Micro Foundations of Export Dynamics

Sanghamitra Das
Indian Statistical Institute, Delhi Centre

Mark Roberts
Pennsylvania State University and the NBER

James Tybout
Pennsylvania State University

April 28, 2000

Abstract

The responses of industrial exports to regime changes are notoriously hard to predict. In this paper we investigate whether the predictability problem traces partly to micro phenomena that are undetectable with aggregated data. We begin by developing a dynamic model of exporting behavior that allows for uncertainty, heterogenous firms, and one-shot entry costs for firms breaking into foreign markets. Then we fit the model to plant-level panel data on Colombian chemical producers. Finally, using the results, we simulate aggregate export trajectories under alternative exchange rate regimes to quantify the effects of entry costs, regime credibility, and heterogeneity. Each proves to matter, but a substantial portion of the randomness in aggregate exports remains unexplained.

Acknowledgments

This paper was partially funded by the World Bank research project "Micro Foundations of Successful Export Promotion," RPO 679-20. The views expressed herein are our own and do not necessarily reflect those of the World Bank. Without implicating them, we would like to thank Behzad Diba, Martin Evans, Michael Keane, Andrew Levin, John Rust, and Ariel Pakes for helpful discussions. We are also grateful to seminar participants at the University of Texas, The University of Toronto, The World Bank, The Board of Governors, George Washington University, and the Bureau of the Census for their comments.

I. Overview

In developing countries, industrial exports are famously unpredictable. Seemingly similar stimuli have given rise to very different export responses in different countries and time periods, leaving analysts to wonder whether the next reform package will generate a surge or a trickle. When foreign exchange reserves are dwindling or domestic demand is slack, the stakes are substantial.

It is not hard to identify micro explanations for the predictability problem. First, a strong export response often means convincing non-exporters to initiate foreign sales. But to break into foreign markets, firms must establish marketing channels, learn bureaucratic procedures, and develop new packaging or product varieties.[1] Exchange rate movements that are viewed as temporary may not induce firms to bear these start-up costs, while credible regime shifts may trigger strong responses. Second, even within narrowly defined industries, firms are quite heterogeneous in terms of their production costs and their product characteristics. For some, small perturbations to the return from exporting may entice them into foreign markets, while for others, dramatic changes in the incentive structure may be needed. The more heterogeneity, the less likely it is that an export boom can be triggered by moving the exchange rate past some critical threshold. Finally, history matters. If a core of exporting firms has already been established, it is less important to induce non-exporters to retool for foreign markets, and export responses may largely reflect volume adjustment among these incumbents.

Unfortunately, we have little sense for which of these factors is quantitatively important,

---

[1]Start-up costs are the focus of the analytical literature on export hysteresis (Baldwin and Krugman, 1989; Dixit, 1989; Krugman, 1989). For evidence on the role of new exporters see Roberts and Tybout (1997b).

or what to look for when assessing potential export responsiveness in a particular context. Early attempts to model export supplies were based on macro data, so they did not quantify the effects of sunk costs, expectations, or firm heterogeneity in shaping export responsiveness.[2] Several more recent studies have used micro panel data to test the hypothesis that sunk costs matter, and have concluded that they clearly do (Roberts and Tybout, 1997a; Sullivan, 1996; Bernard and Jensen, 1999). But these second generation studies have been based on reduced form relationships between current exporting status, exporting history, and exogenous shocks. Thus they have not recovered the deep parameters that link behavior to stochastic processes, and they have not addressed the question of what would happen if regime changes were to stabilize the exchange rate, or change its mean realization. Nor have they shed much light on the role of firm-level distributions of marginal costs or foreign demand conditions. The purpose of the present paper is to estimate a structural model that quantifies the role of each of these factors using plant-level panel data on export revenues and a few other variables.

II. An Empirical Model of Exporting Decisions with Sunk Costs

Previous micro studies of export dynamics have avoided structural models because they are computationally complex. Each period, firms decide whether to export on the basis of observable information about uncertain future market conditions, as well as entry costs, exit costs, and their current exporting status. Hence, for each possible decision, the evaluation of future export profit trajectories involves multi-dimensional integration over realizations on all the

---

[2]Roberts and Tybout (1997a) briefly review studies that use macro data to assess the effects of start-up costs on exports.

state variables.[3] If there are more than several state variables, the problem quickly becomes computationally intractable.

Given these difficulties, we approach structural estimation by, first, keeping the number of state variables in the model small, and second, using techniques recently developed by Rust (1997) that afford new flexibility in estimation. Our key assumptions are the following:

- *The foreign and domestic market for each firm's product are monopolistically competitive.* This eliminates strategic competition, but it ensures that each firm faces a downward-sloping marginal revenue function in each market. It seems like a very reasonable assumption for most manufactured products in a semi-industrialized country.

- *Producers are heterogeneous in terms of their marginal production costs and the foreign demand schedules they face for their products*. Heterogeneity is certainly present in the data, and it is a potentially important explanation for unpredictable export supply responses.

- *Future realizations on the exchange rates, marginal costs, and foreign demand shifters are unknown, but each evolves according to a known Markov process.* Uncertainty matters in sunk-cost models (e.g., Dixit, 1989), so it is critical to allow for it.

- *Plants are risk-neutral and maximize the expected discounted sum of real profits.* This standard assumption eliminates the need to deal with parameters that measure risk aversion.

- *Marginal costs do not respond to output shocks.* This assumption implies that shocks that shift the domestic demand schedule do not affect the optimal level of exports, so it allows us to focus on the export market only. The assumption appears to be reasonable for the industry, country, and time period we will study, since some excess capacity was present.[4]

2.1 Static optimization problem

Let us now specify the plant-level demand and cost functions, and derive the associated

[3] Because the decision to export is a discrete choice, optimal behavior cannot be characterized using Euler equations. Hence integration cannot be avoided.

[4] Estimates of average variable cost functions revealed little dependence on within-plant temporal output fluctuations.

operating profit function.[5] First, conditioned on the real peso/dollar exchange rate (e_t) and a firm-specific stochastic intercept (x_{it}^1), foreign demand is isoelastic, so in logarithms we may write:

$$q_{it} = -\eta_i p_{it} + \psi_1 e_t + x_{it}^1, \tag{1}$$

where q_{it} is the log of the i^{th} firm's exports in year t, $\eta_i > 1$ is the elasticity of foreign demand it faces, p_{it} is the log of the firm's real price of exports in domestic currency, and the random variable x_{it}^1 captures factors that shift the firm-specific demand curve like foreign income fluctuations and changes in the prices of competing products.[6]

Next, let the logarithm of marginal costs be specified as:

$$c_{it} = \psi_2 e_t + x_{it}^2/(1-\eta_i), \tag{2}$$

where x_{it}^2 is a random variable that evolves over time with shocks to real factor prices and technology. (Since $\eta_i > 1$, an increase in x_{it}^2 reduces costs.) The exchange rate appears in our marginal cost expression because it affects imported input prices and real wages. Note that this marginal cost function does not depend upon output levels, so domestic demand shocks do not affect exporting behavior.

Although firms face entry costs when they initiate foreign sales, we assume that they can

---

[5] The static optimization problem developed here is based on Roberts, Sullivan and Tybout (1995) and Clerides, Lach and Tybout (1998).

[6]This demand equation can be derived from a generalized Dixit and Stiglitz (1977) utility function for differentiated products under the assumption that each firm is atomistic in the foreign market (e.g., Bernard, et al, 1999). The coefficient on the exchange rate (ψ) is distinct from the own elasticity of demand (η_i) because exchange rate movements change the foreign currency price of *all* Colombian exports, while changes in *p* only change the price of the individual firm's product.

change the *volume* of their exports without incurring adjustment costs thereafter. Hence the static profit maximization condition yields the optimal q_{it} choice among exporting firms:

$$q_{it}^{*} = \eta_i log\left[1 - \frac{1}{\eta_i}\right] + x_{it}^{1} - \left[\frac{\eta}{1-\eta_i}\right]x_{it}^{2} + \left(\psi_1 - \eta_i \psi_2\right)e_t. \quad (3)$$

Finally, using (1), (2) and (3), the logarithm of current operating profits from exports is linear in x_{it}^1, x_{it}^2 and e_t:

$$\ln\left[\pi(x_{it}^1,\ x_{it}^2,\ e_t)\right] = \pi_{0i} + \left(\psi_1 + (1-\eta_i)\psi_2\right)e_t + x_{it}^1 + x_{it}^2 , \quad (4)$$

where $\pi_{0i} = \left[\frac{\eta_i - 1}{\eta_i} - \ln\left(\frac{\eta_i - 1}{\eta_i}\right)\right]$.

