

Plants and Productivity in International Trade

By ANDREW B. BERNARD, JONATHAN EATON, J. BRADFORD JENSEN, AND SAMUEL KORTUM*

We reconcile trade theory with plant-level export behavior, extending the Ricardian model to accommodate many countries, geographic barriers, and imperfect competition. Our model captures qualitatively basic facts about U.S. plants: (i) productivity dispersion, (ii) higher productivity among exporters, (iii) the small fraction who export, (iv) the small fraction earned from exports among exporting plants, and (v) the size advantage of exporters. Fitting the model to bilateral trade among the United States and 46 major trade partners, we examine the impact of globalization and dollar appreciation on productivity, plant entry and exit, and labor turnover in U.S. manufacturing. (JEL F11, F17, O33)

A new empirical literature has emerged that examines international trade at the level of individual producers. Bernard and Jensen (1995, 1999a), Sofronis Clerides et al. (1998), and Bee Yan Aw et al. (2000), 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 mod-

estly when attention is restricted to producers within a given industry or with similar factor intensity.

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 productiv-

* Bernard: Tuck School of Business at Dartmouth, 100 Tuck Hall, Hanover, NH 03755, and National Bureau of Economic Research (e-mail: andrew.b.bernard@dartmouth.edu); Eaton: Department of Economics, New York University, 269 Mercer Street, 7th Floor, New York, NY 10003, and National Bureau of Economic Research (e-mail: jonathan.eaton@nyu.edu); Jensen: Institute for International Economics, 1750 Massachusetts Avenue, NW, Washington, DC 20036 (e-mail: jbjensen@iie.com); Kortum: Department of Economics, 1035 Heller Hall, University of Minnesota, 271 19th Avenue South, Minneapolis, MN 55455, Federal Reserve Bank of Minneapolis, and National Bureau of Economic Research (e-mail: kortum@econ.umn.edu). We thank Daniel Akerberg, Eli Berman, Zvi Eckstein, Simon Gilchrist, Matthew Mitchell, Nina Pavcnik, and Robert Staiger for helpful comments. Jian Zhang and Jesse Bishop provided excellent research assistance. Eaton and Kortum gratefully acknowledge the support of the National Science Foundation. An earlier version of the paper circulated as NBER Working Paper No. 7688, May 2000. Any opinions expressed are those of the authors and not those of the Bureau of the Census, the NSF, the Federal Reserve Bank of Minneapolis, the Federal Reserve System, or the NBER.

ity 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

The core of our theoretical model is to link the variances and covariances that we observe in productivity, size, and export participation to the single producer-level characteristic of technological efficiency. The most obvious link might seem to be between efficiency and measured productivity, e.g., value added per worker. However, as long as all producers in a country employ inputs in the same proportion at the same cost, under constant returns to scale and either perfect competition or monopolistic competition with a common markup, they would all appear equally productive, in spite of any efficiency differences. With Bertrand competition, however, producers who are more efficient also tend to have a greater cost advantage over their closest competition, set higher markups, and appear more productive.² At the same time, more efficient producers are also likely to have more efficient rivals, charge lower prices, and, with elastic demand, sell more. Finally, more efficient

producers are more likely to beat out rivals in foreign markets.³

A feature of our framework is its empirical tractability. We use it to link the micro- and macro-level data for the manufacturing sector. 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. We estimate these parameters to fit moments of the U.S. plant-level data, and then examine how well our model captures other features of these data. Hence, the framework serves as a bridge between what we know about global trade flows (it is calibrated to fit actual bilateral trade patterns) 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 it can say about how changes in the global economy affect plant entry, exit, exporting, employment, and productivity in manufacturing. In performing these counterfactuals we hold fixed the efficiencies of potential producers around the world. Nevertheless, the two experiments that we perform have a significant impact on aggregate value added per worker in manufacturing. One channel is simply through their impact on the price of intermediates relative to wages, generating substitution of intermediates for labor. A second is through the entry or exit of plants whose efficiency differs from the average. The third is through the reallocation of production across plants with different levels of efficiency.

We first consider the effects of “globalization” in the form of a 5-percent drop in all geographic barriers between countries (resulting in a 15-percent rise in world trade). We find that this move kills off 3.3 percent of U.S. plants. But among the survivors, more than one in 20 of the plants that had previously sold only to the domestic market starts exporting. Since

¹ As in Eaton and Kortum (2002), 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 Paul R. Krugman’s (1979) extension to international trade of the monopolistic competition model introduced by Avinash K. Dixit and Joseph E. Stiglitz (1977). 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 emphasize a *fixed* cost of exporting [see, e.g., Mark J. Roberts and James R. Tybout (1997) and Marc Melitz (forthcoming)]. With only fixed costs, the problem 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., Martin N. Baily et al. (1992), Steven S. Olley and Ariel Pakes (1996), and Eric J. Bartelsman and Phoebus J. Dhrymes (1998). 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, albeit under specific assumptions about technology, demand, and market structure, even when outputs are heterogeneous.

³ Clerides et al. (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 non-exporting plants.

globalization provides the survivors larger markets, and since the survivors were larger to begin with, the decline in manufacturing employment is only 1.3 percent. A drop in the relative price of intermediates, the exit of unproductive firms, and the reallocation of production among survivors lead to a gain in overall manufacturing labor productivity of 4.7 percent.

We then examine 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 3.1 percent and manufacturing employment falls by 13 percent as plants substitute cheaper imported intermediates for labor. Ten percent of plants that initially export drop out of foreign markets, a few of which exit altogether.

Because our model is stylized, the particular numbers generated by these counterfactual simulations should be seen as suggestive more than definitive. Nonetheless, they do illustrate how, even in a very large 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

Our paper is not the only one to explore these issues using a theoretical framework that links measured productivity, size, and export participation to underlying variation in producer efficiency. Melitz (forthcoming) does so assuming a fixed markup (as in Dixit-Stiglitz) and a fixed cost of entry and of exporting. More efficient firms appear more productive because they spread their fixed costs over larger sales and export because they can earn enough abroad to cover the cost of entry. We are agnostic at this point about the empirical relevance of fixed costs, leaving this issue for future work. What

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. José Campa and Linda S. Goldberg (1995) show that imported intermediates are an important link between U.S. producers and the rest of the world. Pierre-Olivier Gourinchas (1999) estimates that changes in the U.S. real exchange rate lead to increased churning in the labor market. Keith Head and John Ries (1999) document the substantial exit and reallocation of production among Canadian producers following tariff reductions under the Free Trade Agreement.

TABLE 1—PLANT-LEVEL EXPORT FACTS

Export status	Percentage of all plants
No exports	79
Some exports	21
Export intensity of exporters (percent)	Percentage of exporting plants
0 to 10	66
10 to 20	16
20 to 30	7.7
30 to 40	4.4
40 to 50	2.4
50 to 60	1.5
60 to 70	1.0
70 to 80	0.6
80 to 90	0.5
90 to 100	0.7

Note: The statistics are calculated from all plants in the 1992 Census of Manufactures.

we show here is that they are not needed to deliver qualitatively the correlations that we observe. Moreover, even without them we can go quite far in explaining quantitatively the U.S. plant-level facts.

