NBER WORKING PAPER SERIES

PLANTS AND PRODUCTIVITY IN INTERNATIONAL TRADE

Andrew B. Bernard
Jonathan Eaton
J. Bradford Jensen
Samuel Kortum

Working Paper 7688
http://www.nber.org/papers/w7688

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2000

We thank Daniel Ackerberg, Eli Berman, Zvi Eckstein, Simon Gilchrist, Nina Pavcnik, and Robert Staiger for helpful comments. Eaton and Kortum gratefully acknowledge the support of the National Science Foundation. The paper was completed while Kortum was a National Fellow at the NBER. Any opinions expressed are those of the authors and not those of the Bureau of the Census, the NSF, or the NBER.

Plants and Productivity in International Trade
Andrew B. Bernard, Jonathan Eaton, J. Bradford Jensen, and Samuel Kortum
NBER Working Paper No. 7688
May 2000
JEL No. F11, F17, O33

ABSTRACT

We reconcile international trade theory with findings of enormous plant-level heterogeneity in exporting and productivity. Our model extends basic Ricardian theory to accommodate many countries, geographic barriers, and imperfect competition. Fitting the model to bilateral trade among the United States and its 46 major trade partners, we see how well it can explain basic facts about U.S. plants: (i) productivity dispersion, (ii) the productivity advantage of exporters, (iii) the small fraction who export, (iv) the small fraction of revenues from exporting among those that do, and (v) the much larger size of exporters. We pick up all these basic qualitative features, and go quite far in matching them quantitatively. We examine counterfactuals to assess the impact of various global shifts on productivity, plant entry and exit, and labor turnover in U.S. manufacturing.

Andrew B. Bernard
Tuck School of Business
Dartmouth College
100 Tuck Hall
Hanover, NH 03755
and NBER
andrew.b.bernard@dartmouth.edu

Jonathan Eaton
Department of Economics
Boston University
270 Bay State Road
Boston, MA 02215
and NBER
jeaton@bu.edu

J. Bradford Jenson
Center for Economics Studies
Bureau of the Census
University of Maryland
College Park, MD 20742

Samuel Kortum
Department of Economics
Boston University
270 Bay State Road
Boston, MA 02215
and NBER
kortum@bu.edu

1 Introduction

A new empirical literature has emerged that examines international trade at the level of individual producers. Bernard and Jensen (1995, 1999a), Clerides, Lach, and Tybout (1998), and Aw, Chung, and Roberts (1998), among others, have uncovered stylized facts about the behavior and relative performance of exporting firms and plants which hold consistently across a number of countries: Most strikingly, exporters are in the minority; they tend to be more productive and larger; yet they usually export only a small fraction of their output. This heterogeneity of performance diminishes only modestly when attention is restricted to producers within a given industry.

International trade theory has not had much to say about these producer-level facts, and in many cases is inconsistent with them. To the extent that empirical implications have been of concern, trade theory has been aimed at understanding aggregate evidence on such topics as the factor content of trade and industry specialization. To understand the effects of trade on micro issues such as plant closings, however, we need a theory that recognizes differences among individual producers within an industry. Moreover, as we elaborate below, such a theory is needed to understand the implications of trade for such aggregate magnitudes as worker productivity.

Our purpose here is to develop a model of international trade that comes to grips with what goes on at the producer level. Such a model requires three crucial elements. First, we need to acknowledge the heterogeneity of plants. To do so we introduce Ricardian differences in technological efficiency across producers and countries. Second, we need to explain the coexistence, even within the same industry, of exporters and purely domestic producers. To capture this fact we introduce costs to exporting through a standard "iceberg" assumption (export costs to a given destination are proportional to production costs). Third, in order for differences in technological efficiency not to be fully absorbed by differences in output prices (thus eliminating differences in measured productivity across plants), we need imperfect competition with variable markups. We take the simplest route of introducing Bertrand competition into the Ricardian framework with a given set of goods.[1]

A novel feature of our theory is to link a plant's underlying technological efficiency to its productivity as normally measured (typically as 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 (or, for that matter, monopolistic competition with a common markup) they would all appear equally productive in terms of value added per worker, in spite of any efficiency differences. 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 markups, measured productivity would say nothing about underlying efficiency. It turns out, however, that the theory implies that producers who are more efficient also tend to have a greater cost advantage over their closest competition, and are thus able to set higher markups. Hence measured productivity provides a signal, albeit a noisy one, about underlying technical efficiency.[2] The link between underlying efficiency and productivity, on the one hand, and underlying efficiency and selection into export markets, on the other, leads our model to predict that exporting plants are more productive.[3]

A novel feature of our empirical approach is to connect the micro and macro level data.

---

[1] As in Eaton and Kortum (2000), specialization emerges endogenously through the exploitation of comparative advantage. An alternative model that also allows for heterogeneity and geographic barriers of the iceberg variety is Krugman's (1979) extension to international trade of the Dixit-Stiglitz (1977) model of monopolistic competition. But this approach delivers the counterfactual implication that every producer exports everywhere. In contrast, in our model a plant exports only when its cost advantage over its competitors around the world overcomes geographic barriers. Other attempts to explain producer heterogeneity in export performance rely on a *fixed* cost of exporting (see, e.g., Roberts and Tybout 1997 and Melitz 1999). The problem here is that a producer would either export nothing or else sell to different countries of the world in proportion to their market sizes. This second implication belies the very small share of exports in the revenues of most exporters.

[2] An extensive literature compares productivity levels across plants. See, e.g., Baily, Hulten, and Campbell (1992), Bartelsman and Dhrymes (1992), and Olley and Pakes (1996). In making such comparisons, it is typically assumed that the plants in question produce a homogeneous output. Our framework shows how such comparisons make sense even when outputs are heterogeneous.

[3] Clerides, Lach, and Tybout (1998) and Bernard and Jensen (1999a) find strong empirical support for this selection mechanism (and little or no empirical support for learning by exporting) in explaining why exporters are more productive than nonexporting plants.

Aggregate production and bilateral trade volumes around the world provide all we need to know about parameters governing geographic barriers, aggregate technology differences, and differences in input costs. The two remaining parameters relate to the heterogeneity of goods in production and in consumption. Reasonable values for these two parameters bring us quite close to fitting the various micro facts as they apply to U.S. manufacturing plants. Hence the framework serves as a bridge between what we know about global trade flows and what we have learned about plant-level export behavior.

Since the model comes to terms with plant-level facts quite well, we go on to ask what changes in the global economy mean for plant entry, exit, and exporting as well as overall productivity and employment in manufacturing. We look at three scenarios.

We first consider the effects of "globalization" in the form of a 5 percent drop in all geographic barriers between countries (resulting in nearly a 40 percent rise in world trade). We find that this move kills off nearly 9 percent of U.S. plants. But among the survivors, one in seven of the plants that had previously sold only to the domestic market starts exporting. Since globalization provides the survivors larger markets, and since the survivors were larger to begin with, the decline in manufacturing employment is less than 3 percent.

We then move in the opposite direction to autarky (raising geographic barriers to eliminate all trade). The number of active U.S. plants rises by 17 percent. But since plants that were exporting lose their overseas markets, nearly as many jobs are destroyed (11 percent of initial employment) as are created (12 percent of initial employment). Reallocation of employment away from exporters to less productive entrants lowers productivity by around 4 percent.

Our final experiment is a decline in U.S. "competitiveness" in the form of an exogenous 10 percent increase in the U.S. relative wage. The number of manufacturing plants falls by 8 percent and manufacturing employment falls by 18 percent as plants substitute cheaper imported intermediates for labor.

Our analysis thus captures how, even in a relatively closed market such as the United

States, changes in the global economy can substantially reshuffle production. This reshuffling in turn can have important implications for overall manufacturing productivity.[4]

We proceed as follows. Section 2 discusses the plant-level facts we seek to explain. In Section 3 we present the theory behind our qualitative explanations, derived in Section 4, for what goes on at the plant level. Section 5 goes on to compare the model's quantitative implications with the plant-level statistics. Section 6 completes the general equilibrium specification of the model required to undertake the counterfactual experiments reported in Section 7. Section 8 concludes.

2 Exporter Facts

Before turning to the theory, we take a closer look at the plant-level statistics about U.S. exporters that our model seeks to explain. These statistics, shown in Table 1, are all calculated from the 1992 U.S. Census of Manufactures (see Appendix B.2).

We first look at the prevalence of exporting among U.S. manufacturing plants. At one extreme, each plant could export the same share of its total output. At the other, a few giant plants would account for all exports. In fact, of the roughly 200,000 plants in the Census, only 21 percent report exporting anything.[5]

While previous work has sought to link trade orientation with industry, it turns out that exporting producers are quite spread out across industries. Figure 1 plots the distribution of industry export intensity: Each of the 458 4-digit manufacturing industries is placed in one of 10 bins according to the percentage of plants in the industry that export. In two-thirds of the industries, the fraction of plants that export lies between 10 and 50 percent. Hence

---

[4] The results of our counterfactual experiments accord well with findings in the literature. Bernard and Jensen (1999b) find productivity gains driven by reallocation among U.S. producers as exporting has increased. Campa and Goldberg (1995) show that imported intermediates are an important link between U.S. producers and the rest of the world. Gourinchas (1999) estimates that changes in the U.S. real exchange rate lead to increased churning in the labor market. Head and Reis (1998) document the substantial exit and reallocation of production among Canadian producers following tariff reductions under the Free Trade Agreement.

[5] Note that the U.S. export share is even lower. For 1992, the OECD (1995) reports that the U.S. manufacturing sector exported about 13 percent of its gross production.

knowing what industry a plant belongs to leaves substantial uncertainty about whether it exports. Industry has less to do with exporting than standard trade models might suggest.

Not only are plants heterogeneous in whether they export, they also differ substantially in measured productivity. Figure 2A plots the distribution across plants of value added per worker (segregating exporters and nonexporters) relative to the overall mean. A substantial number of plants have productivity either less than a fourth or more than four times the average. Again, a plant's industry is a weak predictor of its performance: Figure 2B provides the same distribution only normalizing each plant's productivity by mean productivity in its 4-digit industry. Controlling for industry only marginally tightens the productivity distribution.

While there is substantial heterogeneity in both productivity and export performance, even within industries, Figure 2A brings out the striking association between the two. The exporters' productivity distribution is a substantial shift to the right of the nonexporters' distribution. Figure 2B shows that this association survives even when looking within 4-digit industries. As shown in Table 1, exporters have a 33 percent advantage in labor productivity overall, and a 15 percent advantage relative to nonexporters within the same 4-digit industry. Accounting for differences in capital intensity across plants within industries, the total factor productivity advantage of exporters is 10 percent.

While differences across industries certainly appear in the data, what is surprising is how little industry explains about exporting and productivity. Hence a satisfactory explanation of plant level behavior must go beyond the industry dimension. We consequently pursue an explanation of these facts that, as a first approximation, bypasses industries and goes directly to the plant level.

Table 1 also reports the importance of export markets for the plants that do export. Surprisingly, the vast majority of exporters export less than 10 percent of what they produce. Less than 5 percent of the exporting plants (which also account for about 5 percent of exporters' total output) export more than 50 percent of their production. Even for the minority of plants that do export, domestic sales dominate.

How is it possible for such a small fraction of plants, exporting such a small fraction of what they produce, to account for total exports? An answer is that exporters are much larger. They are almost 5 times the size of nonexporters on average, even when export revenues are excluded from the calculation. While only 21 percent of manufacturing plants report that they export, these plants account for 60 percent of the output of U.S. manufacturing.

An important caveat in considering any of these statistics is that U.S. manufacturing plants as a whole report exports that sum to just over 60 per cent of total U.S. exports of manufactures reported by the OECD. (See Bernard and Jensen (1995) for a discussion of this problem.) We discuss how undercounting could affect our results below.

3 The Basic Model

Our model combines imperfect competition with the Ricardian theory of comparative advantage based on technology differences. We start with Eaton and Kortum (2000) (Henceforth EK), which itself extends the Ricardian model of Dornbusch, Fischer, and Samuelson (1977) to incorporate an arbitrary number of countries separated by geographic barriers.

3.1 A Ricardian Framework

We describe a world of N countries in which each country can produce every good along the interval $[0, 1]$. Efficiency in the production for any good varies across countries. Aside from these Hicks' neutral efficiency differences, production of any good anywhere combines inputs in the same way, although the price of the bundle of inputs may vary across countries.

