

Firm-level Productivity Studies: Illusions and a Solution

Hajime Katayama
The University of Sydney

Shihua Lu
Bates, White Inc.

James R. Tybout
Pennsylvania State University and NBER

April 2005

We gratefully acknowledge the National Science Foundation for support (grant 421-18687W0). We have benefited from useful discussions with Tom Holmes, Marc Melitz, Nina Pavcnik, Mark Roberts, Marc Rysman, Sofronis Clerides, and Johannes Van Biesebroeck.

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 firm-level productivity patterns. But firm-level data on physical quantities of output, capital, and intermediate inputs are typically unobservable. 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 firms' costs and revenues reflect a Bertrand-Nash equilibrium in a differentiated product industry, as in Berry (1994). This allows us to impute each firm's unobserved marginal costs and product appeal from its observed revenues and costs. With these in hand, we calculate each firm'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 appeal 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 paper producers, we study the relation between our welfare-based measures and conventional productivity measures. We find that the two are only weakly correlated with one another. Further, they give contrasting pictures of the relationship between firms' performances and their participation in foreign markets. One reason is that product appeal variation has little effect on standard productivity indices, but it is captured by welfare-based performance measures.

JEL categories: L1, O3, L6

I. Overview

Economists often seek to quantify the effects of a policy or event on the performance of manufacturing firms. Recurrent questions include: How does trade liberalization affect productive 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? And how do entry regulations affect an industry's performance?

To address these issues, many researchers rely on plant- or firm-level productivity analysis. They posit that each establishment's output is a function of the inputs it employs and the productivity level it attains, hereafter indexed by ϕ . Then, using the available output and input measures, they estimate this function and solve for producer- and time-specific approximations to ϕ . Finally, looking across producers and/or through time, they correlate these approximations with things like the extent of foreign ownership, intensity of R&D activity, whether the firms are exporting, rates of effective protection for the firm's product, and whether entry and exit are institutionally constrained.¹

When output and input characteristics are common across plants, and when data on the physical quantities of these variables are available, the productivity approximations that are commonly used makes good sense. Indeed, most of the methodological literature on this approach to analyzing firm or plant-level performance

¹ 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 (2003), and Aw, Chen and Roberts (2000).

presumes that these conditions hold.² But in practice, producer-specific productivity measures are more commonly constructed for differentiated product and/or differentiated input industries, where the characteristics of products and factor inputs vary considerably across producers. Under these circumstances data on physical volumes are usually unavailable, so analysts are forced to 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 imperfect 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 functional relationship between inputs and outputs is precisely estimated, these measures 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.

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 and we incorporate a demand system

² Particular attention has been devoted to the issues of how to estimate the functional relationship and how to decompose production function residuals into "true" productivity and a noise component.

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.

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. It also allows us to construct product-specific measures of consumer and producer surplus, and to relate these measures to policies, events, or managerial decisions.

Our last objective is to demonstrate our methodology using plant-level data on Colombian paper mills, and to compare our welfare-based performance measures with standard measures (Section IV). We find that the two are only weakly correlated with one another. Further, they give contrasting pictures of the relationship between firms' performances and their participation in foreign markets.

II. The Problem with standard performance measures

Consider an industry populated by N single-product, single-plant producers, all of whom have production functions of the form:

$$Q_{jt} = e^{\phi_{jt}} \cdot h(F_{jt}), \quad j = 1, \dots, N. \quad (1)$$

Here Q_{jt} is the physical output of the j^{th} firm in period t , ϕ_{jt} is its “true” productivity level, $h(\cdot)$ is an increasing differentiable function, and F_{jt} is a scalar index of factor usage. More precisely, F_{jt} is a constant returns function of the vector of inputs employed by the firm, $F_{jt} = f(\vec{V}_{jt})$, where $\vec{V}_{jt} = \{V_{jt}^1, V_{jt}^2, \dots, V_{jt}^I\}$ and factors that differ in quality enter \vec{V}_{jt} as distinct inputs.

When Q_{jt} and F_{jt} are observable and the function $h(\cdot)$ is known, the productivity index is retrievable as $\phi_{jt} = \ln Q_{jt} - \ln h(F_{jt})$. But these conditions rarely

prevail, so analysts usually proceed with imperfect information. Specifically, when data on physical output volumes are unavailable, they typically replace Q_{jt} with

$\tilde{Q}_{jt} = R_{jt} / \bar{P}_t$, where R_{jt} is the nominal *value* of the j^{th} firm's output and \bar{P}_t is an

industry-wide output price index. Similarly, when input quantities are unobservable, the convention is to replace them with a deflated measure of expenditure on inputs,

$\tilde{F}_{jt} = \left(\frac{B_{jt}}{\bar{B}_t} \right) F_{jt}$, where \bar{B}_t is a sector-wide input price deflator and B_{jt} is the price of a

unit bundle of inputs for the j^{th} firm.⁴ Thus performance is commonly measured by indices of the form:

$$\tilde{\phi}_{jt} = (\ln R_{jt} - \ln \bar{P}_t) - \ln \hat{h}(\tilde{F}_{jt}), \quad j = 1, \dots, N \quad (2)$$

where $\hat{h}(\cdot)$ is an approximation to the function $h(\cdot)$.

What do performance measures based on equation (2) tell us? Sales revenues depend upon demand conditions and the nature of competition, so we cannot describe the properties of $\tilde{\phi}_{jt}$ without introducing additional assumptions about consumer and producer behavior. Let demand for the j^{th} firm's product be given by the differentiable function:

⁴Presuming cost minimization, $B_{jt} = \min_V \left[\sum_{i=1}^I W_{jt}^i V^i \mid 1 = f(\bar{V}) \right]$ where $\bar{W}_{jt} = \{W_{jt}^1, W_{jt}^2, \dots, W_{jt}^I\}$ is the plant-specific vector of unit factor prices associated with the input vector \bar{V}_{jt} . If some elements of the input vector are measured in physical terms and others are measured in expenditure terms, the expression for \tilde{F}_{jt} is more complicated. We will treat this case in a separate subsection below.

$$Q_{jt} = m^j(\bar{P}_t, \bar{\omega}_t, Y_t), \quad j = 1, \dots, N \quad (3)$$

where $\bar{\omega}_t = \{\omega_{1t}, \omega_{2t}, \dots, \omega_{Nt}\}'$ is a vector of quality or appeal indices for all N products, $\bar{P}_t = \{P_{1t}, P_{2t}, \dots, P_{Nt}\}'$ is the corresponding vector of product prices, and Y_t is an index of total market size.⁵ Further, assume that current prices and product quality/appeal indices are common knowledge, firms are price takers in factor markets, and they are pure Bertrand-Nash price setters in the product market.

Under these assumptions, producer j considers its marginal revenue product at input level F_{jt} to be $\left(1 - \frac{1}{\eta_{jt}}\right) \gamma_{jt} \frac{R_{jt}}{F_{jt}}$, where $\eta_{jt} = \frac{-\partial \ln m^j(\bar{P}_t, \bar{\omega}_t, Y_t)}{\partial \ln P_{jt}}$ is this firm's perceived elasticity of demand and $\gamma_{jt} = \frac{d \ln h(F_{jt})}{d \ln F_{jt}}$ is its returns to scale at the margin. Thus, equating the marginal revenue product of input bundles to their unit prices, producer by producer, we obtain a set of first-order conditions for profit maximization:

$$R_{jt} = \left(\frac{\eta_{jt}}{\eta_{jt} - 1} \right) \left(\frac{B_{jt}}{\gamma_{jt}} \right) F_{jt}, \quad j = 1, \dots, N. \quad (4)$$

⁵ Define $m^{j^{-1}}(P_j | \bar{P}^{-j}, \bar{\omega}, Y)$ to be the inverse demand function for the j^{th} firm, given the vector of prices for all other products, \bar{P}^{-j} , the complete vector of product qualities, and market size. Caplin and Nalebuff (1991) show that a pure Bertrand-Nash equilibrium exists if $m^{j^{-1}}(\cdot | \cdot)$ is convex and diminishing in P_j , so long as cost functions are convex. They also describe sufficient conditions on individual utility functions and the distribution of these utility functions across individuals for this property to obtain, and they demonstrate conditions for uniqueness. In particular, they show that the individual utility functions that underlie logit demand systems satisfy existence and uniqueness conditions.

The key message of equation (4) is that revenue per unit input bundle (R_{jt}/F_{jt}) increases with the mark-up factor $\eta_{jt}/(\eta_{jt} - 1)$ and with factor prices B_{jt} . Substituting (4) into (2), $\tilde{\phi}_{jt}$ may be written as:

$$\tilde{\phi}_{jt} = \ln\left(\frac{F_{jt}}{\hat{h}(\tilde{F}_{jt})}\right) + \ln\left(\frac{\eta_{jt}}{\gamma_{jt}(\eta_{jt} - 1)}\right) + \ln\left(\frac{B_{jt}}{\bar{P}_t}\right), \quad j = 1, \dots, N. \quad (5)$$

Hence, even if $\hat{h}(\cdot)$ properly characterizes returns to scale, measured productivity depends upon scale economies (through the first and second term), demand elasticities (through the second term), and deflated factor prices (through the third term).⁶

Pass-through effects

Does $\tilde{\phi}_{jt}$ nonetheless react to quality and productivity variation in the way that is commonly presumed? That is, does it take on relatively high values when ω_{jt} and/or ϕ_{jt} are relatively large? To simplify our answer to this question, we shall assume for the remainder of this section that analysts are somehow able to correctly estimate the returns-to-scale function $h(\cdot)$ and the input aggregating function $f(\cdot)$.⁷ Further, in order to begin with a simple case, we shall momentarily suppose that factor usage can be precisely

⁶ Hall (1988, 1991) also noted that standard productivity indices are biased when product market competition is imperfect. However, his analysis presumed that inputs and outputs are observable in physical terms, and focused on the problem of estimating the production technology $h(\cdot)$ when finite demand elasticities in product markets drive a wedge between the elasticities of output with respect to the input vector and the vector of factor shares.

