁸To account for sampling error, each bootstrap b replaces m with a different m^b . To account for simulation error, each bootstrap b samples a new set of 500,000 v^b 's, a_n^b 's, and h_n^b 's as described in Section 4.2, thus generating a new $\widehat{m}^b(\boldsymbol{\Theta})$. (Just like m^b , \widehat{m}^b is calculated according to the bins defined from the actual data.) Defining $y^b(\boldsymbol{\Theta}) = m^b - \widehat{m}^b(\boldsymbol{\Theta})$, for each b , we search for

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

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

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

and take the square roots of the diagonal elements as the standard errors. Since we pursue our bootstrapping procedure only to calculate standard errors rather than to perform tests, we do not recenter the moments to account for the initial misfit of our model. Recentering would involve setting

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

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

4.5. Results

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

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

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

We turn to some implications of our parameter estimates.

Our discussion in Section 3.6 foreshadowed our estimate of $\tilde{\theta}$, which lies between the values implied by the slopes in Figure 3C and Figure 4. From equations (31), (29), and (30), 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 (21), our estimate of 2.46 implies that fixed costs dissipate 59 percent of gross profit in any destination.

Our estimate of σ_α implies enormous idiosyncratic variation in a firm's sales across destinations. In particular, the ratio of the 75th to the 25th percentile of the sales shock α is 9.78. In contrast, our estimate of σ_η means much less idiosyncratic variation in the entry shock η , with the ratio of the 75th to 25th percentile equal to 1.58. Given σ_α and σ_η , the variance of sales within a market decreases in ρ , as can be seen from equation (38). Hence the negative estimate of ρ reflects high variation of sales in a market.

A feature of the data is the entry of firms into markets where they sell very little, as seen in Figure 1C. Two features of our estimates reconcile these small sales with a fixed cost of entry. First, our estimate of λ , which is close to 1, means that a firm that is close to the entry cutoff incurs a very small entry cost.²⁹ 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 feature applies to firms that differ systematically in their efficiency, while the second applies to the luck of the draw in individual markets.

²⁹ Arkolakis (2010) found a value around 1, consistent with various observations from several countries.

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 from our model, we use our parameter estimates $\hat{\Theta}$ to simulate a set of artificial firms including nonexporters.³⁰ We then compare four features of these artificial firms with corresponding features of the actual firms.

Entry

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

Sales in a Market

Equation (31) motivates Figure 5A, which plots the simulated (\times 's) and the actual (circles) values of the median and the 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 the variation of each percentile 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 (32) motivates Figure 5B, which plots the median and the 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 datum point for France is below the line, reflecting the role of λ . The model understates noise in these percentiles in markets served by a small number of French firms.

³⁰Here 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 $\bar{u}(s) = \bar{u}_F(s)$.