2.2 Dynamic optimization

If there were no sunk costs associated with becoming an exporter, firms would sell abroad whenever their marginal costs were low enough, or foreign demand and the exchange rate were high enough to generate operating profits in excess of the fixed costs, Γ_F , associated with maintaining a presence in foreign markets: $\pi_i(x_{it}^1, x_{it}^2,\ e_t) > \Gamma_F$. (Fixed costs include dealing with customs, minimum freight and insurance charges, and monitoring foreign standards.) But those firms that are not already exporting must pay an additional start-up cost, Γ_S , to establish distribution channels, learn bureaucratic procedures, and adapt their products and packaging for foreign markets. Defining the binary variable y_t to take a value of one during periods when the firm exports and zero otherwise, and assuming that sunk costs are borne in the first year of exporting, net current profits from exporting, $u(\cdot)$, may be written as:

$$u(\cdot) = \begin{cases} \pi_i(x_{it}^1, x_{it}^2, e_t) - \Gamma_F + \epsilon_{1it} & \text{if } y_{it} = 1 \text{ and } y_{it-1} = 1 \\ \pi_i(x_{it}^1, x_{it}^2, e_t) - \Gamma_F - \Gamma_S + \epsilon_{2it} & \text{if } y_{it} = 1 \text{ and } y_{it-1} = 0 \\ 0 & \text{if } y_{it} = 0 \end{cases} \tag{5}$$

Here we have added the noise components ε_{1it} and ε_{2it} to allow for unobserved transitory shocks that induce managers to deviate from expected profit-maximizing behavior. These variables also pick up transitory variation in the fixed and sunk costs of exporting.

One might object that (5) oversimplifies behavior because it implies that firms completely lose their investment in start-up costs if they are absent from the export market for a single year. However, earlier studies suggest that these investments depreciate very quickly, and that firms which most recently exported two years ago must pay nearly as much to re-enter foreign markets as firms that never exported (Roberts and Tybout, 1997a; Bernard and Jensen, 1996). In light of these findings, and given that more general representations make structural estimation intractable, we consider (5) to be a reasonable abstraction.

To characterize the decision to export, let us write the complete parameter vector of interest as $\theta = (\Lambda, \Gamma_F, \Gamma_S, \sigma_{\varepsilon 1}, \sigma_{\varepsilon 2})$, where $\sigma_{\varepsilon 1}^2 = \text{var}(\varepsilon_{it1})$, $\sigma_{\varepsilon 2}^2 = \text{var}(\varepsilon_{it2})$ and Λ is the vector of parameters that govern the evolution of $\pi_i(e_t, x_{it})$. Further, let us suppose that e_t and the elements of the vector $x_{it} = (x_{it}^1, x_{it}^2)$ each follow first-order Markov processes. Then, suppressing *i* subscripts, we assume that at time *t* each plant chooses the sequence of decision rules over its planning horizon of H years. $Y = \{y_\tau = f_\tau(e_\tau, x_\tau, y_{\tau-1}, \epsilon_\tau, \theta)\}_{\tau=t}^{t+H}$ that maximizes the expected discounted sum of its net profits:

$$\max_{Y} \quad E_t\{\sum_{\tau=t}^{t+H} \delta^{\tau-t} u(e_\tau, x_\tau, y_{\tau-1}, y_\tau, \epsilon_\tau, \theta)\} \tag{6}$$

Here E_t is the expectation operator conditioned on information available at time t, and δ is the discount rate, $0 < \delta < 1$.

This expression is the value function $V_t(e_t, x_t, y_{t-1}, \epsilon_t, \theta)$, which is the unique solution to the Bellman equation:

$$V_t(e_t, x_t, y_{t-1}, \epsilon_t, \theta) = \max_{y_t \in \{0,1\}} [u(e_t, x_t, y_{t-1}, y_t, \epsilon_t, \theta) + \delta EV_t(e_t, x_t, y_t, \epsilon_t, \theta)], \tag{7}$$

where

$$EV_t(e_t, x_t, y_t, \epsilon_t, \theta) = \int_{e_{t+1}} \int_{x_{t+1}} \int_{\epsilon_{t+1}} V_{t+1}(e_{t+1}, x_{t+1}, y_t, \epsilon_{t+1}, \theta) dF(e_{t+1}, x_{t+1}, \epsilon_{t+1} | x_t, \epsilon_t). \tag{8}$$

Under quite weak regularity conditions (Rust, 1995), there exists a unique $V_t(\cdot)$ that solves (7) each period, and the optimal exporting status is given by:

$$f_t(e_t, x_t, y_{t-1}, \epsilon_t, \theta) = \operatorname*{argmax}_{y_t \in \{0,1\}} [u(e_t, x_t, y_{t-1}, y_t, \epsilon_t, \theta) + \delta EV_t(e_t, x_t, y_t, \epsilon_t, \theta)] \tag{9}$$

At any point in time, the i^{th} firm's behavior is determined by its own vector of realizations on the state variables $(e_t, x_{it}, y_{it-1}, \epsilon_{it})$. Using this information, one can calculate *total* export volumes by determining which firms are exporting (from equation 9) and how much each exporter supplies (from equation 3). One can also simulate counterfactual export trajectories by changing the characteristics of the exogenous $\boldsymbol{x}_{it}$ processes or by changing threshold costs. However, each of these exercises requires knowledge of the parameter vector θ. We now turn our attention to estimating these unknowns.

III. Econometric Issues

Taking firms' export revenue trajectories and exogenous characteristics as data, we wish to use the structure described in section II to obtain maximum likelihood estimates of the parameter vector $\theta = (\Lambda, \Gamma_F, \Gamma_S, \sigma_{\varepsilon 1}, \sigma_{\varepsilon 2})$. This is a challenging exercise for several reasons. First, we do not have plant-level information on firms' demand shifters, marginal cost shifters, or prices. Hence we must treat (x_{it}^1, x_{it}^2) as serially correlated, unobserved shocks. This problem is particularly severe for plants that do not always export. For these, we must infer the joint density of marginal cost and demand shifters in non-exporting years from export revenues earned by the same plant in other years. Or, if the plant never exported, inferences must be based solely on other plants in the sample.

Second, because exporting decisions are discrete, there are no closed-form solutions or first-order conditions for the decision rule (9). Hence, to estimate the parameters that characterize dynamic behavior, $(\Gamma_F, \Gamma_S, \sigma_{\varepsilon 1}, \sigma_{\varepsilon 2})$, one must solve the dynamic programming model numerically for each candidate vector and choose the one that yields the best match between predicted and observed y_t trajectories.

Finally, given the number of parameters involved in estimation and the need for several layers of numerical integration, it is not feasible to obtain estimates of all the parameters from a simple maximum likelihood estimator. To handle the complexity we proceed in two stages. First, before solving the dynamic programming model, we estimate the parameters Λ that govern the evolution of $\pi_i(e_t, x_{it})$. Under the assumption that shocks are jointly normal, these parameters imply the density function for the unobserved $\boldsymbol{x_{it}}$ realizations, conditioned on the exchange rate, exogenous firm characteristics, and firms' export revenues. Second, taking this density function as given we construct the likelihood function for observed export market participation patterns and maximize this function over Γ_S, Γ_F, $\sigma_{\varepsilon 1}$ and $\sigma_{\varepsilon 2}$. This approach limits the dimensionality of the maximization problem at the stage which is computationally intensive.[7]

To implement this multi-stage strategy we invoke Rust's (1988) conditional independence assumptions. Specifically, we assume that innovations in the state variables (e_t, x_{it}) are independent of the errors $\varepsilon_{it} = (\varepsilon_{1it}, \varepsilon_{2it})$. This assumption seems reasonable in our context, and it allows us to estimate Λ while ignoring ε_{it}. Further, we assume that ε_{it} is serially uncorrelated so that $EV_t(e_t, x_{it}, y_{it}, \varepsilon_{it}, \theta) = EV_t(e_t, x_{it}, y_{it}, \theta)$.[8]

3.1 First Stage Estimation

To estimate Λ, we first must be specific about the process that generates x_{it}. We posit (and will later test) that each element of x_{it} follows a trend stationary AR1 process:

[7] For details on the computational burden of such models, see Rust (1995).

[8] See Rust (1988) for the proof.

$$x_{it}^{1} = \lambda_{0i} + \lambda_1 x_{it-1}^{1} + \lambda_2 t + \nu_{it} \qquad (10a)$$

$$x_{it}^{2} = \rho_{0i} + \rho_1 x_{it-1}^{2} + \rho_2 t + \omega_{it} \qquad (10b)$$

$$e_t = \gamma_0 + \gamma_1 e_{t-1} + \gamma_2 t + \xi_t \qquad (10c)$$

where all three disturbances are white noise. By construction, ξ is orthogonal to ν and ω; without loss of generality we assume that these latter two disturbances are orthogonal to one another. (If they were not, we could re-state the model in terms of orthogonal shocks by using a Cholesky decomposition of their covariance matrix.) Note that we allow the intercepts of both the demand and marginal cost shifters to be plant specific.

If all three series were observable for each plant it would be a simple matter to estimate the parameters of the three processes and test these assumptions. However, although the real exchange rate is readily observed, the other variables can only be crudely approximated with the available data. Hence, we must recast the model in terms of observables and impute the processes described by (10a) and (10b) indirectly.

