Second Draft

WHY PLANT-LEVEL PRODUCTIVITY STUDIES ARE OFTEN MISLEADING, AND AN ALTERNATIVE APPROACH TO INFERENCE

Hajime Katayama
Pennsylvania State University

Shihua Lu
Charles River Associates

James R. Tybout
Pennsylvania State University and NBER

November, 2002

We gratefully acknowledge the National Science Foundation for support (grant 421-18 687W0). We have benefited from useful discussions with Marc Melitz, Nina Pavcnik, Mark Roberts, Marc Rysman, Johannes Van Biesebroeck and participants in the 2002 Dartmouth Summer Camp on Firms and Trade.

Abstract

Applied economists often wish to measure the effects of policy changes (like trade liberalization) or managerial decisions (like R&D expenditures or exporting) on plant-level productivity patterns. But plant-level data on physical quantities of output, capital, and intermediate inputs are usually unavailable. Therefore, when constructing productivity measures, most analysts proxy these variables with real sales revenues, depreciated capital spending, and real input expenditures. Our first objective is to argue that the resultant productivity indices have little to do with technical efficiency, product quality, or contributions to social welfare. Nonetheless, they are likely to be correlated with policy shocks and managerial decisions in misleading ways.

Our second objective is to develop an alternative approach to inference. We assume plants' costs and revenues reflect a Bertrand-Nash equilibrium in a differentiated product industry, as in Berry (1994). Given the parameters of the demand system, this allows us to impute each plant's unobserved marginal costs and product quality from its observed revenues and costs. With these in hand, we calculate each plant's contribution to consumer and producer surplus. Further, we link these welfare measures to policy and managerial decisions by assuming that marginal costs and product quality indices follow vector autoregressive (VAR) processes, conditioned on policy proxies and/or managerial choice variables. We estimate the demand system parameters and VAR parameters jointly using Bayesian techniques.

Applying our methodology to panel data on Colombian pulp and paper plants, we study the relation between our welfare-based measures and conventional productivity measures. We find that conventional productivity measures are positively correlated with producer surplus because they depend positively upon mark-ups. But the conventional measures are not closely related to product quality measures and they are nearly orthogonal to consumer surplus measures, so from a social welfare standpoint, they are poor characterizations of producer performance. We also examine the response of our welfare-based measures to policy shocks and managerial choices. Specifically, we show that conventional productivity measures imply firms that import their intermediate inputs tend to do worse, while our welfare-based measures suggest they do not.

JEL categories: L1, O3, L6

I. Overview

Economists often seek to quantify the effects of a policy or event on the performance of the manufacturing sector. Recurrent questions include: How much, if at all, does trade liberalization improve efficiency? Do multinational investments cause firms to perform significantly better? How big are the efficiency gains from R&D spending? Are there learning spillovers between firms within an industry? How do entry regulations affect an industry's performance?

To address these issues, many studies rely on plant- or firm-level productivity analysis. They posit a production function of the general form $Q_{jt} = \Phi_{jt} h(V_{jt})$, where Q_{jt} is output of the j^{th} plant in period t, $V_{jt} = \{V_{jt}^i\}_{i \in \{1,\ldots I\}}$ is a corresponding vector of inputs, and Φ_{jt} measures the plant's productivity. Then, using the available output measure (hereafter $\tilde{Q}_{jt}$) and the available input measures (hereafter $\tilde{V}_{jt}$), they construct some estimate $\hat{h}(\cdot)$ for the function $h(\cdot)$ and calculate the associated firm- or plant-specific approximations to Φ_{jt}: $\tilde{\Phi}_{jt} = \dfrac{\tilde{Q}_{jt}}{\hat{h}(\tilde{V}_{jt})}$. These they correlate with things like the extent of foreign ownership, whether the firms are exporting, rates of effective protection for the firm's product, and whether entry and exit are institutionally constrained.

When output and input characteristics are common across plants, and when data on the physical quantities of these variables are available, the use of $\tilde{\Phi}$-type measures (hereafter, "standard performance measures") makes good sense. Indeed, most of the methodological literature on this approach to analyzing firm or plant-level performance presumes that these conditions hold.[1] But in practice, $\tilde{\Phi}$ -type measures are more commonly applied to differentiated product and/or differentiated input industries, where the characteristics of products and factor inputs vary considerably across producers.[2] Under these circumstances data on physical volumes are usually unavailable, so analysts are forced make do with information on the *values* of production, material inputs, and capital stocks.[3] The resulting performance measures are therefore, roughly speaking, indices of revenue per unit input expenditure.

Such measures are viewed as a practical solution to the problem of limited data, and because they are expressed in relative value terms, they are commonly presumed to avoid the problem of comparing heterogeneous goods and factors. Our first objective in this paper is to argue that this benign view is misguided, and that standard performance measures can be very misleading when applied to differentiated product industries (Section II). Even if the function $h(\cdot)$ is precisely estimated, they are contaminated by variation in factor prices and demand elasticities. At worst, they have nothing to do with firms' productive efficiency, product quality or contribution to consumer surplus.

---

[1] Particular attention has been devoted to the issues of how to estimate the function $h(\cdot)$ and how to separate noise from "true" productivity shocks in $\tilde{\Phi}_{jt}$.

[2] A complete list of the relevant studies would take pages. Recent examples include Olley and Pakes (1996), Bahk and Gort (1993), Caves and Barton (1990), Griliches (1986); Aitken and Harrison (1999), Tybout et al (1991); Tybout and Westbrook (1995); Pavcnik (2002); Levinsohn and Petrin (2001), and Aw, Chen and Roberts (2000). Tybout (2000) surveys this type of study for developing countries; Mairesse and Sassenou (1991) survey firm level studies that relate R&D to productivity measures.

[3] On the input side, the typical data set reports the value of intermediate goods purchased, the historical cost of capital stocks, energy usage (sometimes in kilowatt hours, sometimes in value terms), and the number of workers or total hours worked, perhaps broken down by broad skill categories or gender. On the output side it describes sales revenue—sometimes distinguishing exports—and product classification according to standard industrial codes.

Our second objective is to develop an alternative approach to inference (Section III). Specifically, we view firms' costs and revenues as resulting from a Bertrand-Nash equilibrium in a differentiated product industry, as in Berry (1994), and we incorporate the demand system explicitly in the analysis. This allows us to impute the quantities, qualities, marginal costs, and prices of each good from the observed revenues and expenditures. Our approach also yields estimates of the demand system parameters. Thus it enables us to construct product-specific measures of consumer and producer surplus.

Our last objective is to demonstrate our methodology on plant-level panel data from the Colombian manufacturing sector, and to compare our performance measures with standard measures (Section IV). We find, first, that standard performance measures are positively correlated with producer surplus because they depend positively upon mark-ups. (This is probably the reason that they are correlated with producers' survival rates and growth rates.) Second, these standard measures are not closely related to product quality measures and they are nearly orthogonal to consumer surplus measures, so from a social welfare standpoint, they are poor characterizations of producer performance. Third, firms with high marginal costs also tend to produce high quality products, so studies that presume homogenous products and view marginal production costs as an (inverse) index of performance tend to under-appreciate the producers of these goods. Finally, relating firms' performances to whether they engage in international trade, we find that standard measures imply trading firms tend to do worse, while our welfare-based measures suggest they do not.

II. The Problem with standard performance measures

When data on physical output volumes are unavailable, the standard performance measure is typically constructed by replacing Q_{jt} with $\tilde{Q}_{jt} = R_{jt} / \overline{P}_t$, where R_{jt} is the j^{th} plant's nominal sales revenue and $\overline{P}_t$ is an industry-wide output price index. Similarly, if the quantity of the i^{th} input, V_{jt}^i, is unobservable, the convention is to replace it with a deflated measure of expenditure on that input, $\left(\frac{W_{jt}^i}{\overline{W}_t^i}\right) V_{jt}^i$, where $\overline{W}_t^i$ is a sector-wide deflator for factor i and W_{jt}^i is its unobserved price at the j^{th} plant. In this section we shall summarize the shortcomings of the resulting performance index,

$$\log(\tilde{\Phi}_{jt}) \equiv \tilde{\phi}_{jt} = \left(r_{jt} - \overline{p}_t\right) - \log\left[\hat{h}(\tilde{V}_{jt})\right], \tag{1a}$$

lower case letters that are not functions are expressed in natural logarithms.

To be specific about the features of $\tilde{\phi}_{jt}$-type indices, we begin by invoking some simple assumptions about consumer preferences, production technology, and the nature of competition. (Some of these assumptions will be modified later in sections III and IV.) As in Bernard *et al* (2000) and Melitz (2000), let aggregate demand for the j^{th} plant's product be given by $Q_{jt} = \Omega_{jt}^{\eta-1}\left(\frac{P_{jt}}{\overline{P}_t}\right)^{-\eta} \frac{1}{N_t}\left(\frac{Y_t}{\overline{P}_t}\right)$, $j \in \{1, \cdots, N_t\}$, where $\eta > 1$ is the price elasticity of demand, Y_t is industry-wide consumer expenditures on products in the

j^{th} producer's industry, and Ω_{jt} is an index of the quality or appeal of the j^{th} product.[4] Also as in Melitz (2000), let the j^{th} plant's production function be $Q_{jt} = \Phi_{jt} h(V_{jt})$ $= \Phi_{jt} \left[g(V_{jt}) \right]^{\gamma}$, where $g(V_{jt}) = \prod_{i=1}^{I} \left(V_{jt}^{i} \right)^{\alpha_i}$ is a constant returns Cobb-Douglas function, γ measures returns to scale, and as before, $\Phi_{jt} = e^{\phi_{jt}}$ is an index of productive efficiency. Finally, assume that firms are price takers in factor markets but they pure Bertrand-Nash price setters in the product market. Thus, assuming N is large, each behaves as if he faces a demand elasticity of η.

