555
Rev.

DEPARTMENT OF AGRICULTURAL AND RESOURCE ECONOMICS
DIVISION OF AGRICULTURE AND NATURAL RESOURCES
UNIVERSITY OF CALIFORNIA, Berkeley.

WORKING PAPER NO. 555, Rev.

A DYNAMIC MODEL OF OLIGOPOLY IN THE COFFEE EXPORT MARKET

by

Larry S. Karp
and
Jeffrey M. Perloff

California Agricultural Experiment Station
Giannini Foundation of Agricultural Economics
August 24, 1990

A DYNAMIC MODEL OF OLIGOPOLY IN THE COFFEE EXPORT MARKET

by

Larry S. Karp* and Jeffrey M. Perloff**

* Larry S. Karp is an associate professor, Department of Agricultural and Resource Economics, University of California, Berkeley and reader, Department of Economics, University of Southampton

** Jeffrey M. Perloff is a professor, Department of Agricultural and Resource Economics, University of California, Berkeley

We thank Peter Berck, James Chalfant, David Freeman, George Judge, and Ken White for suggestions on econometrics; Drew Fudenberg for advice on theory; participants at seminars at the University of California at Berkeley, BAGEL, and the University of Michigan for comments; Diana Burton for collecting information on coffee; and Max Leavitt for computing help. We especially thank Jim Chalfant for extensive advice concerning importance sampling and other econometric issues. We also thank three anonymous referees and the editor for their comments. Ken White's Shazam was used for all estimations and simulations. Giannini Foundation Paper No. 1026 (for reprint identification only).

A Dynamic Model of Oligopoly in the Coffee Export Market.

Larry S. Karp and Jeffrey M. Perloff

A linear-quadratic, dynamic feedback oligopoly model that nests various market structures is used to estimate the degree of competitiveness and the adjustment paths of the two largest coffee exporters, Brazil and Colombia. Their estimated behavior is relatively competitive. This subgame perfect dynamic model is compared to a standard static oligopoly model and the open-loop model (the dynamic generalization of the standard static model). Both classical and Bayesian tests of open-loop and feedback dynamic models are reported.

Key words: coffee, dynamic oligopoly, econometric, subgame perfect Nash

A DYNAMIC MODEL OF OLIGOPOLY IN THE COFFEE EXPORT MARKET

Are the two largest coffee exporters, Brazil and Colombia, price takers, oligopolists, or in collusion? Casual evidence suggests that major coffee exporters behave noncompetitively and that dynamics play an important role in determining the outcome (Marshall). Since 1959, most exporting and importing countries have participated in a series of International Coffee Agreements (ICAs), which set export quotas (Fischer and Bates and Lien). It is alleged, however, that many countries regularly violate these quotas. In his survey of various ICAs, Gilbert (p. 602) quotes critics who claim that the agreements "are no more than an internationally sanctioned producer's cartel." He concludes that the agreements resulted in higher prices rather than simply more stable prices (p. 604). Greenstone also asserts that the large coffee producers behaved as a cartel. Based on econometric models, de Vries, Akiyama and Duncan, Palm and Vogelvang, and Herrmann argue that the ICAs have resulted in higher prices for member importing countries but lower prices for nonmember importing countries.

We view Brazil and Colombia as a dynamic coffee duopoly facing a fringe with exogenous exports. During the sample period, Brazil and Colombia's share of total world exports averaged 43 percent. Brazil and Colombia act like large "firms" in that each centrally controls exports. The Brazilian Coffee Institute (IBC) controls supply and price; supervises grading, packing, and weighing; and sets quotas within the country. The Colombian Federation of Coffee Growers (FNCC) buys from small farmers, evaluates, blends, grades, cleans, and manages the market through prices and taxes. In the absence of an ICA in 1974, Brazil and Colombia attempted to form an explicit producers' cartel and were later joined by other (smaller) producers. At various times, particularly during the extended ICA negotiations from 1978-1980, the Brazilian and Colombian governments appeared to intervene in the market to maintain stability.

The institutional arrangements in the coffee market constitute circumstantial evidence of an intent by large producers to exert market power. The difficulties of negotiating these agreements and the allegations that many exporting countries fail to comply fully with them suggest that producers do not behave as a perfect cartel. Thus, the hypothesis that the market structure lies between monopoly and competition is plausible. The complexity and inconstancy of coffee marketing institutions and paucity of data make it unreasonable to attempt to estimate the explicit game that producers play. Instead, we estimate a fairly general set of relations that include, as special cases, the equilibrium conditions associated with perfectly competitive, collusive, and Nash-Cournot behavior. Thus, we are able to determine whether the observed outcome is consistent with any of these equilibria.