To begin, note that our assumption of a constant demand elasticity for each firm implies that the ratio of export profits to export revenues is $(\eta_i)^{-1}$. Further, if demand elasticities for a particular plant are the same at home and abroad, this factor of proportionality is also one minus the ratio of total variable costs to total revenue. Hence, assuming that true variable costs are measured up to a zero-mean error by labor and intermediate input costs, we can easily estimate $(\eta_i)^{-1}$ as a firm-specific average of the right-hand-side of this expression across observations. $(\hat{\eta}_i)^{-1} = \frac{1}{T}\sum_{t=1}^{T}\left(1 - \frac{TVC_{it}}{TR_{it}}\right)$.[9]

After converting export revenues to export profits, we use equation (5) and the de-trended, zero mean shifters $\tilde{x}_{it}^1 = x_{it}^1 - (\lambda_{0i} + \lambda_2 t)(1-\lambda_1 L)^{-1}$ and $\tilde{x}_{it}^2 = x_{it}^2 - (\rho_{0i} + \rho_2 t)(1-\rho_1 L)^{-1}$, to re-write equation (4) as:

$$\ln(\pi_{it}) = \beta_{0i} + \beta_1 t + \beta_2 e_{it} + (\tilde{x}_{it}^1 + \tilde{x}_{it}^2) \tag{11}$$

where $\beta_{0i} = \pi_{0i} + \lambda_{0i}(1-\lambda_1)^{-1} + \rho_{0i}(1-\rho_1)^{-1} - \lambda_1\lambda_2(1-\lambda_2)^{-2} - \rho_1\rho_2(1-\rho_2)^{-2}$, $\beta_1 = \lambda_2(1-\lambda_1)^{-1} + \rho_2(1-\rho_1)^{-1}$, $\beta_2 = \psi_1 + (1-\eta^f)\psi_2$, and we control for cross-plant variation in β_{0i} by including a set of dummy variables that distinguish plants according to initial size and four-digit industrial classification. Initial size is measured by initial sales in the domestic market, which serves as a proxy for both marginal costs and product desirability.

Given (10a) and (10b), the elements of the compound disturbance $(\tilde{x}_{it}^1 + \tilde{x}_{it}^2)$ evolve according to: $\tilde{x}_{it} = M\tilde{x}_{it-1} + \upsilon_{it}$, where $\tilde{x}_{it} = (\tilde{x}_{it}^1, \tilde{x}_{it}^2)'$, $E(\upsilon_{it}) = 0$, $E(\upsilon_{it}\upsilon_{it+j}') = 0\ \forall j \neq 0$, $E(\upsilon_{it}\upsilon_{it}') \equiv Q = \begin{bmatrix}\sigma_\nu^2 & 0\\ 0 & \sigma_\omega^2\end{bmatrix}$, and $M = \begin{bmatrix}\lambda_1 & 0\\ 0 & \rho_1\end{bmatrix}$. Hence, under the assumption that each firm begins with a random draw from the steady state distribution of $\tilde{x}_{it}$, $E[\tilde{x}_{it}^1 + \tilde{x}_{it}^2] = 0$ and the covariance matrix for the disturbances in equation (11) is defined by and

$E\left[(\tilde{x}_{it}^1 + \tilde{x}_{it}^2)(\tilde{x}_{it-k}^1 + \tilde{x}_{it-k}^2)\right] = \ell M^{|k|}\Gamma_0 \ell'$, where $\ell = [1,1]$ and $\Gamma_0 = Q(I - M^2)^{-1}$ (e..g., Chow,

[9]Discrepancies between estimated and true elasticities will induce measurement error in profits, and will be reflected in $\varepsilon_{1it}, \varepsilon_{2it}$.

1983). Assuming that all cross-plant correlation in operating export profits is absorbed by the exchange rate and trend term, we use these results to obtain maximum likelihood estimates of the parameter vectors $(\beta_{0i},\beta_1,\beta_2)$, (λ_1,ρ_1) and $(\sigma_v^2,\sigma_\omega^2)$ using maximum likelihood, and thereby to characterize future realizations on exporting profits for any firm, given e_t and its current realization on $(\tilde{x}_{it}^1,\tilde{x}_{it}^2)$.

We fit equation (11) to the unbalanced panel of observations for which exports are positive. But firms self-select into the export market partly on the basis of innovations in the state variables, $(\tilde{x}_{it}^1,\tilde{x}_{it}^2)$, so we need to correct for selectivity bias. Ideally one would combine equation (11) with the likelihood function for export market participation patterns and estimate all parameters in a single stage. Unfortunately, this involves optimizing over too many parameters at the computationally intensive stage of estimation. We therefore crudely control for selection bias by including a Mills ratio in (11). This ratio is based on a simple reduced form Probit equation that explains each plant's probability of exporting as a function of its strictly exogenous characteristics—location, business type, and initial capital stock.

Given estimates of the parameters $(\beta_{0i},\beta_1,\beta_2)$, (λ_1,ρ_1) and $(\sigma_v^2,\sigma_\omega^2)$, we next construct the conditional distributions for $(\tilde{x}_{it}^1,\tilde{x}_{it}^2)$ trajectories that we will need in our second stage estimation. Suppose the i^{th} plant exports in all T years. Then we can infer the vector of (logged) residual export profits, $\pi_i^* = [(\tilde{x}_{i1}^1+\tilde{x}_{i1}^2),(\tilde{x}_{i2}^1+\tilde{x}_{i2}^2),\ldots,(\tilde{x}_{iT}^1+\tilde{x}_{iT}^2)]'$, using our first stage results: $\pi_{it}^* = \tilde{x}_{it}^1+\tilde{x}_{it}^2 = \ln(\pi_{it})-\beta_{0i}-\beta_1 t-\beta_2 e_t$. Further, under the assumption that $(\tilde{x}_{it}^1,\tilde{x}_{it}^2)$ is normally distributed with zero mean, the distribution of the $(\tilde{x}_{it}^1,\tilde{x}_{it}^2)$ trajectory, $\tilde{x}_i = [(\tilde{x}_{i1}^1,\tilde{x}_{i1}^2)',(\tilde{x}_{i2}^1,\tilde{x}_{i2}^2)',\cdots,(\tilde{x}_{iT}^1,\tilde{x}_{iT}^2)']'$, conditioned on the residual export profit trajectory, π_i^*, is $N\left(\Sigma_{12}\Sigma_{22}^{-1}\pi_i^*,\ \Sigma_{11}-\Sigma_{12}\Sigma_{22}^{-1}\Sigma_{12}'\right)$, where $E(\tilde{x}_{it}\tilde{x}_{it+s}') = M^{|s|}\Gamma_0$, $E(\tilde{x}_{it}\pi_{it+s}^*) = M^{|s|}\Gamma_0\ell'$

and $E(\pi^*_{it}\pi^*_{it+s}) = \ell M^{|s|}\Gamma_0\ell'$ define the elements of Σ_{11}, Σ_{12} and Σ_{22}, respectively. Suppressing parameters, we will hereafter denote this density $g(\tilde{x}_i|\pi^*_i)$.

A similar logic yields $g(\tilde{x}_i|\pi^*_i)$ for plants that don't export in all years. When elements of the π^*_i vector are unobserved, we simply delete the columns of Σ_{12} and Σ_{22} that correspond to non-exporting years. The resulting density for $\tilde{x}_i$ uses information on a plant's residual exporting profits in the years when they are observed to extrapolate to those years when they are not, based on (10a) and (10b). When plants *never* export, E($\tilde{x}_i$) = 0 and $E(x_i x_i') = \Sigma_{11}$. This is the only case in which $E(x_i x_i')$ is full rank; otherwise the constraint $\tilde{x}^1_{it} + \tilde{x}^2_{it} = \pi^*_{it}$ applies for at least some years.