We proceed as follows. Section I discusses the plant-level facts we seek to explain. In Section II we present the theory behind our qualitative explanations, derived in Section III, for what happens at the plant level. Section IV goes on to compare the model's quantitative implications with the plant-level statistics. Section V completes the general-equilibrium specification of the model required to undertake the counterfactual experiments reported in Section VI. Section VII concludes.

I. Exporter Facts

Before turning to the theory, we take a closer look at the statistics for U.S. plants that our model seeks to explain. Appendix A, part 2, describes the data from the 1992 U.S. Census of Manufactures from which these statistics are taken.

Table 1 reports that, of the roughly 200,000 plants in the Census, only 21 percent report exporting anything. Even the plants that do export sell mostly at home. Around two-thirds of the exporters sell less than 10 percent of their output abroad. More than half of exports come




FIGURE 1. INDUSTRY EXPORTING INTENSITY

from these plants. Fewer 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.

How can we reconcile the low level of export participation and export intensity by individual plants with the fact that 14 percent of gross U.S. manufacturing production is exported? Part of the answer is that U.S. manufacturing plants as a whole report exports that sum to just over 60 percent of total U.S. exports of manufactures reported by the OECD, an important caveat in considering any of these statistics. (See Bernard and Jensen, 1995, for a discussion of this problem.) The major reason, however, is that the plants that export are much bigger, shipping on average 5.6 times more than nonexporters. Even excluding their exports, plants that export ship 4.8 times as much to the U.S. market than their nonexporting counterparts.

While previous work has sought to link trade orientation with industry, exporting producers are in fact quite spread out across industries. Figure 1 plots the distribution of industry export participation: Each of the 458 4-digit manufacturing industries is placed in one of ten 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, 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 2, exporters have a 33-percent advantage in labor productivity overall, and a 15-percent advantage relative to nonexporters within the same 4-digit industry. While differences




FIGURE 2A. RATIO OF PLANT LABOR PRODUCTIVITY TO OVERALL MEAN

across industries certainly appear in the data, what is surprising is how little industry explains about exporting and productivity.

One might argue that industry is not that informative about exporter status because it is a poor indicator of factor intensity, which is the true determinant of both productivity and export activity. We explore this possibility by allocating plants into 500 bins according to capital intensity (as measured by total assets per worker) and into 500 bins according to the share of payments to nonproduction workers as a share of labor costs, a standard indicator of skill intensity. (Bins were defined so that each contains the same number of plants.) As shown in Table 2, even within these bins the standard deviation of log productivity was nearly as high as in the raw sample. Factor intensity did even less than industry in explaining the productivity advantage of exporters (although each made a modest contribution toward explaining the difference in the raw data). Taking both industry and factor intensity into account took us a bit

further. Assigning plants *within* each 4-digit industry to one of ten factor intensity deciles reduced the productivity advantage of exporters within these bins to 9 percent, using capital intensity, and to 11 percent, using our skill-intensity measure.

Nevertheless, even controlling for industry and factor-intensity differences, substantial heterogeneity in productivity, and a productivity advantage of exporters, remains. Hence a satisfactory explanation of these phenomena must go beyond the industry or factor-intensity dimension (although we concede that these factors are not irrelevant). We consequently pursue an explanation of the plant-level facts that, as an early foray, bypasses industry and factor-intensity differences.

II. The Model

Our model introduces imperfect competition into Eaton and Kortum's (2002) probabilistic formulation of comparative advantage, which




FIGURE 2B. RATIO OF PLANT LABOR PRODUCTIVITY TO 4-DIGIT INDUSTRY MEAN

TABLE 2—PLANT-LEVEL PRODUCTIVITY FACTS

Productivity measure (value added per worker)	Variability (standard deviation of log productivity)	Advantage of exporters (exporter less nonexporter average log productivity, percent)
Unconditional	0.75	33
Within 4-digit industries	0.66	15
Within capital-intensity bins	0.67	20
Within production labor-share bins	0.73	25
Within industries (capital bins)	0.60	9
Within industries (production labor bins)	0.64	11

Notes: The statistics are calculated from all plants in the 1992 Census of Manufactures. The “within” measures subtract the mean value of log productivity for each category. There are 450 4-digit industries, 500 capital-intensity bins (based on total assets per worker), 500 production labor-share bins (based on payments to production workers as a share of total labor cost). When appearing within industries there are 10 capital-intensity bins or 10 production labor-share bins.

itself extends the Ricardian model of Rudiger Dornbusch et al. (1977) to incorporate an arbitrary number N of countries.

As in this earlier literature, there are a con-

tinuum of goods indexed by $j \in [0, 1]$. Demand everywhere combines goods with a constant elasticity of substitution $\sigma > 0$. Hence expenditure on good j in country n , $X_n(j)$, is

$$(1) \quad X_n(j) = x_n \left(\frac{P_n(j)}{p_n} \right)^{1-\sigma},$$

where $P_n(j)$ is the price of good j in country n , x_n denotes total expenditure there, and $p_n = [\int_0^1 P_n(j)^{1-\sigma} dj]^{1/(1-\sigma)}$ is the appropriate price index for country n .

Each country has multiple potential producers of each good with varying levels of technical efficiency. The k th most efficient producer of good j in country i can convert one bundle of inputs into a quantity $Z_{ki}(j)$ of good j at constant returns to scale. Except for this heterogeneity in efficiency, the production technology is identical across producers wherever and whatever they produce.

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 from country i , and we normalize $d_{ii} = 1$ for all i . We impose the plausible “triangle inequality” on the geographic barrier parameters d_{ni} :

$$(2) \quad d_{ni} \leq d_{nk} d_{ki} \quad \forall k,$$

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 .

Inputs are mobile within countries but not between them. We denote the cost of an input bundle in country i by w_i . The k th most efficient producer of good j in country i can thus deliver a unit of the good to country n at a cost:

$$(3) \quad C_{kni}(j) = \left(\frac{w_i}{Z_{ki}(j)} \right) d_{ni}.$$

Eaton and Kortum (2002) (henceforth EK), assume perfect competition. Hence each market n is served only by the lowest-cost supplier of that good to that market. It charges a price equal to this lowest cost, which is:

$$(4) \quad C_{1n}(j) = \min_i \{C_{1ni}(j)\}.$$

But under perfect competition, prices vary in-

versely with efficiency exactly to eliminate any variation in productivity, measured as the value of output per unit input.⁵ Hence we assume a form of imperfect competition in which markups vary across producers to generate variation in measured productivity.⁶

In particular we assume Bertrand competition. As with perfect competition, each market n is still captured by the low-cost supplier of each good j . As in Gene M. Grossman and Elhanan Helpman’s (1991) “quality ladders” model, this supplier is constrained not to charge more than the second-lowest cost of supplying the market, which is:

$$(5) \quad C_{2n}(j) = \min\{C_{2ni^*}(j), \min_{i \neq i^*} \{C_{1ni}(j)\}\},$$

where i^* satisfies $C_{1ni^*}(j) = C_{1n}(j)$. In other words, since i^* is the country with the low-cost supplier, the second-lowest-cost supplier to country n is either: (i) the second-lowest-cost supplier from i^* or else (ii) the low-cost supplier from someplace else.