The standard static assumption is inappropriate where there are large adjustment costs in training, storage, or in capital accumulation. There are two reasons to use a dynamic model for coffee. First, changes in production involve nonlinear costs. A lag of 2 to 5 years exists between planting and first harvest; a tree produces its maximum output between 5 to 10 years of age and bears for up to 30 years. This pattern suggests that average adjustment costs increase with the size of adjustment. Second, Brazil and Colombia maintain large stockpiles. In standard inventory models (e. g., Blinder), the costs of inventory adjustment are assumed to be reasonably approximated by a quadratic function. Costs of adjusting exports, therefore, stem from costly adjustment of production or inventories.[1]

The Model

Our dynamic model is flexible enough to allow for the possibility that Brazil (Firm 1) and Colombia (Firm 2) act like price takers, collude, or export oligopolistic levels between those two extremes. For simplicity, other countries' exports are ignored in this section but are accounted for in the estimates. In period t, firms face the inverse residual linear demand curve

(1) $$p_t = a(t) - bQ_t,$$

where p_t is the real price in period t, Q_t is the combined export of Brazil (q_{1t}) and Colombia (q_{2t}), a(t) includes effects of exogenous variables such as exports of other countries, and b is a positive coefficient.

Each Firm i has constant marginal costs θ_i with respect to contemporaneous exports q_{it} and a quadratic adjustment cost $\theta_i q_{it} + \left(\gamma_i + \frac{\delta}{2} u_{it}\right) u_{it}$, where $u_{i,t} \equiv q_{it} - q_{i,t-1}$ is the change in a firm's export level from period t-1 to period t. In contrast, a static model sets γ_i and δ equal to zero.

Static Model

In most empirical static models of oligopoly, aggregate or firm level data are used to estimate a parameter, which we call v, that reflects the markup of price over marginal cost. Given demand equation (1) for a homogeneous product, Firm i's effective marginal revenue curve (the marginal revenue given the degree of market power actually exercised) is MR_i(v) = p + (1 + $v)p'q_i = p_t - (1 + v_i)bq_i$.

If all firms have a common v and common marginal cost, $MC = MC_i = \theta$, we can rewrite the equilibrium equations for each firm, $MR_i(v) = MC \equiv \theta$, as:

(2) $$p = \theta + (1 + v)bq_1.$$

Estimating (1) and (2), dividing the coefficient on the q_i term in (2) by the estimate of b from (1), and subtracting one gives an estimate of v. The gap between price and marginal cost, p - MC =

$(1 + v)bq_i$, depends on v. If $v = -1$, marginal revenue equals price [$MR = p$] and there is no markup. If $v = 1$, marginal revenue is less than price [$MR = p + p'(2q_i) = p + p'Q$] and the monopoly markup is observed. Intermediate solutions, such as the Nash-Cournot where $v = 0$, are also possible (Carlton and Perloff).

Some authors interpret v as a firm's constant conjectural variation about its rival: $v \equiv dq_j/dq_i$. We prefer the neutral interpretation that v is a measure of market power — the gap between marginal cost and price — so that we do not need to make the behavior assumptions of conjectural variations models. Moreover, the conjectural variation interpretation cannot be used in the feedback dynamic model we now describe.

The Linear-Quadratic Dynamic Model

The game-theoretic literature abounds with very general dynamic models of oligopoly that do not lend themselves to estimation. To make the estimation problem tractable, we restrict our model to a specific family of equilibria indexed by a parameter v as in the static model. We estimate the family of equilibria under the assumption that firms use subgame perfect, Markov (feedback) strategies. That is, firms choose strategies (rules) that determine their exports as a function of the state variables.[2]

For comparison, we also estimate the same family of equilibria under the assumption that firms use open-loop strategies. That is, firms choose, at the initial time, a path that they intend to follow thereafter — they do not expect to revise their decisions after an unexpected shock (such as bad weather) affects production. This failure to anticipate revision is irrational, so the open-loop equilibrium is not subgame perfect. The open-loop model is the dynamic generalization of the static model as explained below.

The open-loop and Markov adjustment paths are identical if firms collude ($v = 1$ and firms are of equal size) or act as price takers ($v = -1$), as shown in Karp and Perloff (1991). In oligopolistic models, such as where firms play Nash within a period, the two models imply different adjustment paths and steady-state export levels. Given sufficient cost information to test overidentifying restrictions, one can empirically discriminate between the two types of behavior (Karp and Perloff 1989). Lacking that information, we estimate both models to determine the sensitivity of the measure of market structure to maintained hypotheses regarding behavior. The two assumptions lead to similar results for coffee.

We estimate the Markov model using a variant of the linear-quadratic game solution (Starr and Ho). A general open-loop model can be estimated (Roberts and Samuelson), but using a linear-quadratic specification enables us to compare easily the open-loop and Markov equilibria and to estimate a market structure parameter.