3.2 Estimating sunk and fixed cost parameters

Once $g(\tilde{x}_i|\pi^*_i)$ is known for all plants, the remaining parameters $(\Gamma_F, \Gamma_S, \sigma_{\varepsilon 1}, \sigma_{\varepsilon 2})$ can be estimated by maximizing the sample likelihood function:

$$L = \prod_{i=1}^{N} \int_{\tilde{x}_i} [\prod_{t=1}^{T} P(y_{it}=0|\beta_{i0}, e_t, t, \tilde{x}_{it}, y_{it-1}=0)^{(1-d_{it})(1-d_{it-1})} P(y_{it}=0|\beta_{i0}, e_t, t, \tilde{x}_{it}, y_{it-1}=1)^{(1-d_{it})(d_{it-1})}$$
$$\cdot P(y_{it}=1|\beta_{i0}, e_t, t, \tilde{x}_{it}, y_{it-1}=1)^{d_{it}d_{it-1}} P(y_{it}=1|\beta_{i0}, e_t, t, \tilde{x}_{it}, y_{it-1}=0)^{d_{it}(1-d_{it-1})}] g(\tilde{x}_i|\pi^*_i)\, d\tilde{x}_i . \qquad (12)$$

where d_{it} is a dummy that takes a value of 1 if the i^{th} firm exports in year t, and a value of zero if it does not. (Parameters common to all plants are suppressed.)[10]

[10]As we mentioned above, the best way to deal with selectivity bias when estimating Λ would be to abandon our multi-stage approach and maximize this likelihood function over the entire parameter vector $\theta = (\Lambda, \Gamma_F, \Gamma_S, \sigma_{\varepsilon 1}, \sigma_{\varepsilon 2})$. As we improve the efficiency of our estimation routines and machines become more powerful, this may become feasible. It is worth noting, however, that if the selectivity problem shows

To integrate over the $\tilde{x}_{i\cdot}$ trajectories, we draw repeatedly from the $g(\tilde{x}_i|\pi_i^*)$ distribution and sum the associated values of the expression in square brackets weighted by $g(\tilde{x}_i|\pi_i^*)$. (The weights are normalized so that for each plant they sum to unity.) We use antethetic sampling to limit the sampling error, and we increase the number of draws on trajectories until the estimates are no longer sensitive to the sampling outcomes.

Calculating the probabilities that appear in the integrand is more cumbersome. To do so we must find the decision rule (9) for each firm at each point in time. In turn, this means we must specify the distribution of the unobserved vector ε_t, and use our information on transition probabilities for the vector of state variables from estimates of equation (10). Let V_{0t}, V_{10t} and V_{11t} respectively be the value from not exporting in period t, from beginning to export in period t after not exporting in period t-1, and from continuing to export in period t after exporting in period t - 1, each exclusive of transitory noise (ε_{it}):

$$\begin{aligned} V_{11it} &= \pi(\beta_{0i}, e_t, t, \tilde{x}_{it}) - \Gamma_F + \delta EV_t(\beta_{0i}, e_t, \tilde{x}_{it}, 1, \theta) \\ V_{10it} &= \pi(\beta_{0i}, e_t, t, \tilde{x}_{it}) - \Gamma_F - \Gamma_S + \delta EV_t(\beta_{0i}, e_t, \tilde{x}_{it}, 1, \theta) \\ V_{0it} &= \delta EV_t(\beta_{0i}, e_t, \tilde{x}_{it}, 0, \theta). \end{aligned} \tag{13}$$

Then suppressing plant (i) subscripts, the probabilities of observing the different exporting states are:

---

up mainly as a biased estimator for the intercept in equations (10a) and (10b), we effectively undo this bias by estimating firms' fixed costs (which shift the profit function) in the second stage.

(14)

$$P[y_t=1|y_{t-1}=0] = P[V_{10t}+\epsilon_{2t} > V_{0t}] = \Phi\left(\frac{V_{10t}-V_{0t}}{\sigma_{\epsilon 2}}\right)$$

$$P[y_t=1|y_{t-1}=1] = P[V_{11t}+\epsilon_{1t} > V_{0t}] = \Phi\left(\frac{V_{11t}-V_{0t}}{\sigma_{\epsilon 1}}\right)$$

$$P[y_t=0|y_{t-1}=0] = P[V_{10t}+\epsilon_{2t} > V_{0t}] = 1 - \Phi\left(\frac{V_{10t}-V_{0t}}{\sigma_{\epsilon 2}}\right)$$

$$P[y_t=0|y_{t-1}=1] = P[V_{11t}+\epsilon_{1t} > V_{0t}] = 1 - \Phi\left(\frac{V_{11t}-V_{0t}}{\sigma_{\epsilon 1}}\right)$$

Clearly, once V_{0t}, V_{10t} and V_{11t} are calculated up to the unknown parameters, maximum likelihood estimation becomes straightforward. However, these expressions are difficult to calculate because they involve the expected value of the period *t+1* value function conditioned on period *t* information. We begin by writing this expectation as:

$$EV_t(e_t, x_t, y_t, \theta) = y_t E_t \max(V_{11t+1}+\epsilon_{1t+1}, V_{0t+1}) + (1-y_t) E_t \max(V_{10t+1}+\epsilon_{2t+1}, V_{0t+1}) \qquad (15)$$

where:

$$E_t \max(V_{11t+1}+\epsilon_{1t+1}, V_{0t+1}) = P[\epsilon_{1t+1} > V_{0t+1} - V_{11t+1}] \cdot [V_{11t+1} + E(\epsilon_{1t+1} | \epsilon_{1t+1} > V_{0t+1} - V_{11t+1})] + P[\epsilon_{1t+1} < V_{0t+1} - V_{11t+1}] \cdot V_{0t+1}$$

$$E_t \max(V_{10t+1}+\epsilon_{2t+1}, V_{0t+1}) = P[\epsilon_{2t+1} > V_{0t+1} - V_{10t+1}] \cdot [V_{10t+1} + E(\epsilon_{2t+1} | \epsilon_{2t+1} > V_{0t+1} - V_{10t+1})] + P[\epsilon_{2t+1} < V_{0t+1} - V_{10t+1}] \cdot V_{0t+1}$$

Then, under the assumption that (ε_{1t}, ε_{2t}) are jointly normally distributed, the conditional expectations above can be expressed as Mills ratios, and the probabilities (conditioned on x_{t+1}) can be obtained from the standard normal distribution function, $\Phi(\cdot)$:

(16*a*) $$E_t \max(V_{11t+1} + \epsilon_{1t+1}, V_{0t+1}) =$$

$$\int_{x_{t+1}} \int_{e_{t+1}} [\Phi\left(\frac{V_{11t+1} - V_{0t+1}}{\sigma_{\epsilon 1}}\right) \cdot \left[V_{11t+1} + \sigma_{\epsilon 1} \phi\left(\frac{V_{0t+1} - V_{11t+1}}{\sigma_{\epsilon 1}}\right) \Big/ \Phi\left(\frac{V_{11t+1} - V_{0t+1}}{\sigma_{\epsilon 1}}\right)\right]$$

$$+ \Phi\left(\frac{V_{0t+1} - V_{11t+1}}{\sigma_{\epsilon 1}}\right) \cdot V_{0t+1}] \, dF(x_{t+1}, e_{t+1} | x_t, e_t)$$

(16*b*) $$E_t \max(V_{10t+1} + \epsilon_{2t+1}, V_{0t+1}) =$$

$$\int_{\tilde{x}_{t+1}} \int_{e_{t+1}} [\ \Phi\left(\frac{V_{10t+1} - V_{0t+1}}{\sigma_{\epsilon 1}}\right) \cdot \left[V_{10t+1} + \sigma_{\epsilon 2} \phi\left(\frac{V_{0t+1} - V_{10t+1}}{\sigma_{\epsilon 2}}\right) \Big/ \Phi\left(\frac{V_{10t+1} - V_{0t+1}}{\sigma_{\epsilon 2}}\right)\right]$$

$$+ \Phi\left(\frac{V_{0t+1} - V_{10t+1}}{\sigma_{\epsilon 2}}\right) \cdot V_{0t+1}] \, dF(\tilde{x}_{t+1}, e_{t+1} | \tilde{x}_t, e_t)$$

The next task is to calculate $EV_{t+1}(e_{t+1}, \tilde{x}_{t+1}, y_{t+1}, \theta)$ so that we can evaluate expressions involving *V*'s. We do so using the backward induction algorithm described by Rust (1995). Specifically, we assume that firms have a finite planning horizon of *H* years, so that $T = t + H$ is the (distant) terminal period.[11] Then, in the terminal year there are no future periods to consider so $EV_T(e_T, \tilde{x}_T, y_T, \theta)$ is set to zero in (13), and each firm's exporting decision maximizes current

[11]So long as *H* is large and the discount rate is positive the resulting solution approximates the solution to the infinite horizon problem.

payoffs, $u(e_T, \tilde{x}_T, \varepsilon_T, y_T, y_{T-1}|\theta)$. Accordingly, for period *T-1,* the expected value function is simply:

$$EV_{T-1}(e_{T-1}, \tilde{x}_{T-1}, y_{T-1}, \theta) = \quad E_{T-1}\Big[\max_{y_T} u(e_T, \tilde{x}_T, y_T, \varepsilon_T | y_{T-1}, \theta)\Big]$$

where expectations are taken over both ε_T and $\tilde{x}_T$ conditioned on $\tilde{x}_{T-1}$, as in (16). Once the expected value function for period *T-1* has been calculated for each possible realization on $\tilde{x}_{T-1}$, the expected value function for period *T-2* can be calculated for each possible realization on $\tilde{x}_{T-2}$:

$$EV_{T-2}(e_{T-2}, \tilde{x}_{T-2}, y_{T-2}, \theta) = E_{T-2}\Big[\max_{y_{T-1}} \big\{u(e_{T-1}, \tilde{x}_{T-1}, y_{T-1}, \varepsilon_{T-1} | y_{T-2}, \theta) + \delta EV_{T-1}(e_{T-1}, \tilde{x}_{T-1}, y_{T-1}, \theta)\big\}\Big]$$

Clearly this calculation can be repeated, backing up one year at a time, until one reaches period *t*. This generates the values needed for V_{10t}, V_{11t}, and V_{0t} which in turn enter the likelihood function.