But, as in the Dixit-Stiglitz (1977) model of monopolistic competition, the low-cost supplier would not want to charge a markup higher than $\bar{m} = \sigma/(\sigma - 1)$ for $\sigma > 1$. (For $\sigma \leq 1$ we set

⁵ To illustrate this point consider a producer with efficiency $Z_1(j)$ that sells only at home. Its measured productivity is $P(j)Z_1(j)$. Under perfect competition its price is $P(j) = w/Z_1(j)$. Hence measured productivity is simply w , which does not vary across producers facing the same input cost regardless of their efficiency. As Tor Jakob Klette and Zvi Griliches (1997) point out, studies that examine productivity at a given producer over time suffer from the same problem unless the value of output is deflated by a producer-specific price index, which is rarely available. Otherwise, an increase in efficiency is masked in any productivity measure by an offsetting drop in price. Of course, looking across countries, efficiency and measured productivity are linked even with perfect competition since countries that are on average more technologically advanced will have higher input costs, particularly wages.

⁶ Consider again, as in the previous footnote, a producer with efficiency $Z_1(j)$ that sells only at home. Under imperfect competition it sets a price $P(j) = M(j)w/Z_1(j)$, where $M(j)$ is the producer-specific markup of price over unit cost. Its measured productivity is therefore $P(j)Z_1(j) = M(j)w$, the cost of an input bundle scaled up by the markup. Variation across producers in $M(j)$ generates variation in measured productivity.

$\bar{m} = \infty$.) Hence the price of good j in market n is:

$$(6) \quad P_n(j) = \min\{C_{2n}(j), \bar{m}C_{1n}(j)\}.$$

From equations (3) and (4), to determine who sells good j to market n we need to know the most efficient way of producing that good in each potential source country i , $Z_{1i}(j)$. From equations (3) and (5), to determine what price the low-cost supplier will charge we also need to know $Z_{2i}(j)$ in the low-cost source in case this potential producer turns out to be the closest competitor in market n . [We do not need to know $Z_{ki}(j)$ for $k > 2$.]

A. A Probabilistic Formulation

To cover all possibilities we thus need to know the highest efficiency $Z_{1i}(j)$ and the next highest $Z_{2i}(j)$ of producing each good j in each country i . Rather than dealing with all these numbers, however, we treat them as realizations of random variables drawn from probability distributions. We can then derive our analytic results in terms of the small number of parameters of these probability distributions. A generalization of the theory of extremes used in EK provides a very convenient family of efficiency distributions which yield tractable distributions of prices and markups along with simple expressions for bilateral trade shares.

EK, needing to concern themselves only with a country's best producers of each good, treat their efficiencies as realizations of a random variable Z_{1i} drawn from the Fréchet distribution. (Our convention is to drop the j index when denoting the random variable as its distribution does not vary with the good.) As we show in Appendix B (available at <http://www.aeaweb.org/aer/contents/>), the analogue to the Fréchet for the joint distribution of Z_{1i} and Z_{2i} is:

$$(7) \quad F_i(z_1, z_2) = \Pr[Z_{1i} \leq z_1, Z_{2i} \leq z_2] \\ = [1 + T_i(z_2^{-\theta} - z_1^{-\theta})]e^{-T_i z_2^{-\theta}},$$

for $0 \leq z_2 \leq z_1$, drawn independently across countries i and goods j .⁷ The distribution may differ by country (through T_i) but we choose units so that it is identical across goods j . The parameter $\theta > 1$ governs the heterogeneity of efficiency, with higher values of θ implying less variability. In a trade context θ determines the scope for gains from trade due to comparative advantage. Given θ , the parameter T_i governs the average level of efficiency in country i . A higher T_i implies on average higher values of Z_{1i} and Z_{2i} (with the distribution of Z_{1i}/Z_{2i} unchanged). In a trade context T_i governs absolute advantage.⁸

We have now introduced all the relevant parameters of the model: (i) geographic barriers d_{ni} , (ii) input costs w_i , (iii) absolute advantage parameters T_i , (iv) the comparative advantage parameter θ , and (v) the elasticity of substitution σ .⁹ We can use expressions (3), (4), and (5) to transform (7) into the joint distribution of the lowest cost C_{1n} and second-lowest cost C_{2n} of supplying some good to country n (Appendix C, on the AER web site, provides the details):

$$(8) \quad G_n(c_1, c_2) = \Pr[C_{1n} \leq c_1, C_{2n} \leq c_2] \\ = 1 - e^{-\Phi_n c_1^\theta} - \Phi_n c_1^\theta e^{-\Phi_n c_2^\theta},$$

for $c_1 \leq c_2$, where:

$$(9) \quad \Phi_n = \sum_{i=1}^N T_i (w_i d_{ni})^{-\theta}.$$

The cost parameter Φ_n distills the parameters of the efficiency distributions, input costs, and trade costs around the world into a single term governing the joint distribution of C_{1n} and C_{2n} ,

⁷ Mathematical Appendices B–J appear on the American Economic Review web site (<http://www.aeaweb.org/aer/contents/>).

⁸ Replacing z_2 with z_1 in (7) yields the marginal distribution of Z_{1i} , the distribution used in EK, where Z_{2i} is irrelevant.

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.

and hence the distribution of prices and markups, in country n .

B. Analytic Results

This framework delivers six key results about prices, markups, and patterns of bilateral trade. In the following section we use these results to link the model with the exporter facts that we described in Section I.

1. The probability π_{ni} that country i is the low-cost supplier to n for any particular good is just i 's contribution to the cost parameter Φ_n ; that is:

$$(10) \quad \pi_{ni} = T_i(w_i d_{ni})^{-\theta} / \Phi_n.$$

Aggregating across goods, π_{ni} becomes the fraction of goods for which country i is the low-cost supplier to country n . As a source i becomes more competitive in market n , through either greater average efficiency T_i , lower input costs w_i , or lower costs of delivery d_{ni} , it exports a wider range of goods there.

2. The distribution $G_n(c_1, c_2)$ applies not only to the first- and second-lowest costs of supplying a good to country n regardless of source, but also to those costs conditional on the nationality of the low-cost supplier. That 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 are more competitive in a market exploit their greater competitiveness by exporting a wider range of products, to the point where entry at the margin equalizes the distribution of costs across source countries.
3. The markup $M_n(j) = P_n(j)/C_{1n}(j)$ is the realization of a random variable M_n drawn from a Pareto distribution truncated at the monopoly markup:

$$(11) \quad H_n(m) = \Pr[M_n \leq m] \\ = \begin{cases} 1 - m^{-\theta} & 1 \leq m < \bar{m} \\ 1 & m \geq \bar{m} \end{cases}.$$

While the distributions of costs differ across destinations, the distribution of markups is the same in any destination. Furthermore, within any destination no source sells at systematically higher markups. Again, greater competitiveness leads to a wider range of exports rather than to higher markups.

4. Assuming $\sigma < 1 + \theta$, the exact price index in country n implied by (1) is:

$$(12) \quad p_n = \gamma \Phi_n^{-1/\theta}.$$

The parameter γ is a function of only the parameters governing the heterogeneity of technology and tastes, θ and σ .¹⁰

5. Since prices in any destination n have the same distribution regardless of source i , the share that country n spends on goods from country i is also the fraction of goods it purchases from there, π_{ni} given in equation (10). That is:

$$(13) \quad \frac{x_{ni}}{x_n} = \pi_{ni},$$

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.