Our feedback and open-loop models include four common models: collusion, price taking, open-loop Nash-Cournot, and Markov Nash-Cournot. Other export paths lying between collusion and price-taking could be produced by other dynamic oligopolistic games. For example, Brazil and Colombia might imperfectly collude. Another possibility is that export levels are chosen subject to political pressures (one group wants to maximize export revenues and another wants to increase labor demand), which causes a deviation from Nash-Cournot or collusive equilibria.[3] Rather than try to model explicitly each of these games, we use an index v that allows for intermediate paths and steady-state exports. This index is the dynamic analog of the price-marginal cost wedge used in static models of oligopoly.

In each period of the dynamic model, Firm i's revenues R_i equal $p_t q_{it}$. Given a discount factor β, Firm i's objective is to maximize its discounted stream of profits,

$$\sum_{t=1}^{\infty} \beta^{t-1}\left[\left(p_t - \theta_i\right)q_{it} - \left(\gamma_i + \frac{\delta}{2}u_{it}\right)u_{it}\right], \tag{3}$$

where $(p_t - \theta_i)q_{it}$ is the contemporaneous profit and the last term is the adjustment cost. With Markov strategies, Firm i chooses changes in exports u_{it} as a function of the current information: its own and its rival's lagged export.

If we define $J_i(\boldsymbol{q}_t; v)$ as the present discounted value of Firm i's program, given the state vector $\boldsymbol{q}_t \equiv (q_{1t}, q_{2t})$ and an index v of market power, Firm i's dynamic programming equation is

$$J_i(\boldsymbol{q}_{t-1}; v) = \max_{u_{i,t}} (p_t - \theta_i)q_{it} - \left(\gamma_i + \frac{\delta}{2}u_{it}\right)u_{it} + \beta J_i(\boldsymbol{q}_t; v). \tag{4}$$

That is, the present discounted value of the stream of future profits as of last period equals the present discounted value of future profits as of this period plus the profits from this period.

The first-order condition corresponding to (4) is

$$p_t = \theta + (1 + v)bq_{it} + \gamma_i + \delta u_i - \beta\left[\frac{\partial J_i(\boldsymbol{q}_t; v)}{\partial q_i} + \frac{\partial J_i(\boldsymbol{q}_t; v)}{\partial q_j}v\right], \tag{5}$$

where $p - (1 + v)bp_{it}$ is marginal revenue and the term in brackets is the discounted shadow value of an extra unit of current exports. Terms in this equation are arranged to emphasize its similarity to the static equilibrium condition (2). The gap between marginal cost and price is the same function of v as in the static model, so v can be interpreted as a market power measure or index of the family of market structures.

Estimation of Dynamic Models

Our objective is to obtain a consistent estimate of the index v of market structure in the subgame perfect, dynamic model. We also estimate δ, the quadratic adjustment parameter. We can allow parameters γ_i, θ_i, and $a_i(t)$ to be firm specific and nonstationary to reflect quality differences, transportation costs, or other firm-specific costs. If $\gamma_i = 0$ (adjustment costs are minimized when there is no adjustment), the open-loop steady-state exports in the collusive, noncooperative Nash-Cournot and price-taking equilibria are identical to the corresponding static equilibria. Because the open-loop and static models have the same equilibria for general cost and revenue functions and not simply quadratic ones, the open-loop equilibrium is the dynamic analog of the static model. The static, open-loop, and feedback models are increasingly complex: the open loop is the dynamic version of the static model and the feedback model replaces the naive strategies of the open loop (and implicitly the static) model with sophisticated Markov strategies.

We now describe the basic estimation approaches (see Appendix). The value of Firm i's program, $J_i(\boldsymbol{q}_{t-1}; v)$, in (4) is a quadratic function. The equilibrium Markov control rules are linear, and the open-loop strategies can be written as linear functions of lagged export:

(6) $$q_t = g(t) + Gq_{t-1},$$

where $\boldsymbol{g}(t)$ is a column vector and $\boldsymbol{G}$ is a 2 x 2 matrix with elements G_{ij} (i, j = 1, 2).[4] In estimating adjustment equation (6), we make no assumptions regarding whether firms have rational expectations about the *exogenous* variables nor do we impose assumptions about whether the inverse demand intercepts and linear costs are constant over time or across firms. Elements of $\boldsymbol{G}$ are used to infer v. A major advantage of our estimation strategy is its simplicity: market structure is logically distinct from the rational expectations hypothesis. One could test rational expectations by including the exogenous information in the state vector (Chow). If that hypothesis is true, there is a loss of efficiency from ignoring it but our estimate is still consistent.