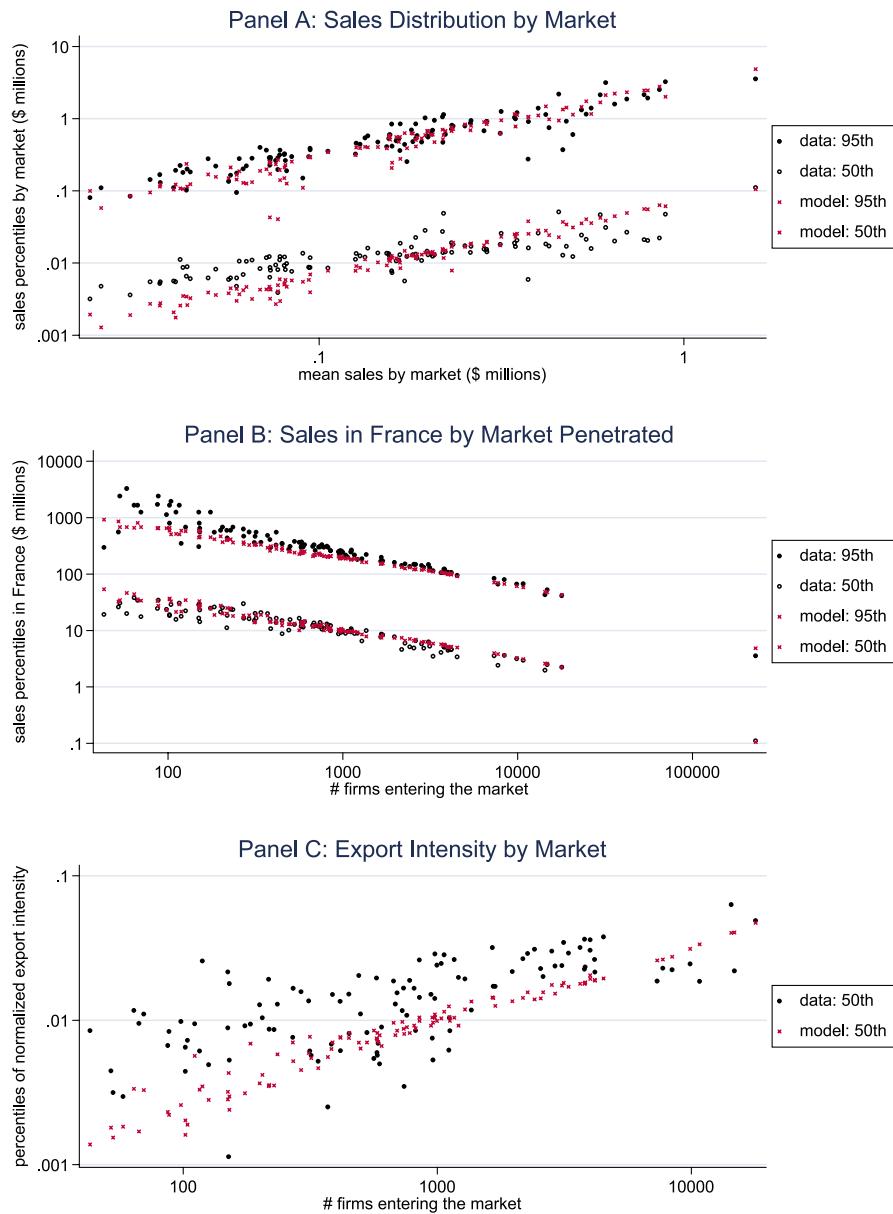


FIGURE 5.—Model versus data.

Export Intensity

Equation (33) motivates Figure 5C, which plots median normalized export intensity in each market against the actual number of French firms selling there. The model picks up the low magnitude of normalized export intensity and how it varies with the number of firms selling in a market. Despite our high estimate of σ_α , however, the model understates the noisiness of the relationship.

4.7. Sources of Variation

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

4.7.1. Variation in Entry

We first calculate the fraction of the variance of entry in each market that can be explained by the cost draw u alone. A natural measure (similar to R^2 in a regression) of the explanatory power of the firm's cost draw for market entry is

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

where V_n^U is the variance of entry and $E[V_n^C(u)]$ is the expected variance of entry given the firm's standardized unit cost u .³¹

We simulated the term $E[V_n^C(u)]$ using the techniques employed in our estimation routine with 230,000 simulated firms and obtained a value of R_n^E for each of our 112 export markets. Looking across markets, the results indicate that we can attribute on average 57 percent (with a standard deviation of 0.01) of the variation in entry in a market to the core efficiency of the firm rather than to its draw of η in that market.

³¹By the law of large numbers, the fraction of French firms selling in n approximates the probability that a French firm will sell in n . Writing this probability as $q_n = N_{nF}/N_{FF}$, the unconditional variance of entry for a randomly chosen French firm is $V_n^U = q_n(1 - q_n)$. 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 $\ln \eta_n$ is distributed $N(0, \sigma_\eta)$, the probability that this condition is satisfied is

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

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

4.7.2. Variation in Sales

Looking at the firms that enter a particular market, how much does variation in u explain sales variation? Consider firm j selling in market n . Inserting (29) into (26), the log of sales is

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

where we have divided sales into four components. Component 4 is common to all firms selling in market n , so it does not contribute to variation in sales there. The first component involves firm j 's idiosyncratic sales shock in market n , while component 3 involves its efficiency shock that applies across all markets. Complicating matters is component 2, which involves both firm j 's idiosyncratic entry shock in market n , $\eta_n(j)$, and its overall efficiency shock, $u(j)$. We deal with this issue by first asking how much of the variation in $\ln X_n(j)$ is due to variation in component 3 and then in the variation in components 2 and 3 together.