6. The share of variable costs in aggregate revenues is $\theta/(1 + \theta)$. This share applies to the set of active producers in any source country i .

Appendices D through H (on the AER web site) provide proofs of results 1 through 4 and 6,

¹⁰ 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 G (on the AER web site). 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.

respectively, while result 5 follows immediately from 1 and 2.

The functional form of the distribution from which efficiencies are drawn is obviously critical to the starkness of some of these results. The Fréchet assumption, like the Cobb-Douglas production function, causes conflicting effects exactly to cancel one another. It might seem surprising, for example, that the distribution of the markup depends only on the parameters θ and σ , and not on levels of technology, factor costs, or geographic barriers. 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, from the perspective of domestic producers who survive, it does, since they now face stiffer competition from abroad. However, an offsetting effect is 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. Under our specification these offsetting effects exactly cancel.

III. Implications for Productivity, Exporting, and Size

How can the model explain the plant-level facts described in Section I? We think of each active producer of some good j in our model as corresponding to a particular plant. At most one plant in each country will produce good j , while a plant may sell good j in several countries. For simplicity we also assume that plants specialize in producing only one good.

We first demonstrate the link between measured productivity and underlying efficiency. We then show why exporting plants tend to have high measured productivity and tend to be big.

A. Efficiency and Measured Productivity

As we discussed above, comparisons of measured productivity across plants reflect only differences in their markups. Hence, in the absence of any connection between markups and effi-

ciency, value-based productivity measures provide information only about monopoly power and not about underlying efficiency. In fact, our model does imply that, on average, plants that are more efficient charge a higher markup. As derived in Appendix I (on the AER web site), conditional on a level of efficiency z_1 , the distribution of the markup M_n is:

$$H_n(m|z_1) = \Pr[M_n \leq m | Z_{1n} = z_1] \\ = \begin{cases} 1 - e^{-\Phi_n w_n^{\theta} z_1^{-\theta} (m^{\theta-1})} & 1 \leq m < \bar{m} \\ 1 & m \geq \bar{m}. \end{cases}$$

A plant with higher efficiency Z_1 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 relative to other producing plants tends to be unusually efficient relative to its latent competitors as well, so charges a higher markup.¹¹

Hence, under imperfect competition, variation in efficiency can generate heterogeneity in measured productivity across plants. As we show next, greater efficiency also makes a producer more likely to export and to be big, explaining the correlations we see in the data.

B. Efficiency and 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

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

¹¹ Looking at the relationship the other way around, how does underlying efficiency Z_1 vary with measured productivity y ? As shown in Appendix J (on the AER web site), the conditional expectation is proportional to y (as long as the markup is less than \bar{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).

But to sell in some other market n requires:

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

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.

Variation in underlying efficiency explains the coexistence of exporting plants and plants that sell only to the domestic market. Plants with higher efficiency are more likely to export and are also more likely to have higher measured productivity. Our model thus captures a key stylized fact: Plants that export appear to be more productive.

C. Efficiency and Size

Our model can also explain why exporting plants tend to have higher domestic sales than plants that don't export. 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 a plant that can charge the Dixit-Stiglitz markup $\bar{m} = \sigma/(\sigma - 1)$, the markup is over a lower unit cost. For a plant whose markup is limited by the costs of potential competitors, the argument is less straightforward. Even though, as we showed above, more efficient plants tend to be further ahead of their rivals, so can charge a higher markup, these rivals, nonetheless, tend to be more efficient themselves, forcing the plant to set a lower price. More formally, from the joint distribution of the lowest and second-lowest cost (8) we can obtain the distribution of the

second-lowest cost (i.e., the price) conditional on the lowest cost:¹²

$$\Pr[C_{2n} \leq c_2 | C_{1n} = c_1] = 1 - e^{-\Phi_n(c_2^0 - c_1^0)}.$$

This distribution is stochastically increasing in c_1 (and hence decreasing in $z_1 = w/c_1$).¹³

IV. Quantification

Our model provides a qualitative explanation of the correlations we observe across plants between measured productivity, exporting, and size through the positive association of each with underlying efficiency. It does not, however, yield tractable closed-form expressions for the moments of the plant-level statistics that we report in Section I. To assess how well our model does quantitatively we take a simulation approach: We first reformulate the model as an algorithm that uses data on aggregate trade shares and expenditures to simulate plant-level

¹² This result is obtained as

$$\begin{aligned} \Pr[C_{2n} \leq c_2 | C_{1n} = c_1] &= \frac{\partial G_n(c_1, c_2)/\partial c_1}{\partial G_{1n}(c_1)/\partial c_1} \\ &= 1 - e^{-\Phi_n(c_2^0 - c_1^0)}, \end{aligned}$$

where $G_{1n}(c_1)$ is the marginal distribution of the lowest cost in country n .

¹³ The following sports analogy might provide intuition into our results on the productivity and size advantage of more efficient plants. Suppose that a running event had been held on many different occasions (with all participants' speeds drawn from the same particular distribution). Imagine trying to assess who, among the winners of each event, likely ran fastest when measurement failures made the winning times unavailable (just as we cannot observe a plant's efficiency Z). Say first that the referees forgot to start the clock at the beginning of the race, but had managed to record the time between when the winner crossed the finish line and the finish of the runner placing second. The winner would probably have been faster the farther ahead she was of the runner placing second (just as in our model a more efficient producer is more likely to be further ahead of its rival, thus able to charge a higher markup). Say instead that the clock had been started properly at the beginning of the race, but that the referees missed recording the winner's time. They had, however, managed to record the time of the runner placing second. The winner would probably have been faster the faster the time of the runner up (just as in our model a more efficient firm is likely to have a more efficient rival, so must charge a lower price).

statistics. We then estimate the two heterogeneity parameters θ and σ to make our simulated data match the actual productivity advantage and size advantage of exporters. Finally, we report how well other moments of our simulated data line up with the remaining facts from Section I.

Each step of a simulation applies to a particular good j as if drawn at random from the continuum. For that good we draw from the efficiency distribution (7) in each country, which together with the w_i 's and d_{ni} 's determines the cost of the low-cost producer from each country supplying each other country. Among these potential suppliers we identify the locations of the active (i.e., lowest-cost) producers for each market. We also determine the second-lowest cost of 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. Doing so repeatedly we build an artificial data set of U.S. plants whose moments we can compare with those of the actual plants in the U.S. Census.

A. Reformulating the Model as an Algorithm

To perform these simulations we need values for the heterogeneity parameters σ (in consumption) and θ (in production), common across countries. It might appear that we also need values for all of the model's numerous parameters for the United States and its trading partners: 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).