These assumptions imply that the j^{th} producer prices according to the mark-up rule $P_{jt}\left(1-\frac{1}{\eta}\right) = C_{jt}$, where $C_{jt} = \frac{b(W_{jt})}{\gamma} Q_{jt}^{\frac{1}{\gamma}-1} e^{\frac{-\phi_{jt}}{\gamma}}$ is the marginal cost, $b(W_{jt}) = \min_V \left[\sum_{i=1}^{I} W_{jt}^{i} V^{i} \mid 1 = g(V) \right] = \alpha_0 \prod_{i=1}^{I} \left(W_{jt}^{i} \right)^{\alpha_i}$ is the minimum cost of a unit bundle of factor services, and $W_{jt} = \{W_{jt}^{i}\}_{i \in \{1,\ldots,I\}}$ is the plant-specific vector of unit factor prices associated with the input vector V_{jt}. Multiplying both sides of the mark-up rule by Q_{jt} yields:

[4] Suppressing time subscripts, the associated utility function is CES in product varieties: $U\left(\left[\sum_j \left(\Omega_j Q_j \right)^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}, Z \right)$, where Z is an outside good.

$$R_{jt}\left(1-\frac{1}{\eta}\right)=\frac{b(W_{jt})}{\gamma}Q_{jt}^{\frac{1}{\gamma}}e^{\frac{-\phi_{jt}}{\gamma}}=\frac{b(W_{jt})}{\gamma}\left(e^{\phi_{jt}}g(V_{jt})^{\gamma}\right)^{\frac{1}{\gamma}}e^{\frac{-\phi_{jt}}{\gamma}}=\left(\frac{b(W_{jt})}{\gamma}\right)g(V_{jt}).$$

where $R_{jt}=P_{jt}Q_{jt}$. Thus, dividing both sides by the price index $\overline{P}_t$, taking logs, and rearranging, we may write the log of deflated revenues as:

$$r_{jt}-\overline{p}_t=\ln\left[g(V_{jt})\right]+\ln\left(\frac{\eta}{\eta-1}\right)+\ln\left[b(W_{jt})\right]-\overline{p}_t-\ln(\gamma). \qquad (2)$$

Subtracting $\ln\left[\hat{h}(\tilde{V}_{jt})\right]$ from both sides of equation (2) yields the standard performance measure (1a). Our concern is not with the estimation of production technologies, so we shall assume that analysts are somehow able to correctly estimate the function $g(\cdot)$ and returns to scale, γ.[5] Then $\ln\left[\hat{h}(\tilde{V}_{jt})\right]=\gamma\ln\left[g(\tilde{V}_{jt})\right]$ and (1a) may be written as:

$$\tilde{\phi}_{jt}=\ln\left[g(V_{jt})\right]-\gamma\ln\left[g(\tilde{V}_{jt})\right]+\ln\left(\frac{\eta}{\eta-1}\right)+\ln\left[b(W_{jt})\right]-\overline{p}_t-\ln(\gamma). \qquad (1b)$$

We are now in a position to state our complaints about $\tilde{\phi}-$type measures. To begin, let us momentarily suppose that all factor inputs can be precisely measured, so $\tilde{V}_{jt}=V_{jt}$ and each element of the factor vector is homogeneous across plants and time. (Think of inputs that differ in quality as distinct inputs.) Then equation (1b) simplifies to:

$$\tilde{\phi}_{jt}=(1-\gamma)\ln\left[g(V_{jt})\right]+\ln\left(\frac{\eta}{\eta-1}\right)+\ln\left[b(W_{jt})\right]-\overline{p}_t-\ln(\gamma), \qquad (1c)$$

[5] Estimation errors introduce another type of problem with $\tilde{\phi}$-type measures, but they do not undo the ones we focus on here.

and remarkably, under constant returns to scale ($\gamma = 1$), $\tilde{\phi}_{jt}$ has nothing to do with productive efficiency or product quality.[6] Rather, it depends exclusively upon demand elasticities and factor prices relative to the output price deflator.

One implication of (1c) is that cross-plant and inter-temporal variation in demand elasticities (η) may induce some misleading conclusions. For example, producers of close substitutes may look relatively inefficient because their demand is relatively elastic. Also, since trade liberalization and exchange rate appreciation tend to reduce the market power of the largest domestic firms—which compete most directly with imports—these shocks will tend to reduce both the average $\tilde{\phi}_{jt}$ value and dispersion in $\tilde{\phi}_{jt}$, even if true productivity remains unaffected.[7] We note, however, that variation in $\tilde{\phi}_{jt}$ due to elasticity effects might contain *some* information on (ω_{jt}, ϕ_{jt}) if we were to make different demand-side assumptions. Specifically, if producers with high ω_{jt} or ϕ_{jt} values were to face relatively low demand elasticities, $\tilde{\phi}_{jt}$ would exhibit positive covariance with these variables.[8]

Other misleading results are likely to come from real factor price variation. For example, in open economies, real exchange rate appreciation often drives up real wages in the tradeable goods industries. But it also tends to increase import penetration rates,

---

[6] Klette and Raknerud (2001) and Bernard *et al* (2000) make the same observation in slightly different contexts.

[7] This dispersion effect is one interpretation for the findings of Caves and Barton (1990).

[8] An inverse relationship between ω_{jt} or ϕ_{jt} and the perceived elasticity of demand arises in a variety of contexts. In the current set-up, one can induce such a relationship by assuming *N* is small. The nested logit demand system we adopt in sections III and IV will also exhibit this property.

so $\tilde{\phi}$ -type indices may falsely create the impression that import competition improves productivity among tradeable goods producers. Interestingly, the market power effect of appreciation we mentioned above works against this real factor price effect. This may help explain why the literature relating trade liberalization and import penetration rates to $\tilde{\phi}_{jt}$ -type measures reports mixed results (Tybout, 2001).

Relaxing the assumption of constant returns to scale makes $\tilde{\phi}_{jt}$ depend upon $\phi_{jt} + \omega_{jt}$, but not necessarily in a desirable way. Specifically, under increasing (decreasing) returns to scale, $\tilde{\phi}_{jt}$ decreases (increases) with $g(V_{jt})$, which is itself an increasing function of $\phi_{jt} + \omega_{jt}$.[9] Most of the productivity studies based on an econometrically estimated $h(\cdot)$ functions find mild *increasing* returns, so this scale effect usually tends to induce a perverse negative correlation between measured performance and $\phi_{jt} + \omega_{jt}$.[10]

---

[9] As Melitz (2000) shows, our production function and CES demand system imply the market clearing condition: $r_{jt} - \bar{p}_t = \left(\frac{\eta - 1}{\eta}\right)\gamma \cdot \ln[g(V_{jt})] + \left(\frac{1}{\eta}\right)[(r_t - \bar{p}) - n_t] + \left(\frac{\eta - 1}{\eta}\right)(\phi_{jt} + \omega_{jt})$. Substituting this expression into the mark-up condition (2), one obtains a reduced-form expression for factor demand: $\ln[g(V_{jt})] = \left[\frac{\eta}{\gamma + \eta \cdot (1-\gamma)}\right]\left\{\frac{1}{\eta}(r_t - \bar{p}_t) + \left(\frac{\eta - 1}{\eta}\right)(\phi_{jt} + \omega_{jt}) + \gamma + \ln\left(\frac{\eta - 1}{\eta}\right) - \ln[b(W_{jt})] + \bar{p}_t\right\}$, so $\mathrm{cov}[\ln g(V), \phi + \omega] > 0$. For future reference, we note that the market clearing condition above is the focus of Melitz's (2000) paper. He uses it as a basis for estimation γ and η, and then as a means to solve for a compound performance index: $(\phi_{jt} + \omega_{jt}) = \left(\frac{\eta - 1}{\eta}\right)^{-1}\left[r_{jt} - \bar{p}_t - \left\{\left(\frac{\eta - 1}{\eta}\right)\gamma \cdot \ln[g(V_{jt})] + \left(\frac{1}{\eta}\right)[(r_t - \bar{p}) - n_t]\right\}\right]$.

[10] If there are *decreasing* returns to scale, the expressions in the previous footnote can be used to describe the limiting case of perfectly competitive product markets by letting η approach infinity. (A unique pure Bertrand-Nash equilibrium does not exist in the limit if returns to scale are constant or increasing.) It is easy to demonstrate that as $\eta \to \infty$, $\tilde{\phi}_{jt} \to \phi_{jt} + \omega_{jt}$.

One other scale effect merits comment. If γ varies across firms, those with large scale economies will appear to be doing relatively poorly, even if the same γ value is used to construct all $\tilde{\phi}_{jt}$ values, so long as firms know their own returns to scale parameters.

Now suppose we relax the assumption that all inputs are accurately measured in physical terms. Let the subset of inputs that are measured in deflated expenditure terms be $E \subseteq \{1,\ldots,I\}$, so that the measured input vector $\tilde{V}_{jt}$ has components $\tilde{V}_{jt}^{i} = \left(\frac{W_{jt}^{i}}{\overline{W}_{t}^{i}}\right) V_{jt}^{i}$ for $i \in E$ and $\tilde{V}_{jt}^{i} = V_{jt}^{i}$ for $i \notin E$. Then, using our expressions for $g(\cdot)$ and $b(\cdot)$ we obtain

$$\ln\left[g(V_{jt})\right] - \gamma \ln\left[g(\tilde{V}_{jt})\right] + \ln\left[b(W_{jt})\right] =$$

$$\ln(\alpha_0) + \sum_{i \notin E} \alpha_i \ln(W_{jt}^{i}) + \gamma \sum_{i \in E} \alpha_i \ln(\overline{W}_{t}^{i}) + (1-\gamma) \sum_{i \in E} \alpha_i \ln(W_{jt}^{i}) + (1-\gamma) \ln\left[g(V_{jt})\right]$$

so using (1b), $\tilde{\phi}_{jt}$ can be written as:

$$\begin{aligned}\tilde{\phi}_{jt} &= \ln(\alpha_0) + \ln\left(\frac{\eta}{\eta-1}\right) + \sum_{i \notin E} \alpha_i \ln\left(\frac{W_{jt}^{i}}{\overline{P}_t}\right) + \gamma \sum_{i \in E} \alpha_i \ln\left(\frac{\overline{W}_{t}^{i}}{\overline{P}_t}\right) \\ &\quad + (1-\gamma)\left(\sum_{i \in E} \alpha_i \ln\left(\frac{W_{jt}^{i}}{\overline{P}_t}\right) + \ln\left[g(V_{jt})\right]\right) - \ln(\gamma)\,.\end{aligned} \tag{1d}$$

All of the properties of (1c) are still present in (1d). But now the role of factor prices is more nuanced. Looking across plants, suppose for the moment that factor prices are unrelated to productive efficiency (ϕ_{jt}) and product quality (ω_{jt}). (The variation in factor prices might come instead from differences in delivery costs, fringe benefit obligations, congestion costs, or other geographic effects like export processing zones.)

Then those plants that pay relatively high prices for the factors that are measured in physical terms, like labor and energy, will appear to be relatively productive. Under increasing (decreasing) returns to scale, $\tilde{\phi}_{jt}$ also falls (rises) with increases in the real price of those factors measured in expenditure terms.

These features of $\tilde{\phi}$-type measures create further scope for spurious inferences. For example, large firms, R&D-intensive firms and multinational firms will tend to look relatively productive because their employees are relatively well-paid, even if their technical efficiency is no better than average.[11] This effect may help explain the robust positive cross-sectional association between R&D and $\tilde{\phi}$-type indices (Mairesse and Sassenou, 1991), as well as the association between foreign ownership and $\tilde{\phi}$-type indices (Blomstrom and Kokko, 1997). Also, under increasing (decreasing) returns to scale, the spurious positive effect of real factor prices on $\tilde{\phi}_{jt}$ induced by real appreciation should be reinforced (weakened) among firms that import their intermediate inputs. Thus, given that multinationals and exporters rely more intensively on imported intermediate goods (Kraay, *et al*, 2001), $\tilde{\phi}$-type indices may falsely imply that they are relatively efficient (inefficient) during periods of real appreciation and heightened import competition. (We will return to this effect in section IV.D below.)