⁷ Estimation errors introduce another type of problem with $\tilde{\phi}$ -type measures, but they do not undo the ones we will focus upon here. Klette and Griliches (1996) discuss the estimation issues that arise when revenue-based output measures are used in place of volume indices.

measured in physical terms (so that $h(\tilde{F}_{jt})$ becomes $h(F_{jt})$), all firms face the same demand elasticity (so that $\eta_{jt} = \eta \quad \forall j, t$), and all firms exhibit constant returns to scale (so that $\gamma_{jt} = 1 \quad \forall j, t$). Under these assumptions, $d \ln h(F_{jt}) = d \ln F_{jt}$ and, with an appropriate choice of units, $\ln\left(\frac{F_{jt}}{h(F_{jt})}\right) = 0$. Hence equation (5) reduces to:

$$\tilde{\phi}_{jt} = \ln\left(\frac{\eta}{\eta-1}\right) + \ln\left(\frac{B_{jt}}{\bar{P}_t}\right), \quad j = 1, \dots, N. \quad (5')$$

By equation (5'), $\tilde{\phi}_{jt}$ varies across firms with factor prices, and is completely *unrelated* to their productive efficiency or product quality. Factor prices matter because firms burdened with high factor costs pass some fraction of them on to consumers as higher output prices, driving up sales revenue per unit input bundle (the *pass-through effect*). In contrast, input-saving productivity shocks *don't* matter because firms respond to them with offsetting reductions in their output prices, leaving their revenues per unit input bundle unaffected.⁸ Similarly, product quality-enhancing technology shocks don't matter because they generate proportionate adjustments in revenues and input usage.⁹

⁸ To see the inverse relationship between output prices and productivity, divide both sides of equation (4) by Q_{jt} and use the constant-returns version of the production function (1), $Q_{jt} = F_{jt}e^{\phi_{jt}}$, to obtain

$\ln P_{jt} + \phi_{jt} = \ln\left(\frac{\eta_{jt}}{\gamma_{jt}(\eta_{jt}-1)}\right) + \ln B_{jt}$. The right-hand side of this expression is exogenous under constant returns and fixed demand elasticities, so shocks to ϕ_{jt} inducing offsetting shocks in $\ln P_{jt}$.

⁹ Shocks to efficiency and/or product quality can, however, affect the cross-firm average $\tilde{\phi}_{jt}$ value if they affect the output price index, \bar{P}_t , as will be discussed shortly.

Klette and Raknerud (2001) and Bernard *et al* (2003) make similar observations in slightly different contexts.

The dependence of $\tilde{\phi}_{jt}$ on factor prices may subvert productivity analysis in a number of ways. For example, the common finding that small and new firms are relatively unproductive may partly reflect the fact that these firms pay relatively low wages and provide few fringe benefits (Baily, et al 1992; Griliches and Ramey, 1996; Aw, Chen and Roberts, 2001).¹⁰ On the other hand, the finding that geographically clustered firms are relatively productive (Henderson, 2001), which is typically attributed to agglomeration economies, may simply reflect high wages and rental costs in urban areas. Likewise, the common tendency to find high $\tilde{\phi}$ indices among R&D-intensive firms (Mairesse and Sassenou, 1991) and among multinational firms (Blomstrom and Kokko, 1997; Aitken and Harrison, 1999) may trace partly to their high unit labor costs.¹¹

Time series variation in $\tilde{\phi}_{jt}$ is less likely to be spurious. As we argued above, firm-specific productivity shocks or quality/appeal shocks leave $\phi_{jt} + \ln P_{it}$ unaffected. But productivity shocks that are common across firms reduce \bar{P}_t , so average $\tilde{\phi}_{jt}$ values do respond to general improvements in efficiency. More precisely, by equations (1) and (2), measured productivity can be written as $\tilde{\phi}_{jt} = \ln\left(\frac{P_{jt}}{\bar{P}_t}\right) + \phi_{jt}$, so cross-plant averages

¹⁰ Evidence that this effect matters comes from Foster, Haltiwanger and Syverson (2003), who study cement manufacturers and therefore are able to measure physical volumes of inputs and outputs with precision. Their ϕ_{jt} estimates do *not* imply that new firms are less productive.

¹¹ Here we are assuming that their productivity measure is constructed using an index of physical labor rather than a measure of expenditures on labor. We will consider the case of expenditure-based labor measures shortly.

of $\tilde{\phi}_{jt}$ will coincide with cross-plant averages of ϕ_{jt} so long as \bar{P}_t is a simple geometric mean of the individual P_{jt} 's that occur in the sample. Even if \bar{P}_t is constructed some other way, cross-firm averages of $\tilde{\phi}_{jt}$ will covary through time with cross firm averages

of values ϕ_{jt} so long as \bar{P}_t indices covary with $\left(\prod_{j=1}^{N_t} P_{jt} \right)^{1/N_t}$.

Of course, output prices generally grow at different rates for different firms, so cross-firm comparisons of growth rates in $\tilde{\phi}$ -type indices can still be misleading. For example, in open economies, real exchange rate appreciation tends to drive down relative output prices among tradeable goods producers. But it also tends to increase import penetration rates, so $\tilde{\phi}$ -type indices may falsely create the impression that import competition hurts productivity among tradeable goods producers. Also, if firms that become exporters or are acquired by multinationals switch to higher quality workers, they are likely to pass some of the associated labor costs on to their consumers.¹² Hence $\tilde{\phi}$ -type indices may falsely imply that these firms are relatively efficient.

Unobserved factor heterogeneity

Thus far we have assumed that each factor is homogeneous across plants, and we have treated inputs with different characteristics—for example, different types of workers—as distinct elements of the V vector. Data are never actually available in sufficient detail to do this, so it is natural to ask how the properties of $\tilde{\phi}$ are affected by

¹² Many empirical studies suggest that this type of adjustment in employment occurs. Girma and Gorg (2004) review the recent literature on multinationals and provide some new evidence. In a study that has been replicated for other countries, Bernard and Jensen (1999) document the relatively high wages paid by exporters in the United States.

unobserved heterogeneity within a particular factor category. The answer is that in cross section, $\tilde{\phi}$ is likely to be positively correlated with ϕ and/or ω because factor prices reflect factor productivity.¹³ But this source of variation in $\tilde{\phi}$ reveals nothing about which firms are doing well in an economic sense. It simply means that firms using high quality inputs get more and/or better output.

Scale effects

Relaxing the assumption of constant returns to scale introduces another possible source of co-variation between $\tilde{\phi}$ and ϕ or ω . Suppose the production function is homogeneous of degree $\gamma > 1$, and the function $h(\cdot)$ correctly captures these scale economies. Then (5) becomes

$$\tilde{\phi}_{jt} = (1 - \gamma) \ln F_{jt} + \ln \left(\frac{B_{jt}}{\bar{P}_t} \right) + \ln \left(\frac{\eta}{\gamma(\eta - 1)} \right), \quad j = 1, \dots, N, \quad (5'')$$

and other things equal, larger firms appear to be relatively unproductive because they charge lower prices (the *scale effect*). More generally, if small firms face increasing returns, but these dissipate beyond some minimum efficient scale, the relation between

size and $\tilde{\phi}$ can be non-monotonic: $\ln \left(\frac{F_{jt}}{h(F_{jt})} \right)$ will fall then rise with F_{jt} , while

$\ln \left(\frac{\eta}{\gamma_{jt}(\eta - 1)} \right)$ will rise with F_{jt} throughout.

¹³ This dependence of measured productivity on unobservable aspects of factor quality is well known (e.g., Griliches and Jorgenson, 1967).

Mark-up effects

Although the assumption of common demand elasticities is often invoked, it is unrealistic in many contexts. Suppose instead that firms with high efficiency (ϕ_{jt}) and/or high quality products (ω_{jt}) face relatively low demand elasticities (η_{jt}) because they enjoy relatively large market shares.¹⁴ Then these “good” firms will charge higher mark-ups and, unless this effect is offset by stronger scale effects, they will appear relatively productive (the *mark-up effect*). Bernard *et al* (2003) demonstrate a similar linkage between firm quality and mark-ups in a contestable markets setting.

Although mark-up effects tend to induce positive covariance between “true” and “measured” productivity, they are also likely to induce spurious variation in $\tilde{\phi}$. For example, producers of close substitutes may look relatively inefficient because their demands are relatively elastic. Also, reductions in institutional barriers to entry may reduce the market power of incumbent firms (e.g., Pakes and McGuire, 1994), making them appear less productive. Similarly, when trade liberalization or exchange rate appreciation reduce the mark-ups 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.¹⁵

Finally, in addition to the pass-through effect captured by the second term in equation 5', factor price shocks induce mark-up effects through the third term. This is

¹⁴ Cournot competition in a homogeneous product industry yields this result, as does Bertrand competition with differentiated products and a Dixit-Stiglitz (constant elasticity of substitution) demand system, when N is small. Berry's (1994) characterization of market equilibrium exhibits the same property.