A reformulation of the model, however, reveals that bilateral trade shares π_{ni} and absorption x_n summarize all we need to know about the country-specific parameters T_i , w_i , and d_{ni} to say who will sell where and at what markup. The reformulation begins by defining transformations of the efficiency terms:

$$U_{1i}(j) = T_i Z_{1i}(j)^{-\theta}$$

$$U_{2i}(j) = T_i Z_{2i}(j)^{-\theta}.$$

Using the efficiency distribution (7), it is easy to show that these transformed efficiencies are realizations of random variables drawn from the parameter-free distributions:

$$(14) \quad \Pr[U_{1i} \leq u_1] = 1 - e^{-u_1}$$

$$\Pr[U_{2i} \leq u_2 | U_{1i} = u_1] = 1 - e^{-u_2 + u_1}.$$

Using equation (3), the transformed efficiencies connect to costs as follows:

$$C_{1ni}(j) = \left(\frac{U_{1i}(j)}{\pi_{ni} \Phi_n} \right)^{1/\theta}$$

$$C_{2ni}(j) = \left(\frac{U_{2i}(j)}{\pi_{ni} \Phi_n} \right)^{1/\theta}.$$

We can now express all of the observables in terms of realizations of the U 's (which are drawn from parameter-free distributions), the π_{ni} [which, via expression (13), can be observed from trade shares], and the parameters θ and σ (for the magnitudes of interest, Φ_n drops out).

The country that sells good j in each market n , which we denote i^* , is given by:

$$(15)$$

$$i^* = \arg \min_i \{C_{1ni}(j)\} = \arg \min_i \left\{ \frac{U_{1i}(j)}{\pi_{ni}} \right\}.$$

Given that a producer from country i^* is the low-cost supplier of good j to market n , its markup there is:

$$(16) \quad M_n(j) = \min \left\{ \frac{C_{2n}(j)}{C_{1n}(j)}, \bar{m} \right\} \\ = \min \left\{ \left(\frac{V_{2n}(j)}{V_{1n}(j)} \right)^{1/\theta}, \bar{m} \right\}$$

where:

$$V_{1n}(j) = \min_i \left\{ \frac{U_{1i}(j)}{\pi_{ni}} \right\} = \frac{U_{1i^*}(j)}{\pi_{ni^*}}$$

$$V_{2n}(j) = \min \left\{ \frac{U_{2i^*}(j)}{\pi_{ni^*}}, \min_{i \neq i^*} \left\{ \frac{U_{1i}(j)}{\pi_{ni}} \right\} \right\}.$$

How much it sells in market n is:

$$(17) \quad X_n(j) = x_n [M_n(j)/\gamma]^{1-\sigma} V_{1n}(j)^{(1-\sigma)/\theta}.$$

For each j we can determine, for each market n , which country i^* serves as the supplier. We use $\Omega_i(j)$ to denote the set of countries which country i supplies. (If the set is empty, country i imports good j .)

Since we are looking at the model's predictions about U.S. plants, we can arbitrarily assign the United States as country 1. In our simulations we treat any j for which $\Omega_1(j)$ is non-empty as a product with a corresponding U.S. plant. For any such product we then calculate, from (17), how much this plant sells in each market and, from (16), its markup in each market. From these expressions we can calculate for an active U.S. plant: (1) whether the plant exports, (2) its total sales around the world, (3) how much it exports, (4) its total production costs, (5) its employment, and (6) its productivity.

The last two calculations force us to take a stand on the inputs to production. We assume that production combines labor, with wage W_i , and intermediates, which are a representative bundle of manufactures with price index p_i given in (12), with labor having a share β in costs. The cost of an input bundle in country i , w_i in equation (3), is therefore:

$$(18) \quad w_i = W_i^\beta p_i^{1-\beta}$$

(where labor units are chosen to eliminate the constant).

For each simulated U.S. plant we calculate the six magnitudes as follows:

1. Whether the plant exports simply involves checking whether $\Omega_1(j)$ contains any element other than 1.
2. Total sales are $X(j) = \sum_{n \in \Omega_1(j)} X_n(j)$, where $X_n(j)$ is calculated according to (17).
3. Total exports are $\sum_{n > 1, n \in \Omega_1(j)} X_n(j)$.
4. Total production costs are $I(j) = \sum_{n \in \Omega_1(j)} [X_n(j)/M_n(j)]$, where $M_n(j)$ is given by (16).
5. Employment, $L(j)$, is proportional to labor cost: $W_1 L(j) = \beta I(j)$.

6. The plant-level productivity measure, value added per worker $v(j)$, is proportional to:

$$(19) \quad v(j)/W_1 \\ = [X(j) - (1 - \beta)I(j)]/[W_1 L(j)].$$

B. Parameterization

For given θ and σ we can calculate each of these statistics using actual data on trade shares π_{ni} , absorption x_n , and the share of intermediates in revenue. We calculate π_{ni} and x_n from 1992 production and trade data in manufactures among the 47 leading U.S. export destinations (including the United States itself).¹⁴ Appendix A, part 1, describes the data. Table 3 lists our choice of partner countries as well as some summary statistics for each of them. The observed share of intermediates in revenues is 0.63. From analytic result 6, intermediate's share in costs is then $0.63(1 + \theta)/\theta$. Given θ , we set $\beta = 1 - 0.63(1 + \theta)/\theta$.¹⁵

We choose the parameters θ and σ to make our artificial data set deliver the same productivity and size advantage of exporters as in the U.S. plant-level data. A smaller value of θ , by generating more heterogeneity in efficiency, will imply a larger productivity advantage of exporters while a larger value of σ , by delivering a larger demand response to price differences, will imply a larger size advantage.

We choose these two particular moments because, unlike variation in productivity and plant size, these moments are invariant to sources of variation that may be required to account fully for the observed heterogeneity in the data. Two such sources of variation are the following:

¹⁴ 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.

¹⁵ The 0.63 figure is calculated from the OECD's STAN database as the value of gross production less value added as a share of gross production for the U.S. manufacturing sector in 1992. The cost notion in our model is variable cost (since we take no particular stand on overhead costs). Since intermediates are more likely to be associated purely with variable costs we use them as the basis for our calibration of β , which represents the share of labor in variable costs.

TABLE 3—AGGREGATE TRADE DATA

No.	Country	Data source	U.S. exports (\$ million)	U.S. percent market share	Imports from ROW (percent of total)	Exports to ROW (percent of total)
1	Arab Emirates	W	1,590	7.9	5.3	36.4
2	Argentina	U	3,498	3.8	2.8	11.9
3	Australia	O	8,570	6.2	2.9	5.9
4	Austria	O	1,785	1.5	5.9	12.5
5	Belgium and Luxembourg	O, U	6,264	4.3	4.2	3.4
6	Brazil	U	5,932	2.9	3.4	7.2
7	Canada	O	83,400	32.9	0.7	0.7
8	Chile	U	2,441	8.8	1.6	2.7
9	China and Hong Kong	U, U	16,200	3.3	1.7	6.2
10	Colombia	U	3,098	12.5	1.8	5.4
11	Denmark	O	1,403	2.5	4.7	7.5
12	Ecuador	U	1,035	14.9	1.2	1.1
13	Egypt	U	1,665	7.1	6.5	22.7
14	Finland	O	914	1.8	2.7	4.9
15	France	O	16,700	2.4	4.0	10.6
16	Germany (unified)	O	23,000	1.8	6.6	7.6
17	Greece	O	804	2.0	4.1	11.9
18	India	U	1,624	1.3	8.4	9.7
19	Indonesia	U	2,846	5.1	0.9	3.9
20	Ireland	U	2,771	8.4	1.8	1.9
21	Israel	U	3,251	10.0	1.5	8.7
22	Italy	O	8,124	1.2	6.4	9.0
23	Japan	O	42,100	1.6	2.5	4.3
24	Korea (South)	O	14,100	5.5	1.4	6.9
25	Kuwait	U	1,471	15.5	3.8	25.2
26	Mexico and Caribbean	O, U, W, U, U	43,700	19.8	2.4	7.0
27	Netherlands	O	9,362	5.1	2.2	4.6
28	New Zealand	O	1,526	7.3	1.1	8.4
29	Nigeria	U	1,012	6.5	1.2	1.0
30	Norway	O	1,779	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	Philippines	U	1,667	4.9	1.1	1.4
34	Portugal	U	807	1.3	1.7	8.0
35	Saudi Arabia	W	7,145	14.4	2.1	8.1
36	Singapore and Malaysia	U, U	15,000	14.1	1.5	6.4
37	Spain	O	5,717	1.9	2.6	7.6
38	Sweden	O	3,403	3.1	3.5	5.0
39	Switzerland	U	4,222	3.2	2.7	5.2
40	South Africa	U	2,106	3.0	1.5	6.1
41	Taiwan	U	14,000	8.1	0.6	1.3
42	Thailand	U	4,094	3.7	3.1	7.3
43	Turkey	U	2,186	2.4	6.0	16.9
44	United Kingdom	O	22,600	3.8	2.0	11.4
45	United States	O	2,520,300	85.0	1.8	3.5
46	USSR (former)	U	2,181	0.4	8.6	9.3
47	Venezuela	U	6,390	17.2	1.4	10.2