Of course, when there is unobserved factor heterogeneity, it is unrealistic to assume that cross-plant and inter-temporal variation in factor prices is orthogonal to ϕ_{jt} and ω_{jt}, as we have been doing thus far. If factors are paid the value of their

[11] Here we are assuming that their productivity measure is constructed using an index of physical labor use rather than a measure of expenditures on labor.

marginal product, and we continue to assume constant returns, we should observe a positive cross-sectional dependence of $\tilde{\phi}_{jt}$ on ϕ_{jt} and ω_{jt} so long as some inputs are measured in physical terms.[12] Similarly, we should observe a positive temporal dependence of $\tilde{\phi}_{jt}$ on ϕ_{jt} and ω_{jt} if some inputs are measured in expenditure terms. But these associations simply tell us that firms using high quality inputs get more and/or better output. They reveal nothing about which firms are doing well in an economic sense, or whether firms are getting better over time.

To all of the above, one might object that $\tilde{\phi}$-type performance measures nonetheless seem to "work." That is, many studies have found that firms with high $\tilde{\phi}_{jt}$ values are more likely to be large or grow, and they are less likely to fail (Baily, Hulten and Campell, 1992; Olley and Pakes, 1996; Aw, Chen and Roberts, 2000; Baldwin and Gorecki, 1991; Lu and Tybout, 1996; Pavcnik, 2002). Does not this finding imply, as the authors of these studies suggest, that high-$\tilde{\phi}_{jt}$ firms are more efficient and/or produce a relatively desirable product? It need not. Success ultimately depends upon profits rather than efficiency or product quality, and firms with low demand elasticities (i.e., large values of $\frac{\eta}{\eta-1}$) tend both to be profitable and to have high $\tilde{\phi}_{jt}$ values, even if their productive efficiency and product quality are unexceptional.

To summarize, when analyzing differentiated product industries, it is a mistake to pretend that sales revenues and input expenditures measure physical outputs and inputs,

[12] This dependence of measured productivity on unobservable aspects of factor quality is well known (e.g., Griliches and Jorgenson, 1967).

respectively. This convention leads to spurious measures of productivity that may have little to do with efficiency or product quality yet tend to be correlated with policy shocks and managerial choices. Furthermore, even when efficiency and product quality are captured in some way, these performance measures do not tell us anything about firms' contributions to welfare.

III. An Alternative Approach to Measuring Performance

How, then, is one to infer something about plants' performances when neither their physical output volumes (Q_{jt}), their physical inputs (V_{jt}), nor their prices (P_{jt}, W_{jt}) are available? If total variable cost and total revenue pairs, $(TC_t, R_t) = \{TC_{jt}, R_{jt}\}_{j=1,\ldots N_t}$, *are* observable for all plants—as is often the case—we suggest viewing (TC_t, R_t) as reflecting equilibrium in a differentiated product industry. Then, by imposing enough structure on producer behavior and the demand system, it is possible to induce a one-to-one mapping between (TC_t, R_t) and (ω_t, C_t), where $(\omega_t, C_t) = \{\omega_{jt}, C_{jt}\}_{j\in 1,\ldots N_t}$ and C_{jt} denotes the j^{th} plant's marginal cost schedule in period *t*. When combined with demand parameters, this mapping allows one to impute plant-specific consumer and producer welfare measures from the observable data, to study the evolution of these welfare measures over time, and to relate them to policy shocks and managerial choices.[13] Below we describe one way to implement this strategy in detail.

A. The demand system, producer behavior, and market equilibrium

To begin, let us mention several constraints on our characterization of market equilibrium. Clearly, if plants produce multiple products, or if the shapes of marginal cost schedules are not somehow restricted, it will be problematic to establish a one-to-one mapping between (ω_t, C_t) and (TC_t, R_t). Also, our assumptions must allow us to map *any* feasible (TC_t, R_t) realization onto a feasible (ω_t, C_t). Thus, for example, we cannot make behavioral assumptions that imply all firms choose identical mark-ups—as we did in the section II—since this would preclude mappings from any set (TC_t, R_t) in which the ratios $\frac{R_{jt}}{TC_{jt}}$, $j \in \{1, \ldots, N_t\}$, differ significantly across firms.[14]

Given these constraints, we assume that each plant produces a single product, and that marginal costs at the j^{th} plant in year t are given by the scalar C_{jt}, regardless of that plant's output level.[15] (We will, however, allow C_{jt} to evolve through time with plant-

---

[13] Without data on factor prices it is impossible to impute productivity measures, $\phi_t = \{\phi_{jt}\}_{j=1,\ldots N_t}$, from observable variables. But these are relevant for welfare only inasmuch as they influence marginal costs, which *are* identified.

[14] The maintained hypothesis that all firms have similar mark-ups is, in our view, the main limitation of Melitz's (2000) approach to performance measurement (see footnote 9).

[15] This is more restrictive than necessary. So long as firms share the same cost function up to a single parameter, one-to-one correspondences are not ruled out.

specific productivity and factor price shocks.) Also, we adopt the nested logit version of Berry's (1994) model of market equilibrium.

1. The nested logit demand system

A brief review of Berry's (1994) model will serve to introduce parameters and their interpretation in the present context. At time *t* let $j \in \{0,1,\dots N_t\}$ index the $N_t + 1$ available varieties, with $j = 0$ corresponding to the "outside" variety. Also, assume that the product varieties can be grouped into $G+1 < N_t + 1$ "nests." In our application the outside product will be a composite imported variety, and the remaining varieties will map one-to-one onto the set of active domestic plants. Also, we will define product nests according to the geographic region in which the plants are located. For example, nest *1* includes the varieties of pulp and paper products manufactured in Bogota, nest *2* includes the varieties manufactured in Medellin or Cali, and so on. (Alternatively, one could distinguish nests within a 4-digit industry by 5-digit product classifications.)

Next, let consumers be indexed by the real number $\ell \in (0, Y_t]$. Each period, each consumer in the market chooses a single unit of the variety that yields her the largest net indirect utility, where variety *j* yields consumer ℓ net utility:

$$u_{\ell jt} = \bar{u}_{jt} + \varsigma_{\ell g_j t} + (1-\sigma)\varepsilon_{\ell jt} \,. \tag{5}$$

Here $g_j \in \{1, \cdots G\}$ denotes the index for the group (nest) that contains the j^{th} variety, and $\bar{u}_{jt}$ is the cross-consumer mean utility delivered by good *j*. The last two terms on the right hand side of (5) are unobserved error components that capture individual taste differences among consumers. The first component, $\varsigma_{\ell g_j t}$, varies across nests but not within them, while $\varepsilon_{\ell jt}$ exhibits within-nest variation. Thus the parameter $0 \le \sigma < 1$

indexes the degree of substitutability among, versus within, the nests.[16] Finally, ε is distributed type-1 extreme value across consumers, given j and t, and $[\varsigma + (1-\sigma)\varepsilon]$ is distributed type-1 extreme value across consumers, given t. This implicitly defines the distribution ς, which is itself a function of σ (Cardell, 1997).

The mean indirect utility delivered by domestic good j in period t depends on both its quality and price:

$$\bar{u}_{jt} = \xi_{jt} - \alpha \ P_{jt} \ , \quad \forall j = 1,..., N, \tag{6}$$

where ξ_{jt} measures product quality, P_{jt} is the price, and α measures the price effect on the mean indirect utility of a given domestic good. Similarly, the mean indirect utility delivered by a unit of the composite imported good is:

$$\bar{u}_{0t} = \xi_{0t} - \gamma \cdot e_t , \tag{7}$$

where e_t is the domestic currency price of a unit of imports, calculated as the product of the tariff-inclusive real price of imports in U.S. dollars and the real effective peso-dollar exchange rate. Since the imported good is a composite, we allow α and γ to differ.

Integrating over consumers yields total demand for each variety as a function of the price vector P_t, the vector of domestic product qualities relative to the quality of the imported good, $\omega_t = \{\omega_{jt} = \xi_{jt} - \xi_{ot}\}_{j \in 1,\dots N_t}$, and the price of the imported good, e_t. Expressed as share of total demand for varieties in the j^{th} product's nest, demand for the j^{th} domestic variety is:

---

[16] As σ goes to zero, within-group correlation of utilities goes to zero, and as σ goes to unity, within-group correlation goes to unity. A more general specification lets σ vary across groups, allowing richer substitution patterns (Berry, 1994; Berry, Levinsohn and Pakes, 1995). This specification has important advantages, but it requires that we observe information about the distinctive features of each group, which makes it infeasible for the present application.

$$S_{j|g_j,t} = \frac{\exp[(\omega_{jt} - \alpha P_{jt} + \gamma e_t)/(1-\sigma)]}{\sum_{k \ni g_k = g_j} \exp[(\omega_{kt} - \alpha P_{kt} + \gamma e_t)/(1-\sigma)]} \quad \text{and}$$

(8)

Also, expressed as a share of total demand for the industry's product, demand for the products in the j^{th} product's nest as is:

$$S_{g_j,t} = \frac{\left\{ \sum_{k \ni g_k = g_j} \exp[(\omega_{kt} - \alpha P_{kt} + \gamma e_t)/(1-\sigma)] \right\}^{(1-\sigma)}}{\sum_{m=1}^{G} \left\{ \sum_{k \ni g_k = m} \exp[(\omega_{kt} - \alpha P_{kt} + \gamma e_t)/(1-\sigma)] \right\}^{(1-\sigma)} + 1}. \quad (9)$$

Finally, consumer surplus can be written as (McFadden, 1981):

$$CS_t = Y_t \cdot E[Max_j(u_{jt})] = y_t \cdot \log\{\exp(\overline{u}_{0t}) + [\sum_{j=1}^{N_t} \exp(\overline{u}_{jt}/(1-\sigma))]^{(1-\sigma)}\} \quad (10)$$

2. Firm-Specific Product Quality

Combining equations (6)-(9) and normalizing the mean utility from the imported good to zero ($\overline{u}_{0t}$ = 0), prices and market shares can be used to impute the relative quality of good *j*:

$$\omega_{jt} \equiv \xi_{jt} - \xi_{0t} = \alpha P_{jt} - \gamma e_t - \sigma \ln(S_{j|g_j,t}) + \ln\left(S_{j|g_j,t} \cdot S_{g_j,t}\right) - \ln(S_{0t}) \quad j = 1,\ldots,N_t \quad (11)$$

This measure of product appeal permits plant-by-plant quality comparisons in differentiated product industries without detailed information about the characteristics and prices of plant-specific inputs and outputs.[17]

[17] Berry, Levinsohn and Pakes (1995) use a similar inversion to study the quality of automobile models.