Within a country there are many potential producers of any good, but only the most efficient ones will ever be in business. Efficient producers of good j in country i can convert one bundle of inputs into a quantity $Z_{1i}(j)$ of good j (where the subscript 1 indicates that Z_1 corresponds to the *most* efficient method for making j in i).

Below we make the input bundle a Cobb-Douglas combination of labor and intermediate inputs, treating the labor share β as independent of j. For now it is more convenient to keep the inputs bundled together, denoting their price index by w_i. It thus costs $w_i/Z_{1i}(j)$ to produce a unit of good j in country i using the cheapest local means.[6]

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 shipping $d_{ni} \geq 1$ units, and we normalize $d_{ii} = 1$ for all i. We impose the plausible "triangle inequality" on the geographic barrier parameters d_{ni}:

$$d_{ni} \leq d_{nk} d_{ki} \quad \forall k, \tag{1}$$

i.e., an upper bound on the cost of moving goods from i to n is the cost of moving them via some third country k.[7]

Taking into account geographic barriers, the lowest cost in country n of obtaining good j from country i is:

$$C_{1ni}(j) = \left(\frac{w_i}{Z_{1i}(j)} \right) d_{ni} \tag{2}$$

while the lowest cost irrespective of source is:

$$C_{1n}(j) = \min_i \{ C_{1ni}(j) \}. \tag{3}$$

To extract any implications for trade we need to specify how efficiencies are distributed across countries. The fact that we must take minima across a large number of sources makes the model potentially cumbersome. To cut through this problem EK take the distribution of

---

[6] By assuming that w_i is the same for any good in a country, we ignore possible factor intensity differences, precluding any Heckscher-Ohlin explanation for trade patterns. We could (and intend to) 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.

[7] Arbitrage across markets should ensure that this restriction holds.

the frontier efficiencies for producing good j in country i to be extremal:[8]

$$F_{1i}(z_1) = \Pr[Z_{1i}(j) \leq z_1] = e^{-T_i z_1^{-\theta}}. \tag{4}$$

We choose units so that the distributions are identical across goods j. Given this normalization, the parameter $T_i > 0$ reflects the overall state of technology in country i, or its absolute advantage, while $\theta > 1$ determines the heterogeneity of relative efficiencies, and governs comparative advantage. A higher T_i for a country means that producers there will on average be more efficient, while a higher θ means that efficiencies are more homogeneous.

Demand everywhere combines goods with a constant elasticity of substitution $\sigma > 0$. Hence expenditure on good j in country n, $X_n(j)$, is:

$$X_n(j) = \alpha(j) X_n \left(\frac{P_n(j)}{P_n} \right)^{1-\sigma}, \tag{5}$$

where $P_n(j)$ is the price of good j in country n, X_n denotes total expenditure there, $\alpha(j)$ effects the size of the market for good j, 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. (We assume that the $\alpha(j)$ are independent of efficiency levels, and normalize $\int_0^1 \alpha(k) dk = 1$.)

We have now introduced all the relevant parameters of the model: (i) geographic barriers d_{ni}, (ii) input costs w_i, (iii) the states of technology T_i, (iv) the elasticity of substitution σ, and (v) the comparative advantage parameter θ. As we show below, we can learn all we need to know about the first three from data on bilateral trade shares. Learning about the last two, which govern the heterogeneity of goods in consumption and production respectively, requires looking at higher moments of the data.[9]

---

[8] Kortum (1997) and Eaton and Kortum (1999) derive this distribution of leading-edge efficiencies from a dynamic model with endogenous innovation. The distribution is extremal because the market selects the best from a large number of possible techniques. If the techniques are random draws from some distribution, then the highest efficiency (suitably normalized) converges to an extremal distribution. The Fréchet distribution (4) 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).

[9] Obviously in general equilibrium input costs w_i are determined endogenously in factor and input markets. This endogeneity turns out not to matter in fitting our model to observed data, but we take it into account in pursuing our counterfactuals below.

To derive the model's implications, we use (2), to move from the distribution of country i's frontier efficiency to the distribution of the lowest cost of supplying a good from country i to country n:

$$G_{1ni}(c_1) = \Pr[C_{1ni}(j) \leq c_1] = \Pr[Z_{1i}(j) \geq d_{ni}w_i/c_1] = 1 - e^{-T_i(w_i d_{ni})^{-\theta} c_1^{\theta}}. \tag{6}$$

Taking the minimum cost across all potential suppliers (3), the distribution of lowest cost in country n (irrespective of source) is:

$$G_{1n}(c_1) = \Pr[C_{1n}(j) \leq c_1] = 1 - \prod_i [1 - G_{1ni}(c_1)] = 1 - e^{-\Phi_n c_1^{\theta}}, \tag{7}$$

where

$$\Phi_n = \sum_{i=1}^{N} T_i (w_i d_{ni})^{-\theta}. \tag{8}$$

The cost parameter Φ_n aggregates states of technology around the world, deflating by input costs and the cost of delivering to country n.

In what follows below we exploit two useful results from EK:

1. The probability π_{ni} that country i is the low cost supplier to n is just its contribution to the cost parameter Φ_n; that is:

   $$\pi_{ni} = T_i (w_i d_{ni})^{-\theta} / \Phi_n. \tag{9}$$

   Aggregating across goods, π_{ni} is then the fraction of goods for which country i is the low-cost supplier to country n.[10]

2. 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.[11] That

---

[10] We obtain this probability by calculating:

$$\pi_{ni} = \Pr\left[C_{1ni}(j) \leq \min_{s \neq i} \{C_{1ns}(j)\}\right] = \int_0^{\infty} \prod_{s \neq i} [1 - G_{1ns}(c_1)]\, dG_{1ni}(c_1).$$

[11] We obtain this result by showing that:

$$G_{1n}(c_1) = \frac{1}{\pi_{ni}} \int_0^{c_1} \prod_{s \neq i} [1 - G_{1ns}(q)]\, dG_{1ni}(q).$$

> is, once transport costs 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 $P_n(j) = C_{1n}(j)$. As we show below, however, with perfect competition (or, for that matter, monopolistic competition when $\sigma > 1$), greater efficiency is reflected in proportionately lower prices. As a consequence, the value of production per unit of input is the same across producers in a country, and hence has nothing to do with differences in efficiency. Since a major feature of our data is enormous variation across plants in measured productivity, we drop perfect competition in favor of a model that delivers heterogeneity in market power.

3.2 Introducing Imperfect Competition

We continue to assume that each market n is captured by the low cost supplier of each good j. Prices, however, are determined by a blend of the Grossman-Helpman (1991) "quality ladders" model (with rungs of random size) and the Dixit-Stiglitz (1977) model of monopolistic competition. As in the quality ladders framework, the low cost supplier to a market n faces latent rivals who can deliver exactly the same good j to that market. The lowest cost among these latent rivals, which we denote $C_{2n}(j)$, places an upper bound on what the low cost supplier can charge. In fact, $C_{2n}(j)$ is the price unless the implied markup exceeds the Dixit-Stiglitz markup $\overline{m}$. Hence:

$$P_n(j) = \min\{C_{2n}(j), \overline{m}C_{1n}(j)\}, \tag{10}$$

where:

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

We must now keep track of not only the lowest cost, but also the second lowest, in supplying each market. The second lowest cost supplier could be either: (i) the second lowest cost supplier from the same country as the low cost supplier, or else (ii) the low cost supplier from some other country. Hence, if country i turns out to be the low cost provider, so that $C_{1n}(j) = C_{1ni}(j)$, then:

$$C_{2n}(j) = \min\left\{\min_{k\neq i}\{C_{1nk}(j)\}, C_{2ni}(j)\right\}, \tag{11}$$

where $C_{2ni}(j)$ is the second-lowest cost of supplying good j to market n from country i.

To learn about $C_{2ni}(j)$ we have to say something not only about the distribution of the best technology in each country, but about the second-best as well. To do so we extend the theory of the distribution of extremes to consider the joint distribution of the first and second best. As we show in Appendix A.1, the analogue to the Fréchet for the joint distribution of the first two order statistics is:

$$F_i(z_1, z_2) = \Pr[Z_{1i}(j) \leq z_1, Z_{2i}(j) \leq z_2] = [1 + T_i(z_2^{-\theta} - z_1^{-\theta})]e^{-T_i z_2^{-\theta}}, \tag{12}$$

for $0 \leq z_2 \leq z_1$, where $Z_{2i}(j)$ denotes the efficiency of the second-best potential producer of good j in country i. (Note that setting $z_1 = z_2$ returns (4).)

We can incorporate (12) with our assumptions about input costs and geographic barriers to obtain the joint distribution of the lowest cost $C_{1ni}(j)$ and second-lowest cost $C_{2ni}(j) = w_i d_{ni}/Z_{2i}(j)$ from potential producers in country i delivering to country n. As shown in Appendix A.2, these distributions in turn, through international competition, generate the joint distribution of the lowest cost $C_{1n}(j)$ and second lowest cost $C_{2n}(j)$ of good j in country n irrespective of source:

$$G_n(c_1, c_2) = \Pr[C_{1n}(j) \leq c_1, C_{2n}(j) \leq c_2] = 1 - e^{-\Phi_n c_1^\theta} - \Phi_n c_1^\theta e^{-\Phi_n c_2^\theta}, \tag{13}$$

for $c_1 \leq c_2$, where Φ_n is the cost parameter given by (8). (Note that letting c_2 approach infinity returns (7).)

Since EK assume perfect competition, the distribution of lowest costs (7) is also their distribution of prices. With imperfect competition goods are sold at a markup $M_n(j) = P_n(j)/C_{1n}(j)$. As we show in Appendix A.3, equations (10) and (13) imply that $M_n(j)$ has a Pareto distribution truncated at the monopoly markup:

$$H_n(m) = \Pr[M_n(j) \leq m] = \begin{cases} 1 - m^{-\theta} & 1 \leq m < \overline{m} \\ 1 & m \geq \overline{m}. \end{cases} \tag{14}$$

While the distribution of costs differ by country, the distribution of markups is the same in any destination. 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. Greater competitiveness leads to larger market share rather than higher markups.

One might have thought that a lowering of geographic barriers, by increasing the number of potential suppliers to a market, would lower markups there. Indeed, it does for domestic producers who survive. However, this effect is exactly offset by the exit of domestic producers who tended to charge the lowest markups. From the perspective of foreign suppliers, a lowering of geographic barriers tends to raise the markup of incumbents (who now have lower costs) but it also leads to entry by marginal foreign suppliers with low markups. Once again the distribution of markups across all active suppliers remains unchanged.

Note that the distribution of the markup depends on only the two heterogeneity parameters θ and σ, with less heterogeneity of either type lowering the markup. If technologies are less heterogeneous (θ high) then markups are typically smaller as leaders tend to have less of a lead over their latent competitors. Alternatively, if different goods are very close substitutes in consumption (σ high) then the markup is truncated by a lower $\overline{m}$ (the Dixit-Stiglitz markup), since individual sellers face more elastic demand.

Since markups have the same distribution everywhere, the cost parameter Φ_n fully captures cross-country variation in the distribution of prices. In particular, assuming $\sigma < 1 + \theta$, the exact price index in country n, appearing in equation (5), is:

$$P_n = \gamma \Phi_n^{-1/\theta}. \tag{15}$$

The parameter γ is a function of only the parameters governing the heterogeneity of technology and tastes, θ and σ.[12] This price index applies not only to what country n buys overall, but to what it buys from any particular source.

Since prices anywhere do not vary systematically according to source, the share country n spends on goods from country i is also the fraction of goods it purchases from there, π_{ni} given in equation (9). That is:

$$\frac{X_{ni}}{X_n} = \pi_{ni}, \tag{16}$$

where X_{ni} is what country n spends on goods from country i and X_n is its total spending. This relationship provides the link between our model and data on aggregate bilateral trade.

4 Implications for Productivity, Exporting, and Size

Before turning to the data themselves, however, we show how this model delivers qualitatively the plant-level facts described in Section 2. We first demonstrate the link between measured productivity and underlying efficiency. We then show why exporting plants tend to be large with high measured productivity.