Each step of the backward induction—and the evaluation of (16a) and (16b)—involves taking expected values over $(e_{t+1}, \tilde{x}_{t+1})$ realizations, given $(e_t, \tilde{x}_t)$. One approach is to discretize the vector $(e_t, \tilde{x}_t)$ and calculate transition probabilities among each vector of values using the estimates of equations 10a-c (Rust, 1988, Das, 1992). The problem with this approach is that it involves a large number of calculations. With 3 state variables and, say, *r* different values per state variable, there are r^3 points in the state space and r^6 transition probabilities. With a reasonably long planning horizon (*T*), approximation errors compound, and it becomes necessary to use very large *r* values. This is the "curse of dimensionality" that made a supercomputer necessary for estimation of Das's (1992) model, which involved 5 state variables.

Recently Rust (1997) has developed an alternative approach that substantially reduces the computational burden. Unlike the method described above, Rust's technique does not treat all possible combinations of values for the discretized state variables as elements of the state space. Rather, the state space is generated with random draws from the joint probability distribution for the vector of state variables, yielding a random Bellman operator. This solves the dimensionality problem because increasing the number of state variables increases the dimension of each point, but it does not increase the number of points over which one integrates.

We adopt Rust's (1997) approach in the present study.[12] First we assume that all $(e_t, \tilde{x}_{it})$ realizations lie in the observed region of support. This ensures boundedness of $(e_t, \tilde{x}_{it})$. Truncating the normal imposes very weak restrictions because almost the entire area under the normal curve lies in this range. Second, we make G uniform draws from within the bounds of $(e_t, \tilde{x}_{it})$.[13] In our case this implies G values of every V_{10}, V_{11}, and V_0. Finally, the transition probabilities $dF\left(e_{t+1}, \tilde{x}_{it+1} | e_t, \tilde{x}_{it}\right)$ in the discrete analog to are normalized to ensure that the random Bellman operator is a contraction mapping. In particular, for each $(e_t, \tilde{x}_{it})$, the imputed probabilities of moving to the various possible $(e_{t+1}, \tilde{x}_{it+1})$ values are scaled to sum to unity.

[12]Rust (1997) also suggests increasing the fineness of the grid for integration over realizations that occur in the new (as opposed to distant) future. We have not yet implemented that aspect of his technique.

[13]For these draws we use the homotopy method to increase the precision of our integrals.

IV. Findings: The Colombian Chemicals Industry

4.1 Descriptive overview of the data

Although our framework should describe any industry in which exporting is potentially profitable for some firms, it will be easiest to identify parameters in those industries which have many exporters, and which exhibit substantial variation in the set of exporters over time. For these reasons, we choose to estimate the model using data on the Colombian chemicals industry for the period 1982 through 1991, which is summarized in table 1 below.

Note that the Colombian peso depreciated substantially in real terms during the sample period, and that chemical exports simultaneously grew. The expansion was partly due to an increase in the number of exporters, and partly due to increases in the magnitude of foreign sales at the typical exporting plant.[14] Colombian chemicals plants produced 35.00 billion pesos worth of exports in 1991, of which 29.94 billion came from plants that were exporting in 1984. So entry by new exporters contributed 5.06 out of the 27.10 expansion. Also, of the 62 plants that existed during the entire sample period, 18 exported in all ten years, 26 never exported, and 18 switched exporting status at least once. So, although there were a number of switches, the data exhibit substantial persistence. This could be due to serial correlation in the plant-specific state variables, $(\tilde{x}_{it}^{1}, \tilde{x}_{it}^{2})$, or it could be due to sunk entry costs, Γ_s, or some combination of both. Our estimates will shed light on the relative importance of these different forces.

[14]The number of Chemical plants remains fixed at 62 during the sample because we have excluded producers who enter or exit to simplify the econometrics, so there is some potential for selectivity bias.

Table 1: Colombian Chemical Producers, Exporters versus Non-exporters*

Year	*Peso Value of Exports* [a]	*Dollar Value of Exports* [b]	*Number of Exporters*	*Number of Entrants*	*Number of Quitters*	*Real Exchange Rate*
1982	6.18	41.71	25	1	1	79.5
1983	8.60	58.05	30	6	1	80.5
1984	7.90	55.32	28	1	3	89.8
1985	11.79	79.58	25	3	6	102.2
1986	14.10	95.17	24	1	2	113.6
1987	15.40	103.95	23	1	2	113.7
1988	21.97	148.30	28	6	1	112.3
1989	20.62	139.19	27	2	3	115.2
1990	27.10	182.93	28	1	0	127.2
1991	35.00	236.25	30	2	0	121.1
Average	16.866	114.045	26.8	2.4	1.9	105.51

*Data describe the 62 Colombian producers of industrial chemicals continually observed over the period 1981-91.

a Billions of 1986 pesos (deflation done using manufacturing-wide wholesale price deflator).

[b] Millions of 1999 dollars (conversion done at the official 1986 exchange rate and brought forward using the U.S. wholesale price deflator).

4.2 First Stage Estimation: Evolution of the State Variables

To estimate the profit function (10), we first need firm-specific estimates of the elasticity of demand to convert export revenues to export profits. Using $(\hat{\eta}_i)^{-1} = \frac{1}{T}\sum_{t=1}^{T}\left(1 - \frac{TVC_{it}}{TR_{it}}\right)$, we obtain the results summarized in figure 1. Note that the magnitudes are generally sensible, and that they imply relatively minor cross-plant variation in π_{i0}, so to simplify our estimator we hereafter ignore this source of heterogeneity.[15] Pooling plants, the average elasticity implies that for each peso of export revenue, a plant typically nets about 28 centavos of operating profits.[16] Figure 2 shows the distribution of average export profits we obtain when we divide plant-specific export revenues by plant-specific demand elasticities. (Plants are sorted from lowest to highest average profit; to roughly convert millions of 1986 pesos to thousands of 1999 dollars, multiply by 6.75.)

Our next task is to estimate $(\beta_{0i}, \beta_1, \beta_2)$, (λ_1, ρ_1) and $(\sigma_\nu^2, \sigma_\omega^2)$ using the techniques described in section 3. As discussed in Appendix 1, specification tests confirm that the compound disturbance is reasonably characterized as be the sum of two AR(1) processes. Accordingly we proceed to our maximum likelihood estimates in table 2 below. Note first that the intercept and dummies describe time-invariant heterogeneity in operating profits that is attributable to domestic sales volume in the initial year. The estimates imply that this type of variation, which proxies for persistent cross-plant differences in market size and marginal production costs, does not appear to be particularly important. Size may nonetheless affect the *fixed* costs of exporting or the sunk costs of breaking into foreign markets; we will explore these possibilities in the following section using our second stage estimator.

---

[15] The expression $\pi_{0i} = \left[(\eta_i - 1)/\eta_i - \ln\left((\eta_i - 1)/\eta_i\right)\right]$ varies from __ at $\eta_i =$ ___ to ___ at $\eta_i =$ ___, so although elasticities are important for converting revenues to profits, they do not induce much cross-plant variation in the profit function intercept. Allowing for plant-specific profit shifters (in addition to our size shifter) would slow down our second stage estimator by a factor of 15; hence we ignore this apparently minor source of variation. To the extent that it matters, it will show up as a source of serial correlation in the disturbance terms $\varepsilon_{1it}, \varepsilon_{2it}$.

[16] Goldberg and Knetter (1999) use data aggregated to the seven-digit industry level to estimate much lower export demand elasticities for manufactured products. The discrepancy between their results and ours probably traces largely to the fact that individual plants *within* a narrowly defined product category export close substitutes for each others' products.

Table 3: Operating Profit Function Parameters

	Coefficient	Standard Error
intercept	-10.189	(6.414)
size dummy 2	-2.383	(1.254)
size dummy 3	-1.580	(1.397)
size dummy 4	-3.045	(1.655)
log exchange rate (β_1)	3.948	(1.430)
trend (β_2)	-0.071	(0.083)
Mills ratio	3.129	(0.904)
ρ_1	-0.594	(0.211)
λ_1	0.957	(0.043)
σ_{ω}^{2}	0.307	(0.147)
σ_{ν}^{2}	0.769	(0.041)
sample size	293 observations	
log-likelihood	-311.690	

In contrast, the exchange rate appears to have an important, positive affect on operating profits from exporting. A one percentage point devaluation translates into nearly a four percent increase in operating profits from exports. The trend term, which picks up secular growth or shrinkage in foreign markets and/or marginal production costs, adds little. We caution, however, that both of these coefficients are identified with only ten years of data, so they may not be representative of longer time periods. If we replace the exchange rate the trend term with annual time dummies (not reported), the fit does not improve significantly, so these variables appear to do a reasonable job of controlling for time effects that are common to all plants, and barring spatial correlation, our assumption that the disturbances in equation (10) are not correlated across plants seems reasonable.

