Plants and Productivity in International Trade: A Ricardian Reconciliation[1]

Andrew B. Bernard***, Jonathan Eaton*, J. Bradford Jensen**,
and Samuel S. Kortum*

*Boston University, **Carnegie-Mellon University, ***Yale School of Management

February 1999

[1]Send correspondence to Department of Economics, Boston University, 270 Bay State Road, Boston, MA 02215, USA; jeaton@bu.edu, kortum@bu.edu. We thank Daniel Ackerberg for helpful comments. Eaton and Kortum gratefully acknowledge the support of the National Science Foundation. Bernard gratefully acknowledges the support of the Frederick Frank Fund for International Research.

Abstract

We attempt to reconcile observations about plant level exporting behavior with international trade theory and the aggregate volume of trade. Our framework extends the basic Ricardian model to accommodate trade among many countries separated by trade frictions in an environment of imperfect competition. We fit the model to data on bilateral trade. We then see how well we can explain several basic facts about U.S. plants: (i) the wide dispersion of measured productivity across plants, (ii) the higher measured productivity of exporters, (iii) the small fraction of active producers who actually export, (iv) the small fraction of revenue earned from exports even by those plants that do export, and (v) the much larger size of plants that export. We pick up all the basic qualitative features, and go quite far in matching relevant magnitudes.

Key words: Ricardian trade, exporting, productivity, Census of Manufactures

JEL classification: F11; F17; O33

1 Introduction

A tradition in trade theory has been to associate goods with industries. Consequently, most empirical work in international trade treats industries as the smallest unit of analysis. Looking inside these industries, however, we observe an enormous variety of behavior and performance. A new empirical literature has emerged examining trade at the level of individual producers. Bernard and Jensen (1995, 1999) and Roberts and Tybout (1997), among others, have uncovered provocative facts about productivity and exporting, e.g., exporters tend to be more productive, most firms do not export, and even those that do typically export only a small fraction of their production. Until now there has been no attempt to reconcile such firm or plant level facts with international trade theory, and with what we know about aggregate bilateral trade flows and their determinants.

This paper attempts a theoretical and empirical integration of behavior at the micro and macro level. Our theory incorporates heterogeneous individual producers in an aggregate Ricardian model of international trade. We deliberately keep the model quite simple to highlight its limitations and power. Our empirics tie together observations on plant export behavior with data on aggregate trade volumes.

The theory is based on Eaton and Kortum (1998b) (Henceforth EK), which itself extends Dornbusch, Fischer, and Samuelson (1977) to incorporate an arbitrary number of countries with potential trade frictions between them. Here we extend it further to allow for imperfect competition. This extension turns out to be essential for understanding the heterogeneity we observe in productivity across individual producers.

We fit the model to data on aggregate bilateral trade. We then see what it predicts about five plant-level facts extracted from the U.S. Census of Manufactures: (i) the wide distribution of productivity across active producers, (ii) the higher productivity of exporters, (iii) the small fraction of active producers who actually export, (iv) the small fraction of revenue earned from

exports even by those plants that do export, and (v) the much larger size of plants that export. An important feature of our analysis is that it ignores industry classifications and attempts to explain features of the population of all manufacturing plants.[1]

Spare as it is, our analytic framework predicts all these features qualitatively and even captures the relevant magnitudes quite well. We model the space of products as a continuum, with each country having a set of potential producers for each product. To capture producer heterogeneity in productivity levels, we treat the efficiencies of these producers as if drawn independently from country specific probability distributions, normalizing units so that these distributions are the same across goods.[2]

Only the lowest cost producer (taking trade frictions as well as factor prices and efficiency into account) ever sells in a country. Because of trade barriers, selling abroad sets a higher hurdle than selling at home. Hence a producer that exports is likely to be more efficient than its nonexporting counterparts (since a producer that beats out its competition in foreign markets is also likely to outshine its domestic counterparts making different goods only for the home market). This generates the higher average productivity of exporters.

In the absence of trade barriers in the model all producers would export and most goods would be imported. The relatively small percentage of output that is exported suggests the need for significant bilateral trade barriers.

Under Bertrand competition, a more efficient producer is likely to charge a lower price than its less efficient counterpart. With elastic demand exporters tend to sell more even domestically than nonexporters, yielding the substantial size advantage of exporters over non-exporters. Elastic demand combined with trade barriers produces the result that exporters will usually sell less abroad than at home.

[1]Our methodology is related to Buldyrev et. al. (1997) who seek to explain features of the distribution of firm size and Klette and Griliches (1998) who seek to explain features of the joint distribution firm size, growth, and R&D investment.

[2]We do recognize that modelling the origins of productivity differences is important. See Kortum (1997) and Eaton and Kortum (1998a) for example.

One potential problem is the link between efficiency (Hicks' neutral) and productivity as it appears in data (typically value added per worker). In fact, as long as all producers in a country employ inputs in the same proportion at the same cost, under perfect competition they would all appear equally productive. In our model, differences in value added per unit of input reflect different markups of price over cost, which emerge through Bertrand competition. In the absence of any link between efficiency and markup, measured productivity would say nothing about underlying efficiency. It turns out, however, that producers who are more efficient also tend to have a greater cost advantage over their closest competition, and are able to set higher markups. These links from underlying efficiency to exporting and size, on the one hand, and from underlying efficiency to measured productivity, on the other, generate the facts that we observe.[3]

We proceed as follows. Section 2 presents the theoretical model. Section 3 then describes how we fit it to bilateral trade data and simulate it to predict what goes on at the plant level. Section 4 then compares our predictions with the facts.

2 Theory

Our model combines the Ricardian theory of comparative advantage driven by technology differences with imperfect competition.

2.1 A Ricardian Framework

We describe a world in which N countries can produce any of a continuum of goods, indexed along the interval $[0, 1]$, but with varying levels of efficiency. Aside from (Hicks' neutral) efficiency differences, production of any good anywhere combines inputs in the same way.