Notes: 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 Robert C. Feenstra et al. (1997). The U.S. market share is a country's imports from the United States relative to its 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).

1. Productivity at the plant level may be observed with error. Hence:

$$\hat{v}(j) = v(j)\xi(j)$$

where $\hat{v}(j)$ is observed value added per worker, $v(j)$ is its actual level, and $\xi(j)$ is a multiplicative error term that is independent of underlying efficiency $Z(j)$. Variation in the realized $\xi(j)$'s thus generates variation in measured productivity that is independent of export status or size. In contrast, variation in $Z(j)$, as governed by the parameter θ , generates variation in measured productivity that correlates with export status and (for $\sigma \neq 1$) with size.

2. Different products may have different weights in demand, so that (1) becomes:

$$(20) \quad X_n(j) = \alpha(j)x_n \left(\frac{P_n(j)}{p_n} \right)^{1-\sigma},$$

where $\alpha(j)$ governs the magnitude of demand anywhere for product j . We treat the $\alpha(j)$ as independent of underlying efficiency $Z(j)$, and hence independent of who produces the good and its price in any market. Plants that produce products with a larger $\alpha(j)$ are, other things equal, larger for reasons that are independent of their underlying efficiency. Variation across the $\alpha(j)$ thus provide a source of variation in plant size that is independent of export status. In contrast, with $\sigma \neq 1$, variation in $Z(j)$ generates variation in plant size that correlates with export status.

Export status correlates only with plant heterogeneity arising from underlying efficiency differences. Hence the productivity and size advantage of exporters arises only from heterogeneity in $Z(j)$ and not from that in α or ξ . Export status serves as an instrument to extract variation in the data relevant for identifying θ and σ .

We implement the algorithm described in the previous subsection as follows. Without assigning any parameter values we draw the $U_{1i}(j)$ and $U_{2i}(j)$ for 47 countries and 1,000,000 j 's from the distributions (14). Using the matrix of π_{ni} we can identify, for each destination n , the

source i^* for each good j using (15). The result generates about 850,000 active U.S. plants (the other goods being imported). For given values of θ and σ (with β chosen to be consistent with θ), we calculate items 1 through 6 above for each U.S. plant. We search over values of θ and σ until our artificial data set delivers the same productivity advantage of exporters (whose value added per worker is on average 33 percent higher than nonexporters') and the same size advantage of exporters (whose domestic shipments are on average 4.8 times higher than nonexporters') as in the 1992 U.S. Census of Manufactures. Once we have found these values of θ and σ , we can calculate analogues to the actual statistics reported in Table 1.

C. The Model's Fit

Searching over parameter values, we find that our simulated data yield the productivity and size advantage of exporters if $\theta = 3.60$ and $\sigma = 3.79$. The estimate of $\theta = 3.60$ based on plant-level data is the same as the lowest of the three estimates from EK, 3.60 (based on trade and wages), as opposed to 8.28 or 12.9 (based on trade and prices).

Table 4 reports how our simulated data, calculated using these parameters, compare with the plant-level export facts computed from the Census data. We also consider how much of the heterogeneity in plants' productivity and size is reflected in the simulated data.

1. **The Fraction Who Export.** A basic prediction of our framework (which 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. One explanation (admittedly favorable to our model) is that a number of plants fail to report exporting. 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. **The Fraction of Revenues from Exports.**

TABLE 4—EXPORT FACTS: SIMULATED VS. ACTUAL

Export status	Percentage of all plants	
	Simulated	Actual
No exports	49	79
Some exports	51	21
Export intensity of exporters (percent)	Percentage of exporting plants	
	Simulated	Actual
0 to 10	76	66
10 to 20	19	16
20 to 30	4.2	7.7
30 to 40	0.0	4.4
40 to 50	0.0	2.4
50 to 60	0.0	1.5
60 to 70	0.0	1.0
70 to 80	0.0	0.6
80 to 90	0.0	0.5
90 to 100	0.0	0.7

Notes: The simulated export facts are based on $\theta = 3.60$ and $\sigma = 3.79$, the parameter values for which the model delivers the observed productivity advantage and size advantage of U.S. exporters. A simulation of the model involves sampling production efficiencies for 1,000,000 goods in 47 countries, calculating the outcome for each good in each market, and averaging over all the outcomes in which a U.S. plant is active. The actual export facts are from Table 1.

Our simulated data match the skewness of the distribution of export intensity among U.S. exporting plants, with most exporters selling only a small fraction abroad. We capture this feature of the data quite nicely despite having ignored it in choosing parameter values.

3. **Variability in Productivity.** The standard deviation of the log of value added per worker is about 0.35 in our simulated data while in the actual data it is 0.75. As discussed in Section IV, subsection B, an explanation for our underprediction that is consistent with our model is that measurement error in the Census data generates much more heterogeneity in the actual data. With this interpretation, heterogeneity in underlying efficiency explains 22 percent of the variance in the log of measured value added per worker. It is obviously somewhat problematic to attribute so much of the variability in productivity to measurement error.

Given σ , a smaller value of θ would allow variability in underlying efficiency to account for more of the variation in measured productivity, but would lead us to overstate the productivity advantage of exporters.

4. **Variability in Size.** The standard deviation of the log of domestic sales is 0.84 in the simulated data and 1.67 in the actual data. As discussed in Section IV, subsection B, an explanation consistent with our model is that variation in demand weights across goods generates additional variability in plant size. With this interpretation, heterogeneity in underlying efficiency explains 25 percent of the variance in log domestic sales. Given θ , a larger value of σ would allow variability in underlying efficiency to account for more of the variation in size but would lead us to overstate the size advantage of exporters.

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.

V. 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.)

Given wages, manufacturing price levels in different countries are connected through trade in intermediates. To take these interactions into account we manipulate equations (12), (9), and (18) to obtain:

$$(21) \quad \mathbf{P}^{-\theta} = \mathbf{\Lambda} \mathbf{P}^{-\theta(1-\beta)},$$

where the n th element of the vector \mathbf{P}^a is p_n^a and the element in the n th row and i th column of the matrix $\mathbf{\Lambda}$ 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 loglinear approximation to (21).