Results

We estimate both dynamic models of the coffee export market from the 1961-62 crop year to the 1983-84 crop year. During that period, Brazil and Colombia's share of total world exports ranged from 32 to 50 percent and averaged 43 percent. Brazil's share was, on average, twice that of Colombia. The share of the rest of South America was 4 percent; the rest of Latin America, 17 percent; Asia, 6 percent; and Africa, 29 percent. Shares of the two largest exporters in Africa were 7 percent (Ivory Coast) and 5 percent (Uganda).

We assume that Brazil and Colombia engage in a dynamic game in which they treat the rest of the world as a fringe with exogenous exports. Most Latin American and some African countries produce Arabica coffee (70 percent of all coffee produced). Most African and Asian countries produce Robustas, which are used mainly in instant coffees. A linear demand curve treating Arabica and Robustas from various countries as imperfect substitutes produced results similar to those reported below.

Quantity data are from *Coffee: World Coffee Situation* (various years) published by the U. S. Department of Agriculture, Foreign Agriculture Service. Coffee price is an average of prices of all coffee traded in New York, the major coffee market. Price data for coffee (New York) and tea (London), the world commodity wholesale price index, and world gross domestic product at constant prices are from the International Monetary Fund, *International Financial Statistics* (various years). In the following, we use real prices, obtained by deflating nominal prices by the world commodity wholesale price index.

Classical Estimates

We estimate inverse demand curve (1) using an instrumental variables technique:[5]

$$\text{Real Price of Coffee} = \underset{(2.25)}{4.401} - \underset{(-4.11)}{.000123}\ \text{World Coffee Exports} + \underset{(0.53)}{.0169}\ \text{World GDP}$$

$$+ \underset{(1.53)}{.559}\ \text{London Price of Tea} + \underset{(1.78)}{.206}\ t - \underset{(-2.59)}{.00380}\ t^2,$$

where t is a time trend (1, 2, ...) and terms in parentheses are t-statistics against the null hypothesis that the coefficient equals zero. The correlation between observed and predicted price is 0.69 and the Durbin-Watson statistic is 1.96. The demand curve plays only a minor role because b, the coefficient on exports, affects only δ and not v.

We estimate the $\boldsymbol{G}$ matrix in adjustment equation (6) using Zellner's seemingly unrelated equations method (table 1). Each country's exports are regressed on its own lagged exports, the other country's lagged exports, a time trend, and a dummy for the 1977-78 freeze in Brazil.[6] Cross-equation symmetry constraints are imposed in table 1: the coefficients on the own lagged exports are equal $[G_{11} = G_{22} \equiv G_1]$ across equations as are the coefficients on the other country's lagged exports $[G_{12} = G_{21} \equiv G_2]$. Based on unrestricted estimates, the F-statistic on these restrictions is 0.64 with 2 and 34 degrees of freedom, so we cannot reject the restrictions. The hypothesis that the coefficients on the lagged exports are zero independently or collectively is rejected by t-tests, F-tests, and likelihood ratio tests.

Based on these estimates of $\boldsymbol{G}$ and assuming $\beta = 0.95$, $v_c^o = \phi_o(G) = -0.84$ and $v_c^f = \phi_f(G) = -0.80$ [where $\phi_o(G)$ is a nonlinear function of G, superscripts o or f indicate open-loop or feedback, and subscript c indicates the classical estimate]. Standard errors of v_c^o and v_c^f, based on a Taylor expansion, are 0.27 and 0.31 (table 2). Thus, we cannot reject the price-taking ($v = -1$)

hypothesis, but we can reject the Nash-Cournot ($\nu = 0$) and less competitive models at the 0.05 level. Adjustment coefficients are the same under both models: $\delta_o^c = \psi_o(G, b) = 0.000037 = \delta_f^c = \psi_f(G, b)$.

For the estimated dynamical system to "make sense," it must have three properties:

- The system is *stable*: $-1 < G_1 + G_2 < 1$ and $-1 < G_1 - G_2 < 1$.
- The *market structure* lies between collusion and price taking:

  $1 > \nu^k = \phi_k(G) > -1$, k = o or f.
- The *adjustment* parameter in each of the models is positive:

  $\delta^k = \psi_k(G, b) > 0$, k = o or f.

Our classical point estimates of the elements of $\boldsymbol{G}$ and our estimates of ν^k and δ^k are consistent with these restrictions.

Bayesian Estimates

Rather than estimate the unconstrained system and hope the point estimates lie in the desired range, we can impose the above three sets of restrictions. Although it would be extremely difficult, if not impossible, to impose such inequalities using a classical approach or to test them, Geweke (1986, 1989) and Chalfant, Gray, and White show how to impose and test inequality restrictions with Bayesian techniques.