3. Market Equilibrium

Given the current group of active producers, product prices are determined by pure strategy Bertrand-Nash competition in the output market. Each period, taking the current prices of all other goods as given, the j^{th} domestic producer chooses P_{jt} to maximize $\pi_{jt} = (P_{jt} - C_{jt}) \cdot S_{j|g_j,t} \cdot S_{g_j,t} \cdot Y_t$, subject to (8) and (9), where C_{jt} is the marginal cost of production at the j^{th} plant. The standard first order conditions imply that in equilibrium, prices satisfy (Berry,1994):

$$P_{jt} = C_{jt} + \frac{(1-\sigma)/\alpha}{1-\sigma \cdot S_{j|g_j,t} - (1-\sigma) \cdot S_{j|g_j,t} \cdot S_{g_j,t}}, \qquad j = 1, \ldots N_t \tag{12}$$

Given (ω_t, C_t), r_t, and the demand parameters (α, γ, σ), a unique set of market shares and prices solves (8), (9) and (12) (Caplin and Nalebuff, 1991). This equilibrium exhibits several intuitive properties. First, firm *j*'s profits (π_{jt}) increase in its own product quality (ω_{jt}), the price of imports (e_t), and market size (Y_t), but decrease in other firms' product quality ($\omega_{kt}\ \forall k \in \{1, \ldots, N_t\}\ k \neq j$). Second, producers with lower marginal cost (C_{jt}) or high product quality (ω_{jt}) realizations also enjoy bigger market shares and profit margins. (Note that big mark-ups imply low demand elasticities, so $\tilde{\phi}_{jt}$ is linked to (c_{jt}, ω_{jt}) through elasticity effects if this demand system characterizes behavior—refer to section II.) Third, at given prices, a reduction in the number of varieties reduces consumer surplus and reduces the own-price elasticity of demand for each remaining variety. Fourth, heightened import competition, as reflected in a reduction in e_t, reduces the prices and market shares of the other goods. Finally, holding the number and characteristics of producers fixed, a reduction in the price of imports increases consumer surplus (6).

B. Imputing product quality and marginal costs

Given the quantity of imports (Q_{0t}) and the demand parameters (α, γ, σ), the conditional mapping from $(\omega_t, C_t \mid e_t)$ to prices and market shares (P_t, S_t) can be restated as a mapping to prices and quantities: $(\omega_t, C_t \mid r_t, Q_{0t}) \rightarrow (P_t, Q_t)$. (That is, Q_{0t} may be used to fix total market size.) In turn, this implies a mapping from $(\omega_t, C_t \mid e_t, Q_{0t})$ to total revenue and total cost pairs, (R_t, TC_t), because $(R_t, TC_t) = \{P_{jt}Q_{jt}, C_{jt}Q_{jt}\}_{j \in \{1,\ldots N_t\}}$. Further, the mapping $(\omega_t, C_t \mid e_t, Q_{0t}) \rightarrow (R_t, TC_t)$ is one-to-one, so it is possible to solve for a unique $(\omega_t, C_t, Q_t, P_t)$, using observable data on (R_t, TC_t), e_t and Q_{0t}. (Appendix I provides a proof.) Therefore, after the demand parameters have been estimated, it is possible to retrieve the necessary ingredients for analysis of producer and consumer surplus. (Details on the construction of these welfare measures are provided in section E below.)

C. Linking product quality and marginal costs to managerial decisions

To complete our model, we add equations that link product quality and market costs to the business environment and managerial decisions. More precisely, we assume that product quality and the log of marginal costs evolve over time according to a vector autoregressive (VAR) process, conditioned on a vector $\boldsymbol{X}_{jt}$ of weakly exogenous variables, including things like R&D expenditures, participation in foreign markets, and the extent of multinational ownership:

$$\omega_{jt} = \omega_0 + \sum_{s=1}^{L} \lambda_s \omega_{j,t-s} + \sum_{s=L+1}^{2L} \lambda_s c_{jt+L-s} + \lambda^x \boldsymbol{X}_{jt} + \varepsilon^{\omega}_{jt}, \tag{13a}$$

$$c_{jt} = c_0 + \sum_{s=1}^{L} \varphi_s c_{j,t-s} + \sum_{s=L+1}^{2L} \varphi_s \omega_{j,t+L-s} + \varphi^x \boldsymbol{X}_{jt} + \varepsilon^{c}_{jt}, \tag{13b}$$

$$j = 1,\cdots,N \;,\quad t = L+1,\cdots,T \,.$$

Accordingly, once the complete vector of parameters has been estimated, it will be possible to trace the welfare effects of any deviation from observed $\boldsymbol{X}_{jt}$ trajectories.

D. Estimation

In addition to providing a basis for policy evaluation, equations (13a) and (13b) help with identification by constraining the way that quality and marginal cost evolve. However, these dynamic relationships bear only obliquely on the demand parameters, and they introduce some new unknowns to be estimated. Prospects for successful maximum likelihood estimation are further dimmed by the irregular shape of the likelihood function for the nested logit (Lahiri and Gao, 2001). Therefore we impose further structure by specifying priors on the demand system parameters (α,γ,σ) and estimating the system (8), (9), (11), (12), (13) using Bayesian techniques.

To summarize this estimation strategy, let us collect all of the parameters we have introduced in the vector $\theta = [\alpha,\gamma,\sigma,\lambda,\varphi,\Sigma]$, where $\Sigma = E\left[\begin{pmatrix}\varepsilon_c \\ \varepsilon_\omega\end{pmatrix}\begin{pmatrix}\varepsilon_c \\ \varepsilon_\omega\end{pmatrix}'\right]$, and define the joint density $f(\theta)$ to describe our priors (to be discussed shortly). Also, let us collect all of the observable data on revenues, costs, imports, the exchange rate, and weakly exogenous firm characteristics in the matrix *D*. Then the posterior distribution for θ is:

$P(\theta \mid D) = \dfrac{f(\theta) \cdot L(D \mid \theta)}{\int_\theta f(\theta) \cdot L(D \mid \theta) d\theta} \propto f(\theta) \cdot L(D \mid \theta)$, where $L(D \mid \theta)$ is the likelihood function based on the pricing rule (12), the market clearing conditions (8) and (9), and the VAR system (13).

Excepting elements of the covariance matrix, Σ, we have no reason to expect that the parameters of our model are correlated. Thus we write the joint prior distribution as a product of our prior marginal densities for the individual parameters:

$$f(\theta) = f_\sigma(\sigma) \cdot f_{\alpha,\gamma}(\alpha, \gamma) \cdot f_{\varphi,\lambda}(\varphi, \lambda) \cdot f_\Sigma(\Sigma).$$

Let us describe each component of $f(\theta)$ in turn. First, the demand system priors we impose are similar to those used by Poirier (1996) and Lahiri and Gao (2001).[18] The underlying utility maximization problem implies that $\sigma \in [0,1]$, so we specify uniform priors on this region of support. [19] Second, we believe the price coefficients α and γ should be positive but we do not know much about their magnitudes, so we specify uniform priors with support [0, 10] for each of these parameters. The remaining parameters describe the VAR (equations 13a and 13b). For the autoregressive parameters we assume joint normality, $f_{\varphi,\lambda}(\varphi, \lambda) = N_{2K}(0_{2K}, 100 \times I_{2K})$, where *K* is one plus the

---

[18] These studies also estimate nested logit models using Bayesian techniques. However, unlike ours, they are concerned with the problem of ill-defined nesting structures.

[19] Restricting σ to be greater or equal to zero reflects our prior knowledge that the products within each nest are at least as good substitutes for each other as those products outside the nest. σ greater than one is not consistent with the underlying assumption of the extreme value distribution of consumer tastes. As σ goes to unity, consumers would purchase only the goods with the highest mean indirect utility in each nest. Restricting σ to be less than or equal to one ensures that all products get consumed.

number of right-hand side variables appearing in each VAR equation. Finally, as is standard in the literature, for the covariance matrix we assume an inverted-Wishart distribution, $f_{\Sigma}(\Sigma) = InvWish(6, 100 \times I_2)$. Overall then, with the exception of σ, we are doing little to constrain the range of plausible realizations on θ.

Closed-form representations of the posterior $P(\theta \mid D)$ are not available; nor is it feasible to make i.i.d. draws directly from $P(\theta \mid D)$. We therefore use a Markov chain Monte Carlo (MCMC) algorithm to generate correlated draws from $P(\theta \mid D)$ and we analyze the moments of the resulting empirical distributions (Gilks, et al, 1996).

The vector θ is relatively large, so we exploit Gibbs sampling techniques to generate our Markov chain. That is, we partition θ into three sub-vectors: $\theta = [\theta_1, \theta_2, \theta_3]$ where $\theta_1 = (\alpha, \gamma, \sigma)$, $\theta_2 = (\lambda, \varphi)$ and $\theta_3 = \Sigma$. Then we update the sub-vectors sequentially by drawing from the full conditional distributions of each in turn. The full conditional distribution of ($\theta_2 \mid \theta_1, \theta_3, D$) is multivariate normal because D and θ_1 imply the (ω, c) trajectories, which contain all of the available information on θ_2. For the same reason, the full conditional distribution of ($\theta_3 \mid \theta_1, \theta_2$, D) is inverted-Wishart. Thus closed-form expressions for the full conditional distributions of θ_2 and θ_3 are easy to construct, and sampling from these distributions is straightforward (Zellner, 1971). However, no simple expression for the full conditional distribution of ($\theta_1 \mid \theta_2, \theta_3, D$) is available, so we use a Metropolis-Hastings sampling algorithm. Appendix 2 provides further details.

E. Constructing Performance Measures

Once we have estimated our posterior distribution, $P(\theta \mid D)$, we solve for the marginal cost and product quality trajectories of each producer in the sample using the expected value of θ.[20] The remaining task is then to translate these trajectories into meaningful performance measures, and to examine the relationship between those measures and the traditional Tornqvist indices described in Part I above.

For the i^{th} producer, we calculate the increment to consumer surplus that it generates each period by evaluating equation (10) with, versus without the i^{th} good:[21]

$$\Delta CS_t^i = \quad y_t \cdot \left[[\sum_j \exp(\bar{u}_{jt} /(1-\sigma))]^{(1-\sigma)} \} - [\sum_{j \neq i} \exp(\bar{u}_{jt} /(1-\sigma))]^{(1-\sigma)} \} \right] \tag{14}$$

Prices and market shares are allowed to adjust to re-establish equilibrium when good *m* is removed. Similarly, we calculate the i^{th} producer's own surplus as $(P_{it} - C_{it})Q_{it}$, and from this we subtract the negative externality this producer imposes on the surplus of other firms. The latter is imputed by evaluating $\sum_{j \neq i} (P_{jt} - C_{jt})Q_{jt}$ with, versus without, the i^{th} producer present, letting prices and market shares adjust to re-establish equilibrium.