We simulate sales of 230,000 firms across our 113 markets and divide the contribution of each component to its sales in each market where it sells. Looking across markets, we find that component 3 itself contributes only 4.8 percent of the variation in $\ln X_{nF}(j)$, and components 2 and 3 together contribute around 39 percent (with very small standard deviations).³²

The dominant role of α is consistent with the shapes of the sales distributions in Figure 2. Even though the Pareto term u must eventually dominate the log-normal α at the very upper tail, such large observations are unlikely in a reasonably sized sample. With our parameter values, the log-normal term is evident in both the curvature and the slope, even among the biggest sellers.³³

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.³⁴

³²In comparison, Munch and Nguyen (2009) found that firm effects drive around 40 percent of the sales variation of Danish sellers across markets.

³³Our estimate of $\tilde{\theta}$ implies that the sales distribution asymptotes to a slope with absolute value of around 0.41, much lower than the values reported in footnote 7 among the top 1 percent of firms selling in a market. Sornette (2000, pp. 80–82) discussed how a log normal with a large variance can mimic a Pareto distribution over a very wide range of the upper tail. Hence α continues to play a role, even among the largest sellers in our sample.

³⁴Our finding that u plays a larger role in entry than in sales conditional on entry is consistent with our higher estimate of σ_α relative to σ_η . A lower value of $\tilde{\theta}$ (implying more sales hetero-

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 . To separate $\tilde{\theta}$ into these components, we turn to data on firm productivity.

A common observation is that exporters are more productive (according to various measures) than the average firm.³⁵ The same is true of our French exporters: The average value added per worker of exporters is 1.22 times the average for all firms. Moreover, value added per worker, like sales in France, tends to rise with the number of markets served, but not with nearly as much regularity.

A reason for this relationship in our model is that a more efficient firm typically both enters more markets and sells more widely in any given market. As its fixed costs are not proportionately higher, larger sales get translated into higher value added relative to inputs used, including those used in fixed costs. An offsetting factor is that more efficient firms enter tougher markets where sales are lower relative to fixed costs. Determining the net effect of efficiency on productivity requires a quantitative assessment.

We calculate productivity among our simulated firms as value added per unit of factor cost, proceeding as follows.³⁶ The value added $V_i(j)$ of firm j from country i is its gross production $Y_i(j) = \sum_n X_{ni}(j)$, less spending on intermediates $I_i(j)$:

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

We calculate intermediate spending as

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

where β is the share of factor costs in variable costs and $E_i(j) = \sum_n E_{ni}(j)$.³⁷ Value added per unit of factor cost is

$$(40) \quad q_i(j) = \frac{V_i(j)}{\beta\bar{m}^{-1}Y_i(j)} = \frac{[\bar{m} - (1 - \beta)] - \bar{m}[E_i(j)/Y_i(j)]}{\beta}.$$

genuity attributable to efficiency) for given σ_α and σ_η would lead us to attribute more variation both in entry and in sales to the firm's underlying efficiency.

³⁵See, for example, Bernard and Jensen (1995), Clerides, Lach, and Tybout (1998), and BEJK (2003).

³⁶In our model, value added per worker and value added per unit of factor cost are proportional since all producers in a country face the same wage and input costs, with labor having the same share.

³⁷We treat all fixed costs as purchased services. See EKK (2008) for a more general treatment.

The only potential source of heterogeneity is in the ratio $E_i(j)/Y_i(j)$.³⁸

The expression for firm productivity (40) depends on the elasticity of substitution σ (through $\bar{m} = \sigma/(\sigma - 1)$) independently of $\tilde{\theta}$. We can thus follow BEJK (2003) and find the σ that makes the productivity advantage of exporters in our simulated data match their productivity advantage in the actual data (1.22). To find the productivity of a firm in the simulated data, we take its sales $X_{ni}(j)$ and fixed cost $E_{ni}(j)$, using (13), in each market where it sells and then sum to get $Y_i(j)$ and $E_i(j)$.³⁹ We then calculate the average $q_i(j)$ for exporters relative to that for all firms. The σ that delivers the same advantage in the simulated data as the actual is $\sigma = 2.98$, implying $\beta = 0.34$. Using our estimate of $\tilde{\theta} = \theta/(\sigma - 1) = 2.46$, the implied value of θ is 4.87.⁴⁰