The parameters of the covariance matrix for the residuals are identified with plant-specific variation, and are estimated with a good deal of precision. Each root is less than one in absolute value, although we cannot reject the null hypothesis that $\lambda_1 = 1$. Of course, we cannot tell from our estimates which root corresponds to marginal costs, $\tilde{x}_{it}^2$, and which to the demand-shifter, $\tilde{x}_{it}^1$. Most likely these error components are correlated and there is no such correspondence—what we have estimated are the roots of the orthogonal processes that are linear combinations of $\tilde{x}_{it}^2$ and $\tilde{x}_{it}^1$.

Finally, note that the Mills ratio has the expected sign and it is quite significant. This implies, as expected, that realizations on the sum of the two latent variables are more favorable than average for firms that are currently exporting.

We estimate the real exchange rate process separately because it is presumably unaffected by plant-specific marginal cost and demand shocks. Ordinary least squares estimates for the years 1967-1992 suggest a structural break around 1981, so a regime switching model is probably necessary to properly characterize this period. Indeed, an interesting issue to pursue is how expectations of regime switches affect decisions to enter or exit the export market. For now, however, we keep things simple by limiting the sample years to 1982-92, which is the period described by our plant-level data. In that sub-period the trend coefficient is very close to zero and quite insignificant, so we set $\gamma_2 = 0$ and obtain the restricted estimates: $\hat{\gamma}_0 = 1.187$, $\hat{\gamma}_1 = 0.752$, $\hat{\sigma}_\xi^2 = 0.250$. The Dickey-Fuller test statistic for stationarity is -3.13, and the critical value is -2.78 at a 90 percent confidence level. So despite the very small sample size we can reject the null hypothesis of a unit root.

4.3 Second Stage Estimates: Sunk Costs and Fixed Costs

Using our estimates of the mark-up parameter and the parameters related to equations 10a-10c, we can now apply the techniques described in section 3.2 to calculate the value function (8) for each plant at each point in time, up to the parameter vector $(\Gamma_F, \Gamma_S, \sigma_{\varepsilon 1}, \sigma_{\varepsilon 2})$. Hence, given any value for this vector, we can characterize the probability of the observed series of export market participation decisions for each producer using (15) and (16). Finally, searching over alternative values of the unknown parameter vector, we find the set that maximizes the likelihood function (15). Because the likelihood function is not globally concave in θ when $\delta > 0$, we start at δ = 0, where global concavity obtains. Then using these θ estimates as starting values, we do a search over δ on the [0,1) interval, incrementing by .05. Given that fixed or sunk costs may systematically vary across firms with their size, we allow the parameters (Γ_F, Γ_S) to be functions of our size dummies.

The parameter estimates are reported in table 4. These are for only a single replication with g=100 simulations and a random grid of size s=100 and are only for illustrative purposes. Several preliminary observations are in order. First, the likelihood function is very flat with respect to the discount rate, so the data do not pin down this parameter. Interestingly, however, as the discount rate rises, the magnitudes of the coefficients systematically falls. This keeps the effect of sunk costs on persistence in exporting status roughly constant. Second, when we re-draw our grid and our simulated trajectories, the parameter estimates vary as much as 30 percent from trial to trial. (Results from repeated trials are not reported.). Hence although the results are qualitatively unaffected, g =100, s =100 is clearly not enough for precise estimation. (Standard errors have not been corrected for this source of noise, nor for the fact that parameters estimated in stage 1 are subject to sampling error.) We are currently estimating the model with values of g up to 300 and values of s up to 1000 but do not have enough estimates with different sets of simulations or discount rates to be able to draw conclusions about the stability of the parameter estimates.

Table 4: Sunk and Fixed Cost Estimates
(grid size = 100, number of simulations =100)

Parameter	***Parameter estimates*** *(standard errors in parentheses)*		
	$\delta = 0.0$	$\delta = 0.4$	$\delta = 0.8$
$\sigma_{\epsilon 1}$	5.61 (3.76)	11.40 (8.02)	16.57 (9.99)
$\sigma_{\epsilon 2}$	542.2 (2,533)	43.47 (31.62)	27.38* (11.63)
Γ_{s_1}	1,354.7 (6,300)	121.38 (84.27)	81.44* (34.47)
Γ_{s_2}	766.4 (3,565)	70.06 (48.77)	50.56* (22.79)
Γ_{s_3}	830.9 (3,846)	80.30 (52.03)	64.22* (26.97)
Γ_{s_4}	291.6 (334.7)	43.15 (29.73)	36.12* (17.87)
Γ_F	-2.41 (3.08)	-2.53 (3.73)	1.78 (1.49)
log-likelihood	119.52	120.12	120.79

Turning to economic implications of our results, several sensible patterns emerge. First, sunk costs appear to systematically fall with plant sizes. The size categories are sorted from smallest to largest so Γ_{s_1} is the entry cost of a plant in the smallest size category. At every discount rate, sunk entry costs for the smallest size category are more than double those for the largest size category, suggesting that big plants are in a better position to step into international markets. This could be due to existing contacts and distribution channels, the types of products these firms produced, or to the mix of people they employ. Second, our estimates for σ_{ε_1} are systematically less than σ_{ε_2}, as our model suggests they should be. The former includes unexplained variation in net operating profits due to fixed costs and imperfectly estimated demand elasticities, while the latter picks up unexplained variation due to both these factors *and* sunk costs. Third, the ratio of sunk costs to σ_{ε_2} is roughly 2.5, which is quite consistent with our earlier results based on reduced-form representations of the decision to export (Roberts and Tybout, 1997a).

As discussed in the introduction, one advantage of structural estimation is that it allows us to calculate the sunk costs of entering foreign markets in currency units. Table 4 implies that at a discount factor of 0.8, the *expected* sunk costs of breaking into export markets vary from 31 to 86 million 1986 pesos, depending upon which plant size category we are describing. However, realized sunk costs vary widely across plants and time, as implied by the estimated standard deviation: σ_{ε_2} = 27. (The reader may wish to benchmark these figures by comparing them to the export profit distribution depicted in figure 2.) Converting back to 1999 dollars, our figures imply that large plants paying 31 million 1986 pesos to break into the export market are paying 209,250 dollars and plants in the smallest category would pay 580,500 dollars, on average.

What do these sunk costs imply about the role of expectations in shaping export response? From equation (13), the payoff to *non*-exporters from entering foreign markets in period t can be written as:

$$V_{10} - V_0 = \pi(\beta_{0i}, e_t, t, \tilde{x}_{it}) - \Gamma_F + \delta\left(EV_t(\beta_{0i}, e_t, \tilde{x}_{it}, y_{it}=1, \theta) - EV_t(\beta_{0i}, e_t, \tilde{x}_{it}, y_{it}=0, \theta)\right) \qquad \textbf{(17)}$$

where the first two right-hand side terms are operating profits net of fixed costs, and the bracketed final term is the option value of being able to export next year without having to pay sunk costs again. (This option value is bounded by $\delta\Gamma_S$.) As discussed in section 3, non-exporting plants do best to enter export markets at the point in time when $V_{10} - V_0$ exceeds the sunk entry costs they face, $\Gamma_S - \varepsilon_{2it}$. And expectations about future export profits affect this exporting decision by influencing the option value component of $V_{10} - V_0$. (Similarly, incumbent exporters decide whether to drop out of the export market partly on the basis of the option value they forego by doing so.) We can therefore gauge the importance of expectations by using our parameter estimates and raw data to compute the components of equation 17, plant by plant.

Figure 3 depicts current operating profits, $\pi(\beta_{0i}, e_t, t, \tilde{x}_{it}) - \Gamma_F$, and the total return from entering, $V_{10} - V_0$, for the 35 plants that were non-exporters in 1991, sorted in order of increasing total returns. Note the option value for these plants, which is the difference between the height of the two lines, constitutes an important fraction of the total return to entering. Hence factors that affected these firms' expectations about their potential future export profits could substantially affect the number of plants participating in foreign markets.

More generally, figure 3 suggests that the set of potential exporters is quite responsive to policy shocks of reasonable magnitude. With entry costs around 60 to 80 million pesos, it appears that roughly 5 firms would choose to enter in 1991.[17] Further, a one-time entry subsidy on the order of 40 million pesos per plant (270,000 US dollars) would induce about half of the non-exporters to initiate foreign sales.[18] Also, given that the elasticity of export profits with

---

[17]This figure is within the range of entry rates actually observed over the 1984-1990 period–refer to Table 1. Entry costs vary across firms in accordance with their size, and we have ignored variation in the pay-off from exporting due to ε_{2it}, so one can only make rough guesses concerning the entry rate using figure 3.

[18]An ongoing entry subsidy would have a smaller effect on the number of exporters because this policy would reduce the cost of repeated entry and exit. This drives up the exit rate as well as the entry rate.

respect to the exchange rate is roughly 4 (table 3), a 25 percent real devaluation would induce many of the non-exporters to jump in, even if the underlying exchange rate process were perceived as unchanged.