---

[20] It would, of course, be possible to also study the distributions for these trajectories that are induced by $p(\theta \mid D)$; we have not pursued this yet.

[21] Ackerberg and Rysman (2001) argue that the nested logit demand system overstates the contribution to consumer surplus provided by each product because it implies very high marginal utility from the first units consumed of each good. To explore the importance of this possible bias, we also look at the rankings of firms in terms of their contribution to consumer surplus as well as the cardinal value of these contributions.

To evaluate these firm-specific welfare contributions, we express them as ratios to firms' reported capital stock to obtain a crude social rate of return on investment. (All other costs of operation are captured by variable costs and will already be netted out of producer surplus.) Obviously we miss pre- and post-sample costs and benefits and our measure of firms' assets will be very crude, but we feel we will come closer to a comprehensive basis for assessment than the standard methodologies. For the sake of comparison, we also calculate the usual Tornqvist measures of total factor productivity under the standard assumptions that deflated revenues measure real output, and deflated expenditures on intermediate goods measure physical intermediate good usage:

$$\tilde{\phi}_{jt} = \ln(R_{jt} / \overline{P}_t) - \sum_{i=1}^{I} \hat{\alpha}^i_{jt} \ln(\tilde{V}^i_{jt}), \tag{15}$$

where $\hat{\alpha}^i_{jt}$ is the share of the i^{th} factor in total costs at firm *j* during period *t*. These indices are normalized to have zero cross-sectional means in each period as suggested by Good, Nadiri and Sickles (1997).

IV. An Application to the Colombian Pulp and Paper Mill Industry

A. The Data

We base our empirical example on panel data describing the Colombian pulp and paper mill industry over the period 1981-1991. These data were originally collected by Colombia's official statistical agency (Departmento Administrativo Nacional de Estadistica) and have been cleaned as described in Roberts (1996). To keep the analysis simple we exclude plants that entered or exited during the sample period, leaving a total of 13 plants over an 11 year period.[22] This naturally creates some selection bias, although the entering and exiting plants were quite small and thus had a minor influence on market shares.

We construct total domestic sales, R_{jt}, as total sales revenue less the value of exports divided by a general wholesale price deflator. To construct total variable costs, TC_{jt}, we first sum payments to labor, intermediate input purchases (net of inventory accumulation), and energy purchases.[23] Then we scale this aggregate by the ratio of total domestic sales to total sales and we divide the result by the same wholesale price deflator we used for output.

Our real exchange rate series, e_t, is taken from Ocampo and Villar (1995), who include an adjustment for tariffs. To impute imports we assume that all imported goods in the relevant industrial classification maintain their same exogenous dollar price during the sample period. Further, we assume that the imported varieties are consumed in fixed proportion to one another, so that they can be treated as a single bundle whose domestic price fluctuates only with the exchange rate. Then, calling the dollar value of imports R_0,

---

[22] Entry and exit would complicate the VAR portion of the likelihood function by creating an unbalanced panel.

[23] Thus we are assuming that capital stocks are exogenous to the firm and all other factors are variable. An equally simple approach would be to assume that capital stocks are perfectly flexible, and to include a rental cost of capital—say 10 percent of the book value—in our total cost measure. The intermediate case in which capital stocks (and perhaps other inputs) are subject to finite adjustment costs is difficult to deal with because it means introducing dynamic optimization into the analysis.

we construct our index of the quantity of imports as $Q_{0t} = \frac{R_{0t}}{e_t}$. The units in which Q_{0t} is measured determine the units in which all domestic varieties are measured.[24]

Finally, the vector of weakly exogenous variables ($\boldsymbol{X}_{jt}$) includes the book value of each plant's initial (1981) capital stock, a trend term, and two dummy variables that summarize plants' participation in foreign markets. The first takes a value of one if the plant was importing some or all of its intermediate inputs in year *t-1*, but *not* exporting any of its output. The second dummy takes a value of one if the plant was both importing some intermediates and exporting some output in year *t-1*. No plant in our sample exported output without importing intermediate inputs, so the omitted category is simply plants that did not buy inputs or sell outputs in international markets in year *t-1*.[25]

B. Posterior Parameter Distributions

Means, standard errors and other summary statistics for our estimated posterior distribution $P(\theta \mid D)$ are reported in Table 1 below. The estimates are constructed using Wooldrich's (2001) correction for persistent unobserved heterogeneity in the disturbance term.[26]

---

[24]Unfortunately, the choice of the units of Q_0 also has implications concerning import volume shares. If we were to halve the imputed quantity of imports, the imputed volume share of imports would also be smaller. This reflects the fact that domestic quantities are not linear in Q_0. An increase in Q_0 does imply bigger domestic quantities but the increase is less than proportional. In practice, we normalize the series of real exchange rate so that in the base year revenue share of imports equals its volume share.

[25] It would have been desirable to also include R&D spending, and to distinguish firms according to whether they were partly owned by foreigners. Unfortunately, this information was not available.

[26] Wooldrich's (2001) correction takes care of initial conditions problem. It amounts to including the initial value of the lagged dependent variables as explanatory variables in all years, and using a standard error components specification for the disturbance. Kraay *et al* (2001) provide further discussion in the context

Overall, the results appear quite well behaved, although the posterior distribution for the price coefficient α is rather diffuse.[27] Given the small sample size we are working with, this is perhaps unsurprising. The mixing parameter σ tends to be close to unity, suggesting that most of the variation in tastes across consumers has to do with region of origin. The VARs for product quality and marginal cost both show a plausible amount of persistence. On the other hand, there is less dynamic interaction among these variables than we expected. (That is, we thought high marginal cost in one year might lead to higher product quality in the next year.) Nonetheless, as we will see shortly, these variables do exhibit contemporaneous covariance across plants because of persistent plant effects. Finally, our coefficients for trends, international transactions, and capital stocks are not estimated with much accuracy. However, the impact of importing intermediate goods on marginal costs and product quality is large and negative, on average. We will return to explore the implications of these international transactions coefficients in section D below.

C. Plant performance measures

Using the posterior means of our demand parameters, we impute relative product qualities (ω_j), marginal costs (C_j), contributions to consumer surplus over total production costs, $\frac{\Delta U_j}{TC_j}$, producer surplus over value of the fixed capital stock, $\frac{\pi_j}{K_j}$,

---

of a similar VAR. Our results indicate that the variance of the random effect is sufficiently small to ignore, so we simply include initial values of the lagged dependent variables.

[27] The standard errors can be misleading because these are not symmetrically distributed random variables. For example, although a standard *t*-test would not reject the null hypothesis that α=0, with 90 percent confidence, the 90 percent confidence intervals for σ lie entirely in the positive domain.

external effects on the producer surplus of other plants over value of the fixed capital stock , $\frac{\Delta \sum_{k \neq j} \pi_k}{K_j}$, and net total surplus created over fixed capital stocks, $\frac{\Delta U_j + \pi_j}{K_j} + \frac{\Delta \sum_{k \neq j} \pi_k}{K_j}$. Then pooling all 11 years of observations on the 13 plants in continuous operation, we obtain the descriptive statistics in table 2. The results imply that the ratio of operating profits to fixed capital is roughly 7 percent, so the average rate of return on fixed capital investment is quite sensible. Domestic products are, on average, somewhat more attractive than the imported goods, but the cross-product standard deviation in relative appeal is substantial. (Robustness tests, not reported, show that it also depends on our $g_\sigma(\cdot)$ specification.) Most social surplus comes from the consumer side rather than the producer side (but see footnote 21). Indeed, the only reason that plants make a positive contribution to welfare is that this consumer surplus effect dominates the negative externality each plant imposes on the others by shrinking their market. (Some plants actually reduce net total welfare—this is a well known possibility in the case of monopolistic competition.)

We next calculated the cross-plant correlations in these variables reported in tables 3.[28] The implications are intriguing. First, relative product quality, ω_j , and marginal production costs, C_j , are weakly correlated ($\rho = 0.148$). Thus on average it costs relatively more to produce a relatively desirable good, and it is probably inappropriate to equate low production costs with superior performance.

[28] We also looked at correlations of firms' rankings in terms of each of these variables. The results are nearly identical to those reported in Table 3, so we do not report them here.

Second, the standard total factor productivity measure ($\tilde{\phi}_{jt}$) is weakly associated with product quality ($\rho =$ 0.113) because of the elasticity effect mentioned in sections II and III. So $\tilde{\phi}_{jt}$ *does* partly capture an aspect of performance that is directly related to welfare. However, product quality is not positively associated with total surplus ($\rho =$ -0.077), so firms that produce high quality goods do not outperform others in terms of their contribution to social welfare. The reason is that high quality goods are more costly to produce, on average.

Third, $\tilde{\phi}_{jt}$ is negatively correlated with marginal costs ($\rho =$-0.109). This association is not strong, but it is robust with respect to the priors we have explored (not reported). Again, this association reflects elasticity effects. Marginal costs vary for two reasons—factor prices and technical efficiency. We cannot sort them out without plant-specific information on factor prices, but if the dominant source of variation is technical efficiency, the tendency of $\tilde{\phi}_{jt}$ to capture ϕ_{jt} is very weak.

Fourth, $\tilde{\phi}_{jt}$ is fairly strongly associated with own-producer surplus, $\frac{\pi_j}{K_j}$ ($\rho =$0.536), and weakly negatively associated with the producer surplus of competing plants ($\rho =$-0.091) . This partly reflects underlying variation in marginal costs because, as just discussed, high marginal costs reduce $\tilde{\phi}_{jt}$ through an elasticity effect on mark-ups and they depress own profits by reducing own market shares (refer to section III.A.3), thereby helping others. (The correlation between own cost shocks and profit rates at competing plants is $\rho =$0.327.) The association between $\tilde{\phi}_{jt}$ and $\frac{\pi_j}{K_j}$ is strengthened by cross-plant variation in product quality, which creates analogous elasticity, mark-up, and market share effects of the opposite sign.

Finally, and most fundamentally, total surplus created over own fixed capital is almost orthogonal to $\tilde{\phi}_{jt}$ $(\rho = 0.038)$. This correlation is the only one that matters if we are exclusively concerned with contributions to social welfare. Taken at face value, it implies that traditional Tornqvist indices—and, we suspect, the entire class of indices discussed in section II—tell us very little about which firms are do well from a social perspective.

D. Linking performance to policy

It is popular to regress performance measures like $\tilde{\phi}_{jt}$ on policy variables or plant characteristics that are considered to respond to policy. For example, variants of $\tilde{\phi}_{jt}$ have often been regressed on measures of exposure to foreign technology, including foreign direct investment in the firm or its industry, and indicators for whether the firm is an exporter. As a final exercise, we demonstrate an alternative exercise using the welfare-based performance measures described in the previous section.