Having determined how prices change, we can easily calculate the change in input cost w_i in each country. Using equation (10) 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 denote each country's expenditure on manufactures for purposes other than as inputs into manufacturing (i.e., final expenditure and expenditure on inputs into nonmanufacturing) by y_n and treat that amount as exogenous. Since, from result 6 in Section II, aggregate costs are a fraction $\theta/(1+\theta)$ of aggregate revenues, the vector of manufacturing absorptions satisfy:

$$(22) \quad \mathbf{X} = \frac{\theta}{1+\theta} (1-\beta) \mathbf{\Pi}' \mathbf{X} + \mathbf{Y},$$

where the n th element of the vector \mathbf{X} is x_n , the representative element of the vector \mathbf{Y} is y_n , and the representative element of the matrix $\mathbf{\Pi}$ is π_{ni} . The first term on the right side of equation (22) represents demand for intermediates in manufacturing while the second term represents other demand for manufactures. We use equation (22) to calculate how a change in $\mathbf{\Pi}$ translates into a change in \mathbf{X} . Together, the changes in $\mathbf{\Pi}$ and \mathbf{X} determine the new values of x_{ni} for each country n and i .

VI. Counterfactuals

We consider two types of aggregate shocks to the world trading regime: (i) a 5-percent worldwide decline in geographic barriers (resulting in 15-percent more world trade), and (ii) a 10-percent exogenous appreciation of the U.S.

wage relative to wages in other countries (leading to a 14-percent decline in U.S. exports). We compare each counterfactual situation to a baseline, holding fixed the efficiency levels of all potential producers.

For each counterfactual we 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.

A. 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:

$$q = \sum_j s(j)q(j),$$

where $s(j) = L(j)/L$ is employment in plant j as a fraction of total manufacturing employment.

Following Lucia Foster et al. (2001), we decompose aggregate productivity growth into the contributions of entering plants (n), exiting plants (x), reallocation among surviving incumbents (c), and productivity gains for continuing

incumbents. Denoting the set of plants of each type as Ω_k , $k = n, x, c$:

$$\begin{aligned}
 (23) \quad & q' - q \\
 &= \sum_{j \in \Omega_c} s(j)[q'(j) - q(j)] \\
 &+ \sum_{j \in \Omega_c} [s'(j) - s(j)][q(j) - q] \\
 &+ \sum_{j \in \Omega_c} [s'(j) - s(j)][q'(j) - q(j)] \\
 &+ \sum_{j \in \Omega_n} s'(j)[q'(j) - q] \\
 &+ \sum_{j \in \Omega_x} s(j)[q - q(j)]
 \end{aligned}$$

where z' denotes the counterfactual value of variable z . The first term is the contribution of productivity changes for continuing plants with initial weights. The second term is the effect of reallocating production among continuing plants given their initial productivity. The third term is the cross-effect of reallocation and productivity changes for continuing plants. The fourth term is the contribution from entry and the fifth from exit.

Even though we are holding a plant's efficiency draw fixed, our counterfactual experiments can affect measured productivity at a continuing plant. Using the cost function, we can write a plant's deflated value added per worker as:

$$(24) \quad q(j) = \frac{1}{\beta} \frac{W}{p} [M^C(j) - (1 - \beta)]$$

where $M^C(j) = X(j)/I(j)$ is the composite markup across all markets, i.e., total revenues over total costs. Note that, across plants, measured productivity varies for no other reason than the markup. In our counterfactuals, which look at the same plant in two different situations, measured productivity can rise either because of an increase in the plant's

markup or because of a fall in manufacturing prices, p .

B. Counterfactual Outcomes

The results of the two counterfactuals are shown in Table 5:

1. Globalization (taking the form of a 5-percent fall in geographic barriers) leads to a 4.7-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 3 percent of U.S. plants exit. Since their productivity averages only 45 percent of the survivors', exit contributes 0.8 percent to the overall productivity gain. As smaller, lower-productivity plants exit, high-productivity exporters expand, leading to an additional 0.2-percent gain. (As they expand, however, they sell to markets where their cost advantage is smaller, hence the covariance term of -0.1 percent.) Net job loss is only 1.3 percent of initial employment, a much lower percent than plant exit. This figure is the net outcome of 1.5-percent gross job creation at plants that expand and 2.8-percent gross job destruction at plants that shrink or close altogether.
2. 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 4.2 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 0.8 percent to the overall productivity gain. Slightly offsetting these gains is the reallocation of production away from the most productive firms (who lose export markets), contributing to a drop of 0.2 percent in value added per worker. Together substitution, reallocation, and exit

TABLE 5—COUNTERFACTUALS

Statistics for U.S. producers	Counterfactual experiment	
	5-percent lower barriers	10-percent higher U.S. relative wage
Productivity decomposition (percent change):		
Aggregate	4.7	4.2
Entrants	0.0	0.0
Exiters	0.8	0.8
Reallocation, continuing plants	0.2	-0.2
Covariance, continuing plants	-0.1	-0.1
Gains, continuing plants	3.9	3.7
Plant exit and entry:		
Number of plants (percent change)	-3.3	-3.1
Relative productivity of exiters (percent)	45	47
Employment share (percent), prior	1.5	1.4
Employment:		
Total employment in manufacturing (percent change)	-1.3	-13
Job creation (percent)	1.5	0.3
Job destruction (percent)	2.8	13
International trade:		
U.S. exports (percent change)	11	-14
U.S. imports (percent change)	17	10
U.S. absorption (percent change)	-1.9	-6.0
World trade (percent change)	15	1.4

Notes: Results are based on simulating the model (with $\theta = 3.60$ and $\sigma = 3.79$) by sampling production efficiencies for 500,000 goods, calculating the outcome for each under the counterfactual and under the baseline, and calculating statistics for the outcomes in which a U.S. plant is ever active. 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 (23), 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. Job creation and job destruction are both shown as a percentage of initial manufacturing employment.

generate a 13-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 each counterfactual. Globalization, as shown in Table 6, generates action among plants initially not exporting: While 6.7 percent of nonexporters are shut down by foreign competition, 5.2 percent take advantage of new export opportunities. Initial productivity is a good indicator of how a nonexporter will fare: 17 percent of those in the lowest quartile exit while only 2.9 percent enter export markets. But none of the plants in the top productivity quar-

tile shuts down, and 12 percent enter export markets.