Within a country there are many potential producers of any good j, but only the most

[3]An extensive literature compares productivity levels across plants. See, e.g., Baily, Hulten, and Campbell (1992) and Olley and Pakes (1996). The second maintains an assumption that the plants in question produce a homogeneous output. Our framework shows how such comparisons make sense even when outputs are heterogeneous.

efficient one will ever be in business. This producer can convert a bundle of inputs into a quantity $Z_{1i}(j)$ of good j. The term $Z_{1i}(j)$ observed across countries i reflects their relative efficiencies in producing good j. If the cost of an input bundle in country i is v_i (which does not vary with j) then it costs $v_i/Z_{1i}(j)$ to produce a unit of good j in country i using the cheapest local means.[4]

Goods can be transported between countries, but at a cost. We make the standard iceberg assumption that delivering one unit of a good in country n requires a shipment of $d_{ni} \geq 1$ units, and we normalize $d_{ii} = 1$ for all i. Hence the lowest cost in country n of obtaining good j from country i is:

$$C_{1ni}(j) = \frac{v_i d_{ni}}{Z_{1i}(j)}.$$

We denote the lowest cost of obtaining good j in country n irrespective of source by:

$$C_{1n}(j) = \min_i \{C_{1ni}(j)\}.$$

Preferences and technologies everywhere combine goods with a constant elasticity of substitution $\sigma > 0$, with good j having weight $\alpha(j)^{1/\sigma}$. Hence expenditure on good j, in country n, $X_n(j)$, is:

$$X_n(j) = \alpha(j) X_n \left(\frac{P_n(j)}{P_n} \right)^{1-\sigma}, \tag{1}$$

where $P_n(j)$ is the price of good j in country n, X_n denotes total expenditure, and $P_n = \left[\int_0^1 \alpha(k) P_n(k)^{1-\sigma} dk \right]^{1/(1-\sigma)}$ is the appropriate price index for country n. (Except in the Cobb-Douglas case, $\alpha(j)$ depends on the units in which j is measured.) Without loss of generality we restrict $\int_0^1 \alpha(j) dj = 1$.

To extract any implications from this model we need to specify how efficiencies are distributed. In general, the fact that we are forced to take minima across a large number of

[4]By assuming that v_i is the same for all goods, we ignore possible factor intensity differences across goods, thus precluding any Heckscher-Ohlin explanation for trade patterns. We could generalize our framework to incorporate differences in factor intensities and in factor endowments as a determinant of specialization, but our goal here is to see how far we get with a purely Ricardian theory.

sources makes the model potentially cumbersome. To cut through this problem EK take the distribution of the frontier efficiencies for producing good j in country i to be extremal:[5]

$$F_{1i}(z_1) \equiv \Pr[Z_{1i}(j) \leq z_1] = e^{-T_i z_1^{-\theta}}. \tag{2}$$

We choose units so that the distributions are identical across goods j (which has the implication that the distributions are independent of the resulting $\alpha(j)$'s). The parameter $T_i > 0$ reflects the average level of efficiency in country i, or its absolute advantage, while $\theta > 1$ determines the variation of efficiencies, which governs comparative advantage.

We can move easily from the distribution of country i's frontier efficiency to the distribution of the lowest cost of delivering a good from country i to country n, which is:

$$G_{1ni}(c_1) = 1 - e^{-\phi_i d_{ni}^{-\theta} c_1^{\theta}}, \tag{3}$$

where $\phi_i = T_i v_i^{-\theta}$ deflates country i's overall efficiency by the cost of paying inputs there. Taking the minimum cost across all potential suppliers, the distribution of lowest cost in country n (irrespective of source), is:

$$G_{1n}(c_1) = 1 - \prod_i [1 - G_{1ni}(c_1)] = 1 - e^{-\Phi_n c_1^{\theta}}. \tag{4}$$

where $\Phi_n \equiv \sum_{i=1}^{N} \phi_i {d_{ni}}^{-\theta}$, which aggregates overall efficiencies across potential suppliers, deflating by input costs and the cost of delivering to country n.

In what follows below we exploit two useful results from EK: (i) The probability π_{ni} that country i is the low cost supplier to n is just its share of Φ_n, that is $\pi_{ni} = \phi_i d_{ni}^{-\theta}/\Phi_n$. (ii) The distribution $G_{1n}(c)$ applies not only to the cost of a good supplied to country n regardless of source, but also to the cost of a good conditional on the source. That is, once transport costs

[5] One explanation for expecting the distribution of leading-edge efficiencies in a country to be extremal is that the market selects the best from a large number of potential producers of any specific good. If the efficiencies of potential producers of a good are thought of as random draws from some distribution, then the highest efficiency (suitably normalized) converges to an extremal distribution. The Fréchet distribution (2) is one of the three limiting functional forms for extremal distributions, and the only one for which heterogeneity does not vanish in the limit (see Billingsley, 1986). Kortum (1997) and Eaton and Kortum (1998a) provide other models leading to the Fréchet distribution.

are taken into account, no exporting country has a systematic cost advantage over any other *in terms of what it actually sells.* Instead, countries that have more advanced technology, lower input costs, or lower costs of delivery to a market, exploit their greater competitiveness there by exporting a wider range of products.

EK assume perfect competition, so that $P_n(j) = C_{1n}(j)$. With perfect competition, however, higher productivity is fully reflected in lower prices, so that the value of production per unit of input across firms in a country is the same irrespective of their efficiency Z_1. Since a major feature of our data is enormous variation across plants in measured productivity (in particular, value added per worker) we are forced to abandon the assumption of perfect competition.