Specifically, we use the estimates in table 1 to quantify the effects of prohibiting firms from becoming exporters and/or importing intermediate goods. It would be straightforward to also prohibit consumers from importing foreign substitutes, but we will not do so in order to focus on these two production-side trade restrictions. Also, for the same reason, we will assume that total domestic demand evolves exactly as it would have in the absence of our policy shock, and that each producer draws the same VAR shocks $(\varepsilon_{jt}^{c}, \varepsilon_{jt}^{\omega})$ that were actually observed.

Under these assumptions we can use our VAR parameters to calculate the paths for (ω_{jt}, c_{jt}) that would have emerged if, beginning in 1982, all international producer trade had been shut down. The cross-plant temporal averages for these variables are graphed in Figure 1. As one could have predicted, since the use of imported intermediate imports reduces marginal costs (Table 1), our hypothetical policy regime results in marginal cost increases. It also results in slight quality increases as firms substitute toward domestic sources. The latter seems counter-intuitive, but it follows from our finding that high quality is weakly associated with high cost.

Substituting our parameter estimates and these counterfactual trajectories for (ω_{jt}, c_{jt}) into equations (8), (9) and (12), we next re-solve for equilibrium each period and calculate the new trajectories for producer and consumer surplus. These are graphed in figure 2. On net, not much happens to producer surplus because prices and costs move in the same direction. (This helps explain why some producers use imported inputs and others do not.) More surprisingly, not much happens to consumer surplus either. The reason is that the higher prices of domestic varieties due to higher marginal costs are accompanied by slight increases in quality, and our demand system estimates imply that consumers care a great deal about quality. We caution that this result appears to depend upon our priors. Other priors led to posterior parameter distributions (not reported) that implied a 10 percent loss in consumer surplus when producers were prohibited from foreign trade.

A very different story would have emerged if we had relied on $\tilde{\phi}_{jt}$-type measures for policy analysis. Fitting an AR(1) like those in table 1 to the Tornqvist index discussed earlier, we find that firms that imported their intermediate goods had significantly *lower* measured productivity.[29] Thus cutting firms off from foreign trade would appear to significantly *improve* performance. This may be an example of the positive link between $\tilde{\phi}_{jt}$ and factor prices that we discussed in section II above.

V. Concluding Remarks

The analysis we have presented here is crude in many ways. We have used a very simple demand system, we have assumed that marginal costs are flat with respect to output, we have ignored producers that were not present for the entire sample period, and we have ruled out any form of forward looking behavior—due either to dynamic pricing games or to capital accumulation. Finally, we have paid no attention to the institutional and technological features of the Colombian pulp and paper industry.

For all of these reasons, we do not wish to argue that the numbers we have presented here are the best that one can do. Rather, our objectives have been to argue that much of the literature on plant-level performance is fundamentally flawed, and to sketch an alternative approach to inference that we feel holds more promise. Significant

[29] We regressed $\tilde{\phi}_{jt}$ on $\tilde{\phi}_{jt-1}$ and the same weakly exogenous variables that appear in the VAR specifications of table 1, making the same correction for unobserved heterogeneity. We obtained a coefficient of –0.181 (standard error 0.063) on our dummy for use of imported intermediates without exporting, and a coefficient of -0.220 (standard error 0.066) on our dummy for use of imported intermediates while simultaneously exporting.

refinements in most of the dimensions mentioned above are possible; we are optimistic that they will enhance the usefulness of our methodology.

Bibliography

Ackerberg, Daniel and Mark Rysman (2002) "Unobserved Product Differentiation in Discrete Choice Models: Estimating Price Elasticities and Welfare Effects," NBER Working Paper No. 8798.

Aitken, Brian, and Ann Harrison (1999) "Do Domestic Firms Benefit from Foreign Direct Investment? Evidence from Panel Data," *American Economic Review*, pp. 605-618

Anderson, Simon, Andre de Palma and Jacques-Francois Thisse (1992) *Discrete Choice Theory of Product Differentiation*. Cambridge: MIT Press.

Aw, Bee-Yan, Xiaomin Chen, and Mark Roberts (2001) "Firm-Level Evidence on Productivity Differentials and Turnover in Taiwanese Manufacturing," *Journal of Development Economics*, pp. 51-86.

Bahk, Byong and Michael Gort (1993) "Decomposing Learning by Doing in New Plants," *Journal of Political Economy*, 101(4), pp. 561-583.

Baily, Martin, Charles Hulten and David Campbell (1992) "Productivity Dynamics in Manufacturing Plants," *Brookings Papers on Economic Activity*: *Microeconomics*, pp. 187-249.

Baldwin, John R. and Paul K. Gorecki (1991) "Entry, Exit, and Productivity Growth," in P. Geroski and J. Schwalbach, eds. *Entry and Market Contestability: An International Comparison.* Oxford and Cambridge: Blackwell, pp. 244-56.

Bernard, Andrew, Jonathan Eaton, J. Bradford Jensen and Samuel Kortum (2000) "Plants and Productivity in International Trade," NBER Working Paper 7688.

Berry, Steven. (1994) "Estimating Discrete-Choice Models of Product Differentiation," *Rand Journal* 25(2), pp. 242-262.

Berry, Steven, James Levinsohn and Ariel Pakes (1995) "Automobile Prices in Market Equilibrium," *Econometrica* 63(4), pp. 841-890.

Berry, Steven, James Levinsohn and Ariel Pakes (1999) "Voluntary Export Restraints in Automobiles," *American Economic Review* 89(3), pp. 400-430.

Blomstrom, Magnus and Ari Kokko (1997). "How Foreign Investment Affects Host Countries," World Bank PRD Working Paper No. WPS 1745.

Caplin, A. and Nalebuff, B. (1991) "Aggregation and Imperfect Competition: On the Existence of Equilibrium," *Econometrica* 59, pp. 25-60.

Cardell, Scott (1997) “Variance Components Structures for the Extreme Value and Logistic Distributions with Application to Models of Heterogeneity,” *Econometric Theory* 13, pp. 185-213.

Caves, Richard, Lau Christiansen and W. Erwin Diewert (1982) “Output, Input and Productivity Using Superlative Index Numbers,” *Economic Journal* 92, pp. 73-96.

Caves, Richard and David Barton (1990) *Efficiency in U.S. Manufacturing Industries*. Cambridge: MIT Press.

Gilks, W. R., S. Richardson and D. J. Spiegelhalter (1996) “Introducing Markov Chain Monte Carlo,” in Gilks, W. R., S. Richardson and D. J. Spiegelhalter, eds., *Markov Chain Monte Carlo in Practice*. London: Chapman and Hall/CRC.

Good, David H., M. Ishaq Nadiri and Robin Sickles. (1997). “Index Number and Factor Demand Approaches to the Estimation of Productivity,” in H. Pesaran and P. Schmidt, *Handbook of Applied Econometrics, Volume 2: Microeconometrics*. Basil-Blackwell.

Griliches, Zvi. (1986). “Productivity, R&D and Basic Research at the Firm Level in the 19870s,” *American Economic Review* 76(1), pp. 141-154.

Jorgenson, Dale and Zvi Griliches (1967). “The Explanation of Productivity Change,” *Review of Economic Studies* 34(3), pp. 249-83.

Klette, Tor and Zvi Griliches (1996) “The Inconsistency of Common Scale Estimators When Output Prices Are Unobserved and Endogenous,” *Journal of Applied Econometrics*, 11(4), pp. 343-61.

Klette, Tor and Arvid Raknerud (2001) “How and Why do Firms Differ?” processed, Department of Economics, University of Oslo.

Kraay, Aart, Isidro Soloaga and James Tybout (2001) “Product Quality, Productive Efficiency, and International Technology Diffusion: Evidence from Plant-Level Panel Data,” World Bank Policy Research Working Paper No. 2759.

Lahiri, Kajal and Jian Gao (2001) “Bayesian Analysis of Nested Logit Model by Markov Chain Monte Carlo,” processed, Department of Economics, University at Albany-SUNY.

Levinsohn, James and Amil Petrin (2001) "Estimating Production Functions using Inputs to Control for Unobservables," processed, The University of Michigan.

Liu, Lili and Tybout, James (1996) “Productivity Growth in Chile and Colombia: The Role of Entry, Exit and Learning,” in Mark Roberts and James Tybout, eds., *Industrial Evolution in Developing Countries*. Oxford: Oxford University Press.

Marshak, J. and W. H. Andrews (1944). “Random Simultaneous Equations and the Theory of Production,” *Econometrica*, 50, pp. 649-670.

Mairesse, Jacques and Mohamed Sassenou (1991). “R&D and Productivity: A Survey of Econometric Studies at the Firm Level.” *STI Review* (OECD, Paris), 1991, pp. 9-43.

McFadden, Daniel (1974) “Conditional Logit Analysis of Qualitative Choice Behavior,” in Paul Zarembka, ed., *Frontiers in Econometrics*. New York: Academic Press, 1974.

Melitz, Marc (2000) “Estimating Productivity in Differentiated Product Industries,” Department of Economics, Harvard University.

Ocampo, Jose Antonio and Leonardo Villar (1995) “Colombian Manufactured Exportes, 1967-91,” in Gerald Helleiner, editor, *Manufacturing for Export in the Developing World: Problems and Possibilities.* New York, N.Y.: Routledge.

Olley, Steven and Ariel Pakes (1996) “The Dynamics of Productivity in the Telecommunications Equipment Industry,” *Econometrica*, 64(6), pp. 1263-1297.

Pavcnik, Nina (2002) “Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants,” *The Review of Economic Studies* 69: 245-276.

Poirier, Dale (1996) “A Bayesian Analysis of Nested Logit Models,” *Journal of Econometrics* 75, pp. 163-181.

Roberts, Mark (1996) “Colombia, 1977-85: Producer Turnover, Margins, and Trade Exposure,” in Mark Roberts and James Tybout, editors, *Industrial Evolution in Developing Countries*. New York: Oxford University Press.

Tybout, James, Jaime de Melo and Vittorio Corbo (1991) “The Effects of Trade Reforms on Scale and Technical Efficiency: New Evidence from Chile,” *Journal of International Economics* 31, pp. 231-250.

Tybout, James (2000) “Manufacturing Firms in Developing Countries: How Well do They do, and Why?” *Journal of Economic literature* 38(1), pp. 11-44.

Tybout, James (forthcoming) “Plant- and Firm-Level Evidence on the ‘New’ Trade Theories,” in James Harrigan (ed.) *Handbook of International Trade*, Oxford: Basil-Blackwell.

Tybout, James and M. Daniel Westbrook (1995). "Trade Liberalization and Dimensions of Efficiency Change in Mexican Manufacturing Industries," *Journal of International Economics* 39, pp. 53-78.

Wooldridge, Jeffrey (2000) “The Initial Conditions Problem in Dynamic, Nonlinear Panel Data Models with Unobserved Heterogeneity,” unpublished working paper, Department of Economics, Michigan State University.

Zellner, Arnold (1971) *An Introduction to Bayesian Inference in Econometrics*. New York: Wiley.