4.1 Efficiency and Measured Productivity

To simplify things, consider a plant producing good j for only the home market (we therefore drop the i subscript in this subsection). Defining the number of input bundles used to make good j as $I(j)$, efficiency is $Z_1(j) = Y(j)/I(j)$, where $Y(j)$ is physical output. Comparing this efficiency measure across plants producing different goods is not meaningful since it depends

---

[12] Specifically, γ is $\left[\frac{1+\theta-\sigma+(\sigma-1)\bar{m}^{-\theta}}{1+\theta-\sigma}\Gamma\left(\frac{1+2\theta-\sigma}{\theta}\right)\right]^{1/(1-\sigma)}$, as shown in Appendix A.4. The restriction on σ and θ ensures that goods are sufficiently heterogeneous in consumption relative to their heterogeneity in production so that buyers do not concentrate their purchases on a few low-price goods. As long as we obey this parameter restriction, γ is irrelevant for anything that we do empirically.

on the units in which output is measured. But in any case, available measures of plant productivity y are almost always based on the value of output, $y(j) = P(j)Y(j)/I(j)$.[13]

Under perfect competition, $P(j) = w/Z_1(j)$, so that $y(j) = w$. Measured productivity is thus the same for all plants facing common input prices, regardless of their relative efficiency. With perfect competition the value-based measure does not capture any differences in efficiency across producers.[14]

Under imperfect competition, however, $P(j) = M(j)w/Z_1(j)$, so that measured productivity is:

$$y(j) = M(j)w, \tag{17}$$

the cost of an input bundle scaled up by the producer-specific markup. Differences across producers in measures of productivity can now emerge, but these simply reflect differences in markups. In the absence of any connection between markups and efficiency, value-based productivity measures provide information only about monopoly power.

In fact, our model does imply that, on average, plants that are more efficient charge a higher markup. As derived in Appendix A.5, conditional on a level of efficiency z_1, the distribution of the markup $M(j)$ is:

$$H(m|z_1) = \Pr[M(j) \leq m | Z_1(j) = z_1] = \begin{cases} 1 - e^{-\Phi w^{\theta} z_1^{-\theta}(m^{\theta}-1)} & 1 \leq m < \overline{m} \\ 1 & m \geq \overline{m}. \end{cases}$$

A plant with higher efficiency Z_1 than another is likely to have a higher markup (its distribution of M stochastically dominates the other's) and hence higher measured productivity. The reason is that a plant that is unusually efficient tends to be unusually efficient relative to its latent competitors as well. Plants that are particularly advanced can typically charge higher

---

[13] As noted by Klette and Griliches (1995), producer-level price indices are essentially nonexistent. Consequently an observed increase in a given plant's productivity (as typically measured) over time could reflect either increased efficiency or an increase in its relative price. In our comparisons of productivity levels at a given time across producers making different goods efficiency and relative price are also confounded. Plant-level deflators, even if they existed, could not solve this problem.

[14] Of course, looking across countries, efficiency and measured productivity are linked, since countries that are on average more technologically advanced will have higher input costs, particularly wages.

markups.[15]

Hence, under imperfect competition, variation in efficiency can capture heterogeneity in measured productivity across plants. We still have to show why greater productivity is associated with exporting and why exporting in turn correlates with size.[16]

4.2 Efficiency and Exporting

Having shown that a more efficient plant is likely to have higher measured productivity, we now ask what efficiency implies for exporting.

Consider the best potential producer of good j from country i facing potential competitors from abroad with efficiencies $Z_{1k}(j)$ for $k \neq i$. In order to sell at home its efficiency $Z_{1i}(j)$ must satisfy

$$Z_{1i}(j) \geq Z_{1k}(j)\frac{w_i}{w_k d_{ik}} \; \forall k \neq i. \tag{18}$$

But to sell in some other market n requires:

$$Z_{1i}(j) \geq Z_{1k}(j)\frac{w_i d_{ni}}{w_k d_{nk}} \; \forall k \neq i.$$

The triangle inequality implies that $d_{nk} \leq d_{ni}d_{ik}$ or that $w_i d_{ni}/(w_k d_{nk}) \geq w_i/(w_k d_{ik})$. Hence exporting anywhere imposes a higher efficiency hurdle than selling only at home. While any plant good enough to sell abroad will also sell at home, only a fraction of those selling domestically will succeed in exporting anywhere.

Plants with higher efficiency are more likely to export and are also more likely to have higher measured productivity. Variation in underlying efficiency induces correlation between

---

[15] Looking at the relationship the other way around, how does underlying efficiency Z_1 vary with measured productivity y? As shown in Appendix A.6, the conditional expectation is proportional to y (as long as the markup is less than $\overline{m}$). Hence, a plant appearing to be 2 percent more productive than another is, *on average*, 2 percent more efficient (unless it is charging the monopoly markup, in which case expected efficiency is even greater).

[16] For simplicity, we have demonstrated the relationship between efficiency and measured productivity ignoring exports. Since exporting plants typically charge different markups in different markets, the relationship between measured productivity and efficiency is more complicated. Nevertheless, higher underlying efficiency leads to a higher markup in any given market. Our simulations take these complications into acccount, as we discuss below.

exporting and measured productivity: More efficient plants tend to have a greater cost advantage over their rivals in any market where they can sell, and they can typically sell in more markets.

Our model thus captures a key stylized fact: Plants that export appear to be more productive. It also explains the coexistence of exporting plants and plants that sell only to the domestic market, and why plants don't export everything.

4.3 Efficiency and Size

Our model can also explain why exporting plants tend to be larger than plants selling only to the domestic market. Obviously exporting plants are larger because they sell to more markets. But why should we expect them to sell more at home?

The reason is that greater efficiency not only raises the probability of exporting, it will also likely result in a lower domestic price. For elasticities of substitution $\sigma > 1$, lower prices translate into more spending.

Greater efficiency leads to lower prices for either of two reasons. For plants that can charge the Dixit-Stiglitz markup $\overline{m} = \sigma/(\sigma - 1)$, the markup is over a lower unit cost. For plants whose markups are limited by the costs of potential competitors, the argument is less straightforward. From the joint distribution of the lowest and second-lowest cost (13) we can obtain the distribution of the second lowest cost (i.e. the price) conditional on the lowest cost:

$$G_{2n}(c_2|c_1) = \Pr[C_{2n} \leq c_2 | C_{1n} = c_1] = \frac{\partial G_n(c_1, c_2)/\partial c_1}{\partial G_{1n}(c_1)/\partial c_1} = 1 - e^{-\Phi_n(c_2^\theta - c_1^\theta)}.$$

This distribution is stochastically increasing in c_1 (and hence decreasing in $z_1 = w/c_1$). Even though, as we showed above, plants with lower costs typically charge higher markups, their prices nevertheless tend to be lower. Hence, for $\sigma > 1$, they earn more revenues.

5 Predicting Plant-Level Magnitudes from Aggregate Trade

Our model explains the correlations we observe between measured productivity, exporting, and size through the positive association of each with underlying efficiency. While we capture the directions of these correlations, a greater challenge is to explain their magnitude.

5.1 Parameterizing the Model

To make any quantitative predictions we need to assign values to the model's parameters: For the United States and all its trading partners these are the country-specific parameters w_i (the cost of an input bundle) and T_i (the state of technology), and, for each country pair, d_{ni} (the geographic barrier). We also need to know the heterogeneity parameters σ (in consumption) and θ (in production) common across countries. (We set the labor share in total cost, β, to 0.21, as described in EK. The preference parameters $\alpha(j)$ do not matter for the statistics we consider.)

It may seem that there are so many countries, and hence so many country-specific parameters, that we could fit whatever we want. It turns out, however, that the π_{ni}'s, given by (9), summarize all we need to know about the country specific parameters T_i, w_i, and d_{ni}.[17] The π_{ni}'s can be observed directly as bilateral trade shares, using (16). We calculate trade shares from 1992 production and trade data in manufactures among the 47 leading U.S. export destinations (including the United States itself). Even though our goal is to learn what the model has to say about U.S. plants, we need to consider all bilateral trade relationships. Whether a U.S. producer exports to France, for example, depends, among other things, on its ability to edge out a German rival. Appendix B.1 describes the data. Table 2 lists our choice of partner countries as well as some summary statistics for each of them.

For the remaining two parameters we use two different estimates of θ, the heterogeneity of

---

[17] Country i's ability to compete in market n depends on a combination of its input costs w_i, its state of technology T_i, and the geographic barrier from i to n, d_{ni}. The π_{ni}'s summarize how these various factors combine to determine i's competiveness in n. How π_{ni} breaks down into its components is irrelevant from the point of view of producers in country n.

goods in production, taken from EK, 8.28 (the preferred estimate from that paper) and 3.60. We then experiment with various values of σ, the heterogeneity of goods in consumption. We use 1 and 4 for each value of θ, along with 6 and 8 for $\theta = 8.28$ (obeying the theoretical restriction $\sigma < \theta + 1$).

5.2 Simulating the Model

We extract the model's quantitative implications through simulation. Details of the simulation methodology are given in Appendix C. Here, we just describe the basic setup.

Each step of a simulation involves sampling a good j at random from the continuum. For that good we draw from the efficiency distribution (12) in each country, which together with the w_i's and d_{ni}'s determines the potential cost of each country supplying each other country (Appendix C explains how we actually carry out these steps given what we know about the parameters). From these potential costs we identify the actual locations of the active (i.e., low cost) producers for each market.[18] We also determine the second lowest cost in supplying each market, governing the markup there. If it turns out that the United States has an active producer of good j, which we interpret as a U.S. plant, we determine whether it exports, calculate its price markup in each market where it sells, determine its revenue in those markets, and calculate its measured productivity.

As noted above, exporting complicates the expression for measured productivity since the same producer sets different markups in each market. Consider a U.S. plant producing good j and selling it to a set of markets Ω. The generalization of expression (17), representing total revenues per input bundle, is:

$$y(j) = M^C(j)w = \left[\frac{\sum_{n \in \Omega} X_n(j)}{\sum_{n \in \Omega} \frac{X_n(j)}{wM_n(j)}} \right],$$

---

[18] The number of active producers can range from 1, supplying good j to the whole world, to 47, with each producer supplying only the local market. Even though our goal is to learn what the model has to say about U.S. plant behavior, we need to consider what goes on in the entire set of 47 countries. Whether a U.S. producer exports to France, for example, depends, among other things, on its ability to edge out a German rival.

where the composite markup of the U.S. plant $M^C(j)$ is total revenues over total costs. The plant-level productivity measure we are trying to match is value added per worker. Since we assume that production combines intermediates and labor, with labor having a share β, value added per worker is:

$$v(j) = \frac{W}{\beta}\left[\frac{y(j)}{w} - (1-\beta)\right], \tag{19}$$

where W is the U.S. wage.[19]

We generate aggregate predictions by taking 500,000 samples from the continuum of goods. About 75,000 of these turn out to be goods that the United States imports, leaving us with a data set for around 425,000 simulated U.S. plants. From this data set we calculate the simulated analogues of the actual statistics reported in Table 1: (i) the fraction of U.S. plants that export, (ii) the revenues of exporters relative to nonexporters, (iii) the distribution of export intensity for exporters, (iv) the heterogeneity of productivity, and (v) the productivity advantage of exporters. We now turn to what happens.

5.3 Simulation Results

Table 3, giving each plant equal weight in the statistics, and Table 4, weighting by shipments, compare the predictions from our simulations (under different values of θ and σ) to the actual statistics from the 1992 U.S. Census of Manufacturing. We now turn to our results fact by fact.

1. **The Fraction who Export.** A basic prediction of our framework (which, as shown in Appendix C, does not rely on our estimates of θ or σ) is the fraction of plants that export at all. Our model's prediction that 51 percent of plants export is substantially above the 21 percent of plants that report exporting anything in 1992. When we weight by shipments, however, our predictions are quite close to the actual fraction of 60 percent for

---

[19] Since intermediates are a share $1-\beta$ of total cost and the ratio of total revenues to total cost is $y(j)/w$, value added relative to total cost is $y(j)/w - (1-\beta)$. Furthermore, the ratio of total costs to labor costs is $1/\beta$ and hence the ratio of total cost to labor is W/β. Multiplying value added relative to total costs by total costs relative to labor gives (19).

low to moderate values of the elasticity of substitution, σ. An explanation (admittedly favorable to the model) for why weighting by size improves our predictive ability is that a large number of small plants erroneously report no export activity. (Recall that total exports reported by manufacturing plants in the Census survey constitute just over 60 percent of total aggregate U.S. manufacturing exports as measured by OECD.)