A loss of competitiveness, as shown in Table 7, leads to the exit of 6.1 percent of the plants originally producing only for the U.S. market. Fewer than 1 percent of exporters shut down, but 10 percent do stop exporting. Breaking down these statistics by a plant's initial position in the productivity distribution, nearly 15 percent of low-productivity nonexporting plants exit while none exit from the top three quartiles of the productivity distribution. The results for exporters are similar: Among the low-productivity group 21 percent stop exporting, of which almost 3 percent actually exit. For the highest-productivity exporters

TABLE 6—PLANT-LEVEL TRANSITIONS: 5-PERCENT LOWER GEOGRAPHIC BARRIERS

	Before geographic barriers fall	After geographic barriers fall		
	Percent of plants	Exit (percent)	Domestic (percent)	Export (percent)
All plants:				
Domestic	49	6.7	88	5.2
Export	51	0.1	0.9	99
By productivity quartile:				
Lowest quartile				
Domestic	77	17	80	2.9
Export	23	0.6	1.5	98
Second quartile				
Domestic	66	0.1	95	4.6
Export	34	0.0	1.6	98
Third quartile				
Domestic	30	0	93	7.2
Export	70	0	0.9	99
Highest quartile				
Domestic	21	0	88	12
Export	79	0	0.4	99

Notes: Results are based on simulating the model (with $\theta = 3.60$ and $\sigma = 3.79$) by sampling production efficiencies for 500,000 goods, calculating the outcome for each under the counterfactual and under the baseline, and calculating statistics for the outcomes in which a U.S. plant is ever active. 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 shut down ("Exit"), produce only for the domestic market ("Domestic"), or continue to export ("Export").

fewer than 4 percent stop exporting. For either counterfactual, we see a striking heterogeneity of outcomes from aggregate shocks.

VII. 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 parameterized 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 highly 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

TABLE 7—PLANT-LEVEL TRANSITIONS: 10-PERCENT HIGHER U.S. RELATIVE WAGE

	Before appreciation of the U.S. wage	After appreciation of the U.S. wage		
	Percent of plants	Exit (percent)	Domestic (percent)	Export (percent)
All plants:				
Domestic	49	6.1	94	0.0
Export	51	0.3	10	90
By productivity quartile:				
Lowest quartile				
Domestic	77	15	85	0.0
Export	23	2.6	18	79
Second quartile				
Domestic	66	0.0	100	0.0
Export	34	0.0	18	82
Third quartile				
Domestic	31	0	100	0
Export	69	0	11	89
Highest quartile				
Domestic	21	0	100	0
Export	79	0	3.9	96

Notes: Results are based on simulating the model (with $\theta = 3.60$ and $\sigma = 3.79$) by sampling production efficiencies for 500,000 goods, calculating the outcome for each under the counterfactual and under the baseline, and calculating statistics for the outcomes in which a U.S. plant is ever active. Every pair of numbers in the first column sum to 100 percent. The last three numbers in every row also sum to 100 percent (except, due to rounding). Following the wage appreciation, an active U.S. plant will either shut down (“Exit”), produce only for the domestic market (“Domestic”), or continue to export (“Export”).

important features of the world. We ignore possible differences across industries in relevant parameter values. We treat labor as homogeneous and have implicitly lumped capital together with intermediates. We assume the absence of any internal trade barriers. 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.

APPENDIX A: DATA

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.

1. Aggregate Trade Data

We chose our sample of countries as follows (the 47 countries or regions are listed in Table 3). 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 Robert C. Feenstra et al. (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 the 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 3.

We get home purchases x_{nm} 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 3 suggest that undercounting is not a serious problem.

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 (excluding very small plants not mailed a census form). It provides data on their value of shipments, production and nonproduction employment, salaries and wages, value added, capital stock, ownership structure, and exports. The plant export measure is the reported value of exports, specifically "the value of products shipped for export [including] direct exports and products shipped to exporters or other wholesalers for export." As some indirect exports are not included in this measure, we do

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

REFERENCES

- Aw, Bee Yan; Chung, Sukkyun and Roberts, Mark J.** "Productivity and Turnover in the Export Market: Micro Evidence from Taiwan and South Korea." *World Bank Economic Review*, January 2000, 14(1), pp. 65–90.
- Baily, Martin N.; Hulten, Charles R. and Campbell, David.** "Productivity Dynamics in Manufacturing Plants." *Brookings Papers on Economic Activity: Microeconomics*, 1992, Spec. Iss., pp. 187–267.
- Bartelsman, Eric J. and Dhrymes, Phoebus J.** "Productivity Dynamics: U.S. Manufacturing Plants, 1972–1986." *Journal of Productivity Analysis*, January 1998, 9(1), pp. 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, pp. 67–119.
- . "Exceptional Exporter Performance: Cause, Effect, or Both?" *Journal of International Economics*, February 1999a, 47(1), pp. 1–25.
- . "Exporting and Productivity." National Bureau of Economic Research (Cambridge, MA) Working Paper No. 7135, May 1999b.
- Campa, José and Goldberg, Linda S.** "Investment, Exchange Rates and External Exposure." *Journal of International Economics*, May 1995, 38(3–4), pp. 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*, August 1998, 113(3), pp. 903–47.
- Dixit, Avinash K. and Stiglitz, Joseph E.** "Monopolistic Competition and Optimum Product Diversity." *American Economic Review*, June 1977, 67(3), pp. 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*, December 1977, 67(5), pp. 823–39.
- Eaton, Jonathan and Kortum, Samuel.** "Technology, Geography, and Trade." *Econometrica*, September 2002, 70(5), pp. 1741–80.
- Feenstra, Robert C.; Lipsey, Robert E. and Bowen, Henry P.** "World Trade Flows, 1970–1992, with Production and Tariff Data." National Bureau of Economic Research (Cambridge, MA) Working Paper No. 5910, January 1997.
- Foster, Lucia; Haltiwanger, John and Krizan, C. J.** "Aggregate Productivity Growth: Lessons from Microeconomic Evidence," in Charles R. Holtén, Edwin R. Dean, and Michael J. Harper, eds., *New developments in productivity analysis*. Chicago: University of Chicago Press, 2001.
- Gourinchas, Pierre-Olivier.** "Exchange Rates and Jobs: What Do We Learn from Job Flows?" in Ben S. Bernanke and Julio J. Rotemberg, eds., *National bureau of economic research macroeconomics annual 1998*. Cambridge, MA: MIT Press, 1999, pp. 153–208.
- Grossman, Gene M. and Helpman, Elhanan.** "Quality Ladders in the Theory of Growth." *Review of Economic Studies*, January 1991, 58(1), pp. 43–61.
- Head, Keith and Ries, John.** "Rationalization Effects of Tariff Reductions." *Journal of International Economics*, April 1999, 47(2), pp. 295–320.
- Klette, Tor Jakob and Griliches, Zvi.** "The Inconsistency of Common Scale Estimators When Output Prices are Unobserved and Endogenous." *Journal of Applied Econometrics*, July–August 1996, 11(4), pp. 343–61.
- Krugman, Paul R.** "Increasing Returns, Monopolistic Competition, and International Trade." *Journal of International Economics*, November 1979, 9(4), pp. 469–79.
- Melitz, Marc.** "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica* (forthcoming).
- Olley, Steven S. and Pakes, Ariel.** "The Dynamics of Productivity in the Telecommunications Industry." *Econometrica*, November 1996, 64(6), pp. 1263–97.
- Organization for Economic Cooperation and Development (OECD).** *The OECD STAN database*. Paris: OECD, 1995.
- Roberts, Mark J. and Tybout, James R.** "The Decision to Export in Colombia: An Empirical Model of Entry with Sunk Costs." *American Economic Review*, September 1997, 87(4), pp. 545–64.
- United Nations Industrial Development Organization (UNIDO).** *Industrial statistics database*. Vienna: UNIDO, 1999.
- World Bank.** *World tables*. Baltimore, MD: Johns Hopkins University Press, 1995.