Appendix 1: Inferring Qualities and Quantities from Revenues and Costs

This appendix demonstrates that in each period, given the quantity of imports, Q_{0t}, there is a one-to-one mapping from (TC_t, R_t) to (Q_t, C_t). Using equation (11) of the text, it then follows immediately that there exists a one-to-one mapping from (TC_t, R_t) to (ω_t, C_t), given the price of imported goods, e_t.

We shall hereafter drop *t* subscripts to reduce clutter. Note first that our assumption of flat marginal cost schedules implies that total variable costs at the j^{th} plant are $TC_j = Q_j C_j$. Thus the within-group market share of the j^{th} firm is:

$S_{j|g_j} = \frac{TC_j}{C_j} \Big/ Q_{g_j}^{tot}$ where total output from the j^{th} firm's group is

$Q_{g_j}^{tot} = \sum_{k \ni g_k = g_j} \frac{TC_k}{C_k}$ and total domestic output is $Q^{tot} = \sum_{g=1,G} Q_g^{tot}$. Note also that total revenues at the j^{th} plant are $R_j = P_j Q_j$, so the j^{th} firm's price-cost markup may be expressed in terms of observable variables as $m_j = \frac{R_j}{TC_j} - 1$, and if marginal cost is known, price can be calculated as $P_j = (m_j + 1)C_j$.

Substituting these market share and price expressions into the pricing rule (15) and solving for marginal cost, we obtain:

$$C_j = \sigma \cdot \left(\frac{TC_j}{Q_{g_j}^{tot}} \right) + (1-\sigma) \cdot \left(\frac{TC_j}{Q^{tot} + Q_0} + \frac{1}{\alpha \cdot m_j} \right). \qquad \text{(A1.1)}$$

This expression defines the unobservable C_j as a monotonic decreasing function of Q_{g_j}, given data on TC_j, m_j, Q^{tot} and Q_0. Thus, once the nest subtotals are known, each firm's marginal costs are implied by (A1.1). With these marginal costs, prices can be retrieved from $P_j = (m_j + 1)C_j$. In turn, these imply quantities $Q_j = R_j / P_j$, and market shares follow immediately. Substituting prices and market shares into (11) yields the vector of product qualities, ω.

To solve for the nest quantity subtotals, note that $Q_{g_j}^{tot} = \sum_{k \ni g_k = g_j} \frac{TC_k}{C_k}$.

Substituting the marginal cost expression (A1.1) into this sum, and dividing both sides by $Q_{g_j}^{tot}$, one obtains:

$$1 = \sum_{k \ni g_k = g} \left[\frac{1}{\sigma + (1-\sigma) \cdot Q_g^{tot} \cdot \left(\left(Q^{tot} + Q_0 \right)^{-1} + \left(\alpha \cdot TC_k \cdot m_k \right)^{-1} \right)} \right], \qquad \text{(A1.2)}$$

$$g = 1, \ldots G.$$

The right-hand side of (A1.2) is a monotonic negative function of Q_g^{tot} with value $n_g \sigma^{-1} > 1$ at $Q_g^{tot} = 0$ and limit 0 as $Q_g^{tot} \to \infty$, where n_g is the number of producers in nest *g*. Thus, for all $g \in \{1, \ldots G\}$, $Q_0 > 0$, and $Q^{tot} \geq 0$, equation (A1.2) has a unique, positive root: $Q_g^{tot} = f_g(Q^{tot} \mid Q^0)$.

Finally, we will show that $Q^{tot} - \sum_{g=1,G} f_g(Q^{tot} \mid Q_0) = 0$ has a unique positive root for any given $Q_0 > 0$. The existence of at least one root follows from the fact that

$\sum_{g=1,G} f_g(Q^{tot} \mid Q_0)$ is continuous in Q^{tot} , $\lim_{Q^{tot} \to 0}\left[Q^{tot} - \sum_{g=1,G} f_g(Q^{tot} \mid Q_0)\right] < 0$ and

$\lim_{Q^{tot} \to \infty}\left[Q^{tot} - \sum_{g=1,G} f_g(Q^{tot} \mid Q_0)\right] > 0$. Uniqueness follows from the fact that the

right-hand side of:

$$\frac{Q^{tot}}{Q^{tot} + Q^0} = \sum_{g=1,G} \frac{f_g(Q^{tot} \mid Q^0)}{Q^{tot} + Q^0} = \sum_{g=1,G} s_g(Q^{tot} \mid Q^0) \tag{A1.3}$$

is a continuous decreasing function of Q^{tot}. This can be seen by restating (A1.2) as:

$$1 = \sum_{k \ni g_k = g}\left[\frac{1}{\sigma + (1-\sigma)\cdot s_g \cdot \left(1 + \left(Q^{tot} + Q_0\right)\left(\alpha \cdot TC_k \cdot m_k\right)^{-1}\right)}\right],\ g = 1,\ldots G,$$ which implies

that s_g falls with Q^{tot}, $g = 1,\ldots G$.

Appendix 2: The Gibbs Sampler

Because it is not feasible to sample independent draws from the density $P(\theta \mid D) \propto f(\theta) \cdot L(D \mid \theta)$, we use Markov chain Monte Carlo (MCMC) techniques. The idea is to draw a sequence of realizations on θ from some Markov process, $\{\theta^{(1)}, \theta^{(2)}, \cdots, \theta^{(i)}\}$, with elements whose unconditional distributions converge to $P(\theta \mid D)$ as $i \to \infty$. After discarding the early draws to eliminate the effects of the starting values, one can approximate the posterior moments of θ by constructing their sample counterparts from the chain.

The mostly commonly used MCMC algorithm is the Gibbs sampler. It generates a Markov chain by breaking the parameter vector into sub-vectors with full conditional distributions that *can* be sampled from, then using these conditional distributions to update the sub-vectors sequentially (Gilks, et al, 1996). We exploit Gibbs sampling techniques by breaking θ into 3 sub-vectors: $\theta_1 = (\alpha, \gamma, \sigma)$, $\theta_2 = (\lambda, \phi)$, and $\theta_3 = vec(\Sigma)$. These we update according to the following algorithm:

Step 0: Set the initial values $\theta^{(0)} = (\theta_1^{(0)}, \theta_2^{(0)}, \theta_3^{(0)})$, and $i = 0$.

Step 1: Draw $\theta^{(i+1)}$ as follows:

a) Draw $\theta_1^{(i+1)} \sim \pi_1(\theta_1 \mid \theta_2^{(i)}, \theta_3^{(i)}, D)$

b) Draw $\theta_2^{(i+1)} \sim \pi_2(\theta_2 \mid \theta_1^{(i+1)}, \theta_3^{(i)}, D)$

c) Draw $\theta_3^{(i+1)} \sim \pi_3(\theta_3 \mid \theta_1^{(i+1)}, \theta_2^{(i+1)}, D)$

Step 2: Set $i = i + 1$, and go to step 1.

The distribution $\pi_1(\theta_1 \mid \theta_2, \theta_3, D)$ is the most difficult to construct. It is proportional to $L(D \mid \theta_1, \theta_2, \theta_3) f_{\theta_1}(\theta_1)$ where $L(D \mid \theta_1, \theta_2, \theta_3)$ is the likelihood function based on (8), (9), (11), (12) and (13); and $f_{\theta_1}(\theta_1)$ is the prior distribution defined in the text. But $L(D \mid \theta_1, \theta_2, \theta_3) f_{\theta_1}(\theta_1)$ does not have a closed form expression, so we draw θ_1 using the random-walk Metropolis algorithm with a normal proposal density. The performance of the random-walk Metropolis algorithm depends crucially on the variance –covariance matrix of the proposal density. If the variance-covariance matrix is too big, then nearly all proposed moves will be accepted (high acceptance) but the random walk will move around the parameter space very slowly (slow mixing). On the other hand, if the variance-covariance matrix is too small, then an excessively large fraction of proposed moves will be rejected (low acceptance), although those draws that are accepted will move the chain by large increments. To balance these two effects, the convention is to choose the variance-covariance matrix in such a way that the empirical overall acceptance rate is around between 0.15 and 0.5. For more details, see Gilks, *et al*, (1996, chapter 7). We experimented until this condition was satisfied.

To describe $\pi_2(\theta_2 \mid \theta_1, \theta_3, D)$ and $\pi_3(\theta_3 \mid \theta_1, \theta_2, D)$, let us rewrite (13a) and (13b) as $Y_{jt} = B' Z_{jt} + \varepsilon_{jt}$ where $Y_{jt} = (\omega_{jt}, c_{jt})'$, $Z_{jt} = (1, \omega_{jt-1}, c_{jt-1}, X'_{jt})'$, $\varepsilon_{jt} = (\varepsilon^{\omega}_{jt}, \varepsilon^{c}_{jt})$, and $B' = \begin{pmatrix} \omega_0 & \lambda_\omega & \lambda_C & \lambda'_X \\ c_0 & \phi_\omega & \phi_C & \phi'_X \end{pmatrix}$. Also, stacking observations, let us define:

$$
\begin{aligned}
Y &= [Y_{12} \quad \cdots \quad Y_{1T} \quad \cdots \quad \cdots \quad Y_{N2} \quad \cdots \quad Y_{NT}]', \\
X &= [X'_{12} \quad \cdots \quad X'_{1T} \quad \cdots \quad \cdots \quad X'_{N2} \quad \cdots \quad X'_{NT}]' \\
Z &= [Z_{12} \quad \cdots \quad Z_{1T} \quad \cdots \quad \cdots \quad Z_{N2} \quad \cdots \quad Z_{NT}]', \\
U &= [U_{12} \quad \cdots \quad U_{1T} \quad \cdots \quad \cdots \quad U_{N2} \quad \cdots \quad U_{NT}]'
\end{aligned}
$$

Then, we can write the VAR system as $Y = ZB + U$. Further, conditional on θ_1, our one-to-one mapping from (TC_t , $R_t | Q_{0t}, r_t$) to (ω_t, C_t) allows us to infer (Y, Z) from (θ_1, D). Thus the construction of $\pi_2(\theta_2 | \theta_1, \theta_3, D)$ and $\pi_3(\theta_3 | \theta_1, \theta_2, D)$ is a standard exercise (Zellner, 1971).

Specifically, the likelihood-based full conditional distribution of θ_2, given (θ_1, θ_3, D), is normal with mean $\left((Z'Z)^{-1} Z' \otimes I_2\right) \cdot vec(Y')$ and variance $(Z'Z)^{-1} \otimes \Sigma$. The full conditional posterior distribution for θ_2 efficiently blends this information with our priors. We have assumed that θ_2 has prior distribution $N(u_0, V_0)$, so $\pi_2(\theta_2 | \theta_1, \theta_3, D)$ is multivariate normal with mean $u_n = V_n \left[\left(Z' \otimes \Sigma^{-1}\right) vec(Y') + V_0^{-1} u_0\right]$ and variance $V_n = \left[\left((Z'Z) \otimes \Sigma^{-1}\right) + V_0^{-1}\right]^{-1}$.