¹⁵ This dispersion effect is one interpretation for the findings of Caves and Barton (1990).

because positive shocks to input prices have the same effect on marginal costs as negative productivity shocks. For example, wage increases drive up firms' marginal costs, inducing them to reduce their mark-up and to appear less productive. And exchange-rate depreciation has similar effects on firms that import their intermediate inputs.

Measuring some factor in expenditure terms

Our discussion to this point presumes that all factors inputs can be measured in physical terms, but this is hardly realistic. With the possible exception of firms that process primary goods (e.g., dairies, grain mills, petroleum refineries, cement plants), data on intermediate inputs for manufacturers are almost always expressed in expenditure terms. It is even rarer to find physical capital measured in terms of numbers of machines of each type and vintage. What are the properties of $\tilde{\phi}$ when input expenditures are used as proxies for input usage?

Let some subset $E \subseteq \{1, \dots, I\}$ of inputs be measured in deflated expenditure terms,

so that the measured input vector \tilde{V}_{jt} has components $\tilde{V}_{jt}^i = \left(\frac{W_{jt}^i}{\bar{W}_t^i} \right) V_{jt}^i$ for $i \in E$ and

$\tilde{V}_{jt}^i = V_{jt}^i$ for $i \notin E$. Further, let the input aggregator function be Cobb-Douglas:

$F_{jt} = f(\tilde{V}_{jt}) = \prod_{i=1}^I (V_{jt}^i)^{\mu_i}$. Then measured productivity may be written as:

$$\tilde{\phi}_{jt} = (1 - \gamma) \left[\ln F_{jt} + \mu^E \ln \left(\frac{B_{jt}^E}{\bar{B}_t^E} \right) \right] + \ln \left(\frac{(\bar{B}_t^E)^{\mu^E} (B_{jt}^{NE})^{1 - \mu^E}}{\bar{P}_t} \right) + \ln \left(\frac{\eta_{jt}}{\gamma(\eta_{jt} - 1)} \right), \quad (5''')$$

where $\mu^E = \sum_{i \in E} \mu_i$, $B_{jt}^E \propto \left(\prod_{i \in E} (W_{jt}^i)^{\mu_i} \right)^{1/\mu^E}$, $B_{jt}^{NE} \propto \left(\prod_{i \notin E} (W_{jt}^i)^{\mu_i} \right)^{1/(1-\mu^E)}$, and

$$\bar{B}_t^E \propto \left(\prod_{i \in E} (\bar{W}_t^i)^{\mu_i} \right)^{1/\mu^E}.$$

Equation (5''') shows that increases in the prices of the E factors do not generate a pass-through effect on $\tilde{\phi}$. That is, B_{jt}^E does not appear in the second term of equation 5'''. The reason is that idiosyncratic shocks to the prices of E factors induce increases in measured input usage so, if demand elasticities are fixed and constant returns prevail, revenues and costs adjust in proportion to one another, leaving measured productivity unaffected. Of course, output and mark-ups are likely to fall with marginal cost increases, as discussed in the previous subsection. So regardless of whether factors are measured in physical or expenditure terms, shocks to their prices should create scale effects and mark-up effects on $\tilde{\phi}$.

But don't we get sensible stories from $\tilde{\phi}$?

Given all of the sources of spurious variation in $\tilde{\phi}$ -type performance measures, it is unsurprising that robust findings on their relation to policy proxies can be elusive.¹⁶ Nonetheless, some stylized facts concerning these measures suggest that they do capture some key features of firms' performance. In particular, many studies have found that firms with high $\tilde{\phi}$ values are more likely to be large or grow, and they are less likely to

¹⁶ For example, Tybout (2003) describes the wide range of findings on productivity and various trade policy proxies.

fail (Baily, Hulten and Campell, 1992; Olley and Pakes, 1996; Aw, Chen and Roberts, 2001; Baldwin and Gorecki, 1991; Liu and Tybout, 1996; Pavcnik, 2002). Doesn't this literature imply, as the authors of these studies suggest, that high- $\tilde{\phi}$ firms are more efficient and/or produce relatively appealing products? 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 $\tilde{\phi}$ values, even if their productive efficiency and product appeal are unexceptional.

III. An Alternative Approach to Measuring Performance

Thus far we have argued that, when analyzing differentiated-product industries, it is unwise to pretend that revenues and input expenditures measure physical inputs and outputs. How, then, is one to infer something about plants' performances when their product characteristics, physical output volumes (Q_{jt}) and prices (P_{jt}) are unobservable? We are aware of two partial solutions in the literature. The first, due to Bernard *et al* (2003), is to assume that each product can be produced by multiple firms. Then, further assuming that firms engage in pure Bertrand competition (limit pricing) and that all products share the same elasticity of demand, relatively productive firms exhibit relatively large mark-ups. It follows that revenue per unit input bundle can be used as a proxy for physical productivity. This approach constitutes an elegant resuscitation of traditional productivity measures, but it hinges on the presumptions that (1) in each product category, at least two firms stand ready to supply perfect substitutes; (2) demand

elasticities do not vary across products; and (3) input vectors (\bar{V}_{jt}) are measurable in physical units.

The second approach is due to Melitz (2000). He also assumes that the input vector is observable in physical units and demand elasticities are the same for all firms. But unlike Bernard et al (2003), but he treats each firm's product as distinct. Then, building on Klette and Griliches (1996), he shows that the residuals from a revenue function can be used to make inferences about $\phi_{jt} + \omega_{jt}$.

Our approach, like Melitz's, treats each firm's product as distinct. However we relax the assumptions that all firms face the same demand elasticity and that the input vector is observable in physical units. Further, we show how to measure ϕ_{jt} and ω_{jt} separately for each producer, and we propose an alternative approach to evaluating firms' performances based on these indices.

The logic of our approach is straightforward. Suppose that the vector of total costs $T\bar{C}_t = [TC_{1t}, TC_{2t}, \dots, TC_{Nt}]'$ and the vector of total revenues $\bar{R}_t = [R_{1t}, R_{2t}, \dots, R_{Nt}]$ can be observed for all N producers in a differentiated product industry. Then, so long as these data reflect product market equilibrium, they can be used in conjunction with the demand functions (3) and the first-order conditions for profit maximization (4) to impute the unobservable vectors of marginal costs $\bar{C}_t = [C_{1t}, C_{2t}, \dots, C_{Nt}]$ and product appeal indices ($\bar{\omega}_t$), as well as product prices (\bar{P}_t) and quantities (\bar{Q}_t). Further, once this mapping is established, it can be used to calculate plant-specific consumer and producer

welfare measures, to study the evolution of these welfare measures over time, and to relate them to policy shocks and managerial choices.¹⁷

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

Several conditions must be satisfied in order to implement this strategy. First, one's assumptions concerning consumer and producer behavior must imply a unique set of equilibrium prices and quantities, (\bar{P}_t, \bar{Q}_t) , at each $(\bar{\omega}_t, \bar{C}_t)$, given observable control variables. (Sufficient conditions are described in Caplin and Nalebuff, 1991.) Second, each possible pair of output and marginal cost values, (Q_{jt}, C_{jt}) , must map onto a unique total cost, TC_{jt} . Thus it is generally necessary to impose some structure on the marginal cost function.

These requirements rule out non-parametric approaches and some flexible functional forms, but it is not difficult to find a reasonable set of assumptions that suits our purposes. In the remainder of this section we describe one approach to inference that seems to work well. It is based on adopt Berry's (1994) representation of market equilibrium with McFadden's (1974) nested-logit demand functions.¹⁸

Suppressing t subscripts for now, suppose that each domestic producer $j \in \{1, \dots, N\}$ is associated with a single, distinct product variety. Also, let imported goods constitute one additional composite variety (indexed by $j = 0$) available to domestic

¹⁷ Without data on factor prices it is impossible to impute productivity measures, $\phi_t = \{\phi_{1t}, \phi_{2t}, \dots, \phi_{Nt}\}$, from observable variables. But these are relevant for welfare only inasmuch as they influence marginal production costs, which *are* identified.

¹⁸ A simple logit demand system would also work; we use the nested logit for added generality.

consumers. Finally, assume that product varieties can be grouped into $G+1 < N+1$ nests, with products in different nests presumed to be poorer substitutes for one another than products within the same nest.¹⁹ Then, invoking further assumptions about the distribution of tastes and the form of consumers' indirect utility functions, it is possible to write demand for the j^{th} product as (McFadden, 1974):

$$Q_j = m^j(\bar{P}, \bar{\omega}, Y \mid \sigma, \alpha_d, \alpha_f) = s_{j|g} \cdot s_g \cdot Y, \quad j = 1, \dots, N \quad (3')$$

where Y is the total number of consumers active in the market,

$$s_{j|g_j} = \frac{\exp[(\omega_j - \alpha_d P_j + \alpha_f P_0)/(1 - \sigma)]}{\sum_{k \ni g_k = g_j} \exp[(\omega_k - \alpha_d P_k + \alpha_f P_0)/(1 - \sigma)]}, \quad j = 1, \dots, N$$

is the quantity share of product j in total demand for goods from nest g_j , and

$$s_{g_j} = \frac{\left\{ \sum_{k \ni g_k = g_j} \exp[(\omega_k - \alpha_d P_k + \alpha_f P_0)/(1 - \sigma)] \right\}^{(1 - \sigma)}}{\sum_{m=1}^G \left\{ \sum_{k \ni g_k = m} \exp[(\omega_k - \alpha_d P_k + \alpha_f P_0)/(1 - \sigma)] \right\}^{(1 - \sigma)} + 1}.$$

is the quantity share of nest g_j in total consumption.