4.9. The Price Index Effect

With values for the individual parameters θ and σ , we can return to equation (14) and ask how much better off buyers are in a larger market. Taking

³⁸An implication of (40) is that the distribution of productivity of sellers in supplying any given market is invariant to any market-specific feature such as size, location, or trade openness. To see this result, replace $E_i(j)/Y_i(j)$ in (40) with $E_{ni}(j)/X_{ni}(j)$. Using (13) for the numerator and (26) for the denominator, exploiting (30), we can write this ratio in terms of $v_{ni}(j)$ as

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

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

³⁹We calibrate β from data on the share of manufacturing value added in gross production. Denoting the value-added share as s^V , averaging across data from the United Nations Industrial Development Organization (UNIDO (1999)) gives us $s^V = 0.36$. Taking into account profits and fixed costs, we calculate

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

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

⁴⁰While the model here is different, footnote 11 suggests that the parameter θ plays a similar role here to its role in Eaton and Kortum (2002) and in BEJK (2003). Our estimate of $\theta = 4.87$ falls between the estimates of 8.28 and 3.60, respectively, from those papers. Our estimate of $\sigma = 2.98$ is not far below the estimate of 3.79 from BEJK (2003). Using Chaney (2008) as a basic framework, Crozet and Koenig (2010) estimated θ and σ (along with a parameter δ that reflects the effect of distance D on trade costs) using French firm-level data on exports to countries contiguous to France. They employed a three-step procedure, estimating (i) a probit equation for firm entry, retrieving $\delta\theta$ from the coefficient $\ln D$, (ii) an equation that estimates firm sales, retrieving $\delta(\sigma - 1)$ from the coefficient $\ln D$, and (iii) a relationship between productivity and sales, retrieving $[\theta - (\sigma - 1)]$. Performing this procedure for 34 industries, they obtained estimates of θ ranging from 1.34 to 4.43 (mean 2.59) and estimates of σ ranging from 1.11 to 3.63 (mean 1.72).

E_{ni} 's as given, the elasticity of P_n with respect to X_n is $1/\theta - 1/(\sigma - 1) = -0.30$: Doubling market size leads to a 30 percent lower manufacturing price index.⁴¹

5. GENERAL EQUILIBRIUM AND COUNTERFACTUALS

We now use our framework to examine how counterfactual aggregate shocks in policy and the environment affect individual firms. To do so, we need to consider how such changes 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:

A1. 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 .

A2. Intermediates are as in [Eaton and Kortum \(2002\)](#). Consistent with Section 4.8, manufacturing inputs are a Cobb–Douglas combination of labor and intermediates, where intermediates have the price index P_i given in (14). Hence an input bundle costs

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

where β is the labor share.

A3. Nonmanufacturing is as in [Alvarez and Lucas \(2007\)](#). Final output, which is nontraded, is a Cobb–Douglas combination of manufactures and nonmanufactures, with manufactures having a share γ . Labor is the only input to nonmanufactures. Hence the price of final output in country i is proportional to $P_i^\gamma W_i^{1-\gamma}$.

A4. Fixed costs pay for labor in the destination. We thus decompose the country-specific component of the entry cost $E_{ni} = W_n F_{ni}$, where F_{ni} reflects the labor required for entry for a seller from i in n .⁴²

A5. Each country i 's manufacturing deficit D_i and overall trade deficit D_i^A are held at their 1986 values (relative to world GDP).

⁴¹Recall that our analysis in footnote 16 suggested that entry costs for French firms rise with market size with an elasticity of 0.35, attenuating this effect. Assuming that this elasticity applies to the entry costs E_{ni} for all sources i , a calculation using (14) and (16) still leaves us with an elasticity of $(1 - 0.35) * (-0.30) = -0.20$.

⁴²Combined with our treatment of fixed costs as intermediates in our analysis of firm-level productivity, A4 implies that these workers are outsourced manufacturing labor.