2.2 Adding Imperfect Competition

The most natural alternative to perfect competition is Bertrand competition: As with perfect competition, each market n is captured by the lowest cost producer. But this source charges the cost of the next lowest cost supplier, which we denote $C_{2n}(j)$, unless the resulting markup exceeds the monopoly markup:

$$\overline{m} = \begin{cases} \sigma/(\sigma-1) & \sigma > 1 \\ \infty & \sigma \leq 1. \end{cases}$$

Hence:

$$P_n(j) = \min\{C_{2n}(j), \overline{m}C_{1n}(j)\}.$$

We must now keep track not only of the lowest cost supplier, but also the second lowest. This second lowest supplier could be either: (i) the second lowest cost producer in the same country as the lowest cost provider, or else (ii) the lowest cost provider from some other country. Hence, if country i turns out to be the low cost provider for country n:

$$\begin{aligned} C_{1n}(j) &= C_{1ni}(j) \\ C_{2n}(j) &= \min\left\{\min_{k\neq i}\{C_{1nk}(j)\}, C_{2ni}(j)\right\}. \end{aligned} \tag{5}$$

The second most efficient potential producer of good j in country i can make $Z_{2i}(j)$ units with a bundle of inputs. The analogue to the Fréchet for the limiting joint distribution of the first and second best is:

$$F_i(z_1, z_2) \equiv \Pr(Z_{1i}(j) \leq z_1, Z_{2i}(j) \leq z_2) = [1 + T_i(z_2^{-\theta} - z_1^{-\theta})]e^{-T_i z_2^{-\theta}},$$

for $0 \leq z_2 \leq z_1$. (Note that setting $z_1 = z_2$ returns (2).)

With trade, the resulting joint distribution of first and second costs in country n turns out to be:

$$G_n(c_1, c_2) \equiv \Pr(C_{1n}(j) \leq c_1, C_{2n}(j) \leq c_2) = 1 - e^{-\Phi_n c_1^{\theta}} - \Phi_n c_1^{\theta} e^{-\Phi_n c_2^{\theta}},$$

for $c_1 \leq c_2$. (Note that letting c_2 approach infinity returns (4).)

Since EK assume perfect competition, the distribution of lowest costs (4) is also their distribution of prices. In our analysis goods are sold at a markup $M_n(j) \equiv P_n(j)/C_{1n}(j)$. These have a Pareto distribution truncated at the monopoly markup:

$$H_n(m) = \begin{cases} 1 - m^{-\theta} & 1 \leq m < \overline{m} \\ 1 & m \geq \overline{m}. \end{cases} \tag{6}$$

Note that $\Pr[M_n(j) = \overline{m}] = \overline{m}^{-\theta}$. While the distribution of costs (and hence prices) differ by country, the distribution of markups does not depend on n,and thus is the same everywhere. Furthermore, since the distribution of costs (including transportation) of goods actually sold is the same whether or not we condition on source, so is the distribution of markups. That is, no country sells at systematically higher markups than another.

The distribution of prices under imperfect competition is complicated since it depends on both the distribution of costs and the distribution of the markup. Exploiting the result that the markup $M_n(j)$ is independent of second lowest cost $C_{2n}(j)$ we obtain an expression for the exact price index in country n:

$$P_n = \gamma \Phi_n^{-1/\theta}, \tag{7}$$

where

$$\gamma \equiv \left[\frac{1+(\sigma-1)\bar{m}^{-\theta}}{1+\theta-\sigma}\Gamma\left(\frac{1+2\theta-\sigma}{\theta}\right)\right]^{1/(1-\sigma)}$$

Two important properties of the price index are that: (i) It is only valid for $\sigma < 1+\theta$, as can be seen from the definition of γ. (ii) Variation in the price index across countries is completely captured in the term $\Phi_n^{-1/\theta}$. As long as we obey the parameter restriction that $\sigma < 1+\theta$, the complicated constant γ is irrelevant for anything we do empirically.

2.3 Efficiency vs. Measured Productivity

We now turn to the implications of these assumptions for observations on productivity. To simplify things we first consider a nonexporting producer. Defining $E_i(j)$ as total expenditure on inputs employed in making good j in country i, a natural measure of the real physical inputs used is $I_i(j) = E_i(j)/v_i$. If good j is actually produced in country i, physical productivity or efficiency is then $Z_{1i}(j) = Y_i(j)/I_i(j)$, where $Y_i(j)$ is the quantity produced. Comparing this productivity measure across commodities only makes sense given our particular normalization. Since any actual data measure goods in arbitrary units, such comparisons are meaningless.

In fact, output is usually measured in units of value. Measured productivity is then $y_i(j) = P_i(j)Y_i(j)/I_i(j)$. Under perfect competition, $P_i(j) = v_i/Z_{1i}(j)$, so that $y_i(j) = v_i$, which is the same for all j, in striking contrast with the data. With perfect competition the value-based measure does not capture efficiency differences across producers, but only the overall level of efficiency of the economy, as reflected in input costs, such as wages.

Under monopolistic competition, however, $P_i(j) = M_i(j)v_i/Z_{1i}(j)$, so that measured productivity is given by:

$$y_i(j) = M_i(j)v_i. \tag{8}$$

Differences across producers in measures of productivity can now emerge, but these reflect differences in markups, not efficiency. In the absence of any connection between the two, value-based productivity measures provide information only about differences across producers

in monopoly power, not in relative efficiency.

All is not lost, however. In fact, our model does imply that, on average, firms that are more efficient can charge a higher markup. Conditional on efficiency z, the distribution of the markup $M_i(j)$ is:

$$H_i(m|z) \equiv \Pr[M_i(j) \leq m | Z_{1i}(j) = z] = \begin{cases} 1 - e^{-\Phi_i v_i^\theta z^{-\theta}(m^\theta - 1)} & 1 \leq m < \overline{m} \\ 1 & m \geq \overline{m}. \end{cases}$$

Note that $\Pr[M_i(j) = \overline{m} | Z_{1i}(j) = z] = e^{-\Phi_i v_i^\theta z^{-\theta}(\overline{m}^\theta - 1)}$. A producer with higher efficiency than another is likely to have a higher markup (its distribution of M stochastically dominates the other's) and hence higher measured productivity.

Looking at the relationship the other way around, how does underlying efficiency Z_1 vary with measured productivity y? The conditional expectation is proportional to y (as long as the markup is less than $\overline{m}$). Hence, a firm appearing to be 2 per cent more productive than another is, *on average* , 2 per cent more efficient (unless it is already charging the monopoly markup, in which case expected efficiency is even greater).