The parameters α_d and α_f measure consumers' sensitivity to product prices and the price of the outside good, and the parameter $\sigma \in [0,1)$ measures the degree to which products within a nest substitute for one another—larger σ values imply greater

¹⁹ The geographic location of plants, or their 5-digit industrial classifications, can be use a basis for nesting when product features are unobservable. We will use the former nesting strategy in section IV below.

substitutability.²⁰ Also, the quality/appeal index for the imported variety is normalized to zero ($\omega_0 = 0$), so that all measured qualities of domestic varieties are expressed as deviations from the quality of imported goods.

On the supply side, assume that the N domestic producers engage in Bertrand-Nash competition with one another, given the exogenous price of the “outside” variety (P_0). Then the profit maximization condition (4) becomes (Berry, 1994):

$$P_j = C_j + \frac{(1 - \sigma) / \alpha_d}{1 - \sigma \cdot S_{j|g_j} - (1 - \sigma) \cdot S_{j|g_j} \cdot S_{g_j}}, \quad j = 1, \dots, N \quad (4')$$

Given the demand parameters, $(\alpha_d, \alpha_f, \sigma)$, the market size ($Y = \sum_{j=0}^N Q_j$), and the price of the imported good (P_0), equations (3') and (4') describe $2 \cdot N$ relationships among the $4 \cdot N$ unknowns, $(\bar{P}, \bar{Q}, \bar{\omega}, \bar{C})$. A similar set of $2 \cdot N$ relationships is implied by (3') and (4') if total market size (Y) is unobserved but the quantity of imports (Q_{0t}) is available. Hence it is possible to establish a unique mapping from the observables, (P_0, Q_0, TC, \bar{R}) , to the unobservables, $(\bar{P}, \bar{Q}, \bar{\omega}, \bar{C})$, if we exploit (3'), (4') and $2 \cdot N$ more equations (see Appendix 1):

$$R_j = P_j \cdot Q_j \quad j = 1, \dots, N \quad (6)$$

$$TC_j = C_j \cdot Q_j \quad j = 1, \dots, N \quad (7)$$

Equation (6) simply equates total sales revenue to the product of price and quantity, and (7) equates total operating costs to the product of marginal cost and quantity. The latter is

²⁰ As σ increases, products within a nest become closer substitutes. 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.

less innocuous than the former, in that it presumes constant returns to scale and no adjustment costs. However the same assumptions are embodied in Tornqvist indices, and more general formulations are feasible.²¹

B. The evolution of product quality and marginal costs

It remains to link product quality and market costs to the business environment and managerial decisions. To this end 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 \mathbf{X}_{jt} of weakly exogenous performance determinants. Depending upon the context, these might include things like R&D expenditures, participation in foreign markets, and the extent of multinational ownership:

$$\omega_{jt} = \omega_{j0} + \sum_{s=1}^L \lambda_s \omega_{j,t-s} + \sum_{s=L+1}^{2L} \lambda_s \ln C_{jt+L-s} + \lambda^x \mathbf{X}_{jt} + \varepsilon_{jt}^{\omega}, \quad (8a)$$

$$\ln C_{jt} = c_{j0} + \sum_{s=1}^L \varphi_s \ln C_{j,t-s} + \sum_{s=L+1}^{2L} \varphi_s \omega_{j,t+L-s} + \varphi^x \mathbf{X}_{jt} + \varepsilon_{jt}^c, \quad (8b)$$

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

²¹ For example, constant returns can be relaxed by positing an upward or downward-sloping cost function that is common across plants up to a single parameter that captures idiosyncratic efficiency and/or factor price effects. An alternative treatment of adjustment costs is to assume capital stocks are fixed from period to period, and that TC_{jt} reflects the costs of labor, materials and energy. Minor adjustments to the algorithm described in Appendix I become necessary because marginal costs are not flat with respect to output. 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.

Here we assume that the errors are jointly normal, $(\varepsilon^\omega, \varepsilon^c) \sim i.i.d.N(0, \Sigma)$, where

$$\Sigma = E \left[\begin{pmatrix} \varepsilon_c \\ \varepsilon_\omega \end{pmatrix} \begin{pmatrix} \varepsilon_c \\ \varepsilon_\omega \end{pmatrix}' \right],$$

and that they are orthogonal to the unobservable plant effects (ω_{j0}, c_{j0}) .

Once the complete vector of parameters has been estimated, it is possible to simulate the effects of X_{jt} on marginal cost and product quality trajectories. In turn, these trajectories can be used to infer the effects of X_{jt} on welfare by re-solving for product market equilibria and the associated consumer and producer surpluses under alternative assumptions about X_{jt} paths. Specifics of this exercise are discussed in section IV D below.

C. Estimation

Without the VAR system (8a) and (8b), the demand parameters $(\alpha_d, \alpha_f, \sigma)$ are not identified. A different mapping from $(P_0, Q_0, T\bar{C}, \bar{R})$ to $(\bar{P}_t, \bar{Q}_t, \bar{\omega}_t, \bar{C}_t)$ exists for each feasible set of $(\alpha_d, \alpha_f, \sigma)$ values, and without more structure, each is equally likely.

Equations (8a) and (8b) help with identification by constraining the shapes of the cross-sectional $(\bar{\omega}_t, \bar{C}_t)$ distributions and the way that individual (ω_{jt}, C_{jt}) pairs evolve through time. However, these constraints 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 unknown parameters and estimating the system (3'), (4'), (6), (7), (8a) and (8b) using Bayesian techniques.

To summarize this estimation strategy, let us collect all of the parameters we have introduced in the vector $\theta = [\alpha_d, \alpha_f, \sigma, \lambda, \varphi, \Sigma]$, and define the density $p(\theta)$ to describe our priors, which we will discuss 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:

$$\pi(\theta | D) = \frac{p(\theta) \cdot L(D | \theta)}{\int_{\theta} p(\theta) \cdot L(D | \theta) d\theta} \propto p(\theta) \cdot L(D | \theta), \text{ where } L(D | \theta) \text{ is the likelihood}$$

function based on (3'), (4'), (6), (7), (8a) and (8b).

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:

$$p(\theta) = p_{\sigma}(\sigma) \cdot p_{\alpha_d, \alpha_f}(\alpha_d, \alpha_f) \cdot p_{\varphi, \lambda}(\varphi, \lambda) \cdot p_{\Sigma}(\Sigma).$$

The demand system priors we impose are similar to those used by Poirier (1996) and Lahiri and Gao (2001).²² The underlying utility maximization problem implies that $\sigma \in [0,1]$, so we specify uniform priors on this region of support.²³ Similarly, although the price coefficients α_d and α_f should be positive, we do not know much about their magnitudes, so we specify uniform priors with support $[0, 10]$ for each of these

²² These studies also estimate nested logit models using Bayesian techniques. However, unlike ours, they are concerned with the problem of ill-defined nesting structures.

²³ 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.

parameters. (Experimentation suggests that our choice of priors for the price coefficients has little effect on the analysis, but it is necessary to impose priors with finite variance on at least one of these coefficients to deal with the under-identification problem.) The remaining parameters describe the VAR (equation 8). For the autoregressive coefficients we assume joint normality, $p_{\phi,\lambda}(\phi,\lambda) = N_{2L}(0_{2L}, 100 \times I_{2L})$, where L is one plus the 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, $p_{\Sigma}(\Sigma) = InvWish(6, 100 \times I_2)$. Overall then, we are imposing some structure on the demand parameter posteriors, but we are doing little to constrain the range of plausible realizations on the VAR parameters.

Closed-form representations of the posterior $\pi(\theta | D)$ are not available; nor is it feasible to make *i.i.d.* draws directly from $\pi(\theta | D)$. We therefore use a Markov chain Monte Carlo (MCMC) algorithm to generate correlated draws from $\pi(\theta | D)$ and we analyze the moments of the resulting empirical distributions (Gilks, et al, 1996). The vector θ is relatively large, so we partition θ into three sub-vectors: $\theta = [\theta_1, \theta_2, \theta_3]$, where $\theta_1 = (\alpha_d, \alpha_f, \sigma)$, $\theta_2 = (\lambda, \varphi)$ and $\theta_3 = \Sigma$. Then we update the sub-vectors sequentially using Gibbs sampling. Appendix 2 provides further details.

E. Constructing Performance Measures

Once we have estimated our posterior distribution, $\pi(\theta | D)$, we solve for the marginal cost and product quality trajectories of each producer in the sample using the

mean of θ .²⁴ 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 consumer surplus with, versus without the i^{th} good:²⁵

$$\Delta CS_t^i = y_t \cdot \left[\sum_j \exp(\bar{u}_{jt}/(1-\sigma))^{(1-\sigma)} - \sum_{j \neq i} \exp(\bar{u}_{jt}/(1-\sigma))^{(1-\sigma)} \right] \quad (9)$$

where $\bar{u}_{jt} = \begin{cases} \omega_{jt} - \alpha_d P_{jt} & j \in \{1, \dots, N\} \\ -\alpha_f P_{0t} & j = 0 \end{cases}$ measures the cross-consumer mean utility

delivered by product j . Prices and market shares are allowed to adjust to re-establish equilibrium when good i 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.