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

$$(41) \quad Y_i = \sum_{n=1}^N \pi_{ni} X_n.$$

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

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

$$(42) \quad Y_i = \frac{\gamma\sigma(W_i L_i + D_i^A) - \sigma D_i}{1 + (\sigma - 1)(\beta - \gamma/\theta)}.$$

Since $X_n = Y_n + D_n$,

$$(43) \quad X_n = \frac{\gamma\sigma(W_n L_n + D_n^A) - (\sigma - 1)(1 - \beta + \gamma/\theta)D_n}{1 + (\sigma - 1)(\beta - \gamma/\theta)}.$$

From expression (20), we can write

$$(44) \quad \pi_{ni} = \frac{T_i (W_i^\beta P_i^{1-\beta} d_{ni})^{-\theta} (\sigma F_{ni})^{-[\theta-(\sigma-1)]/(\sigma-1)}}{\sum_{k=1}^N T_k (W_k^\beta P_k^{1-\beta} d_{nk})^{-\theta} (\sigma F_{nk})^{-[\theta-(\sigma-1)]/(\sigma-1)}}.$$

Substituting (42), (43), and (44) into (41) gives us a set of equations that determine wages W_i given prices P_i . From expression (14), we have

$$(45) \quad P_n = \bar{m} \kappa_1^{-1/\theta} \left[\sum_{i=1}^N T_i (W_i^\beta P_i^{1-\beta} d_{ni})^{-\theta} (\sigma F_{ni})^{-[\theta-(\sigma-1)]/(\sigma-1)} \right]^{-1/\theta} \\ \times \left(\frac{X_n}{W_n} \right)^{(1/\theta)-1/(\sigma-1)},$$

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

5.2. Calculating Counterfactual Outcomes

We apply the method used in Dekle, Eaton, and Kortum (2008; henceforth DEK) to equations (41) and (45) to calculate counterfactuals.⁴³ 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

$$(46) \quad Y'_i = \sum_{n=1}^N \pi'_{ni} X'_n.$$

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 \hat{W}_i and prices \hat{P}_i using (42), (43), and (44) as

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

where sticking these three equations into (46) yields a set of equations involving \hat{W}_i 's for given \hat{P}_i 's. From (45), we can get a set of equations that involve \hat{P}_i 's for given \hat{W}_i 's:

$$(47) \quad \hat{P}_n = \left[\sum_{i=1}^N \pi_{ni} \hat{W}_i^{-\beta\theta} \hat{P}_i^{-(1-\beta)\theta} \hat{d}_{ni}^{-\theta} \hat{F}_{ni}^{-[\theta-(\sigma-1)]/(\sigma-1)} \right]^{-1/\theta} \left(\frac{\hat{X}_n}{\hat{W}_n} \right)^{(1/\theta)-1/(\sigma-1)}.$$

We implement counterfactual simulations for our 113 countries in 1986, aggregating the rest of the world (ROW) into a 114th country. 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 take Y_i^L and country-specific γ_i 's from data on GDP, manufacturing production, and trade. Appendix F describes our sources of data and our procedures for assembling them to execute the counterfactual.

⁴³DEK (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 the rest of the world).

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 the numéraire, as $\widehat{X}_{nF} = \widehat{\pi}_{nF}\widehat{X}_n/\widehat{P}_F$ and (ii) the change in the number of French firms selling there, using (27), as $\widehat{N}_{nF} = (\widehat{\pi}_{ni}\widehat{X}_n)/(\widehat{W}_n\widehat{F}_{nF})$.

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

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

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

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

- (b) In Step 4, we use our baseline values X_{nF} and N_{nF} to calculate baseline $\bar{u}_n(s)$'s, and use 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 (28) and (36).

- (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.

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

5.3. Counterfactual: Implications of Globalization

We consider a 10 percent drop in trade barriers, that is, $\widehat{d}_{ni} = 1/(1.1)$ for $i \neq n$ and $\widehat{d}_{nn} = \widehat{F}_{ni} = 1$. This change roughly replicates the increase in French import share over the decade following 1986.⁴⁴ Table III shows the aggregate

⁴⁴Using time-series data from the Organization for Economic Cooperation and Development's (OECD's) STAN data base, 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 through 2000. By 1997, this share had risen for all 16 countries,

TABLE III
AGGREGATE OUTCOMES OF COUNTERFACTUAL EXPERIMENT

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

(Continues)

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.