2. **Productivity.** Table 3 reports the simulated standard deviation of the logarithm (multiplied by 100 to yield a percentage) of value added per worker (19). Note that the model implies substantial heterogeneity, but not as much as appears in the data. The simulated variation in productivity is smaller with less heterogeneity either in consumption (σ higher) or in production (θ higher). More homogeneity of either type forces plants to compress their markups toward the competitive level, reducing differences in measured productivity. Table 3 also shows the difference between exporters and nonexporters in the mean of log value-added per worker (multiplied by 100 to approximate a percentage difference). Our simulations straddle the actual exporter productivity advantage of 33 per cent, with more heterogeneity (lower σ or θ) implying a greater differential.

3. **Exporters' Size Advantage.** The next two rows of Table 3 show the average size of exporters relative to nonexporters. Even excluding export revenues, as long as demand is elastic ($\sigma > 1$), the model captures the fact that exporters are bigger. We overpredict the size differences in the upper ranges of σ allowed by our estimates of θ.

4. **The Fraction of Revenues from Exports.** For all parameter values we predict that the vast majority of exporters earn less than 10 percent 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, fewer than 5 percent 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 the elasticity of substitution

σ is large. We do about as well weighting by size.

In summary, our model not only picks up the qualitative features of the plant-level data, parameterizing the model with aggregate trade data we can go quite far in fitting the quantitative magnitudes. Our inability to fit all the facts perfectly is not surprising since we are trying to match 15 moments of the plant-level data with only two free parameters (σ and θ). Standing by our preferred estimate of $\theta = 8.28$ from EK, we are down to one.

The main trade off in choosing σ is capturing the size advantage of exporters (requiring high σ) versus capturing more of the heterogeneity in productivity (favoring low σ). In turning to the counterfactuals below we compromise with a value of $\sigma = 6$.

6 General Equilibrium

We have been able to infer the connection between aggregate trade flows and plant level facts from the model taking input costs and trade patterns as given. But in using the model to infer the effects of exogenous changes in the global environment, we need to specify how these magnitudes respond.

To close the model in the simplest way, we assume that there is a tradeable nonmanufactured good which can serve as our numeraire. Each country n produces this good competitively with labor productivity W_n. The manufacturing sector in country n therefore faces an elastic supply of labor at wage W_n. (EK describe other ways of closing the model.)

Recall that inputs are a Cobb-Douglas combination of labor and intermediates. We treat intermediates as representative of manufactures generally (both in terms of their prices and how they aggregate). Their price index in country i is thus simply P_i, so that the input cost index w_i is proportional to $W_i^{\beta} P_i^{1-\beta}$.

Given wages, manufacturing price levels in different countries are connected through trade in intermediates. To take these interactions into account we manipulate equations (15) and

(8) to obtain:

$$P^{-\theta} = \Lambda P^{-\theta(1-\beta)}, \tag{20}$$

where the nth element of the vector P^x is P_n^x and the element in the nth row and ith column of the matrix Λ is proportional to $T_i W_i^{-\theta\beta} d_{ni}^{-\theta}$. We solve for the endogenous response of prices to the exogenous shocks considered in our counterfactuals using a log-linear approximation to (20).

Having determined how prices change, we can easily calculate the change in input cost w in each country. Using equation (9) we can then calculate changes in the market share π_{ni} of any country i in any other country n. The remaining step is to calculate changes in manufacturing absorption in each country.

We take each country's aggregate GDP Y_n as exogenous and assume that a fixed fraction γ_n is devoted to manufactured final goods. Given imperfect competition, we can show that aggregate costs are a fraction $\theta/(1+\theta)$ of aggregate revenues. It follows that the vector of manufacturing absorptions satisfy:

$$X = \frac{\theta}{1+\theta}(1-\beta)\Pi' X + \gamma Y, \tag{21}$$

where the nth element of the vector X is X_n, the representative element of the vector γY is $\gamma_n Y_n$, and the representative element of the matrix Π is π_{ni}. (The first term on the right side of equation (21) represents demand for intermediates while the second term represents final demand for manufactures.) We use equation (21) to calculate how a change in Π translates into a change in X. Together, the changes in Π and X determine the new values of X_{ni} for each country n and i.

7 Counterfactuals

We consider three types of aggregate shocks to the world trading regime: (i) a 5 percent worldwide decline in geographic barriers (resulting in 39 percent more world trade), (ii) a rise in geographic barriers so extreme that trade shuts down, and (iii) a 10 percent exogenous appreciation of the U.S. wage relative to wages in other countries. We compare each counterfactual situation to a baseline, holding fixed the efficiency levels of all potential producers.

For each counterfactual ask: (i) How much entry and exit occurs, both in and out of production and in and out of exporting? (ii) What happens to a conventional measure of overall U.S. manufacturing productivity and what are the contributions of entry, exit, and reallocation among surviving incumbents? (iii) What happens to total employment, job creation, and job destruction in manufacturing?

Before turning to the results themselves, we explain our productivity measure and its components.

7.1 Productivity Accounting

In assessing the impacts of our counterfactuals on measured productivity we look at total manufacturing value added divided by manufacturing employment. Previously we considered productivity at a given moment across a given set of plants facing the same input prices. We now have to account for the role of entry, exit, reallocation, and changes in input costs and prices.

Starting at the plant level, we modify (19) by defining $q(j) = v(j)/P$ to take account of changes in the manufacturing price level. (Since from now on we consider only U.S. plants we drop the subscript i.) Aggregating across plants, overall manufacturing productivity q is:

$$q = \sum_j \frac{L(j)}{L} q(j),$$

where $L(j)$ is employment in plant j.

Following Baily, Hulten, and Campbell (1992), we decompose aggregate productivity growth into the contributions of entering plants (n), exiting plants (x), reallocation among surviving incumbents (s), and productivity gains for surviving incumbents. Denoting the set of plants of each type as $\Omega_k,\ k = n, x, s$:

$$\begin{aligned} q' - q &= \frac{L'_n}{L'}\left(q'_n - q'_s\right) - \frac{L_x}{L}\left(q_x - q_s\right) \\ &\quad + \sum_{j\in\Omega_s}\left(\frac{L'(j)}{L'_s} - \frac{L(j)}{L_s}\right) q'(j) + \sum_{j\in\Omega_s}\frac{L(j)}{L_s}\left(q'(j) - q(j)\right), \end{aligned} \tag{22}$$

where $L_k = \sum_{j\in\Omega_k} L(j)$, $q_k = \sum_{j\in\Omega_k} \frac{L(j)}{L_k} q(j)$, and x' denotes the counterfactual value of variable x. The first term represents the productivity contribution from entrants whose productivity levels differ on average from that of surviving incumbents. The second term represents the corresponding productivity contribution from plant exit. The third is the contribution from reallocation across incumbent survivors. The fourth is the contribution of productivity changes within the incumbent survivors. (We will present each term as a percentage of initial productivity q.)

It might seem that since we are holding a plant's efficiency draw fixed, our various counterfactuals should leave an incumbent survivor's measured productivity unaffected, so that the last term in (22) would equal zero. But this is not the case. Incorporating cost minimizing behavior, a plant's deflated value added per worker is:

$$q(j) = Z_1(j)\frac{P^C(j)}{P}\left[1 - \frac{(1-\beta)}{M^C(j)}\right]\left(\frac{W}{P}\right)^{1-\beta}\left(\frac{1-\beta}{\beta}\right)^{1-\beta}, \tag{23}$$

where the composite price $P^C(j)$ is the plant's revenues from around the world over its total physical output. In our experiments an incumbent survivor's measured labor productivity can rise for any of three reasons: (i) the plant's own output price $P^C(j)$ rises relative to the price of manufactures P, (ii) the composite markup $M^C(j)$ rises, raising the share of value added in gross production, or (iii) the price of intermediates P falls relative to the wage W, leading to use of more intermediates per worker. The last effect turns out to be the dominant one quantitatively, as our experiments generate substantial changes in the ratio of intermediates prices to the wage.

7.2 Counterfactual Outcomes

The results of all three counterfactuals are shown in Table 5:

1. Globalization (taking the form of a 5 percent fall in geographic barriers) leads to more than a 4 percent increase in our productivity measure. The main factor is the gains within surviving plants driven by the decline in the price of intermediates (as cheaper imports replace domestically produced inputs). But the reallocation of production is also important. Over 8 percent of U.S. plants exit. Since their productivity averages only 70 percent of the survivors', exit contributes 1.5 percent to overall productivity. As smaller, lower-productivity plants exit, high-productivity exporters expand. Hence net job loss is only 2.7 per cent of initial employment, a much lower percent than plant loss. This figure is the net outcome of 4.6 percent gross job creation and 7.3 percent gross job destruction.

2. Not surprisingly, going to autarky moves things the other way. Measured productivity falls by almost 9 percent, primarily due to the increase in intermediates prices facing incumbents. In addition, inefficient U.S. producers enter to produce the goods that had been imported, dragging down overall productivity by 1.5 percent. The reallocation of production away from productive exporters knocks 2.4 percentage points from overall productivity growth. Although the number of U.S. plants rises by 17 percent, manufacturing employment rises by less than 2 percent: The jobs destroyed in exporting plants losing their overseas markets (nearly 11 percent of initial employment) almost offsets the job-creation from relatively small entrants (over 12 percent of employment).

3. A loss in U.S. "competitiveness" (taking the form of a 10 percent rise in the U.S. wage relative to wages elsewhere) actually pushes measured U.S. manufacturing productivity up by 4 percent. The primary reason is that imports keep intermediates prices from rising by as much as the wage, so that plants substitute intermediates for workers.

> Exit by unproductive domestic producers contributes an additional 1.5 percent to the overall productivity gain. Together substitution and exit generate an 18 percent fall in manufacturing employment.

To show what kind of churning goes on at the plant level Tables 6 and 7 illustrate transitions in and out of production and in and out of exporting for the first and third counterfactuals, respectively.

Table 6 shows what happens as a consequence of globalization. The action is among plants initially not exporting: While 17 percent of nonexporters are shut down by foreign competition, 12 percent take advantage of new export opportunities. Initial productivity is a good indicator of how a nonexporter will fare: Almost a third of those in the lowest quartile exit while only 7 percent enter export markets. But none of the plants in the top productivity quartile shut down, and over one third enter export markets.

Turning to Table 7, we see the result of a 10 percent increase in the U.S. wage. Of the U.S. plants originally producing for only the U.S. market, 16 percent exit. Only 1 per cent of exporters shut down, but 24 percent do stop exporting. Breaking down these statistics by a plant's initial position in the productivity distribution, nearly one third of low productivity nonexporting plants exit while none exit from the top half of the productivity distribution. For either counterfactual, we see a striking heterogeneity of outcomes from aggregate shocks.

8 Conclusion

Recent plant-level findings pose challenges to standard trade theory. Most notably, plants that export are scattered across industries, even exporters earn most of their revenues domestically, and productivity differs dramatically across plants within an industry. We reconcile what goes on at the plant level with a fully articulated and calibrated model of international trade. Our framework captures the stylized facts qualitatively, and goes quite far in matching data on U.S. manufacturing plants. The framework points to the importance of export costs in segmenting markets, and of efficiency differences across producers in generating heterogeneity in market power, measured productivity, and the ability to overcome geographic barriers.

Although foreign markets are small in plants' revenues, the international economy nonetheless plays an important role in determining which producers are in business and which are good enough to export. Simulations of counterfactuals illustrate the potentially diverse impact at the plant level of aggregate policy shifts. Lower trade barriers, for example, tend to nudge out low-productivity plants while enabling the more productive to sell more abroad. Even though the number of U.S. plants fall there is little net job destruction (but substantial job turnover). Aggregate productivity rises as employment shifts from low productivity plants driven out by import competition to high productivity plants turning toward export markets.

Our model captures very parsimoniously the remarkable heterogeneity of plant-level experience. To achieve this parsimony it omits many important features of the world. We ignore possible differences across industries in relevant parameter values. We treat labor as homogeneous and have simply lumped capital together with intermediates. We ignore dynamics entirely. In principle one could extend our approach to incorporate these features (and such extensions remain topics for future research). But our theory has already gone much further than previous work in bridging the gap between macro and micro-level trade data.