To evaluate these firm-specific welfare contributions, we express them as ratios to firms' total costs. The larger this ratio, the larger the amount of surplus created per unit

²⁴ It would, of course, be possible to also study the distributions for these trajectories that are induced by $\pi(\theta | D)$; we have not pursued this yet.

²⁵ Akerberg 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. Thus our results may over-emphasize consumer surplus relative to producer surplus.

expenditure on inputs. For the sake of comparison, we also calculate traditional Tornqvist measures of total factor productivity under the standard assumptions that deflated revenues measure real output, deflated expenditures on intermediate goods measure physical intermediate good usage, and deflated book values of capital measure physical capital stocks:

$$\tilde{\phi}_{jt} = \ln(R_{jt} / \bar{P}_t) - \sum_{i=1}^I \hat{\mu}_{jt}^i \ln(\tilde{V}_{jt}^i), \quad (10)$$

where $\hat{\mu}_{jt}^i$ is the share of the i^{th} factor in total costs at firm j during period t . For comparison purposes, we consider two variants of this measure—one in which labor is measured in physical units ($\tilde{\phi}_{1jt}$) and one in which labor is measured in terms of the wage bill, including fringe benefits ($\tilde{\phi}_{2jt}$). Differences between these two indices should reveal the importance of mark-up effects related to labor costs, as discussed in section II.

IV. An Application to the Colombian Paper Mill Industry

A. The Data

We base our empirical example on panel data describing the 12 Colombian paper mills that operated continuously during 1981-1991.²⁶ These data were originally collected by Colombia's official statistical agency (Departamento Administrativo Nacional de Estadística) and they have been cleaned as described in Roberts (1996).²⁷ We choose the

²⁶ Lu (1999) provides an overview of the Colombian pulp and paper industry. See also Departamento Nacional de Planeación (2005).

²⁷ We exclude plants that entered or exited during the sample period to avoid complicating the VAR portion of the likelihood function. This naturally creates some selection bias, although the entering and exiting plants were quite small and thus had a minor influence on market.

Colombian paper industry because it has a relatively small number of producers who are differentiated from one another by geographic location and by modest differences in their product characteristics. Thus its producers enjoy some market power, and given the transport costs for paper, their locations imply a natural nesting structure for our demand system.²⁸ Finally, because many Colombian mills import pulp from Chile and the United States, and many export paper to Latin American destinations, this industry allows us to study the relation between plants' performances and their participation in international markets.

Using these data 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 costs, TC_{jt} , we first sum payments to labor, intermediate input purchases net of inventory accumulation, energy purchases, and capital costs. (The latter are measured as 10 percent of the book value of firms' capital stocks.) Then we scale this aggregate by the fraction of sales that go to domestic consumers and we divide the result by the same wholesale price deflator we used for output. This definition of total cost implies that all inputs are variable, as do standard Tornqvist indices.

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

²⁸ Specifically, we assign each domestic producer to one of three regional nests: central (including Bogotá); coastal (including Medellín and Cartagena), and other (including Cali).

price fluctuates only with the exchange rate. Then, calling the period t dollar value of imports R_{0t} , 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 and effectively fix the size of the market.²⁹

Our vector of weakly exogenous performance determinants (X_{jt}) includes four variables. The first is a trend captures secular movements in input costs, productive efficiency, and product appeal. The second is a real exchange rate, which is a potentially important source of fluctuation in input prices. The last two variables characterize plants' participation in foreign markets. (Our data do not allow us to explore the effects of multinational ownership, advertising, or R&D expenditures.) One is a dummy that takes a value of one if the plant was importing some or all of its intermediate inputs in year t , but *not* exporting any of its output. The other is a dummy that takes a value of one if the plant was both importing some intermediates and exporting some output in year t . No plant in our sample exported output without importing intermediate inputs, so the omitted category is plants that did not buy inputs *or* sell outputs in international markets in year t .

B. Posterior Parameter Distributions

Means, standard errors and other summary statistics for our estimated posterior distribution $\pi(\theta | D)$ are reported in Table 1. These estimates are constructed using

²⁹Unfortunately, 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.

Wooldridge's (forthcoming) correction for the unobserved plant effects in the VAR, c_{j0} and ω_{j0} .³⁰

Most marginal posterior distributions have much smaller variances than the associated priors, implying that the data provide the main basis for identification. However, for the reasons discussed in section IIIC above, the data are not informative about the domestic price coefficient, α_d . The posterior distribution for this parameter is nearly coincident with our prior distribution, regardless of what priors we choose. Fortunately, experimentation confirms that the posterior distributions for all other parameters except α_f proved to be insensitive to our α_d prior.

Product quality and marginal cost both show a moderate amount of persistence, as evidenced by the posteriors for coefficients on lagged dependent variables (λ_2 and ϕ_2). There is not much dynamic interaction among these variables (λ_3 and ϕ_3), but their innovations show modest negative correlation.³¹ Macro variables do not exhibit strong relationships with quality or marginal costs, but it is worth noting that domestic product qualities trend weakly upward relative to imported products, and they tend to fall with real exchange rate depreciation. Also, marginal costs trend downward, and tend to rise

³⁰ Wooldridge's (forthcoming) correction takes care of the 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 of a similar VAR. Our preliminary results indicated that the variances of the random effects are sufficiently small to ignore, so the specification reported in table 1 simply includes initial values of the lagged dependent variables.

³¹ The posterior mean of Σ implies that the correlation between marginal cost innovations and product quality innovations is -0.347. Correlation between plant effects for marginal costs and product quality is implied by the coefficients on initial cost and product appeal realizations (g^w and g^c).

with the exchange rate. The latter is presumably because many firms import their pulp. Finally, the dummy coefficients suggest that firms engaged in exporting tend to produce higher quality products and that firms that import some of their inputs (weakly) tend to have lower marginal costs. 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 next 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 total production costs, $\frac{\Pi_j}{TC_j}$, external

effects on the producer surplus of other plants over total production costs, $\frac{\Delta \sum_{k \neq j} \Pi_k}{TC_j}$, and

net total surplus created over total production costs, $\frac{\Delta U_j + \Pi_j}{TC_j} + \frac{\Delta \sum_{k \neq j} \Pi_k}{TC_j}$. Then,

pooling all observations, we calculate the cross-plant correlations in these variables reported in tables 2.³² To contrast our performance measures with traditional productivity measures, we include $\tilde{\phi}_1$ and $\tilde{\phi}_2$ in our analysis.

Several intriguing findings emerge concerning the relation between $\tilde{\phi}$ -type measures, marginal costs, and product quality. First, the standard total factor productivity

³² 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.

measures ($\tilde{\phi}_{1jt}$ and $\tilde{\phi}_{2jt}$) are negatively associated with marginal costs ($\rho = -0.473$, $\rho = -0.736$), reflecting the mark-up effect mentioned in sections II and III. Second, however, $\tilde{\phi}_{jt}$ completely fails to capture cross-firm variation in product quality or appeal ($\rho = -0.098$, $\rho = -0.051$).

Other correlations speak to the relationship between $\tilde{\phi}$ -type measures and welfare-based measures. Because firms with low marginal costs have relatively large mark-ups and relatively large operating profits, $\tilde{\phi}_1$ and $\tilde{\phi}_2$ are strongly correlated with producer surplus ($\rho = 0.620$, $\rho = 0.844$). And because low marginal cost firms also have relatively low prices, $\tilde{\phi}$ -type performance measures are somewhat correlated with consumer surplus as well ($\rho = 0.066$, $\rho = 0.301$). However, firms with low prices tend to take surplus away from others, so the correlation between $\tilde{\phi}$ -type measures and *net* contributions to social welfare is quite modest ($\rho = 0.144$, $\rho = 0.219$). That is, under our assumptions about producer and consumer behavior, traditional Tornqvist indices tell us almost nothing about which firms are doing well from a social perspective. Finally, $\tilde{\phi}_{2jt}$ is more closely correlated with welfare-based measures than $\tilde{\phi}_{1jt}$, most likely because the former does not reflect spurious variation due to labor cost pass-through effects (section II).

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 welfare 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 $(\varepsilon_{jt}^c, \varepsilon_{jt}^w)$ that were actually observed.

Under these assumptions we can use our VAR parameters to calculate the paths for $(\omega_{jt}, \ln C_{jt})$ that would have emerged if, beginning in 1982, all international producer trade had been shut down. The cross-plant temporal averages for these variables simply reflect the posterior distributions discussed in section B above (Figure 1). Firms that export and/or use imported intermediate imports tend to have low marginal costs; but the tendency for exporters to produce high-quality goods is largely offset by the tendency for importers to produce lower quality goods.

Substituting our parameter estimates and these counterfactual trajectories for $(\omega_{jt}, \ln C_{jt})$ into equations (3') and (4'), 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, the effect on marginal cost of shutting down trade proves more important than the effect on product appeal. That is, the higher marginal costs induced by a prohibition on producer trade translate into lower producer surplus, higher output prices, and therefore, lower consumer surplus. Thus, although our results are not statistically strong, and we haven't dealt with endogeneity issues, they suggest that firms participating in foreign markets contribute more to welfare. One possible reason is that low cost firms and high quality firms self-select into exporting, as suggested by Clerides et al (1998) and Bernard et al (1999).