Our Bayesian prior is the product of a conventional uninformative distribution and an indicator function that equals 1 where the inequality constraints are satisfied and 0 elsewhere. The posterior distribution is calculated using Monte Carlo numerical integration with importance sampling. Given a quadratic (absolute difference) loss function, estimates of the parameters consistent with the restrictions are obtained by calculating the mean (median) of the coefficient estimates for all replications in which the constraints are satisfied. Indeed, we obtain the full posterior distributions of ν_b^o and ν_b^f.

Table 2 summarizes results of the classical and Bayesian estimates. The ν_a^k based on an absolute difference loss function (medians) are close to the classical point estimates. The ν_b^k based on a quadratic loss function (means) are about 0.2 higher than the classical estimates. Standard deviations on ν_b^k are slightly greater than the Taylor approximations for the classical estimates.

To estimate the probability that the restrictions hold, we calculate the (importance weighted) proportion of Monte Carlo replications satisfying the restrictions. Stability conditions are virtually always met (table 2). All three sets of conditions hold in approximately three-quarters of the replications, so imposing the restrictions seems reasonable (table 2). Because the restriction that δ is positive holds in three-quarters of the cases (the odds in favor of a positive δ are 3 to 1), the data indicate there is dynamic adjustment.[7]

The assumption in this Bayesian approach that the original error terms are normal can be relaxed by using a bootstrapping approach as reported in table 2.[8] The bootstrap estimates show slightly higher standard deviations corresponding to the mean ν_s^k estimates (0.43 and 0.44), a lower probability of rejecting due to $\nu_s^k \leq -1$ (19 percent), and a higher probability of rejecting due to $\nu_s^k < 1$ (11 percent).

Market Structure

Bayesian estimates provide an entire posterior distribution of market parameter ν^k. Some of the information from the importance sampling histograms is summarized in table 3, which shows the probability that ν^k lies within certain ranges. The probability that ν_b^k lies between -1 (price taking) and 0 (Nash-Cournot) is greater than 90 percent. There is a slightly higher probability (nearly 50 percent) that v_b^k lies between the classical estimate and 0 than between -1 and the classical estimate (over 40 percent). Two-thirds of the distribution lies below v_b^k. In the feedback model, the posterior odds ratio that the market structure lies between price taking and Nash-Cournot rather than between Nash-Cournot and collusive is 12.9.

Results for the bootstrap model are similar; however, the bootstrap distribution has thicker tails. The bootstrap posterior odds ratio that the market structure is more competitive than Nash-Cournot versus the opposite is 8.3 — less than the odds using the importance sampling estimates but still very high. Based on either the classical or Bayesian approaches, the Brazil-Colombia exports are close to price taking. The probability that they are at least as noncompetitive as Nash-Cournot is no greater than 11 percent according to all our estimating approaches.

Simulations

We can simulate adjustment paths and steady states based on the above estimates.[9] For a given v, the open-loop steady-state export and the corresponding static export are equal. For a constant $v \in (-1, 1)$, the feedback games are more competitive than open-loop games in the sense that steady-state exports are greater for a given v (see Reynolds for the intuition for the Nash-Cournot model). Based on simulations, differences in steady states between the open-loop and feedback models is maximized at values close to those estimated, -0.8. In absolute terms, however, these differences are small.

In the feedback model based on the unconstrained classical estimator v_c^f, steady-state exports are 6 percent lower than in the corresponding price-taker model (table 4). In the open-loop and static models based on the classical estimator v_c^o, steady-state exports are 7 percent below those of the price-taker model. Thus, the feedback model is closer to price taking than is the open-loop model. This quantity difference leads to large price effects because the demand elasticity evaluated at the mean price and quantity is about 0.2.

The dynamic models based on unconstrained classic estimates deviate from the price-taker level by only about half that of those based on the Bayesian constrained estimates. Using the constrained importance-sampling Bayesian quadratic loss estimates (bootstrapping), the feedback steady-state export is 13 (12) percentage points below the price-taker level and the open-loop steady-state export is 15 (14) percentage points below (table 4).

The steady states vary in the two estimated dynamic models for two reasons. First, for either model, the lower the estimated v, the closer its steady-state export is to price taking. As table 4 shows, the steady state of either the feedback or open-loop model based on v_c^o is 2 percent higher than that based on v_c^f. The open-loop model's v_c^o = -0.84 is slightly closer to price taking than is the feedback model's v_c^f = -0.80.

Second, and more than offsetting the first effect, in the open-loop model there is more collusive behavior for any given v. As table 4 shows, the feedback model (v_c^f) predicts a steady state 6 percentage points below price taking, whereas the open-loop model predicts a steady state 9 percentage points below. As a result, the less competitive open-loop model must have a lower estimated v to be consistent with the data.