TABLE III—*Continued*

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

(Continues)

TABLE III—*Continued*

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

^aBelgium includes Luxembourg.

general equilibrium consequences of this change: (i) the change in the real wage $\widehat{\omega}_n = (\widehat{W}_n/\widehat{P}_n)^{\gamma_n}$, (ii) the change in the relative wage $\widehat{W}_n/\widehat{W}_F$, (iii) the change in the sales of French firms to each market \widehat{X}_{nF} , and (iv) the change

TABLE IV
COUNTERFACTUALS: FIRM TOTALS^a

	Baseline	Counterfactual	
		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

^aCounterfactual simulation of a 10% decline in trade costs.

in the number of French firms selling to each market \hat{N}_{nF} . Lower trade barriers raise the real wage in every country, typically by less than 5 percent.⁴⁵ Relative wages move quite a bit more, capturing terms of trade effects from globalization.⁴⁶ The results that matter at the firm level are French sales and the number of French firms active in each market. While French sales decline by 5 percent in the home market, exports increase substantially, with a maximum 80 percent increase in Japan.⁴⁷ The number of French exporters increases roughly in parallel with French exports.⁴⁸

Table IV 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,

⁴⁵There are 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 entrepôt trade or (for Liberia) ships. These anomalies have little consequence for our overall results.

⁴⁶A 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.

⁴⁷A 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 that captures the change in French cost advantage relative to domestic producers in n is $x_1 = \pi_{nn}[\beta \ln(\hat{W}_n/\hat{W}_F) - \ln d_{nF}]$ with a predicted elasticity of θ (we ignore changes in the relative value of the manufacturing price index, since they are small). The variable that captures the percentage change in n 's absorption is $x_2 = \ln \hat{W}_n$ 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.

⁴⁸We can compare these counterfactual outcomes with the actual change in the number of French sellers in each market between 1986 and 1992. We tend to overstate the increase and to underestimate heterogeneity across markets. The correlation between the counterfactual change and the change in the data is nevertheless 0.40. Our counterfactual, treating all parameters as given except for the iceberg trade costs, obviously misses much of what happened in those 6 years.

TABLE V
COUNTERFACTUALS: FIRM ENTRY AND EXIT BY INITIAL SIZE

Initial Size Interval (percentile)	All Firms				Exporters	
	Baseline No. of Firms	Counterfactual		Baseline No. of Firms	Counterfactual	
		Change From Baseline	Change (%)		Change From Baseline	Change (%)
Not active	0	1118	—	0	1118	—
0 to 10	23,140	-11,551	-49.9	767	15	2.0
10 to 20	23,140	-5702	-24.6	141	78	55.1
20 to 30	23,140	-3759	-16.2	181	192	106.1
30 to 40	23,140	-2486	-10.7	357	357	100.0
40 to 50	23,140	-1704	-7.4	742	614	82.8
50 to 60	23,138	-1141	-4.9	1392	904	65.0
60 to 70	23,142	-726	-3.1	2450	1343	54.8
70 to 80	23,140	-405	-1.8	4286	1829	42.7
80 to 90	23,140	-195	-0.8	7677	2290	29.8
90 to 99	20,826	-38	-0.2	12,807	1915	15.0
99 to 100	2314	0	0.0	2169	62	2.8
Total	231,402	-26,589		32,969	10,716	

competition from imports drives almost 27,000 firms out of business, although almost 11,000 firms start exporting.

Tables V and VI decompose these changes into the contributions of firms of different baseline size, with Table V 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 VI decomposes sales revenues. All of the increase is in the top decile, and most of that in the top percentile. For every other decile, sales decline. Almost two-thirds of the increase in export revenue is from the top percentile, although lower deciles experience much higher percentage increases in their export revenues.

Comparing the numbers in Tables V and VI 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.⁴⁹

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

⁴⁹The 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.