Similarly, using the mapping $(\theta_1, \theta_2, D) \to (B, Y, Z)$, we may write the likelihood-based full conditional estimator of Σ, given (θ_1, θ_2, D), as $\frac{1}{N(T-1)} \sum_{i=1}^{N} \sum_{t=2}^{T} (Y_{it} - Z_{it} B)(Y_{it} - Z_{it} B)'$. When multiplied by N(T-1), this estimator has a Wishart distribution with N(T-1) degrees of freedom. Thus, given that we have assumed Σ has prior distribution $InvWish\left(m_0, G_0^{-1}\right)$, the full conditional posterior distribution for $\theta_3 = vec(\Sigma)$, i.e., $\pi_3(\theta_3 | \theta_1, \theta_2, D)$, is the vector version of a $Inv\,Wish(m_n, G_n^{-1})$ distribution, where $m_n = m_0 + N(T-1)$ and $G_n^{-1} = G_0^{-1} + (Y - ZB)'(Y - ZB)$.

TABLE 1: POSTERIOR PARAMETER DISTRIBUTIONS

	Mean	Std. Error	Median	Skewness	5%	95%
			Demand System			
α (prior: $\alpha \sim U[0,10]$)	4.194	2.845	3.696	0.375	0.437	9.198
σ (prior: $\sigma \sim U[0,1]$)	0.970	0.023	0.975	-1.029	0.927	0.997
γ (prior: $\gamma \sim U[0,10]$)	1.086	0.779	0.997	0.463	0.070	2.423
			Product Quality VAR			
λ_1 (constant)	-0.436	3.186	-0.229	-0.199	-5.876	4.608
λ_2 (ω_{it-1})	0.469	0.376	0.468	0.008	-0.152	1.086
λ_3 (c_{it-1})	0.053	0.237	0.056	-0.032	-0.343	0.437
λ_4 (trend)	-0.021	0.046	-0.020	-0.219	-0.098	0.052
λ_5 (initial capital stock)	-0.009	0.101	-0.007	-0.133	-0.179	0.153
λ_6 (exported, t-1)	0.029	0.374	0.022	0.059	-0.578	0.648
λ_7 (imported intermediates, t-1)	-0.040	0.385	-0.041	0.030	-0.679	0.595
$\vartheta_1^{\omega}(\omega_{i1})$	0.788	1.373	0.781	0.051	-1.467	3.051
$\vartheta_2^{\omega}(c_{i1})$	-0.192	0.382	-0.157	-0.511	-0.897	0.353
			Log Marginal Cost VAR			
ϕ_1 (constant)	0.291	3.211	0.069	0.286	-4.684	5.896
ϕ_2 (c_{it-1})	0.551	0.244	0.554	-0.060	0.152	0.946
ϕ_3 (ω_{it-1})	0.314	0.383	0.311	0.020	-0.306	0.942
ϕ_4 (trend)	0.001	0.045	0.001	0.080	-0.071	0.076
ϕ_5 (initial capital stock)	-0.003	0.105	-0.004	0.062	-0.170	0.169
ϕ_6 (exported, t-1)	-0.146	0.387	-0.145	-0.042	-0.788	0.490
ϕ_7 (imported intermediates, t-1)	-0.098	0.403	-0.104	0.029	-0.764	0.557
$\vartheta_1^{C}(c_{i1})$	-0.799	1.382	0.255	0.435	-0.267	0.954
$\vartheta_1^{C}(\omega_{i1})$	0.286	0.384	-0.788	-0.064	-3.129	1.433
			Covariance Matrix			
Σ_{11}	0.943	0.127	0.931	0.532	0.755	1.168
Σ_{12}	0.010	0.093	0.010	0.009	-0.143	0.164
Σ_{22}	1.004	0.137	0.992	0.535	0.802	1.245

TABLE 2:

DESCRIPTIVE STATISTICS ON PERFORMANCE MEASURES (Colombian Pulp Mills)

	$\frac{\Delta U_j}{K_j}$	$\frac{\pi_j}{K_j}$	$\frac{\Delta\sum_{k\neq j}\pi_k}{K_j}$	$\frac{\Delta U_j+\pi_j}{K_j}+\frac{\Delta\sum_{k\neq j}\pi_k}{K_j}$	c_j	ω_j
Mean	0.065	0.071	-0.221	-0.085	-3.198	1.868
Median	0.004	0.067	-0.156	-0.057	-3.252	1.886
Std. Dev.	0.134	0.041	0.222	0.147	0.496	0.279
Skewness	12.516	2.082	4.752	7.266	-0.208	-0.026

TABLE 3:

CORRELATIONS OF PERFORMANCE MEASURES (Colombian Pulp Mills)

	$\frac{\Delta U_j}{K_j}$	$\frac{\pi_j}{K_j}$	$\frac{\Delta\sum_{k\neq j}\pi_k}{K_j}$	$\frac{\Delta U_j+\pi_j}{K_j}+\frac{\Delta\sum_{k\neq j}\pi_k}{K_j}$	c_j	ω_j	$\tilde{\phi}_j$
$\frac{\Delta U_j}{K_j}$	1.000	0.262	-0.713	-0.088	-0.225	0.235	0.028
$\frac{\pi_j}{K_j}$		1.000	-0.414	-0.104	-0.466	0.060	0.536
$\frac{\Delta\sum_{k\neq j}\pi_k}{K_j}$			1.000	0.738	0.327	-0.077	-0.091
$\frac{\Delta U_j+\pi_j}{K_j}+\frac{\Delta\sum_{k\neq j}\pi_k}{K_j}$				1.000	0.157	0.116	0.038
c_j					1.000	0.148	-0.109
ω_j						1.000	0.113
$\tilde{\phi}_j$							1.000

FIGURE 1:
EFFECTS OF PRODUCER TRADE ON QUALITY AND MARGINAL COST TRAJECTORIES

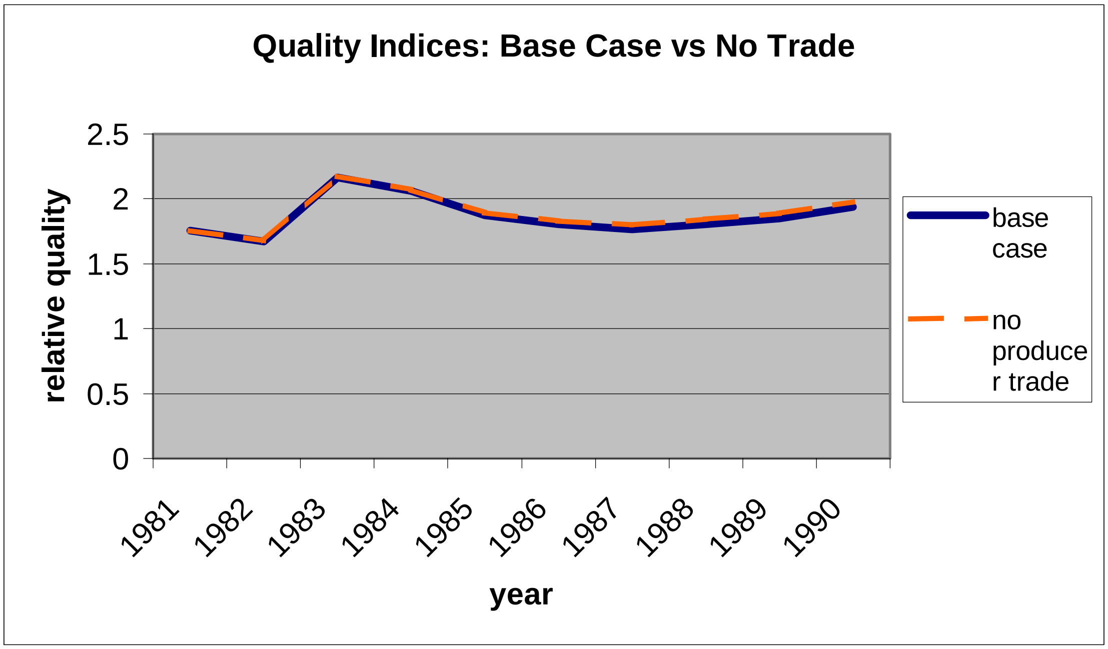

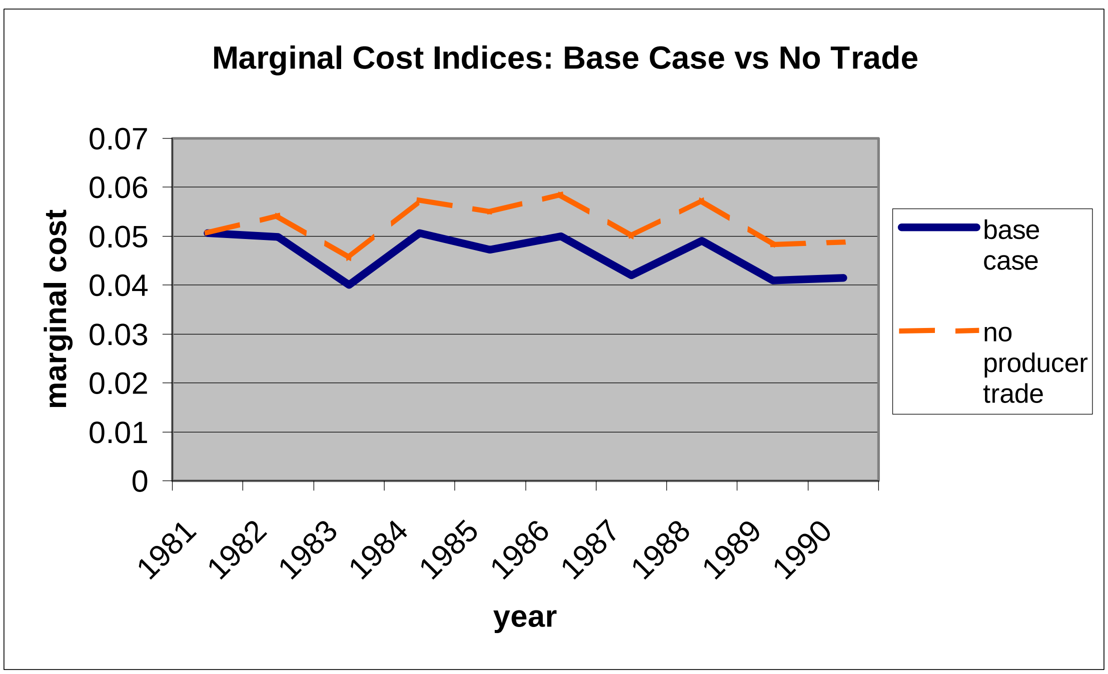

FIGURE 2: EFFECTS OF PRODUCER TRADE ON WELFARE MEASURES