Adjustment paths vary across models. As figure 1 (where exports of both countries are set to zero in year zero) shows for the classical estimates, the combined exports of the two countries in the feedback model reaches a higher steady-state level than in the open-loop model; but, in both models, exports reach their steady state level after only three years.

To illustrate a more realistic adjustment path, we solved the dynamic game using estimated v and δ and marginal costs obtained by model calibration (assuming the steady-state exports equal the average exports in the sample period). In each period, we added the residuals plus the coefficient (from the estimated adjustment equation) times the Brazilian freeze dummies to the simulated q_t. These simulations exclude demand-side shocks and also ignore the upward trend in the demand curve, which may explain why the simulations shown in figure 2 are low toward the end of the period. Even ignoring the demand trend, though, figure 2 shows the model is capable of emulating reality reasonably well.

Summary and Conclusions

Based on both classical and Bayesian estimates of a dynamic Markov model, Brazil and Colombia compete vigorously with each other in the coffee export market. Steady-state exports based on the classical (Bayesian) estimate are 6 (12 or 13) percent less than if both Brazil and Colombia were price takers. From the Bayesian analysis, we are reasonably confident that the behavior of Brazil and Colombia is closer to price taking than to collusion.

Our results from a subgame perfect dynamic oligopoly model differ only slightly from those based on the much simpler open-loop dynamic model. If these results can be replicated with other data sets, open-loop estimates can be used to approximate Markov strategies without major bias.

References

Akiyama, T., and R. C. Duncan. *Analysis of the World Coffee Market*, World Bank Staff Commodity Working Paper No. 1, Washington, D. C., 1982.

Bates, Robert H., and Da-Hsiang Donald Lien. "On the Operations of the International Coffee Agreement." *International Organ.* 3(1985):553-559.

Blinder, Alan S. "Inventories and Sticky Prices: More on the Microfoundations of Macroeconomics." *American Economic Review.* 72(1982):334-349.

Carlton, Dennis W., and Jeffrey M. Perloff. *Modern Industrial Organization.* Glenview, IL: Scott, Foresman and Company, 1990.

Chalfant, James A., Richard S. Gray, and Kenneth J. White. "Evaluating Prior Beliefs in a Demand System." *Amer. J. Agr. Econ.* 73(1991):476-490.

Chow, Gregory C. *Econometric Analysis by Control Methods.* New York: John Wiley and Sons, Inc., 1981.

de Vries, Joseph. *Structure and Prospects of the World Coffee Economy*, World Bank Staff Working Paper No. 208, Washington, D. C., 1975.

Fischer, Bart S., *The International Coffee Agreement, A Study in Coffee Diplomacy.* New York: Praeger, 1972.

Freedman, David A., and Stephen C. Peters. "Bootstrapping an Econometric Model: Some Empirical Results." *J. Bus. and Econ. Stat.* 10(1984):150-158.

Fudenberg, Drew, and Eric Maskin. "The Folk Theorem in Repeated Games with Discounting with Incomplete Information." *Econometrica* 54(1986):533-54.

Geweke, John. "Exact Inference in the Inequality Constrained Normal Linear-Regression Model." *J. Applied Econometrics.* 26(1986):295-321.

Geweke, John. "Bayesian Inference in Econometric Models Using Monte Carlo Integration." *Econometrica*. 57(1989):1317-1339.

Gilbert, C. L. "International Commodity Agreements: Design and Performance." *World Devel.* 15(1987):591-616.

Greenstone, W. D. "The Coffee Cartel: Manipulation in the Public Interest." *J. Futures Markets*. 1(1981):1-18.

Herrmann, Roland. "Free Riders and the Redistributive Effects of International Commodity Agreements: The Case of Coffee." *J. Pol. Modeling*, 8(1986):597-621.

Karp, Larry S., and Jeffrey M. Perloff. "Open-Loop and Feedback Models in Dynamic Oligopoly." University of California, Department of Agricultural and Resource Economics, Working Paper, Berkeley, 1991.

Karp, Larry S., and Jeffrey M. Perloff. "Dynamic Oligopoly in the Rice Export Market." *Rev. Econ. and Stat.* 71(1989):462-470.

Marshall, C. F. *The World Coffee Trade*. Cambridge: Woodhead-Faulkner, 1983.

Palm, F. C., and E. Vogelvang. "A Short-Run Econometric Analysis of the International Coffee Market." *European Rev. of Agr. Econ.* 13(1986):451-476.

Reynolds, Stanley S. "Capacity Investment, Preemption and Commitment in an Infinite Horizon Model." *Intern. Econ. Rev.* 28(1987):69-88.

Roberts, Mark J., and Larry Samuelson. "An Empirical Analysis of Dynamic, Nonprice Competition in an Oligopolistic Industry." *Rand J. Econ.* 19(1988):200-220.