A rather different story would have emerged if we had relied on $\tilde{\phi}$ -type measures for policy analysis. Fitting a conditional AR(1) like those in Table 1 with $\tilde{\phi}$ as the dependent variable, we find that firms that exported had *lower* measured productivity, significantly so for the case of expenditure-based labor measures, $\tilde{\phi}_2$ (refer to Table 3). One interpretation is that exporters produce relatively high-quality goods, and in doing so they incur relatively high marginal costs. Since quality is not captured by $\tilde{\phi}$ -type measures, but high marginal cost shows up as low productivity, exporters look relatively bad. Finally, those plants that imported their intermediate goods do neither better nor worse, and exchange rate devaluation drags down productivity, suggesting that the mark-up effect associated with imported pulp and chemicals is important.

To demonstrate the implications of these estimated $\tilde{\phi}$ -type trajectories, we once again cut plants off from foreign trade. Now isolation appears to *improve* their performance, largely because of the negative correlation between exporting and $\tilde{\phi}$ (Figure 3).

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 little attention to the institutional and technological features of the Colombian 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 many findings in the literature on plant-level performance may be spurious, to sketch an alternative approach to inference that we feel holds more promise, and to contrast our methodology with the standard approach by applying both to the same data. Significant refinements in most of the dimensions mentioned above are possible; we are optimistic that they will enhance the usefulness of our methodology.

Bibliography

- Akerberg, 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 and J. Bradford Jensen (1999) "Exceptional Exporter Performance: Cause, Effect, or Both?" *Journal of International Economics*, 1999, Vol. 47, pp. 1-25.
- Bernard, Andrew, Jonathan Eaton, J. Bradford Jensen and Samuel Kortum (2003) "Plants and Productivity in International Trade," *American Economic Review* 93(4), pp. 1268-1290.
- 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.
- _____ (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, Andrew and Barry Nalebuff (1991) "Aggregation and Imperfect Competition: On the Existence of Equilibrium," *Econometrica* 59, pp. 25-60.
- Caves, Douglas, 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.
- Clerides, Sofronis, Saul Lach and James Tybout (1998) "Is Learning by Exporting Important? Micro Dynamic Evidence from Colombia, Mexico and Morocco." *Quarterly Journal of Economics*, pp. 903-947.
- Departamento Nacional Planeacion (2005) "Analysis de Cadenas Productivas: Cadena Pulpa, Papel e Industria Grafica," www.dnp.gov.co/archivos/documentos/DDE_Desarrollo_Emp_Industria/Imprenta.pdf, downloaded April 19, 2005.
- Foster, Lucia, John Haltiwanger and Chad Syverson (2003) "Reallocation, Firm Turnover and Efficiency: Selection on Productivity or Profitability?" University of Chicago, Graduate School of Business Working Paper.
- Gilks, Walter R., Sylvia Richardson and David 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.
- Girma, Sourafel and Holger Gorg (2004) "Evaluating the Foreign-Ownership Wage Premium using a Matching and Difference-in-Differences Approach," University of Nottingham, Department of Economics.
- Griliches, Zvi. (1986). "Productivity, R&D and Basic Research at the Firm Level in the 19870s," *American Economic Review* 76(1), pp. 141-154.
- Griliches, Zvi and Haim Ragev (1995) "Firm Productivity in Israeli Industry: 1979-1988," *Journal of Econometrics*, 65(1), pp.175-203.
- Hall, Robert (1988) "The Relation between Price and Marginal Cost in U.S. Industry," *Journal of Political Economy*, pp. 921-947.
- Hall, Robert (1990) "Invariance Properties of Solow's Productivity Residual," in Peter Diamond, ed., *Growth, Productivity and Employment*. Cambridge: MIT Press.
- Henderson, Vernon (2001) "Marshall's Scale Economies," processed, Department of Economics, Brown University.

- 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 (2002) "Bayesian Analysis of Nested Logit Model by Markov Chain Monte Carlo," *Journal of Econometrics*, 111(1), pp. 103-133.
- Levinsohn, James and Amil Petrin (2003) "Estimating Production Functions using Inputs to Control for Unobservables," *Review of Economic Studies*.
- 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.
- Lu, Shihua (1999). *Modeling Industrial Evolution in an Import Competing Industry, with an Application to the Pulp and Paper Industry in Colombia*. Ph.D. dissertation, Georgetown University.
- Marshall, 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.
- Pakes, Ariel and Paul McGuire (1994). "Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model," *RAND Journal of Economics*, 25(4), pp. 555-589.
- 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 (2000) "Manufacturing Firms in Developing Countries: How Well do They do, and Why?" *Journal of Economic Literature* 38(1), pp. 11-44.
- _____ (2003) "Plant- and Firm-Level Evidence on the 'New' Trade Theories," in E. Kwan Choi and James Harrigan (eds.) *Handbook of International Trade*, Oxford: Basil-Blackwell.
- 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 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 (forthcoming) "Simple Solutions to the Initial Conditions Problem for Dynamic, Nonlinear Panel Data Models with Unobserved Heterogeneity," *Journal of Applied Econometrics*.
- 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, given (P_0, Q_0, TC, \bar{R}) , the $4 \cdot N$ unknowns $(\bar{P}, \bar{Q}, \bar{\omega}, \bar{C})$ are uniquely determined by (3'), (4'), (6) and (7) of the text. It also sketches an algorithm for finding these unknowns. The mapping is done period by period, so we shall hereafter drop t subscripts to reduce clutter.

First, by equation (7), total costs at the j^{th} plant are $TC_j = Q_j C_j$, so the within-group market share of the j^{th} firm is: $S_{j|g_j} = \frac{TC_j}{C_j} / Q_{g_j}^{\text{tot}}$ where total output from the j^{th} plant's group is $Q_{g_j}^{\text{tot}} = \sum_{k \ni g_k = g_j} \frac{TC_k}{C_k}$ and total domestic output is $Q^{\text{tot}} = \sum_{g=1, G} Q_g^{\text{tot}}$.

Also, by equation (6), total revenues at the j^{th} plant are $R_j = P_j Q_j$, so the j^{th} plant's price-cost markup may be expressed in terms of observable variables as $m_j = \frac{R_j}{TC_j} - 1$, and once its marginal cost is known, its price can be calculated as $P_j = (m_j + 1)C_j$.

Substituting these market share and price expressions into the pricing rule (4') and solving for marginal cost, we obtain:

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

This expression defines the unobservable C_j as a monotonic decreasing function of $Q_{g_j}^{\text{tot}}$, given data on TC_j , m_j , Q^{tot} and Q_0 . Thus, once the nest quantity subtotals are known, each firm's marginal costs are implied by (A1.1). With these marginal costs,

prices can be retrieved from $P_j = (m_j + 1)C_j$, and using $Q_j = R_j / P_j$, market shares follow immediately. Finally, once prices and market shares are known, the vector of product qualities can be found by substituting into:

$$\omega_j = \alpha_d P_j - \alpha_f P_0 - \sigma \ln(S_{j|g_j}) + \ln(S_{j|g_j} \cdot S_{g_j}) - \ln(S_0) \quad j = 1, \dots, N, \quad (\text{A1.2})$$

which follows from the expressions in the text for $s_{j|g_j}$ and s_{g_j} .³³ (Here

$$S_0 = 1 - \sum_{g=1}^G S_g \text{ is the market share of the imported variety.})$$

The key, then, is to solve for the nest quantity subtotals. Substituting the marginal cost expression (A1.1) into $Q_{g_j}^{tot} = \sum_{k \ni g_k = g_j} \frac{TC_k}{C_k}$ and dividing both sides by $Q_{g_j}^{tot}$, one

obtains:

$$1 = \sum_{k \ni g_k = g} \left[\frac{1}{\sigma + (1 - \sigma) \cdot Q_g^{tot} \cdot \left((Q^{tot} + Q_0)^{-1} + (\alpha_d \cdot TC_k \cdot m_k)^{-1} \right)} \right], \quad (\text{A1.3})$$

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

The right-hand side of this expression 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 $Q_g^{tot} \rightarrow \infty$, where n_g is the number of producers in nest g . Thus, for all $g \in \{1, \dots, G\}$, $Q_0 > 0$, and $Q^{tot} \geq 0$, equation (A1.3)

³³ Berry, Levinsohn and Pakes (1995) use a similar inversion to study the quality of automobile models.

has a unique, positive root: $Q_g^{tot} = f_g(Q^{tot} | Q^0)$, which can be found using a bisection algorithm at any Q^{tot} .

It remains to show that $Q^{tot} - \sum_{g=1,G} f_g(Q^{tot} | Q_0) = 0$ has a unique positive root

for any given $Q_0 > 0$. The existence of at least one root follows from the fact that

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

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

share of domestic quantities in total quantity,

$$\frac{Q^{tot}}{Q^{tot} + Q^0} = \sum_{g=1,G} \frac{f_g(Q^{tot} | Q^0)}{Q^{tot} + Q^0} = \sum_{g=1,G} s_g(Q^{tot} | Q^0),$$

is a continuous decreasing function of Q^{tot} . This can be seen by restating (A1.3) as:

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

that s_g falls with Q^{tot} , $g = 1, \dots, G$. Again, a bracketing and bisection algorithm suffices to generate numerical solutions.