TABLE VI
COUNTERFACTUALS: FIRM GROWTH BY INITIAL SIZE

Initial Size Interval (percentile)	Baseline (\$ millions)	Total Sales		Exports		
		Counterfactual		Baseline (\$ millions)	Counterfactual	
		Change From Baseline	Change (%)		Change From Baseline	Change (%)
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	1793	-476	-26.6	6	18	307.8
50 to 60	3299	-712	-21.6	18	48	269.7
60 to 70	6188	-1043	-16.9	58	130	223.0
70 to 80	12,548	-1506	-12.0	206	391	189.5
80 to 90	31,268	-1951	-6.2	1085	1501	138.4
90 to 99	148,676	4029	2.7	16,080	11,943	74.3
99 to 100	230,718	18,703	8.1	56,301	20,486	36.4
Total	436,144	16,442		73,758	34,534	

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 (2010) would predict, sales by firms with an initially smaller presence grow substantially more than those at the top.⁵⁰

6. CONCLUSION

A literature that documents the superior performance of exporters, notably Bernard and Jensen (1995), inspired a new generation of trade theories that incorporate producer heterogeneity. These theories, in turn, deliver predictions beyond the facts that motivated them. To account for the data presented here,

⁵⁰We can examine another counterfactual analytically: a uniform change \hat{F} in labor required for entry. Proportional changes in \hat{F} cancel out in the expression for π'_{ni} , so that the only action is in equation (47) for the manufacturing price index:

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

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

which break down firms' foreign sales into individual export destinations, these theories are confronted with a formidable challenge.

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

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

REFERENCES

- ALVAREZ, F., AND R. E. LUCAS (2007): "A General Equilibrium Analysis of the Eaton–Kortum Model of International Trade," *Journal of Monetary Economics*, 54, 1726–1768. [1486]
- ARKOLAKIS, K. (2010): "Market Penetration Costs and the New Consumers Margin in International Trade," *Journal of Political Economy*, 118, 1151–1199. [1454,1463,1479,1495]
- AW, B. Y., S. CHUNG, AND M. J. ROBERTS (2000): "Productivity and Turnover in the Export Market: Micro-Level Evidence From the Republic of Korea and Taiwan (China)," *World Bank Economic Review*, 14, 65–90. [1453]
- BERNARD, A. B., AND J. B. JENSEN (1995): "Exporters, Jobs, and Wages in U.S. Manufacturing: 1976–1987," *Brookings Papers on Economic Activity: Microeconomics*, 1995, 67–119. [1453,1484, 1495]
- BERNARD, A. B., J. EATON, J. B. JENSEN, AND S. KORTUM (2003): "Plants and Productivity in International Trade," *American Economic Review*, 93, 1268–1290. [1453,1462,1484,1485]
- BISCOUPR, P., AND F. KRAMARZ (2007): "Employment, Skill Structure, and International Trade: Firm-Level Evidence for France," *Journal of International Economics*, 72, 22–51. [1455]
- CHANAY, T. (2008): "Distorted Gravity: Heterogeneous Firms, Market Structure, and the Geography of International Trade," *American Economic Review*, 98, 1707–1721. [1454,1462,1485]
- CLERIDES, S., S. LACH, AND J. R. TYBOUT (1998): "Is Learning by Exporting Important? Micro-Dynamic Evidence From Colombia, Mexico, and Morocco," *Quarterly Journal of Economics*, 113, 903–947. [1484]
- CROZET, M., AND P. KOENIG (2010): "Structural Gravity Equations With Intensive and Extensive Margins," *Canadian Journal of Economics*, 43, 41–62. [1485]
- DEKLE, R., J. EATON, AND S. KORTUM (2008): "Global Rebalancing With Gravity," *IMF Staff Papers*, 55, 511–540. [1488]
- DIXIT, A., AND J. E. STIGLITZ (1977): "Monopolistic Competition and Optimum Product Diversity," *American Economic Review*, 67, 297–308. [1464]
- EATON, J., AND S. KORTUM (1999): "International Technology Diffusion: Theory and Measurement," *International Economic Review*, 40, 537–570. [1462]
- (2002): "Technology, Geography, and Trade," *Econometrica*, 70, 1741–1780. [1453,1462, 1485,1486]
- EATON, J., S. KORTUM, AND F. 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.) [1455]