References

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

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

Bartelsman, Eric J. and Dhrymes, Phoebus J. "Productivity Dynamics: U.S. Manufacturing Plants, 1972-1986." *Journal of Productivity Analysis 9* (January 1998): 5-34.

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 1999a): 1-25.

Bernard, Andrew B. and Jensen, J. Bradford. "Exporting and Productivity" *NBER Working Paper No. 7135* (1999b).

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

Campa, José and Goldberg, Linda S. "Investment, Exchange Rates and External Exposure." *Journal of International Economics 38* (May 1995): 297-320.

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

Dixit, Avinash and Stiglitz, Joseph. "Monopolistic Competition and Optimum Product Diversity." *American Economic Review 67* (June 1977): 297-308.

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." *International Economic Review 40* (August 1999): 537-570.

Eaton, Jonathan and Kortum, Samuel. "Technology, Geography, and Trade." mimeo, Boston University (2000).

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).

Gourinchas, Pierre-Olivier. "Exchange Rates and Jobs: What Do We Learn from Job Flows?" in *NBER Macroeconomics Annual 1998*, Ben S. Bernanke and Julio J. Rotemberg (eds.), The MIT Press, 1999.

Grossman, G.M. and Helpman, E. "Quality Ladders in the Theory of Growth." *Review of Economic Studies 58* (1991): 43-61.

Head, Keith and Ries, John. "Rationalization Effects of Tariff Reductions." mimeo, University of British Columbia (1998).

Klette, Tor Jakob, and Griliches, Zvi. "The Inconsistency of Common Scale Estimators When Output Prices are Unobserved and Endogenous." *Journal of Applied Econometrics 11* (1996): 343-361.

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

Krugman, Paul. “Increasing Returns, Monopolistic Competition, and International Trade.” *Journal of International Economics 9* (November 1979): 469-479.

Melitz, Marc. “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity.” mimeo, University of Michigan (1999).

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 Colombia: An Empirical Model of Entry with Sunk Costs.” *American Economic Review 87* (1997): 545-564.

UNIDO,*Industrial Statistics Database,* 1999.

World Bank, *World Tables.* Baltimore: Johns Hopkins University Press, 1995.

A Mathematical Appendix

In this appendix we derive some of the results used in the paper, beginning with those from Section 3.2.

A.1 The Joint Distribution of Efficiency for the Best and Second-Best

The Fréchet distribution $e^{-Tz^{-\theta}}$ has two useful properties for our purposes: (i) For a positive constant λ, if $Z_1(j)$ has a Fréchet distribution, then so to does $\lambda Z_1(j)$ and (ii) If $Z_{1a}(j)$ and $Z_{1b}(j)$ are drawn independently from Fréchet distributions with a common parameter θ (but possibly different parameters T) then $Z_1(j) = \max\{Z_{1a}(j), Z_{1b}(j)\}$ itself has a Fréchet distribution. We now demonstrate that the analogues of these properties hold for our proposed generalization (12) of the Fréchet distribution to cover the joint distribution of the maximum $Z_1(j)$ and the next highest $Z_2(j)$:

$$F(z_1, z_2; T) = \Pr[Z_1(j) \leq z_1, Z_2(j) \leq z_2] = [1 + T(z_2^{-\theta} - z_1^{-\theta})]e^{-Tz_2^{-\theta}},$$

for $0 \leq z_2 \leq z_1$. (To facilitate the derivations below, we have made explicit that T parameterizes F.)

The first property of the joint distribution is verified as follows:

$$\begin{aligned}
\Pr[\lambda Z_1(j) \leq z_1, \lambda Z_2(j) \leq z_2] &= F(z_1/\lambda, z_2/\lambda; T) \\
&= \{1 + T[(z_2/\lambda)^{-\theta} - (z_1/\lambda)^{-\theta}]\}^{-T(z_2/\lambda)^{-\theta}} \\
&= F(z_1, z_2; T'),
\end{aligned}$$

where $T' = \lambda^{\theta}T$.

To verify the second property, we take the pair $(Z_{1a}(j), Z_{2a}(j))$ to be drawn from the distribution $F(z_1, z_2; T_a)$ and the pair $(Z_{1b}(j), Z_{2b}(j))$ to be drawn independently from $F(z_1, z_2; T_b)$. Furthermore, we define

$$Z_1(j) = \max\{Z_{1a}(j), Z_{2a}(j), Z_{1b}(j), Z_{2b}(j)\}$$

and

$$Z_2(j) = \max 2\{Z_{1a}(j), Z_{2a}(j), Z_{1b}(j), Z_{2b}(j)\},$$

where $\max 2$ denotes the second highest from the set. For $z_1 > z_2$, the event

$$[Z_1(j) \leq z_1, Z_2(j) \leq z_2]$$

can be broken down exhaustively into three mutually exclusive events:

$$[Z_{1a}(j) \leq z_2, Z_{1b}(j) \leq z_2],$$

$$[z_2 < Z_{1a}(j) \leq z_1, Z_{2a}(j) \leq z_2, Z_{1b}(j) \leq z_2],$$

and

$$[z_2 < Z_{1b}(j) \leq z_1, Z_{2b}(j) \leq z_2, Z_{1a}(j) \leq z_2].$$

Applying this breakdown:

$$\begin{aligned}
\Pr[Z_1(j) \leq z_1, Z_2(j) \leq z_2] &= F(z_2, z_2; T_a)F(z_2, z_2; T_b) \\
&\quad +[F(z_1, z_2; T_a) - F(z_2, z_2; T_a)]F(z_2, z_2; T_b) \\
&\quad +[F(z_1, z_2; T_b) - F(z_2, z_2; T_b)]F(z_2, z_2; T_a) \\
&= [1 + T_a(z_2^{-\theta} - z_1^{-\theta}) + T_b(z_2^{-\theta} - z_1^{-\theta})]e^{-T_a z_2^{-\theta}} e^{-T_b z_2^{-\theta}} \\
&= F(z_1, z_2; T_a + T_b),
\end{aligned}$$

which is the property we sought.

A.2 The Joint Distribution of Lowest and Next-Lowest Cost

We now turn from the distribution of efficiencies in each source country (12) to the closely related distribution of the costs of providing goods to each destination. It is convenient to work with the complementary distribution (with inequalities reversed):

$$G_{ni}^c(c_1, c_2) = \Pr[C_{1ni}(j) \geq c_1, C_{2ni}(j) \geq c_2]$$

$$
\begin{aligned}
&= \Pr[Z_{1i}(j) \le w_i d_{ni}/c_1, Z_{2i}(j) \le w_i d_{ni}/c_2] \\
&= F_i(w_i d_{ni}/c_1, w_i d_{ni}/c_2) \\
&= F(c_1^{-1}, c_2^{-1}; T_i[w_i d_{ni})]^{-\theta}),
\end{aligned}
$$

where, as above, we denote $F_i(z_1, z_2)$ by $F(z_1, z_2; T_i)$.

Now consider the complementary distribution for the lowest and next-lowest costs for goods available in destination n, without regard to the source. For $c_1 \le c_2$:

$$
\begin{aligned}
G_n^c(c_1, c_2) &= \Pr[C_{1n}(j) \ge c_1, C_{2n}(j) \ge c_2] \\
&= \prod_{i=1}^{N} G_{ni}^c(c_2, c_2) + \sum_{i=1}^{N}[G_{ni}^c(c_1, c_2) - G_{ni}^c(c_2, c_2)] \prod_{k \ne i} G_{nk}^c(c_2, c_2) \\
&= \prod_{i=1}^{N} e^{-T_i(w_i d_{ni})^{-\theta} c_2^\theta} + \sum_{i=1}^{N} T_i(w_i d_{ni})^{-\theta}(c_2^\theta - c_1^\theta) e^{-T_i(w_i d_{ni})^{-\theta} c_2^\theta} \prod_{k \ne i} e^{-T_k(w_k d_{nk})^{-\theta} c_2^\theta} \\
&= e^{-\Phi_n c_2^\theta} + e^{-\Phi_n c_2^\theta}(c_2^\theta - c_1^\theta) \sum_{i=1}^{N} T_i(w_i d_{ni})^{-\theta} \\
&= F(c_1^{-1}, c_2^{-1}; \Phi_n),
\end{aligned}
$$

where Φ_n is the cost parameter defined in (8).

We obtain the cost distribution (13) from the relationship,

$$
G_n(c_1, c_2) = 1 - G_n^c(0, c_2) - G_n^c(c_1, c_1) + G_n^c(c_1, c_2).
$$

A.3 The Distribution of the Markup

Defining $M_n'(j) = C_{2n}(j)/C_{1n}(j)$, the markup is simply $M_n(j) = \min\{M_n'(j), \bar{m}\}$. First, consider the distribution of $M_n'(j)$ given $C_{2n}(j) = c_2 \ge 0$. For any $m' \ge 1$ we have:

$$
\begin{aligned}
\Pr[M_n'(j) \le m' | C_{2n}(j) = c_2] &= \Pr[c_2/m' \le C_{1n}(j) \le c_2 | C_{2n}(j) = c_2] \\
&= \frac{\int_{(c_2/m')}^{c_2} g_n(c_1, c_2) dc_1}{\int_0^{c_2} g_n(c_1, c_2) dc_1} \\
&= \frac{c_2^\theta - (c_2/m')^\theta}{c_2^\theta} \\
&= 1 - m'^{-\theta},
\end{aligned}
$$

where $g_n(c_1, c_2)$ is the joint density function corresponding to the distribution (13). Since conditional on $C_{2n}(j) = c_2$ the distribution of $M'_n(j)$ is Pareto and does not depend on c_2, it follows that the unconditional distribution of $M'_n(j)$ is also Pareto. The distribution of the markup $H(m) = \Pr[M_n(j) \leq m]$ is therefore Pareto, but truncated at $\bar{m}$.

A.4 The Price Index

The main step in deriving an expression for the price index P_n is obtaining an expression for the expectation, $E[P_n(j)^{1-\sigma}]$. The two are closely linked since:

$$P_n^{1-\sigma} = E[P_n^{1-\sigma}] = E[\int_0^1 \alpha(k) P_n(k)^{1-\sigma} dk] = E[\alpha(j) P_n(j)^{1-\sigma}] = E[P_n(j)^{1-\sigma}],$$

where the last equality follows from our assumption that $\alpha(j)$ is independent of efficiency.

¿From A.3, $M'_n(j) = C_{2n}(j)/C_{1n}(j)$ has a Pareto distribution and is independent of $C_{2n}(j)$. Assuming $\sigma < 1 + \theta$, we therefore have:

$$\begin{aligned} E[P_n(j)^{1-\sigma}] &= \int_1^\infty E[P_n(j)^{1-\sigma} | M'_n(j) = m'] \theta m'^{-(\theta+1)} dm' \\ &= \int_1^{\bar{m}} E[C_{2n}(j)^{1-\sigma}] \theta m'^{-(\theta+1)} dm' + \int_{\bar{m}}^\infty E[(\bar{m} C_{2n}(j)/m')^{1-\sigma}] \theta m'^{-(\theta+1)} dm' \\ &= E[C_{2n}(j)^{1-\sigma}][(1 - \bar{m}^{-\theta}) + \bar{m}^{-\theta} \frac{\theta}{1 + \theta - \sigma}]. \end{aligned}$$

¿From (13) we can derive the density function for $C_{2n}(j)$, which we denote $g_{2n}(c_2)$. We can then calculate

$$E[C_{2n}(j)^{1-\sigma}] = \int_0^\infty c_2^{1-\sigma} g_{2n}(c_2) dc_2 = \Phi_n^{-(1-\sigma)/\theta} \int_0^\infty x^{(1-\sigma+\theta)/\theta} e^{-x} dx = \Phi_n^{-(1-\sigma)/\theta} \Gamma(\frac{1-\sigma+2\theta}{\theta}).$$

Combining these results:

$$P_n^{1-\sigma} = \Gamma(\frac{1-\sigma+2\theta}{\theta})(1 + \frac{\sigma - 1}{\theta - (\sigma - 1)} \bar{m}^{-\theta}) \Phi_n^{-(1-\sigma)/\theta}.$$

Raising both sides of the equation to the power $1/(1-\sigma)$ yields (15).