For producers who export, these results on measured productivity and efficiency go through as long as productivity is measured as domestic sales relative to inputs employed to make output for the domestic market. Unfortunately, data do not usually segregate inputs according to the destination of the output they make. Since firms will typically charge different markups in different markets, the relationship between measured productivity and efficiency is more complicated, as we show below. Nevertheless, as we verify in our simulations, there is a strong tendency for higher underlying efficiency to lead to higher markups in all markets, and to higher measured productivity.

2.4 Efficiency, Who Exports, and the Correlation Between Exporting and Productivity

How hard is it for a producer in country i with efficiency z to penetrate country n's market? First, it has to be the most efficient potential producer in country i, which, from (2), it is

with probability $\exp(-T_i z^{-\theta})$. A tougher test is actually selling at home, which means beating out foreign competition as well. This it does with probability $\exp[-\Phi_i v_i^{\theta} z^{-\theta}]$. In general, it can sell in country n with probability $\exp[-\Phi_n (v_i d_{ni})^{\theta} z^{-\theta}]$. Selling abroad is more demanding than selling at home under the plausible "triangle inequality" restriction on trade frictions $d_{nk} \leq d_{ni} d_{ik}$ for all countries k: that is, it is never cheaper to move goods from k to n through i than to move them there directly. Under this restriction any firm efficient enough to export will also sell at home, with exporters typically having higher levels of efficiency than those that sell only at home.[6]

Hence, firms with higher efficiency are more likely to export. As we saw above, they are also more likely to have higher measured productivity. Hence variation in underlying efficiency (which is not usually measured) induces correlation between exporting and measured productivity. Our model thus captures a key stylized fact: Firms that export appear to be more productive. While we capture this feature qualitatively, a greater challenge is to fit the distribution of measured productivity and how it differs between exporters and nonexporters. We now turn to this task.

3 Deriving Plant-Level Predictions From Aggregate Trade

We now fit the model to aggregate trade data to garner predictions about the following producer-level observations: (i) the distribution of measured productivity across active producers, (ii) the shift in this distribution between exporters and nonexporters, (iii) the fraction of active producers who export, (iv) the distribution across producers of the share of total revenue earned from exports, and (v) the difference between exporters and nonexporters in average sales to their home market.

To fit the model, we exploit a simple relationship that it delivers between parameters and observations on aggregate bilateral trade flows. The probability π_{ni} that country n will

[6]This will cause a problem in the data as there do exist some plants that export a large majority of their output.

purchase a particular good from country i is just the probability that country i is the lowest cost supplier to country n, as discussed above. Aggregating across goods, this probability becomes the share of country $i's$ goods in country $n's$ purchases. Hence

$$\frac{X_{ni}}{X_n} = \pi_{ni} = \frac{\phi_i d_{ni}^{-\theta}}{\Phi_n}, \tag{9}$$

where X_{ni} is country $n's$ expenditure on goods from country i.

For simulation purposes we can fully parameterize the model by inserting observations on bilateral trade shares into (9) and enlisting estimates of θ and σ. We can then simulate repeated draws from the appropriate efficiency distributions in each different country. For each draw, we identify the location of the active producers, determine whether they export, calculate their price markups in each market in which they are active, determine their revenues in each market, and calculate their measured productivity. Aggregating across draws we obtain the model's predictions for our producer level facts.

3.1 Who Exports?

The first step of our simulations is to draw a sample of lowest costs for producing in each country. We first define $U_{1ni}(j) \equiv \phi_n \left[C_{1ni}(j)\right]^{\theta}$. In particular, $U_{1ii} = T_i \left[Z_{1i}(j)\right]^{-\theta}$, representing the efficiency for good j relative to the country mean, has the exponential distribution $1 - e^{-u_1}$ (which does not depend on any parameters of the model). For each of our N countries we draw U_{1ii} independently from the unit exponential distribution. Manipulating (9) appropriately, we then calculate, for each n:

$$U_{1ni}(j) = \frac{\phi_n d_{ni}^{\theta}}{\phi_i} U_{1ii}(j) = \frac{X_{nn}}{X_{ni}} U_{1ii}(j).$$

Country n buys from country i if and only if $U_{1ni}(j) \leq U_{1nk}(j)$ for all k. Hence for this draw we can identify the active producers in each country and where they sell. Note that we can do so using data only on trade shares, without reference to θ or σ. Repeated draws give a prediction about the fraction of active producers who export.

3.2 What Do They Charge?

We now need to draw from the distribution of second lowest cost conditional on our draw for lowest cost. We define $U_{2ni}(j) \equiv \phi_n \left[C_{2ni}(j)\right]^\theta$. The distribution of $U_{2ii}(j)$ conditional on $U_{1ii}(j) = u_1$ is again "parameter free" and equal to $1 - e^{-u_2+u_1}$. Whenever country i turns out to be an active producer, conditional on its realization of $U_{1ii}(j)$, we draw $U_{2ii}(j)$. For each country n we then calculate:

$$U_{2ni}(j) = \frac{\phi_n d_{ni}^\theta}{\phi_i} U_{2ii}(j) = \frac{X_{nn}}{X_{ni}} U_{2ii}(j).$$

Suppose i turns out to be the low cost supplier to n, that is, $U_{1ni}(j) \leq U_{1nk}(j)$ for all k. We then define, parallel to (5):

$$\begin{aligned} U_{1n}(j) &= U_{1ni}(j) \\ U_{2n}(j) &= \min\left\{\min_{k \neq i}\{U_{1nk}(j)\}, U_{2ni}(j)\right\}. \end{aligned}$$

We can then calculate markups as:

$$M_n(j) = \min\left\{\left(\frac{U_{2n}(j)}{U_{1n}(j)}\right)^{1/\theta}, \overline{m}\right\}.$$

While learning about who exports and where requires only trade data, learning about markups forces us to take a stand on θ and σ (since $\bar{m}$ depends on σ).