Our estimates also imply a role for history. Figure 4 depicts the probability of being an exporter in 1991 for each plant in the sample, presuming that it exported in 1990, and contrasts this with the probability of being an exporter in 1991 given that it *didn't* export in 1990. (For each plant, one of these probabilities is counter-factual.) Plants are once again sorted in order of increasing value from exporting in the current period.[19] The gap between the two lines, dubbed the "hysteresis band" in the theoretical literature, can be as large as .70, and is substantial for most firms.[20] So for a set of realizations on the exogenous state variables, and a given set of expectations about the future, export market participation patterns may vary dramatically across time or countries, depending upon inherited exporting patterns.

There is one critical qualification to these observations. Among the ten plants with the highest pay-off from being an exporter, the probability of being an exporter is unity *regardless* of whether they were exporting last period. Figure 5, which redoes figure 3 for *all* 62 plants, shows why. These top ten plants expect profits so high that they are certain to jump back in the market immediately if something bumps them out. In fact, operating profits (and export sales) among the top ten plants dwarf the operating profits (and export sales) of the remaining plants.

The important implication for policy is that, while expectations and history matter a great deal to most plants, they are nearly irrelevant to the producers who supply most of the industry's

[19]Here we use the value of staying in export markets for incumbent exporters and the value of *getting* in for others. If all firms faced the same entry costs, each probability would be a monotonic function of total value.

[20]This figure is quite similar to the maximum width of the hysteresis band we calculated in earlier work using a reduced-form approximation to the value function (Roberts and Tybout, 1997).

exports. So policies designed to induce export market participation are likely to work if the objective is to increase the number of exporting plants. But the effect of induced entry on the *volume* of exports is likely to be minor. In the short run, most export response will come from volume adjustments among the large incumbent suppliers.

5. Concluding remarks

Let us summarize what we have learned about the effects of history, heterogeneity and expectations on export dynamics in the Colombian chemicals industry. First, sunk entry costs are substantial, and vary systematically across plants with different levels of output in the domestic market. Accordingly, for many plants expectations and history are important determinants of exporting status. Second, however, cross-plant heterogeneity in demand and marginal costs has created three basic categories of plants: majors exporters, plants that export on a small scale, and non-exporters. Entry and exit moves plants between the second and third categories, so while subsidies for new exporters might well induce net entry, the impact on aggregate chemical exports would probably be quite limited. This fact has not come through in earlier plant-level studies of export hysteresis because those studies have focused on the decision to export rather than on export volumes. Finally, much of the cross-plant and inter-temporal variation in exporting behavior remains unexplained by our model. We have squeezed a good deal of information out of a few key variables, but further progress in modeling export dynamics may require better information on the determinants of plant-specific demand and cost shocks.

The next step in this research project is to use the estimated parameters from both the first and second stage to simulate the effects of changes in the economic environment on the number of exporters and aggregate export flows. For example, changes in the exchange rate regime can be simulated by allowing different parameters in the stochastic process generating the exchange rate. The model can then be used to trace how responsive aggregate export flows are to changes in the exchange rate and how these responses are affected by the level of sunk entry costs. Examination of these issues is underway.

Finally, one qualifier about the generality of results is important. While we believe the model we have developed provides a general framework for measuring and analyzing the sources of export market dynamics, there is no reason to expect that one explanation will apply to all markets. In particular, both sunk entry costs and the degree of micro-level profit heterogeneity will likely vary significantly across industries and, thus, the sources of export dynamics for the Colombian industrial chemicals industry need not be true for other industries. The theoretical literature on export hysteresis was first motivated by puzzling patterns of U.S./Japanese automobile trade in the mid 1980s. That context was very different from Colombian chemicals trade because the producers were each potentially large contributors to total sales, and because entry costs were inflated by the need for distribution networks and advertising. Hence, application of the methodology developed herein to autos and other sectors with similar market structures might yield dramatically different results.

References

Arellano, Manuel and Stephen Bond. 1991. "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations." *Review of Economic Studies*, 58, pp. 277-297.

Baldwin, Richard E. and Paul Krugman. 1989. "Persistent Trade Effects of Large Exchange Rate Changes." *Quarterly Journal of Economics* 104(4), 635-54.

Bernard, Andrew, Jonathan Eaton, J. Bradford Jensen and Samuel Kortum. 1999. "Plants and Productivity in International Trade: A Ricardian Reconciliation." Processed, Boston University.

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

Breitung, Jorg and Wolfgang Meyer. 1994. "Testing for Unit Roots in Panel Data: Are Wages on Different Bargaining Levels Co-integrated?" *Applied Economics*, 26, 353-361.

Chow, Gregory. 1983. *Econometrics*. New York: McGraw-Hill.

Clerides, S., S. Lach, and J. Tybout.1998. "Is Learning by Exporting Important? Micro Dynamic Evidence from Colombia, Mexico and Morocco," *Quarterly Journal of Economics*, August, pp. 903-947.

Das, S. 1992. "A Micro-Econometric Model of Capital Utilization and Retirement: The Case of the U.S. Cement Industry," *Review of Economic Studies*, 59, 277-297.

Dixit, Avinash K. 1989. "Entry and Exit Decisions Under Uncertainty." *Journal of Political Economy* 97(3), 620-38.

Goldberg, Penelopi and Micheal Knetter.1999. "Measuring the Intensity of Competition in Export Markets," *Journal of International Economics* 47(1), pp. 27-60.

Krugman, Paul R. 1989. *Exchange Rate Instability*. Cambridge, MA: The MIT Press.

Roberts, Mark J. and James R. Tybout. 1997a. "The Decision to Export in Colombia: An Empirical Model of Entry with Sunk Cost," *American Economic Review*, September, pp. 545-564.

Roberts, Mark and James Tybout. 1997b. *What Makes Exports Boom?* The World Bank. Directions in Development Monograph.

Roberts, Mark, Theresa.A. Sullivan and James.R. Tybout. 1995. "Micro Foundations of Export Booms," Pennsylvania State University, Department of Economics Working Paper.

Rust, John. 1997. “Using Randomization to Break the Curse of Dimensionality,” *Econometrica* 65(3), May, pp. 487-516.

Rust, John. 1988. “Maximum Likelihood Estimation of Discrete Control Processes,” *SIAM Journal of Control and Optimization*, 26(5), 1006-1023.

Rust, John. 1995. “Numerical Dynamic Programming in Economics,” in *Handbook of Computational Economics*, edited by H. Amman, D. Kendrik and J. Rust.

Figure 1: Elasticities of Demand

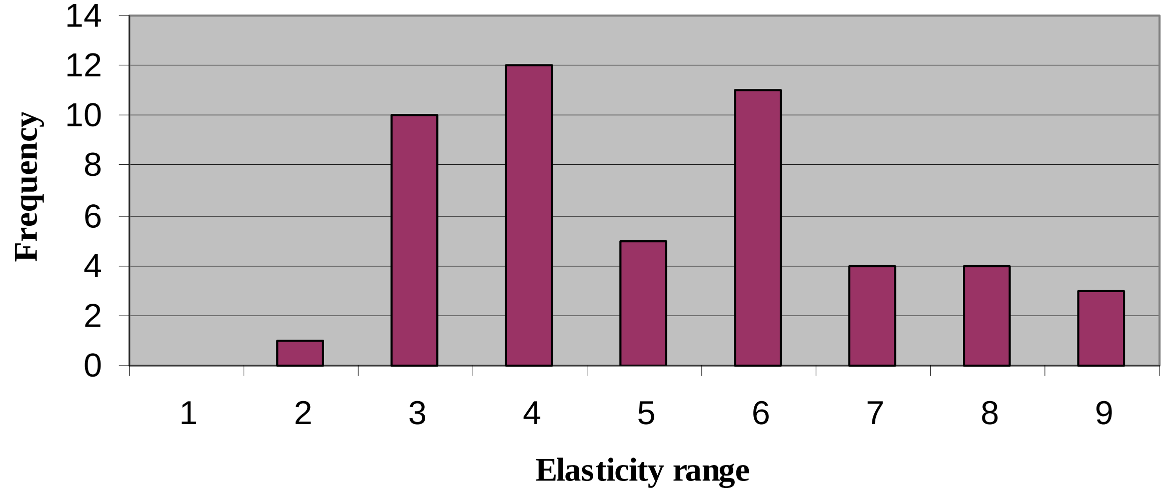

Note: Bars indicate the frequency of estimated elasticities between 2 and 3, 3 and 4, *et cetera*. The right-most bar is an open-ended category.