A.5 The Markup Conditional on Efficiency

Consider the distribution of $M'_n(j) = C_{2n}(j)/C_{1n}(j)$ conditional on $C_{1n}(j) = c_1 \geq 0$. For any $m' \geq 1$ we have:

$$\begin{aligned}
\Pr[M'_n(j) \leq m' | C_{1n}(j) = c_1] &= \Pr[c_1 \leq C_{2n}(j) \leq m'c_1 | C_{1n}(j) = c_1] \\
&= \frac{\int_{c_1}^{m'c_1} g_n(c_1, c_2)dc_2}{\int_{c_1}^{\infty} g_n(c_1, c_2)dc_2} \\
&= \frac{e^{-\Phi_n c_1^\theta} - e^{-\Phi_n (m'c_1^\theta}}{e^{-\Phi_n c_1^\theta}} \\
&= 1 - e^{-\Phi_n c_1^\theta (m'^\theta - 1)}.
\end{aligned}$$

Suppose that good j is supplied by a producer from country n. Then, $C_{1n}(j) = w_n/Z_{1n}(j)$ so that conditioning on $C_{1n}(j) = c_1$ is the same as conditioning on $Z_{1n}(j) = w_n/\bar{c}_1 = z_1$. In other words, $c_1 = w_n/z_1$. Thus, for $1 \leq m \leq \bar{m}$ we get,

$$\Pr[M_n(j) \leq m | Z_{1n}(j) = z_1] = 1 - e^{-\Phi_n (w_n/z_1)^\theta (m^\theta - 1)}.$$

A.6 Efficiency Conditional on the Markup

Suppose that we could observe $M'_n(j) = C_{2n}(j)/C_{1n}(j)$. Consider the distribution of $C_{1n}(j)$ conditional on $M'_n(j) = m'$:

$$\begin{aligned}
\Pr[C_{1n}(j) \leq c_1 | M'_n(j) = m'] &= \Pr[C_{2n}(j) \leq m'c_1 | M'_n(j) = m'] \\
&= \Pr[C_{2n}(j) \leq m'c_1] \\
&= G_{2n}(m'c_1) \\
&= G_n(m'c_1, m'c_1),
\end{aligned}$$

where we have used the result above about the independence of $M_n(j)$ and $C_{2n}(j)$. It follows that

$$\Pr[C_{1n}(j) \leq c_1 | M'_n(j) = m'] = \Pr[C_{1n}(j)/m' \leq c_1 | M'_n(j) = 1],$$

so that a shift up in M' is equivalent to a shift down in costs by the same factor.

Consider goods j_a and j_b that are each supplied to country n by local producers. Let $M_n(j_a) = m_a$ and $M_n(j_b) = m_b$, and suppose $m_a = \lambda m_b$ for $\lambda > 1$. Ignoring exports, according to (17), we will observe the productivity of the j_a producer exceeding that of the j_b producer by the same factor λ. If $m_a < \bar{m}$ then $E[C_{1n}(j_a)] = E[C_{1n}(j_b)]/\lambda$ and hence $E[Z_{1n}(j_a)] = \lambda E[Z_{1n}(j_b)]$. If $m_a = \bar{m}$, then typically $M_n'(j_a) > m_a$ and so $E[Z_{1n}(j_a)] > \lambda E[Z_{1n}(j_b)]$.

B Data Appendix

Our empirical work combines macro-level observations on bilateral trade and production in manufacturing with micro-level statistics calculated from observations of individual U.S. manufacturing establishments. We describe each in turn.

B.1 Aggregate Trade Data

We chose our sample of countries as follows (the 47 countries or regions are listed in Table 2). We started with the 52 countries that import the most from the United States. To avoid problems of entrepot trade we combined Hong Kong with China and Singapore with Malaysia. A remaining anomaly is the large U.S. market share in manufacturing absorption of a number of countries in the Caribbean Basin (Costa Rica, the Dominican Republic, Guatemala, and Panama). U.S. exports to these countries turn out to be dominated by apparel and textile products. This trade is essentially legislated by preferential trading agreements (the Caribbean Basin Initiative and Special Access Program 807A of the U.S. Harmonized Tariff) which give U.S. manufacturers a strong incentive to outsource the production of apparel from fabric formed and cut in the United States. These programs grossly inflate the U.S. share in these countries' absorption of manufactures. We deal with the problem by consolidating Caribbean Basin Countries with Mexico, whose size swamps the influence of apparel trade governed by these statutes. (Dealing with this phenomenon properly in our framework would require pursuing an industry level analysis.)

Bilateral trade (X_{ni}) among these countries (in millions of U.S. Dollars) is from Feenstra, Lipsey, an Bowen (1997). Starting with the file WBEA92.ASC, we aggregate over all manufacturing industries.

Data on 1992 gross production in manufacturing in millions of U.S. Dollars came from three sources. When possible we used the data published by the OECD (1995). If that was unavailable we used gross production data from UNIDO (1999). In a few cases, we resorted to value added in manufacturing from World Bank (1995), scaling up the numbers by the factor 2.745 to make them consistent with gross production. Some basic statistics, as well as additional information on our data source for each country, are in Table 2.

We get home purchases X_{nn} by subtracting total exports of manufactures from 1992 gross manufacturing production. Total manufacturing absorption is $X_n = \sum_{i=1}^{47} X_{ni}$, where X_{ni} is the imports by country n of manufactures produced in country i. There is some undercounting since we do not have all the countries of the world. The last two columns of Table 2 suggest undercounting is not a serious problem.

B.2 Plant-Level Data

We extract our plant-level facts from the 1992 U.S. Census of Manufactures in the Longitudinal Research Database of the Bureau of the Census (see Bernard and Jensen, 1999a). The 1992 Census includes over 200,000 plants. It provides 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, specifically "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 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.

The construction of total factor productivity for each plant (as shown in Table 1) is described in Bernard and Jensen (1999b).

C Simulation Methodology

For each good j that we sample we perform the following:

1. **Who Sells Where?** We first draw a sample of lowest costs for producing good j in each country. The cost distribution $G_{1ii}(c_1) = \Pr[C_{1ii}(j) \leq c_1]$ of the most efficient local provider is given by (6) after setting $n = i$. We actually work with a simple transformation of costs, chosen so that the distribution of transformed costs is free of any parameters. In particular the term:

$$U_{1ii}(j) = T_i w_i^{-\theta} \left[C_{1ii}(j)\right]^{\theta}$$

has a unit exponential distribution: $\Pr[U_{1ii}(j) \leq u_1] = 1 - e^{-u_1}$. For each of our N countries we draw $U_{1ii}(j)$ independently from this distribution. Next we define, for $n \neq i$:

$$U_{1ni}(j) = T_n w_n^{-\theta} C_{1ni}(j)^{\theta} = T_n w_n^{-\theta} [d_{ni} C_{1ii}(j)]^{\theta}.$$

>From (9) and (16) we can calculate $U_{1ni}(j)$ from $U_{1ii}(j)$ and aggregate trade shares since:

$$U_{1ni}(j) = \frac{T_n w_n^{-\theta}}{T_i (w_i d_{ni})^{-\theta}} U_{1ii}(j) = \frac{\pi_{nn}}{\pi_{ni}} U_{1ii}(j) = \frac{X_{nn}}{X_{ni}} U_{1ii}(j).$$

These transformed costs tell us where each country is buying good j: Country n buys from country i if and only if $U_{1ni}(j) \leq U_{1nk}(j)$ for all k (since $U_{1ni}(j) \leq U_{1nk}(j)$ if and only if $C_{1ni}(j) \leq C_{1nk}(j)$). Hence, for this draw of good j, 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 the heterogeneity parameters θ or σ.

2. **Markups.** To get predictions about markups we must also draw from the distribution of second lowest cost conditional on our draw for lowest cost. In parallel with our

transformation of lowest cost, our transformation of second lowest cost is $U_{2ni}(j) = T_n w_n^{-\theta} \left[C_{2ni}(j)\right]^{\theta}$. It turns out that the distribution of $U_{2ii}(j)$ conditional on $U_{1ii}(j) = u_1$ is again "parameter free": $\Pr[U_{2ii}(j) \leq u_2 | U_{1ii}(j) = u_1] = 1 - e^{-u_2 + u_1}$. Whenever country i turns out to be an active producer, conditional on its realization of $U_{1ii}(j)$, we draw $U_{2ii}(j)$ from this shifted unit exponential distribution. We then calculate:

$$U_{2ni}(j) = \frac{T_n w_n^{-\theta}}{T_i (w_i d_{ni})^{-\theta}} U_{2ii}(j) = \frac{\pi_{nn}}{\pi_{ni}} U_{2ii}(j) = \frac{X_{nn}}{X_{ni}} U_{2ii}(j).$$

Suppose i turns out to be the low cost supplier to n (that is, $U_{1ni}(j) \leq U_{1nk}(j)$ for all k) so that $T_n w_n^{-\theta} C_{1n}(j)^{\theta} = U_{1n}(j) = U_{1ni}(j)$. We define, parallel to (11):

$$U_{2n}(j) = T_n w_n^{-\theta} C_{2n}(j)^{\theta} = \min\left\{ \min_{k \neq i}\{U_{1nk}(j)\}, U_{2ni}(j) \right\}.$$

We can then calculate the markup on this good in each country n 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 the heterogeneity parameters θ and σ (since the Dixit-Stiglitz markup $\bar{m}$ depends on σ).

3. **Revenues.** Knowing the markups producers charge in each destination, we can use expression (5) to figure out revenues in each country. Using expressions (9), (15), and (16) to replace unknown parameters with aggregate trade data, the revenue earned by a producer of good j from country i who has captured country n's market is:

$$X_n(j) = \alpha(j) X_n \left[M_n(j)/\gamma\right]^{1-\sigma} \left(U_{1ni}(j) \frac{X_n}{X_{nn}}\right)^{(1-\sigma)/\theta}. \qquad (24)$$

We can use this expression for $n = i$ to get its domestic sales (up to the unknown parameter $\alpha(j)$). We then calculate $x_{ni}(j) = X_n(j)/X_i(j)$, (which no longer depends on either $\alpha(j)$ or γ), setting $x_{ni}(j) = 0$ in markets that the producer from i fails to penetrate. We then manipulate the sum of $x_{ni}(j)$ across all foreign destinations $n \neq i$ to get the plant's export revenue as a fraction of total revenue.

4. **Productivity.** The generalization of expression (17) for an active producer in country i, representing total revenues relative to an index of inputs, is:

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

   which can be substituted into the expression for value added per worker (19).

With a collection of simulation steps in hand, we can simulate statistics for U.S. plants. We simulate the fraction of U.S. plants that export as the number of draws that yield a U.S. exporter relative to the number that yield an active U.S. producer. We simulate the revenues of U.S. exporters relative to nonexporters as the ratio of the average of expression (24) for draws yielding an exporting plant relative to this average for draws yielding a U.S. plant serving only the local market. (Averaging eliminates the heterogeneity that would be generated by $\alpha(j)$; taking ratios eliminates γ.) We obtain the distribution of export intensity by counting the number of draws that yield U.S. plants in each export intensity bin as a fraction of the number of draws yielding a U.S. exporter. We calculate heterogeneity of productivity as the standard deviation of the logarithm of (19) across all draws with an active U.S. producer. Finally, we simulate the productivity advantage of exporters as the difference in the average of the logarithm of (19) between draws yielding a U.S. exporter and draws yielding a U.S. plant serving only the local market. (Taking the standard deviation or the difference of logs eliminates the term W_i/β.)

Table 1: Plant-Level Facts

Exporter shares	Percentage of all plants	Percentage of total output
	21	60

Productivity	Standard deviation of log productivity (%)	Exporter less nonexporter avg. log productivity (%)
Labor productivity (LP)	76	33
LP, within industries	66	15
TFP, within industries	40	10

Exporter size advantage	Ratio of average U.S sales	Ratio of average total sales
	4.8	5.6

Export intensity (%)	Percentage of all exporters	Percentage of total output of exporters
0 to 10	66	58
10 to 20	16	19
20 to 30	7.7	9.6
30 to 40	4.4	6.1
40 to 50	2.4	2.4
50 to 60	1.5	2.2
60 to 70	1.0	1.0
70 to 80	0.6	1.0
80 to 90	0.5	0.4
90 to 100	0.7	0.3

The statistics are calculated from all plants in the 1992 Census of Manufactures. Labor productivity (LP) is measured as value added per worker. The construction of total factor productivity (TFP)–shipments not accounted for by labor, capital, or materials, and using estimated industry-specific output elasticities–is described in Bernard and Jensen (1999b). Heterogeneity is the standard deviation of the logarithm of productivity (LP or TFP), multiplied by 100. The productivity advantage of exporters is the difference (multiplied by 100) in the mean logarithm of productivity between exporting and nonexporting plants. Within industries indicates that we subtract (from the log of productivity for each plant) average log productivity of the appropriate 4-digit industry. The size advantage of exporters is the average shipments of exporting plants relative to the average for nonexporting plants, presented as a simple ratio.