3.3 How Much Do They Sell?

Knowing the markups producers charge in each destination, we can use expression (1) to figure out revenues in each country. Using expression (9) and (7) to replace unknown parameters with aggregate trade data, revenues earned by a producer in country i who has captured the market for good j in country n are:

$$X_n(j) = \alpha(j) X_n \left[M_n(j)/\gamma\right]^{1-\sigma} \left(U_{1ii}(j) \frac{X_n}{X_{ni}}\right)^{(1-\sigma)/\theta}. \tag{10}$$

Assuming the producer also sells at home, we define $x_{ni}(j) \equiv X_n(j)/X_i(j)$, (which no longer depends on either $U_{1ii}(j)$, $\alpha(j)$, or γ). The $x_{ni}(j)$ yield a prediction about the distribution

of revenues from exports as a share of the total. Equation (10), as it applies to home sales, predicts (v) above, how much exporters sell at home on average relative to their nonexporting peers.

3.4 How Productive Are They?

As noted above, exporting complicates the expression for measured productivity since the same producer sets different markups in each market. The generalization of expression (8) is:

$$y_i(j) = \sum_{n=1}^{N} \left[\frac{x_{ni}(j)/M_n(j)}{\sum_{k=1}^{N} x_{ki}(j)/M_k(j)} \right] v_i M_n(j),$$

where we set $x_{ni}(j) = 0$ for markets that the producer fails to penetrate. ¿From this expression we garner predictions about the distribution of measured productivity and how it differs between exporters and nonexporters.

4 How Well Do We Fit?

We implement this methodology as follows: We take bilateral trade in manufactures among the OECD in 1990 from EK. We plug these data, along with estimates of θ and σ, into the framework to obtain predictions about productivity and exporting for U.S. producers. We use two alternative estimates of θ from EK and experiment with various values of σ.

Interpreting producers as plants or establishments, we then compare our predictions to measures obtained from the 1992 U.S. Census of Manufacturing and information about each plant's revenues from exporting. (See Bernard and Jensen (1999) for a full description.) How well do our predictions match the facts?

4.1 Who Exports

A basic prediction of our framework (which, recall, does not rely on our estimates of θ or σ) is the fraction of plants that export at all. At one extreme, each plant could export the same share of its total output. At the other, plants completely specialize, with roughly the

fraction of plants exporting representing the fraction of total production exported. Our model's prediction, that 32 per cent of plants export, is very far from either extreme (see Table 1). In fact, only 21 per cent of plants exported in 1992. Here, we seem to be off partly because of the large number of tiny plants which almost never export. If plants with under 20 employees are excluded from the calculation, then the fraction of plants that export rises to over 30 per cent.

4.2 How Much They Export

For the plants that do export, what fraction of revenues flow from the export market? Table 1 compares our predictions, under various assumptions about θ and σ, with the facts. We take the two values of θ estimated by EK and try σ's of 1, 3, and 6 (for the smaller value of θ, the value of $\sigma = 6$ is dropped since it violates the conditions for a well behaved price index). Note that across alternative parameter values we predict that the vast majority of exporters earn less than 10 per cent of their revenues from exports, in line with the actual data. Our predictions are also in line with the fact that few plants specialize in exporting. In the data, less than 5 per cent of exporters earn more than half of their revenues from exports. Our model picks up this general pattern, in fact underestimating the fraction of export-oriented plants when σ is large. For no parameter values, however, can we explain the roughly 1 per cent of exporting plants that earn more than 80 per cent from exports.

4.3 How Productive They Are

To connect our theoretical productivity measure y to our data on value added per worker y^V we assume a Cobb-Douglas production function in labor and material inputs, with labor having a share $\beta = .21$ (See EK for an explanation.) Simple manipulation yields:

$$\ln y^V(j) = \ln(w/\beta) + \ln\left[y(j)/v - (1-\beta)\right].$$

In Table 1 we show the difference (multiplied by 100) between exporters and nonexporters in the mean of log value-added per worker. Our simulations straddle the actual exporter productivity advantage of 33 per cent. The simulated differential in productivity is smaller with larger values of either σ or θ.

Figure 1 shows the distribution of productivity implied by the model and how it shifts between exporters and nonexporters (with $\theta = 3.6$ and $\sigma = 1$). The two upward sloping lines represent log productivity at each percentile (from 1 to 100) for exporters and nonexporters. The hump-shaped curve shows the productivity advantage of exporters at different points in the distribution. Figure 2 shows the results from the 1992 Census of Manufactures.

There is a strong resemblance between the two figures in several respects. (i) Throughout the respective productivity distributions, exporters are more productive than nonexporters. (ii) The productivity distributions overlap substantially so that a significant fraction of nonexporters are more productive than many exporters. There are, however, three important deviations between the simulated and actual figures: (i) The model does not capture the long tail at the low end of the productivity distribution. (ii) The model finds that the productivity advantage of exporters is maximized in the middle of the distribution whereas in the data it is at the low end of the distribution where exporters have the greatest advantage over nonexporters. (iii) The model tends either to overstate the productivity advantage of exporters (as shown in Figure 1, with $\theta = 3.6$) or to understate the amount of heterogeneity in measured productivity (not shown, if $\theta = 8.3$).

4.4 How Big They Are

The last two rows of Table 1 show the size advantage of exporters relative to nonexporters. Even excluding export revenues, the model predicts that exporters will be bigger as long as demand is elastic. For the parameter values shown, however, we consistently underestimate the size advantage of exporters. Apparently, factors not captured by our model are important

in explaining why exporters are so much larger.

5 Conclusion

We have presented a first cut at matching the predictions of a fully articulated model of trade with plant-level facts. Our framework captures all the relevant stylized facts qualitatively, although it is a bit off in predicting some magnitudes. That we do not fit things exactly is not surprising given the spareness of the model and the richness of the data. Our current analysis omits many important features of the world. We ignore, for example, possible differences across industries in relevant parameter values. We treat labor as homogeneous and ignore capital altogether. We ignore the implications of multi-product plants. In principle one could extend our approach to incorporate such features and such extensions remain topics for future research. Nevertheless, at this point we think we learn most about the ability of our model to grapple with plant-level facts by examining both in their most naked form.