- (2008): “An Anatomy of International Trade: Evidence From French Firms,” Working Paper 14610, NBER. [1484]
- (2011): “Supplement to ‘An Anatomy of International Trade: Evidence From French Firms,’ *Econometrica Supplemental Material*, 79, http://www.econometricsociety.org/ecta/Supmat/8318_miscellaneous.pdf; http://www.econometricsociety.org/ecta/Supmat/8318_data_and_programs-1.zip; http://www.econometricsociety.org/ecta/Supmat/8318_data_and_programs-2.zip. [1455]
- GABAIX, X. (1999): “Zipf’s Law for Cities: An Explanation,” *Quarterly Journal of Economics*, 94, 739–767. [1462]
- GABAIX, X., AND R. IBRAGIMOV (2011): “Rank-1/2: A Simple Way to Improve the OLS Estimation of Tail Exponents.” *Journal of Business & Economic Statistics*, 29, 24–39. [1458]
- GOFFE, W. L. (1996): “SIMANN: A Global Optimization Algorithm Using Simulated Annealing,” *Studies in Nonlinear Dynamics and Econometrics*, 1, Algorithm 1. Available at <http://www.bepress.com/snnde/vol1/iss3/algorithm1>. [1478]
- GOFFE, W. L., G. D. FERRIER, AND J. ROGERS (1994): “Global Optimization of Statistical Functions With Simulated Annealing,” *Journal of Econometrics*, 60, 65–99. [1478]
- GORIÉROUX, C., AND A. MONFORT (1994): *Simulation Based Econometric Methods*. CORE Lecture Series. Louvain-la-Neuve: Université catholique de Louvain. [1475]
- HELPMAN, E., M. J. MELITZ, AND Y. RUBINSTEIN (2008): “Estimating Trade Flows: Trading Partners and Trading Volumes,” *Quarterly Journal of Economics*, 123, 441–487. [1476]
- HELPMAN, E., M. J. MELITZ, AND S. R. YEAPLE (2004): “Exports vs. FDI With Heterogeneous Firms,” *American Economic Review*, 94, 300–316. [1454,1462]
- HOROWITZ, J. (2001): “The Bootstrap,” in *Handbook of Econometrics*, Vol. 5, ed. by J. J. Heckman and E. E. Leamer. Amsterdam: Elsevier Science B.V., 3159–3228, Chapter 52. [1478]
- KORTUM, S. (1997): “Research, Patenting, and Technological Change,” *Econometrica*, 65, 1389–1419. [1462]
- KRUGMAN, P. R. (1980): “Scale Economies, Product Differentiation, and the Patterns of Trade,” *American Economic Review*, 70, 950–959. [1466]
- LUTTMER, E. (2007): “Selection, Growth, and the Size Distribution of Firms,” *Quarterly Journal of Economics*, 122, 1103–1144. [1462]
- MELITZ, M. J. (2003): “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, 71, 1695–1725. [1453,1454,1462,1463,1496]
- MUNCH, J. R., AND D. X. NGUYEN (2009): “Decomposing Firm-Level Sales Variation,” Working Paper 2009-05, EPRU. [1483]
- OECD (1995): The OECD STAN Database. [1489]
- PEDERSEN, N. K. (2009): “Essays in International Trade,” Ph.D. Thesis, Northwestern University, Evanston, IL. [1470]
- RICARDO, D. (1821): *On The Principles of Political Economy and Taxation*. London: Murray. [1486]
- ROBERTS, M., AND J. R. TYBOUT (1997): “The Decision to Export in Colombia: An Empirical Model of Entry With Sunk Costs,” *American Economic Review*, 87, 545–564. [1454]
- SIMON, H. A. (1955): “On a Class of Skew Distribution Functions,” *Biometrika*, 42, 425–440. [1462]
- SIMON, H. A., AND C. P. BONINI (1958): “The Size Distribution of Business Firms,” *American Economic Review*, 98, 607–617. [1458]
- SORNETTE, D. (2000): *Critical Phenomena in Natural Sciences*. Berlin: Springer. [1483]
- UNIDO (1999): Industrial Statistics Database. [1485]

Dept. of Economics, Pennsylvania State University, 608 Kern Graduate Building, University Park, PA 16802, U.S.A.; jxe22@psu.edu,

Dept. of Economics, University of Chicago, 1126 East 59th Street, Chicago, IL 60637, U.S.A.; kortum@uchicago.edu,

and

CREST (ENSAE), 15 Boulevard Gabriel Péri, 92245, Malakoff, France;
kramarz@ensae.fr.

Manuscript received December, 2008; final revision received December, 2010.