Table 2: Aggregate Trade Data (continued on the following page)

#	Country	Data Source	U.S. Exports ($ mill.)	U.S. % Market Share	Imports from ROW (% of tot.)	Exports to ROW (% of tot.)
1	Arab Emirates	W	1590	7.9	5.3	36.4
2	Argentina	U	3498	3.8	2.8	11.9
3	Australia	O	8570	6.2	2.9	5.9
4	Austria	O	1785	1.5	5.9	12.5
5	Belgium & Luxembourg	O,U	6264	4.3	4.2	3.4
6	Brazil	U	5932	2.9	3.4	7.2
7	Canada	O	83400	32.9	0.7	0.7
8	Chile	U	2441	8.8	1.6	2.7
9	China & Hong Kong	U,U	16200	3.3	1.7	6.2
10	Colombia	U	3098	12.5	1.8	5.4
11	Denmark	O	1403	2.5	4.7	7.5
12	Ecuador	U	1035	14.9	1.2	1.1
13	Egypt	U	1665	7.1	6.5	22.7
14	Finland	O	914	1.8	2.7	4.9
15	France	O	16700	2.4	4.0	10.6
16	Germany (Unified)	O	23000	1.8	6.6	7.6
17	Greece	O	804	2.0	4.1	11.9
18	India	U	1624	1.3	8.4	9.7
19	Indonesia	U	2846	5.1	0.9	3.9
20	Ireland	U	2771	8.4	1.8	1.9
21	Israel	U	3251	10.0	1.5	8.7
22	Italy	O	8124	1.2	6.4	9.0
23	Japan	O	42100	1.6	2.5	4.3
24	Korea (South)	O	14100	5.5	1.4	6.9
25	Kuwait	U	1471	15.5	3.8	25.2
26	Mexico & Caribbean	O,U,W,U,U	43700	19.8	2.4	7.0
27	Netherlands	O	9362	5.1	2.2	4.6
28	New Zealand	O	1526	7.3	1.1	8.4
29	Nigeria	U	1012	6.5	1.2	1.0
30	Norway	O	1779	3.2	3.0	8.5
31	Paraguay	W	807	11.5	0.9	3.9
32	Peru	U	858	5.8	2.8	2.1
33	Phillipines	U	1667	4.9	1.1	1.4
34	Portugal	U	807	1.3	1.7	8.0

Explanatory notes are on the following page.

#	Country	Data Source	U.S. Exports ($ mill.)	U.S. % Market Share	Imports from ROW (% of tot.)	Exports to ROW (% of tot.)
35	Saudi Arabia	W	7145	14.4	2.1	8.1
36	Singapore & Malaysia	U,U	15000	14.1	1.5	6.4
37	Spain	O	5717	1.9	2.6	7.6
38	Sweden	O	3403	3.1	3.5	5.0
39	Switzerland	U	4222	3.2	2.7	5.2
40	South Africa	U	2106	3.0	1.5	6.1
41	Taiwan	U	14000	8.1	0.6	1.3
42	Thailand	U	4094	3.7	3.1	7.3
43	Turkey	U	2186	2.4	6.0	16.9
44	United Kingdom	O	22600	3.8	2.0	11.4
45	United States	O	2520300	85.0	1.8	3.5
46	USSR (Former)	U	2181	0.4	8.6	9.3
47	Venezuela	U	6390	17.2	1.4	10.2

The Caribbean Basin countries are Costa Rica, Dominican Republic, Guatemala, and Panama. All data are for 1992 and cover the manufacturing sector. Data on bilateral exports and imports (as measured by the importer) are from Feenstra, Lipsey, and Bowen (1997). The U.S. market share is a country's imports from the United States relative to it absorption of manufactures. Absorption is defined as gross manufacturing production minus total manufactured exports plus manufactured imports from the other countries in the sample. The data sources for gross manufacturing production (in order of our preference for using them) are: OECD (O), UNIDO (U), and World Bank (W). (In using UNIDO data: for Argentina we took the (weighted) geometric mean of the 1990 and 1993 figure, for Thailand we took the geometric mean of 1991 and 1993 figure, and for the former USSR we took the 1990 figure.) The World Bank provides only value added data, which we multiply by 2.745 (the average ratio of gross production to value added for 39 of the countries). The United States' imports from itself are defined as gross manufacturing production less all exports. Imports from ROW are reported as imports from countries not in the sample as a percentage of all imports (exports to ROW are defined in a parallel fashion).

Table 3: Simulated Exporting Facts

	$\theta = 3.60$		$\theta = 8.28$				Census
	$\sigma = 1$	$\sigma = 4$	$\sigma = 1$	$\sigma = 4$	$\sigma = 6$	$\sigma = 8$	Data
	Fraction of Plants that Export (%):						
	51	51	51	51	51	51	21
	Heterogeneity of Productivity (%):						
	48	31	28	26	22	18	75
	Productivity Advantage of Exporters (%):						
	50	30	29	27	22	17	33
	Size Advantage of Exporters (ratio of means):						
U.S. Sales	1.0	6.9	1.0	1.4	2.4	6.4	4.8
Total Sales	1.3	8.1	1.3	1.8	3.0	7.5	5.6
Export Intensity (%)	Fraction of Exporting Plants (%):						
0 to 10	49	80	49	55	64	79	66
10 to 20	26	18	26	25	25	19	16
20 to 30	9.2	1.4	9.5	8.5	6.9	1.3	7.7
30 to 40	5.0	0.0	4.9	4.5	3.8	0.0	4.4
40 to 50	3.2	0.0	3.2	3.4	0.0	0.0	2.4
50 to 60	2.5	0.0	2.4	3.1	0.0	0.0	1.5
60 to 70	2.5	0.0	2.5	0.1	0.0	0.0	1.0
70 to 80	2.5	0.0	2.5	0.0	0.0	0.0	0.6
80 to 90	0.0	0.0	0.0	0.0	0.0	0.0	0.5
90 to 100	0.6	0.6	0.5	0.6	0.6	0.6	0.7

The statistics in the first six columns are based on simulating the model (with $\theta = 8.28$ and $\sigma = 6$) by sampling 500,000 goods. The last column is calculated from all plants in the 1992 Census of Manufactures. Export intensity is exports as a percentage of total shipments. The productivity advantage of exporters is the difference (multiplied by 100) in the mean logarithm of value added per worker between exporting and nonexporting plants. The heterogeneity of productivity is the standard deviation across all plants of the logarithm of value added per worker (multiplied by 100). The size advantage of exporters is the average shipments of exporting plants relative to the average for nonexporting plants.

Table 4: Simulated Exporting Facts (Size-Weighted)

	$\theta = 3.60$		$\theta = 8.28$				Census
	$\sigma = 1$	$\sigma = 4$	$\sigma = 1$	$\sigma = 4$	$\sigma = 6$	$\sigma = 8$	Data
	Fraction of Output due to Plants that Export (%):						
	58	90	58	66	76	89	60
Export Intensity (%)	Fraction of Exporters' Output (%):						
0 to 10	40	25	40	40	35	26	58
10 to 20	23	63	23	23	22	65	19
20 to 30	9.3	12	9.5	9.5	15	8.7	9.6
30 to 40	5.8	0.0	5.8	6.8	27	0.0	6.1
40 to 50	4.4	0.0	4.4	7.8	0.2	0.0	2.4
50 to 60	4.1	0.0	4.1	13	0.0	0.0	2.2
60 to 70	5.6	0.0	5.5	0.5	0.0	0.0	1.0
70 to 80	7.5	0.0	7.5	0.0	0.0	0.0	1.0
80 to 90	0.0	0.0	0.0	0.0	0.0	0.0	0.4
90 to 100	0.0	0.0	0.0	0.0	0.0	0.0	0.3

The statistics in the first six columns are based on simulating the model (with $\theta = 8.28$ and $\sigma = 6$) by sampling 500,000 goods. The first row of statistics is the output of exporting plants as a percentage of the output of all plants. The remaining rows are the output of plants with a particular export intensity (exports relative to total shipments) as a percentage of the output of all exporting plants.

Table 5: Counterfactuals

Statistics for U.S. Producers	Counterfactual Experiment		
	5 % Lower Barriers	Autarky	10 % Higher U.S. rel. wage
Productivity Decomposition:			
Aggregate (% change)	4.3	-8.8	4.0
Entrants (% change)	0.0	-1.5	0.0
Exiters (% change)	1.5	0.0	1.5
Reallocation among survivors (% change)	0.3	-2.4	-0.6
Gains within survivors (% change)	2.5	-4.9	3.1
Plant Exit and Entry:			
Number of plants (% change)	-8.3	17	-8.3
Relative productivity of exiters (%)	70	–	70
Employment share (%), prior	4.9	–	5.1
Relative productivity of entrants (%)	–	81	–
Employment share (%), post	–	8.5	–
Employment:			
Total employment in mfg. (% change)	-2.7	1.4	-18
Job creation (%)	4.6	12	0.9
Job destruction (%)	7.3	11	19
International Trade:			
U.S. exports (% change)	28	-100	-31
U.S. imports (% change)	42	-100	21
U.S. absorption (% change)	-5.6	3.2	-17
World trade (% change)	39	-100	3.6

Results are based on simulating the model (with $\theta = 8.28$ and $\sigma = 6$) by sampling 500,000 goods, comparing each outcome under the counterfactual and under the baseline. Aggregate productivity is manufacturing value added deflated by the manufacturing price level and divided by manufacturing employment. The next four rows correspond to the decomposition of equation (22), shown as percentages of the initial level of productivity q. Relative productivity of exiters is calculated (prior to the counterfactual) as the employment-weighted average productivity of plants that would eventually exit divided by the employment-weighted average productivity of plants that would survive the counterfactual. Relative productivity of entrants is calculated in a parallel fashion but using post-counterfactual productivity. Job creation and job destruction are both shown as a percentage of initial manufacturing employment.

Table 6: Plant Level Transitions: 5 % Lower Geographic Barriers

Before Geographic Barriers Fall		After Geographic Barriers Fall		
	% of Plants	Exit (%)	Domestic (%)	Export (%)
All Plants:				
Domestic	49	17	71	12
Export	51	0	2	98
By Productivity Quartile:				
Lowest quartile				
Domestic	77	30	63	7
Export	23	1	3	96
Second quartile				
Domestic	66	14	76	10
Export	34	0	3	97
Third quartile				
Domestic	41	0	82	18
Export	59	0	2	98
Highest quartile				
Domestic	11	0	68	32
Export	89	0	0	100

Results are based on simulating the model (with $\theta = 8.28$ and $\sigma = 6$) by sampling 500,000 goods, comparing each outcome under the counterfactual and under the baseline. Every pair of numbers in the first column sum to 100 percent. The last three numbers in every row also sums to 100 percent (except, due to rounding). Following the decline in geographic barriers, a U.S. plant will either shutdown ("Exit"), produce only for the domestic market ("Domestic"), or continue to export ("Export").

Table 7: Plant Level Transitions: 10 % Higher U.S. Relative Wage

Before Appreciation of the U.S. Wage		After Appreciation of the U.S. Wage		
	% of Plants	Exit (%)	Domestic (%)	Export (%)
All Plants:				
Domestic	49	16	84	0
Export	51	1	24	76
By Productivity Quartile:				
Lowest quartile				
Domestic	77	30	70	0
Export	23	6	38	56
Second quartile				
Domestic	66	13	87	0
Export	34	1	37	62
Third quartile				
Domestic	41	0	100	0
Export	59	0	30	70
Highest quartile				
Domestic	11	0	100	0
Export	89	0	11	89

Results are based on simulating the model (with $\theta = 8.28$ and $\sigma = 6$) by sampling 500,000 goods, comparing each outcome under the counterfactual and under the baseline. Every pair of numbers in the first column sum to 100 percent. The last three numbers in every row also sums to 100 percent (except, due to rounding). Following the wage appreciation, an active U.S. plant will either shutdown ("Exit"), produce only for the domestic market ("Domestic"), or continue to export ("Export").