References

Baily, Martin N., Hulten, Charles, and Campbell, David. "Productivity Dynamics in Manufacturing Plants." *Brookings Papers on Economic Activity: Microeconomics* (1992): 187-267.

Bernard, Andrew B. and Jensen, J. Bradford. "Exporters, Jobs, and Wages in U.S. Manufacturing: 1976-1987." *Brookings Papers on Economic Activity: Microeconomics* (1995): 67-119.

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

Billingsley, Patrick. *Probability and Measure,* 2nd edition. New York: John Wiley & Sons, 1986.

Buldyrev, Sergey V., Luis A. Amaral, Shlomo Havlin, Heiko Leschhorn, Philipp Maass, Michael A. Salinger, H. Eugene Stanley, and Michael H.R. Stanley. "Scaling Behavior in Economics: II. Modeling of Company Growth." *Journal Physique I France* 7 (April 1997): 635-650.

Dornbusch, Rudiger, Fischer, Stanley and Samuelson, Paul A. "Comparative Advantage, Trade, and Payments in a Ricardian Model with a Continuum of Goods." *American Economic Review* 67 (December 1977): 823-839.

Eaton, Jonathan and Kortum, Samuel. "International Technology Diffusion: Theory and Measurement." forthcoming in *International Economic Review* (1998a).

Eaton, Jonathan and Kortum, Samuel. "Technology, Geography, and Trade." mimeo, Boston University (1998b).

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

Klette, Tor Jakob, and Griliches, Zvi. "Empirical Patterns of Firm Growth and R&D Investment: A Quality Ladder Model Interpretation." *NBER Working Paper No. 6753* (1998).

Kortum, Samuel. "Research, Patenting and Technological Change." *Econometrica* 65 (November 1997): 1389-1419.

Maskus, Keith. "Comparing International Trade Data and Product and National Characteristics Data for the Analysis of Trade Models." In International Economic Transactions: *Issues in Measurement and Empirical Research National Bureau of Economic Research Studies in Income and Wealth, Volume 55,* edited by Peter Hooper and J. David Richardson. University of Chicago Press, 1991.

OECD. "The OECD STAN Database." 1995.

Olley, Steven and Pakes, Ariel. "The Dynamics of Productivity in the Telecommunications Industry." *Econometrica* 64 (November 1996): 1263-1297.

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

A Mathematical Appendix

To be added.

B Data Appendix

B.1 Aggregate Trade Data

Our bilateral trade data for 1990 cover the manufacturing sectors of 19 OECD countries (Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Greece, Italy, Japan, Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, United Kingdom, and United States). Country i's imports from home are gross manufacturing production less manufacturing exports. Its total manufacturing expenditures are home purchases plus imports from everywhere else. These measures are reported by the STAN database in local currencies (OECD, 1995). We calculate imports from each of the other 18 countries, as a fraction of total manufactured imports, from the United Nations–Statistics Canada bilateral merchandise trade data by 4-digit SITC, as described in Feenstra, Lipsey, and Bowen (1997).[7] All import measures are c.i.f. Since we always normalize imports either by home sales or by total expenditures, no exchange rate translation is required.

B.2 Plant Data

The data on individual U.S. manufacturing plants come from the 1992 Census of Manufactures in the Longitudinal Research Database of the Bureau of the Census. The 1992 Census includes over 200,000 plants. It includes data on their value of shipments, production and nonproduction employment, salaries and wages, value-added, capital stock, ownership structure, and direct exports. The plant export measure is the reported value of direct exports which is reported as "The value of products shipped for export [including] direct exports and products shipped to exporters or other wholesalers for export." As indirect exports are not

[7] We used the concordance of Maskus (1991) to determine those SITC codes corresponding to manufactures. Using Feenstra et al.'s concordance made virtually no difference.

included in this measure, we find systematic undercounting of total exports as measured by the Census. See Bernard and Jensen (1995) for a more detailed analysis of undercounting.

Table 1: Simulations of Exporting Facts

	$\theta = 8.28$			$\theta = 3.60$		Actual
	$\sigma = 1$	$\sigma = 3$	$\sigma = 6$	$\sigma = 1$	$\sigma = 3$	
	Per cent of Plants that Export:					
	32	32	32	32	32	21
Per cent of Revenues from Exports	Per cent of Exporting Plants					
0 to 10	69	69	82	68	82	66
10 to 20	13	15	13	13	12	16
20 to 30	4.9	5.4	5.5	4.8	6.0	7.7
30 to 40	3.5	3.4	0.1	3.5	0.5	4.4
40 to 50	3.0	3.3	0.0	3.3	0.0	2.4
50 to 60	2.6	3.6	0.0	2.8	0.0	1.5
60 to 70	3.0	0.7	0.0	3.3	0.0	1.0
70 to 80	1.0	0.0	0.0	0.9	0.0	0.6
80 to 90	0.0	0.0	0.0	0.0	0.0	0.5
90 to 100	0.0	0.0	0.0	0.0	0.0	0.7
	Percentage Productivity Advantage of Exporters:					
	32	32	21	54	38	33
	Percentage Size Advantage of Exporters:					
U.S. Sales	0.0	19	75	0.0	70	141
Total Sales	19	35	83	20	79	158

The statistics in the first 5 columns are averages from simulating the model 50,000 times. The last column (Actual) is calculated from the over 200,000 plants in the 1992 Census of Manufactures. In both the simulations and the actual, the percentage productivity advantage for exporters is calculated as the difference (multiplied by 100) in the mean logarithm of value added per worker between exporters and nonexporting producers. Likewise the percentage size advantage is the difference (multiplied by 100) in the mean logarithm of revenues between exporters and nonexporting producers (either excluding or including revenues from the export market).