Figure 2: Average Export Revenues and Export Profits (plants sorted by export profits)

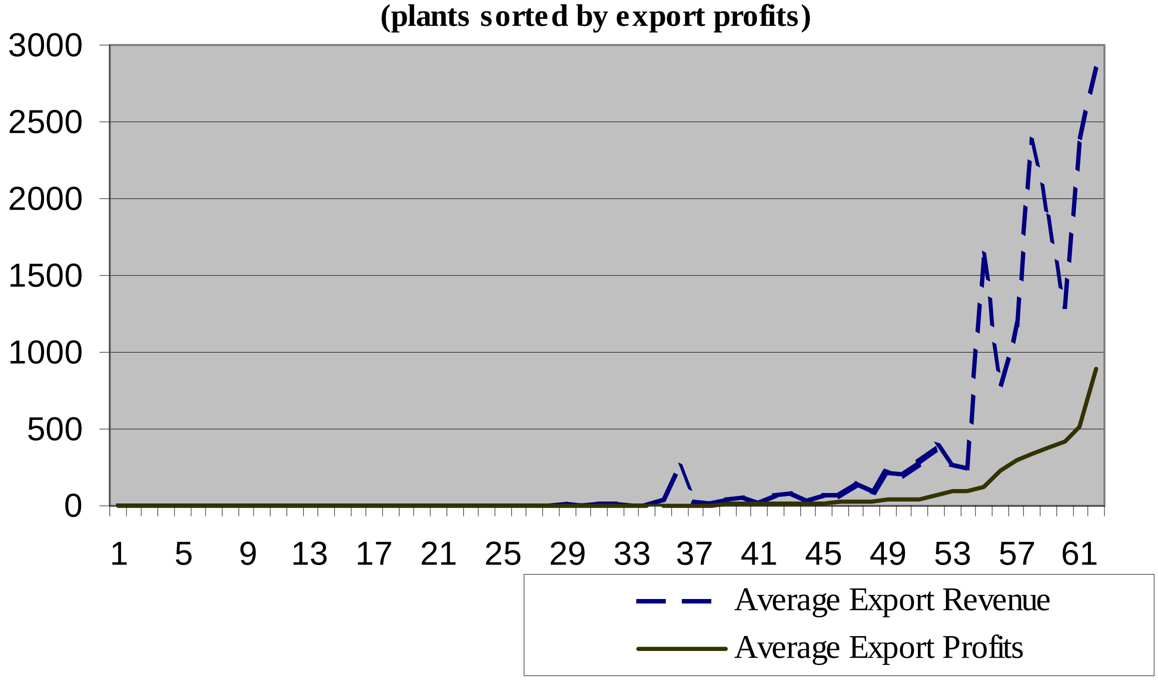

Units: Millions of 1986 pesos. (Multiply by 6.75 to get thousands of 1999 dollars.)

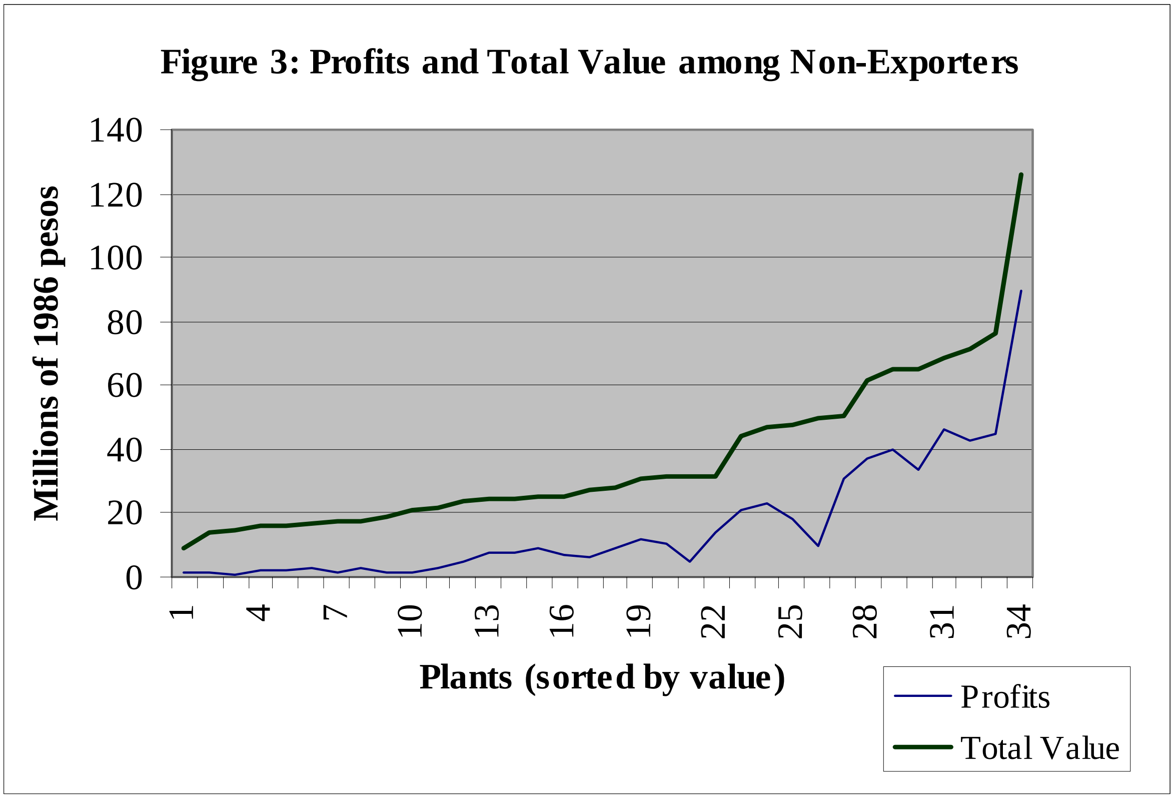
Figure 3: Profits and Total Value among Non-Exporters
140
120
100
80
60
40
20
0
Millions of 1986 pesos
1
4
7
10
13
16
19
22
25
28
31
34
Plants (sorted by value)
Profits
Total Value

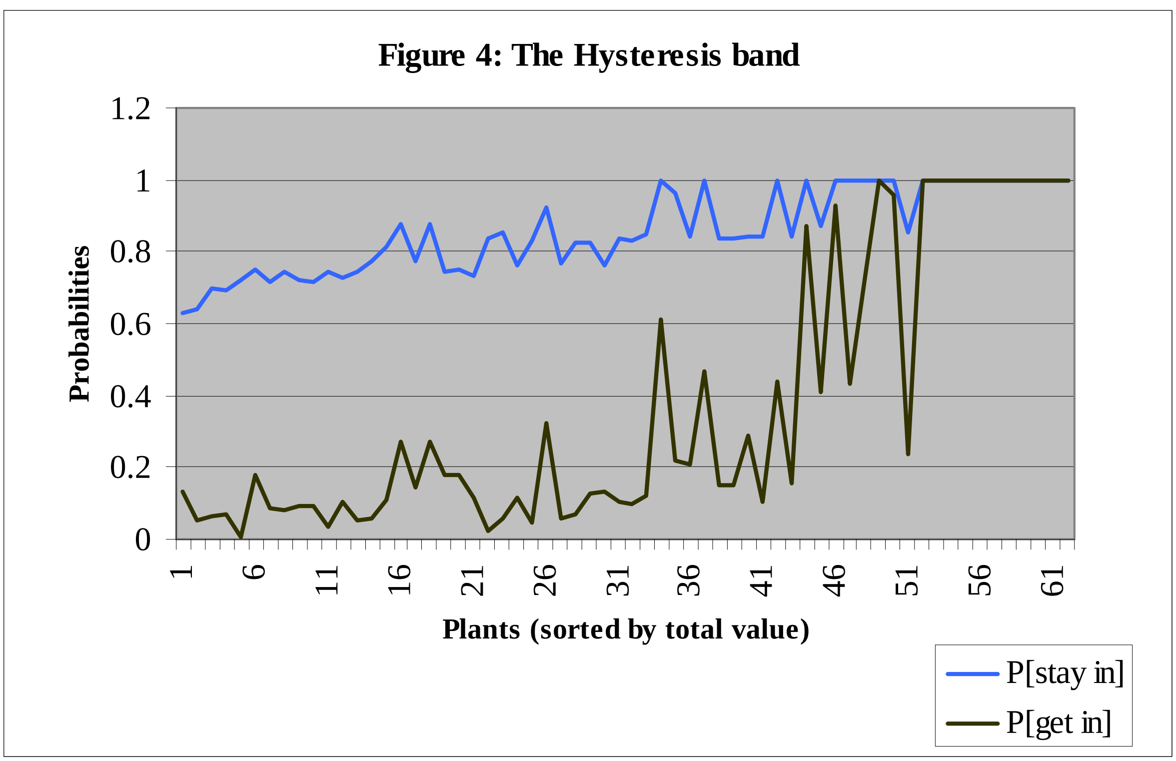
Figure 4: The Hysteresis band
1.2
1
0.8
0.6
0.4
0.2
0
Probabilities
1
6
11
16
21
26
31
36
41
46
51
56
61
Plants (sorted by total value)
P[stay in]
P[get in]

Figure 5: Profits and Total Value, All Plants