Figure 1: Industry Exporting Intensity

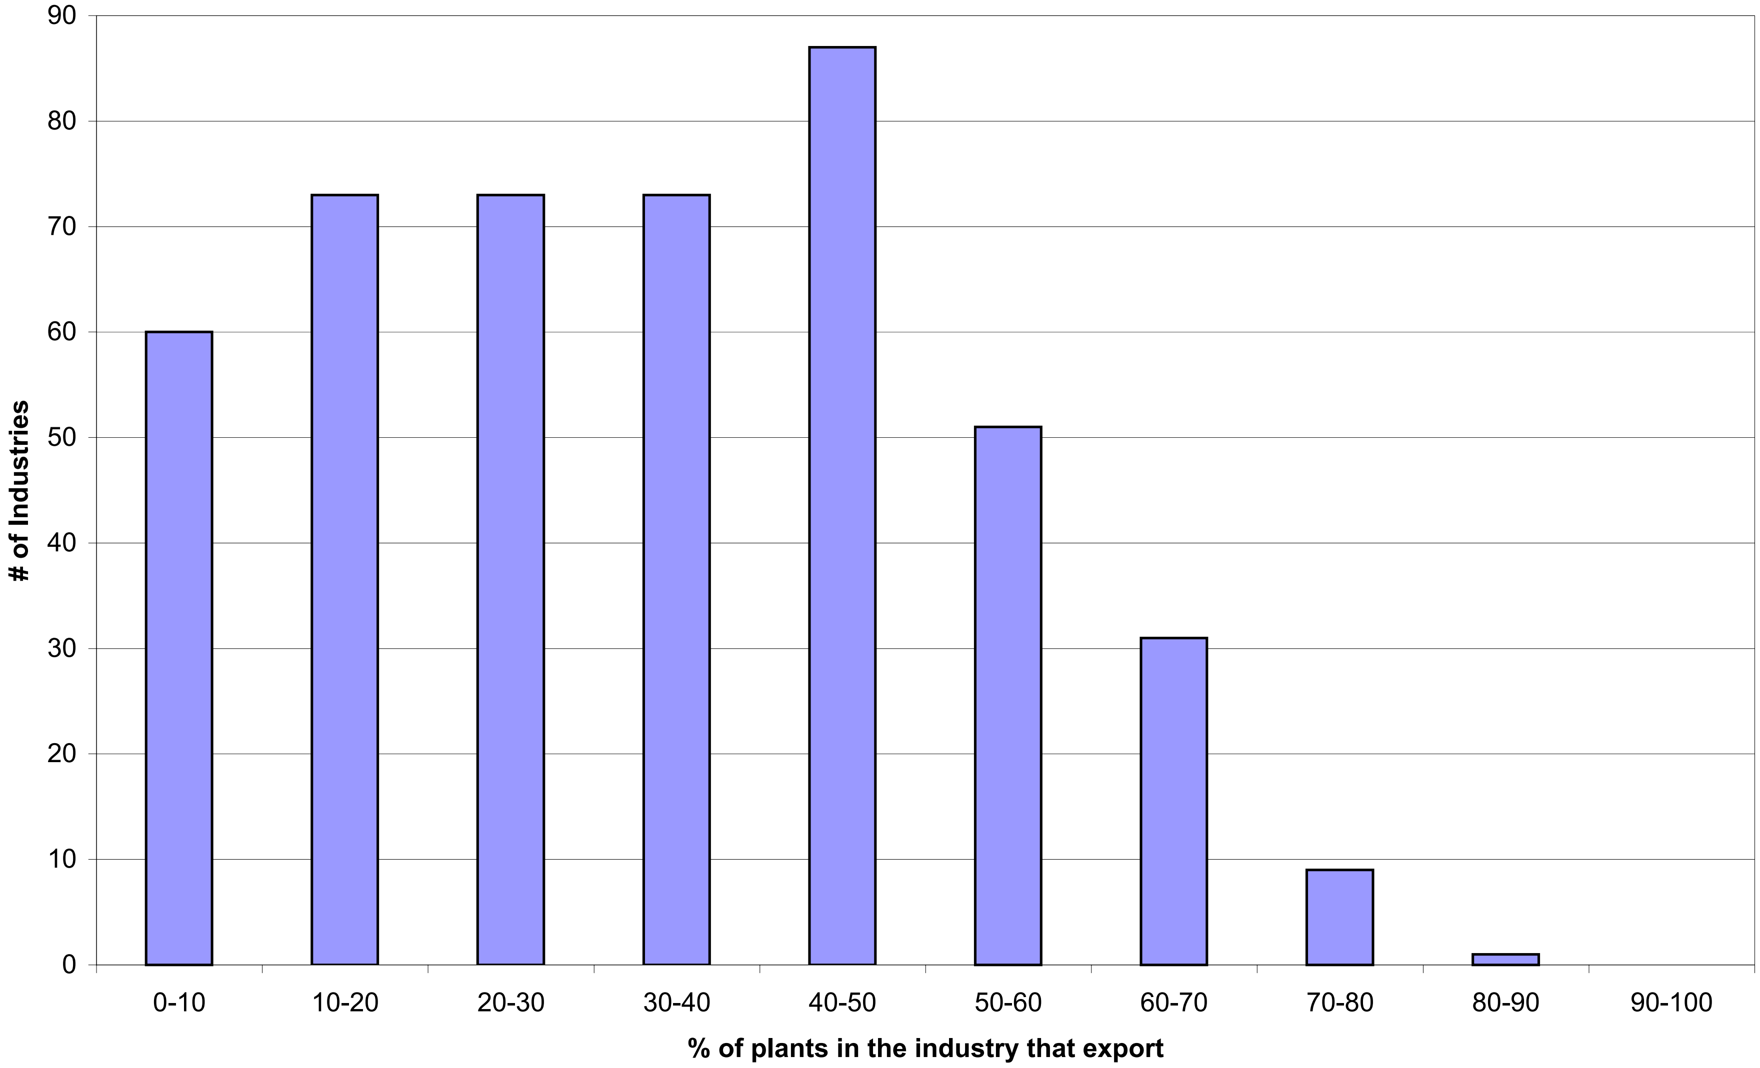

Figure 2a: Ratio of Plant Labor Productivity to Overall Mean

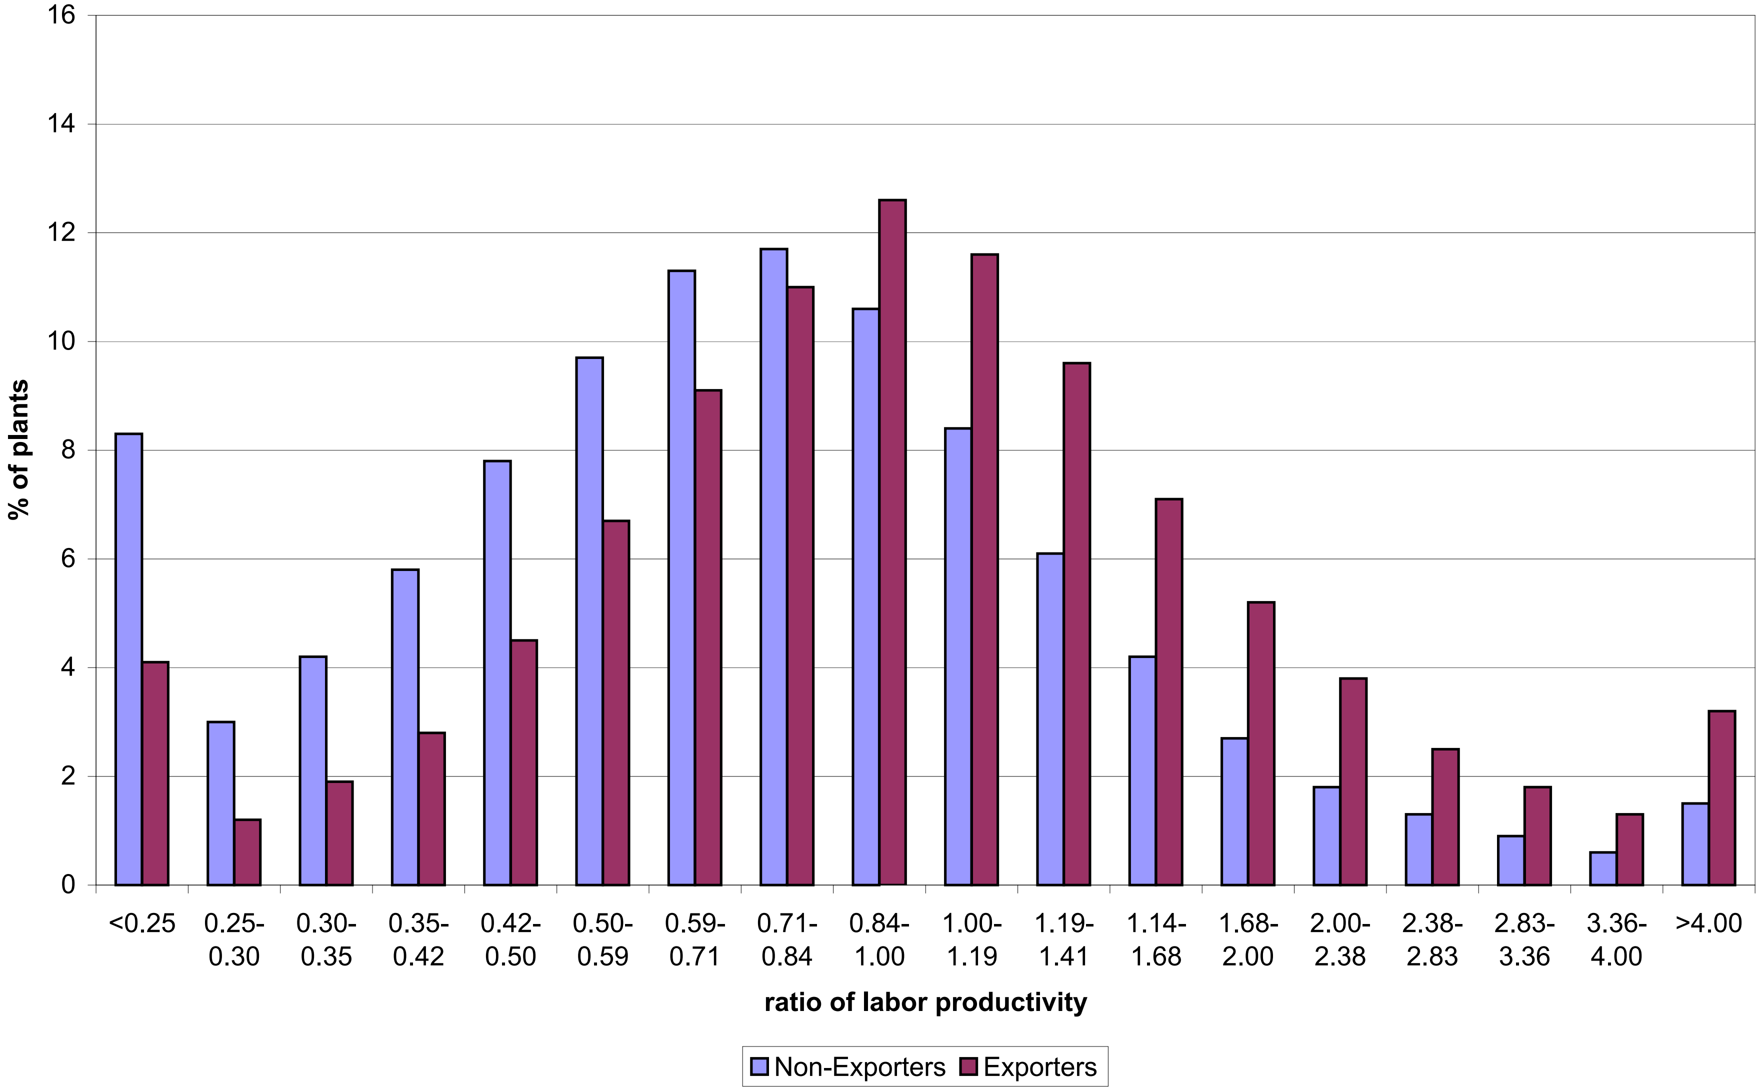

Figure 2b: Ratio of Plant Labor Productivity to 4-digit Industry Mean