Figure 1: Simulated Distribution of Log Labor Productivity

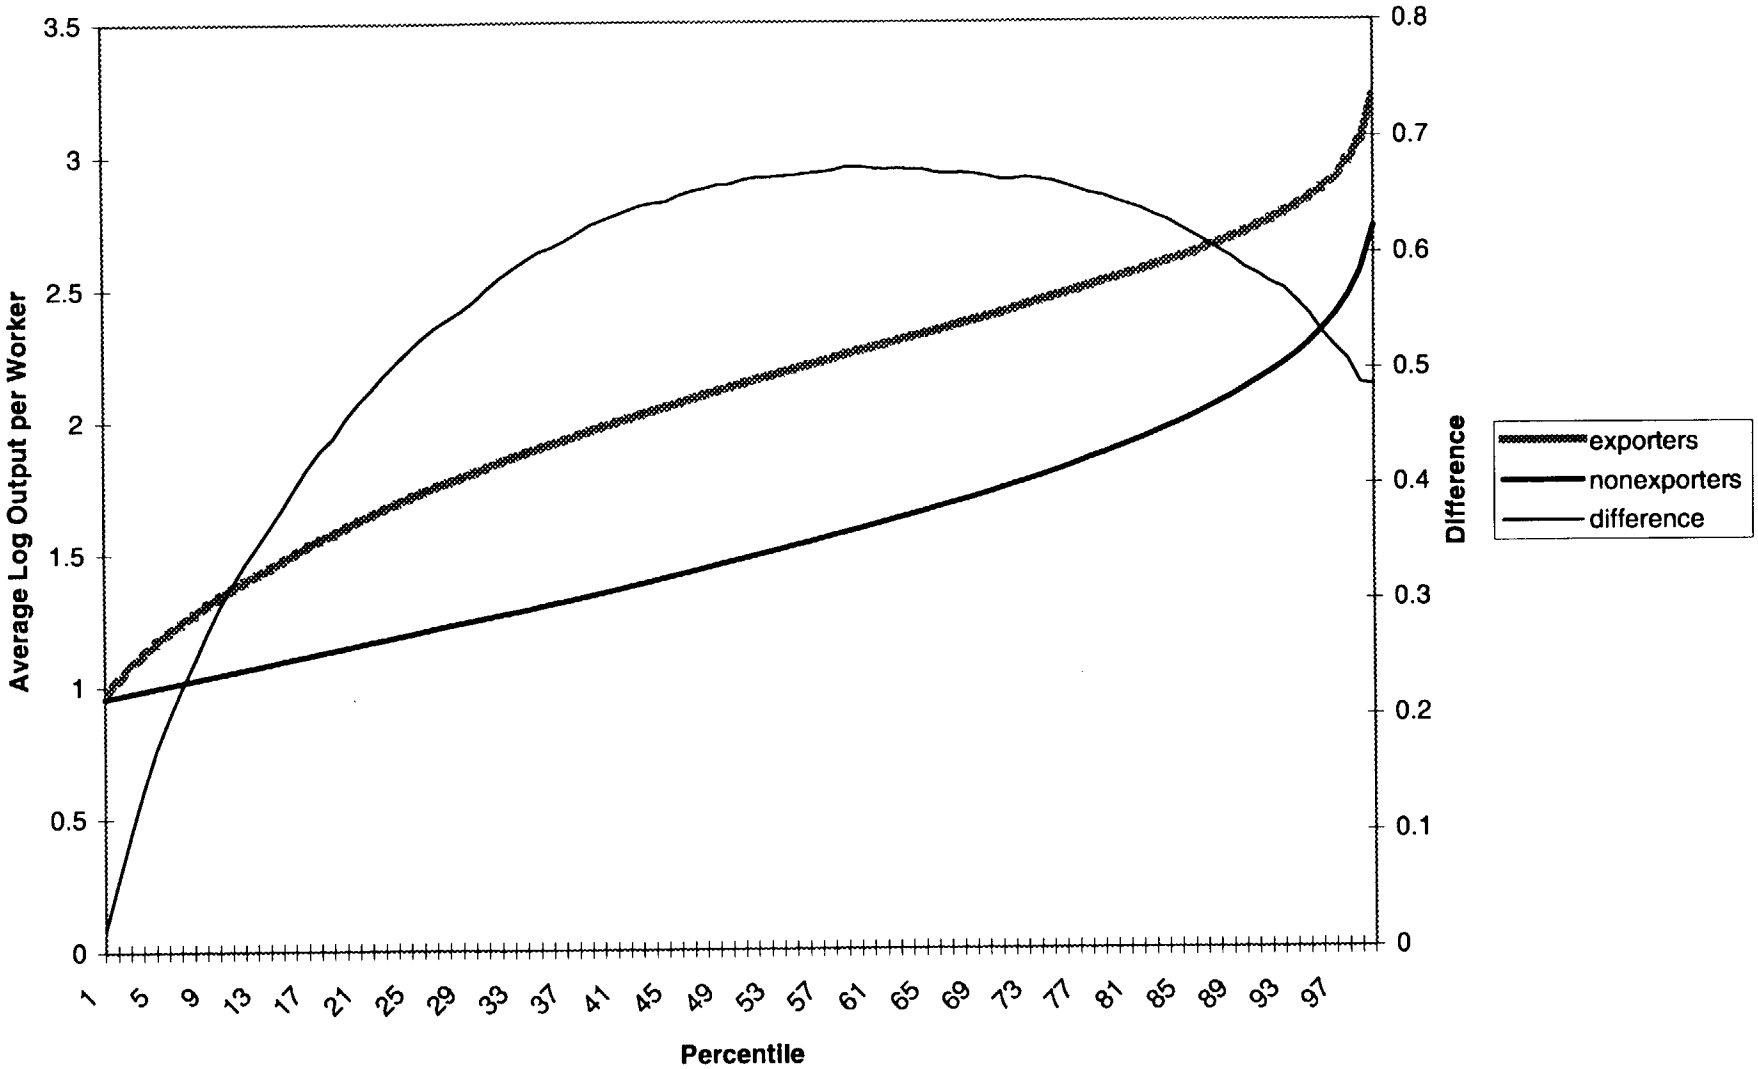