Appendix 2: The Gibbs Sampler

Because it is not feasible to sample independent draws from the density $\pi(\theta | D) \propto p(\theta) \cdot L(D | \theta)$, we use Markov chain Monte Carlo (MCMC) techniques. The idea is to draw a sequence of realizations on θ from some Markov process, $\{\theta^{(1)}, \theta^{(2)}, \dots, \theta^{(i)}\}$, with elements whose unconditional distributions converge to $\pi(\theta | D)$ as $i \rightarrow \infty$. 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: $\theta_1 = (\alpha_d, \alpha_f, \sigma)$, $\theta_2 = (\lambda, \phi)$, and $\theta_3 = \text{vec}(\Sigma)$. 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 $\theta_1^{(i+1)} \sim \pi_1(\theta_1 | \theta_2^{(i)}, \theta_3^{(i)}, D)$

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

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

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

The distribution $\pi_1(\theta_1 | \theta_2, \theta_3, D)$ is the most difficult to construct. It is proportional to $L(D | \theta_1, \theta_2, \theta_3)p_{\theta_1}(\theta_1)$ where $L(D | \theta_1, \theta_2, \theta_3)$ is the likelihood function based on (3'), (4'), (6) and (7); and $p_{\theta_1}(\theta_1)$ is the prior distribution defined in the text. But $L(D | \theta_1, \theta_2, \theta_3)f_{\theta_1}(\theta_1)$ 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 $\pi_2(\theta_2 | \theta_1, \theta_3, D)$ and $\pi_3(\theta_3 | \theta_1, \theta_2, D)$, rewrite (11a) and (11b) as

$$Y_{jt} = \beta'Z_{jt} + \varepsilon_{jt} \text{ where } Y_{jt} = (\omega_{jt}, c_{jt})', Z_{jt} = (\mathbf{1}, \omega_{jt-1}, c_{jt-1}, X'_{jt})', \varepsilon_{jt} = (\varepsilon_{jt}^{\omega}, \varepsilon_{jt}^c),$$

$$\text{and } \beta' = \begin{pmatrix} \omega_0 & \lambda_{\omega} & \lambda_C & \lambda'_X \\ c_0 & \phi_{\omega} & \phi_C & \phi'_X \end{pmatrix}. \text{ Also, stacking observations, let us define:}$$

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

Then, we can write the VAR system as $Y = Z\beta + U$. Further, conditional on θ_1 , our one-to-one mapping from $(T\bar{C}_t, \bar{R}_t | Q_{0t}, P_{0t})$ to $(\bar{\omega}_t, \bar{C}_t)$ 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 $\left((Z'Z)^{-1}Z' \otimes I_2\right) \cdot \text{vec}(Y')$ 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 $u_n = V_n \left[\left(Z' \otimes \Sigma^{-1} \right) \text{vec}(Y') + V_0^{-1} u_0 \right]$ and variance $V_n = \left[\left((Z'Z) \otimes \Sigma^{-1} \right) + V_0^{-1} \right]^{-1}$.

Similarly, using the mapping $(\theta_1, \theta_2, D) \rightarrow (\beta, Y, Z)$, we may write the likelihood-based full conditional estimator of Σ , given (θ_1, θ_2, D) , as

$$\frac{1}{N(T-1)} \sum_{i=1}^N \sum_{t=2}^T (Y_{it} - Z_{it}\beta)(Y_{it} - Z_{it}\beta)'. \text{ When multiplied by } N(T-1), \text{ this estimator has a}$$

Wishart distribution with $N(T-1)$ degrees of freedom. Thus, given that we have assumed Σ has prior distribution $\text{InvWish}(m_0, G_0^{-1})$, the full conditional posterior distribution for

$\theta_3 = \text{vec}(\Sigma)$, i.e., $\pi_3(\theta_3 | \theta_1, \theta_2, D)$, is the vector version of a $\text{InvWish}(m_n, G_n^{-1})$

distribution, where $m_n = m_0 + N(T-1)$ and $G_n^{-1} = G_0^{-1} + (Y - Z\beta)'(Y - Z\beta)$.

TABLE 1: POSTERIOR PARAMETER DISTRIBUTIONS

	Prior	Mean	Std. Dev.	Z-stat.	Median	5%	95%
<i>Demand System</i>							
α_d	$U[0,10]$	5.320	2.720	1.956	5.380	1.093	9.504
α_f	$U[0,10]$	2.100	0.871	2.411	2.018	0.809	3.972
σ	$U[0,1]$	0.893	0.047	19.055	0.897	0.808	0.966
<i>Product Quality VAR</i>							
λ_1 (constant)	$N(0,10)$	0.039	0.208	0.187	0.038	-0.301	0.380
λ_2 (ω_{it-1})	$N(0,10)$	0.307	0.104	2.939	0.303	0.141	0.483
λ_3 ($\ln C_{it-1}$)	$N(0,10)$	-0.024	0.035	-0.692	-0.023	-0.083	0.033
λ_4 (trend)	$N(0,10)$	0.029	0.060	0.473	0.029	-0.070	0.127
λ_5 (exchange rate)	$N(0,10)$	-1.839	0.958	-1.918	-1.738	-3.813	-0.413
λ_6 (exporter)	$N(0,10)$	0.036	0.017	2.091	0.036	0.007	0.064
λ_7 (importer)	$N(0,10)$	-0.019	0.057	-0.331	-0.018	-0.112	0.074
$\mathfrak{G}_1^\omega(\omega_{it})$	$N(0,10)$	0.690	0.116	5.956	0.693	0.494	0.875
$\mathfrak{G}_2^\omega(\ln C_{it})$	$N(0,10)$	0.052	0.062	0.835	0.051	-0.048	0.156
<i>Log Marginal Cost VAR</i>							
ϕ_1 (constant)	$N(0,10)$	0.131	0.472	0.278	0.138	-0.653	0.900
ϕ_2 ($\ln C_{it-1}$)	$N(0,10)$	0.437	0.079	5.544	0.437	0.308	0.567
ϕ_3 (ω_{it-1})	$N(0,10)$	0.162	0.202	0.801	0.162	-0.168	0.496
ϕ_4 (trend)	$N(0,10)$	-0.222	0.139	-1.596	-0.220	-0.454	0.004
ϕ_5 (exchange rate)	$N(0,10)$	0.456	0.663	0.687	0.454	-0.629	1.546
ϕ_6 (exporter)	$N(0,10)$	-0.020	0.037	-0.538	-0.020	-0.081	0.041
ϕ_7 (importer)	$N(0,10)$	-0.122	0.132	-0.929	-0.123	-0.339	0.092
$\mathfrak{G}_1^C(\ln C_{it})$	$N(0,10)$	0.482	0.136	3.532	0.482	0.258	0.707
$\mathfrak{G}_1^C(\omega_{it})$	$N(0,10)$	-0.101	0.230	-0.438	-0.100	-0.479	0.275
<i>Covariance Matrix</i>							
Σ_{11}	$\Sigma \sim$ Inverted Wishart (3,1)	0.032	0.005	5.979	0.031	0.024	0.041
Σ_{12}		-0.026	0.011	-2.255	-0.025	-0.046	-0.009
Σ_{22}		0.175	0.025	6.991	0.172	0.138	0.219

TABLE 2:

CORRELATIONS OF PERFORMANCE MEASURES (Colombian Paper Mills)*

	$\frac{\Delta U_j}{TC_j}$	$\frac{\Pi_j}{TC_j}$	$\frac{\Delta \sum_{k \neq j} \Pi_k}{TC_j}$	$\frac{\Delta U_j + \Pi_j}{TC_j}$ $\frac{\Delta \sum_{k \neq j} \Pi_k}{TC_j}$	$\ln C_j$	ω_j	$\tilde{\phi}_{1j}$	$\tilde{\phi}_{2j}$
$\frac{\Delta U_j}{TC_j}$	1	0.639	-0.756	0.741	-0.544	0.622	0.066	0.301
$\frac{\Pi_j}{TC_j}$		1	-0.805	0.402	-0.818	0.147	0.62	0.844
$\frac{\Delta \sum_{k \neq j} \Pi_k}{TC_j}$			1	-0.187	0.825	-0.141	-0.239	-0.525
$\frac{\Delta U_j + \Pi_j}{TC_j} + \frac{\Delta \sum_{k \neq j} \Pi_k}{TC_j}$				1	-0.139	0.746	0.144	0.219
$\ln C_j$					1	-0.011	-0.493	-0.736
ω_j						1	-0.098	-0.051
$\tilde{\phi}_{1j}$							1	0.929
$\tilde{\phi}_{2j}$								1

* All variables are purged of annual time effects. Productivity normalizations are based on Caves et al (1982). $\tilde{\phi}_{1j}$ is based on the number of workers; $\tilde{\phi}_{2j}$ is based on the cost of labor.