Starr, A. W., and Y. C. Ho. "Nonzero Sum Differential Games." *Journal of Optimization Theory and Applications*. 3(1969): 184-206.

Appendix: Estimating the Dynamic Model

The equilibrium conditions of the dynamic games can be written in matrix notation. Let $\boldsymbol{e}_i$ be the i^{th} unit vector, $\boldsymbol{e}$ equal the column vector of 1's, $\boldsymbol{S}_i = \boldsymbol{e}_i\boldsymbol{e}_i'\delta$, and $\boldsymbol{K}_i = b(\boldsymbol{e}\boldsymbol{e}_i' + \boldsymbol{e}_i\boldsymbol{e}')$. The dynamic programming equation for the linear-quadratic model is

$$J_i(\boldsymbol{q}_{t-1};v) = \max_{u_{it}} \left[\left(a\boldsymbol{e}_i'\boldsymbol{q}_t - \frac{1}{2}\boldsymbol{q}_t'\boldsymbol{K}_i\boldsymbol{q}_t - \frac{1}{2}\boldsymbol{u}_t'\boldsymbol{S}_i\boldsymbol{u}_t\right) + \beta J_i(\boldsymbol{q}_t;v)\right],$$

where $J_i(\cdot)$ is a quadratic function, $\gamma_i = 0$ for simplicity, and the equilibrium Markov control rules are linear as shown in (6). If the open-loop equilibrium sequence is expressed as a function of the current state, the open-loop equilibrium can be written in "feedback form" as in (6). Firms revise their plans if something unexpected happens, but these revisions are unanticipated. Define $\boldsymbol{v}_i$ as a 2 x 1 column vector with 1 in the i^{th} position and v in the other position. Given β, an estimated matrix $\boldsymbol{G}$, and an estimated demand slope b, the open-loop v and δ satisfy

(A.1)
$$\boldsymbol{K}_i\boldsymbol{v}_i = \left[\boldsymbol{G}^{-1}(\boldsymbol{I} - \boldsymbol{G})(\boldsymbol{I} - \beta\,\boldsymbol{G})\right]'\boldsymbol{e}_i\delta\,,$$

where $\boldsymbol{K}_i$ is of rank 2 so that the solution to (A.1) is unique. The derivation of (A.1) does not require an assumption of symmetry.

To estimate v and δ in the feedback case, we define the vectors

$$\boldsymbol{w}_i = \left[\boldsymbol{I} - \beta(\boldsymbol{G}' \otimes \boldsymbol{G}')\right]^{-1}\left[(\boldsymbol{G}' \otimes \boldsymbol{G}')\,(\mathrm{vec}(\boldsymbol{K}_i))\right],$$

$$\boldsymbol{x}_i = \left[\boldsymbol{I} - \beta(\boldsymbol{G}'\otimes\boldsymbol{G}')\right]^{-1}\left[\boldsymbol{G}'\otimes\boldsymbol{G}') - (\boldsymbol{I}\otimes\boldsymbol{G}') - (\boldsymbol{G}'\otimes\boldsymbol{I}) + \boldsymbol{I}\right][\mathrm{vec}(\boldsymbol{e}_i\boldsymbol{e}_i')],$$

where vec(Z) is a vector of the stacked columns of the matrix $\boldsymbol{Z}$. The inverse vec operation is then used to "rematricize" w_i and x_i to obtain the 2 x 2 matrices W_i and X_i. W_i is linear in Firm i's demand coefficients and X_i depends only on β and $\boldsymbol{G}$. If agents use feedback strategies, v and δ must satisfy

$$\text{(A.2)} \qquad \left[K_i + \beta\, W_i + (e_i\, e_i' + \beta\, X_i)\, \delta_i\right] v_i = G^{-1}\, e_i \delta_i \equiv y_i^* \delta_i\,.$$

Given symmetry, the left side of (A.2) is of rank 2 and the estimate of v is independent of b. Defining matrix A^i such that $bA^i \equiv K_i + \beta W_i$ and matrix $B^i \equiv e_i e_i' + \beta X_i$, A^i and B^i depend only on β and G. To recover v and δ, rewrite the i^{th} and the k^{th} ($k \neq i$) equation of (A.2) as

$$\text{(A.3)} \qquad b\left(A_{ii} + v \sum_{j \neq i} A_{ij}\right) + \left(B_{ii} + v \sum_{j \neq i} B_{ij}\right)\delta = y_{ii}^* \delta\,,$$

$$\text{(A.4)} \qquad b\left(A_{ki} + v \sum_{j \neq i} A_{kj}\right) + \left(B_{ki} + v \sum_{j \neq i} B_{kj}\right)\delta = y_{ik}^* \delta\,,$$

where subscripts designate the element of A^i, B^i, and y, and the superscript i is suppressed. Equation (A.4) can be solved for δ as a linear function of b and a nonlinear function of v. Substituting this function into (A.3) gives a quadratic in v that is independent of b. Hence a symmetric v can be estimated with knowledge of only β and G. Although there are two solutions to (A.3), extensive simulation experiments show that one value is close to the open-loop value and that the other is infeasible ($v < -1$, $v > 1$, or $\delta < 0$). Therefore, in practice it is easy to choose the correct root.