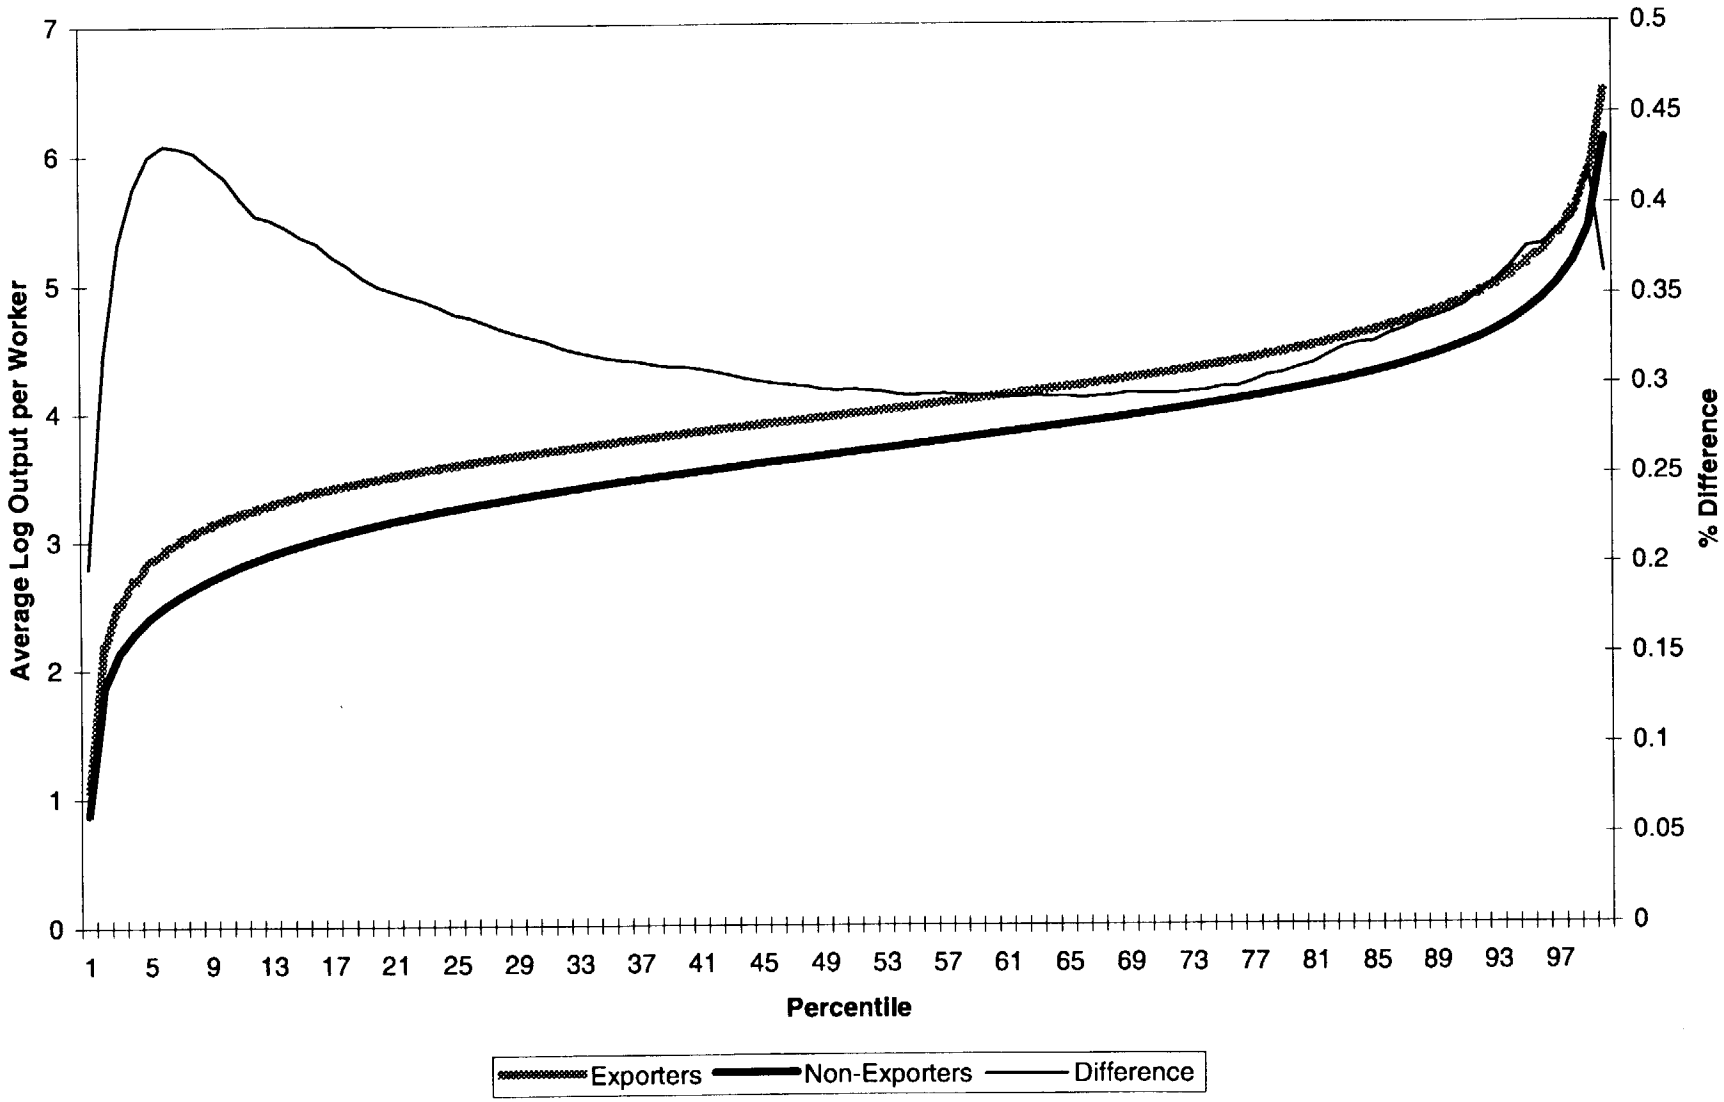
Figure 2: Actual Distribution of Log Plant Labor Productivity - 1992