TABLE 3: VAR BASED ON STANDARD PRODUCTIVITY MEASURES

	Dependent variable: $\tilde{\phi}_1$	Dependent variable: $\tilde{\phi}_2$
Intercept	0.395 (0.113)	0.288 (0.061)
Trend	0.026 (0.024)	0.007 (0.013)
Exporter Dummy	-0.071 (0.080)	-0.106 (0.048)
Importer Dummy	0.099 (0.072)	-0.049 (0.043)
Exchange rate	-1.152 (0.436)	-0.280 (0.236)
$\tilde{\phi}_{t-1}$	0.649 (0.069)	0.587 (0.075)
$\tilde{\phi}_1$	0.153 (0.054)	0.171 (0.064)

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

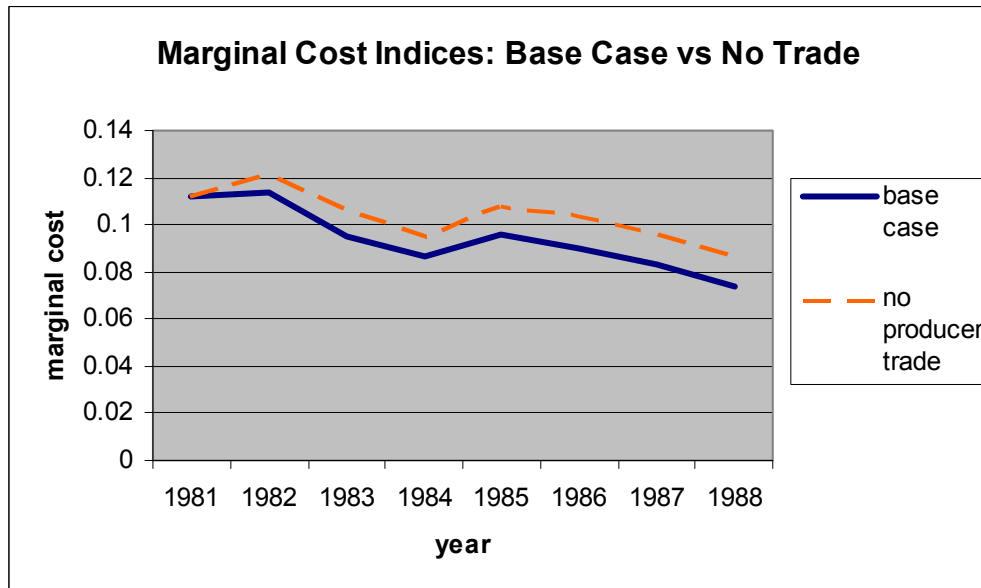
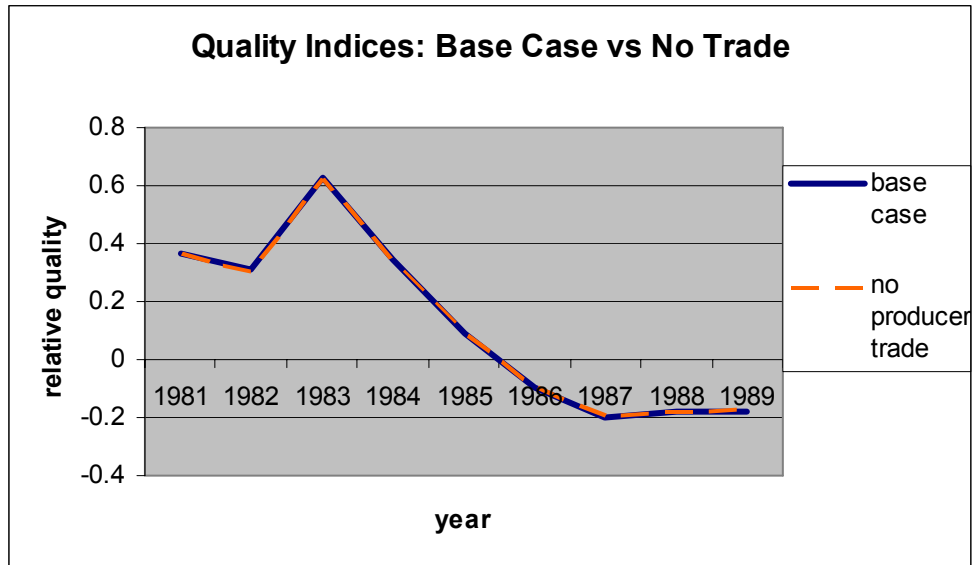
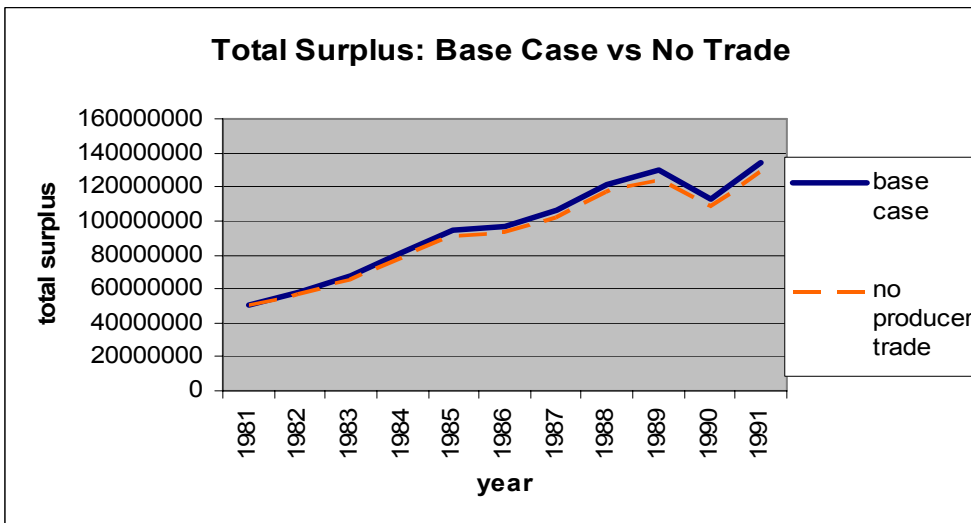
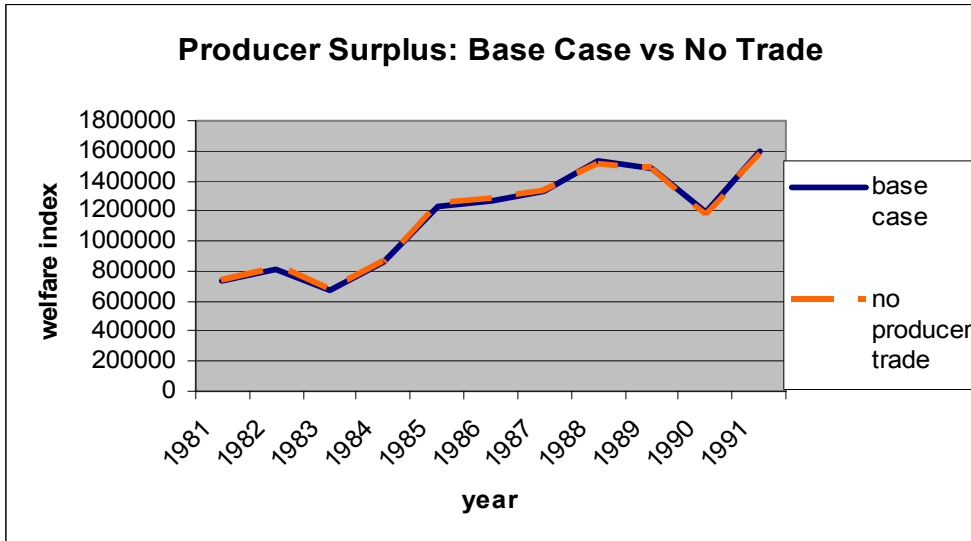
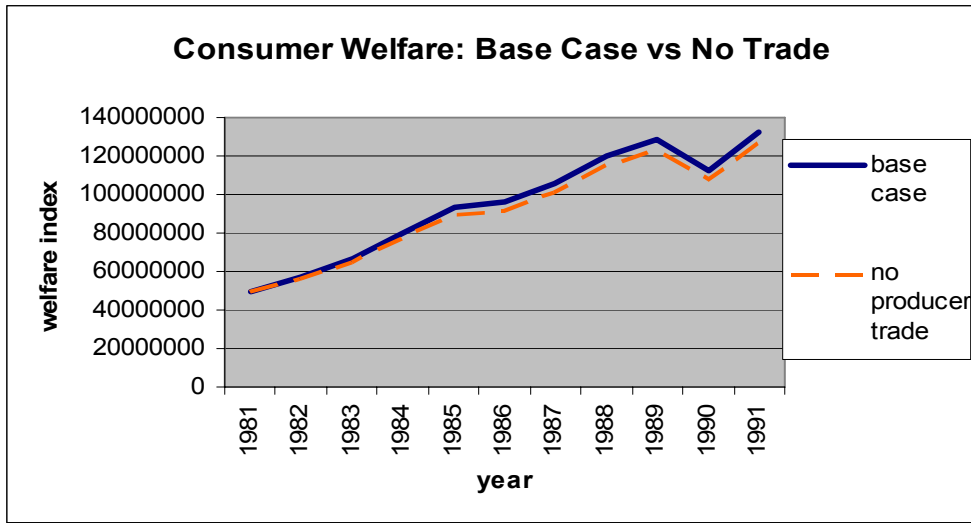


FIGURE 2: EFFECTS OF PRODUCER TRADE ON WELFARE MEASURES



**FIGURE 3:
EFFECTS OF PRODUCER TRADE ON TRADITIONAL PERFORMANCE MEASURES**