Footnotes

1. Our technique could be generalized to allow for adjustment costs in production and inventories. However, because reliable inventory data are not available, we assume the estimated adjustment cost is a proxy for both the cost of adjusting output and the cost of inventories.

2. A combination of strategies is a subgame perfect Nash equilibrium if it is a Nash equilibrium for the game and the relevant strategies are a Nash equilibrium for every subgame (subperiod within the game).

3. Further, the game-theoretic folk theorem (Fudenberg and Maskin) justifies any outcome that is less competitive than Nash.

4. Infinite horizon games may have many subgame perfect equilibria. We obtain a unique equilibrium by taking the equilibrium strategies that result from the game with a finite horizon T and letting $T \to \infty$.

5. Instrumental variables are time; time squared; time cubed; the freeze dummy; the London price of tea; the world gross domestic product; and Brazilian and Colombian rain, temperatures, gross domestic products, and populations.

6. To test the effects of the ICA agreements we tried the following experiments: 1) Dummies for the time of ICA meetings were included; 2) a one-year lag or lead of these meeting dummies were tried; and 3) a continuous variable measuring time from the previous meeting was used. Estimated coefficients were not statistically different from zero at the 0.05 level individually or collectively. We also divided the period in half and estimated separate adjustment equations for the two subperiods (to capture a fundamental structural change, possibly due to changes in the ICA). Based on a Chow test, we could not reject the hypothesis of no structural change.

7. We also estimated the corresponding static model where marginal cost is a function of the freeze and a time trend (but not of dynamic adjustments). In the static optimization equation, the real price was regressed on quantity, a dummy for the 1977-78 Brazilian freeze, and a time trend. The estimated $v < -1$ is outside the plausible range.

8. With lagged endogenous variables but no autocorrelation (as indicated by a Durbin's h test), we can bootstrap by choosing rows (left- and right-hand-side variables) of the original data (Freedman and Peters).

9. In the simulations, we choose marginal cost parameters for Brazil and Colombia so the steady-state export under the classically estimated feedback model equals the average export for the sample period. Because the constant marginal costs are not identified in our estimation procedure, we set the marginal cost of Brazil to zero and searched for the marginal cost of Colombia and the residual demand intercept producing the average steady-state outputs for the two countries. Subtracting b times the average output of the rest of the world's producers from the estimated demand curve intercept approximately equals the resulting residual demand intercept.

Table 1. Adjustment Equations: Regression of Exports on Lagged Exports

	Brazil	Colombia
Constant	12,986.0	6,967.9
	(4.99)[a]	(4.28)
Brazilian freeze 1977-78	-9,980.7	843.7
	(-4.66)	(0.74)
Time (1, 2, ...)	22.4	124.8
	(0.30)	(2.59)
Lagged Brazil exports	0.302	-0.192
	(2.27)	(-2.42)
Lagged Colombia exports	-0.192	0.302
	(-2.42)	(2.27)
R^2	0.57	0.74
Durbin-Watson	2.23	1.57
Durbin's h	-0.72	1.34

[a] Figures in parentheses are t-statistics against the null hypothesis that the coefficient equals zero.

Table 2. Classical and Bayesian Inequality-Constrained Estimates

	v^o	v^f
Classical estimates		
v_c^k (unrestricted)	-0.84	-0.80
σ (Taylor approximation)	0.27	0.31

Bayesian inequality	Importance sampling (T = 5,000)		Bootstrapping (T = 2,000)	
constrained estimates	v^o	v^f	v^o	v^f
Quadratic loss				
v_b^k (mean)	-0.65	-0.62	-0.68	-0.63
σ	0.35	0.36	0.43	0.44
Precision of mean of v_b^k ($\sigma/\sqrt{T}$)	0.0059	0.0060	0.0096	0.0097
Absolute loss				
v_a^k (median)	-0.76	-0.73	-0.86	-0.81
σ	0.37	0.37	0.46	0.47
Reject because (%)				
Unstable	0.002	0.002	0.0	0.0
$\delta^k \leq 0$	25.2	23.6	18.7	18.7
$v^k < -1$	24.9	25.4	18.7	18.7
$v^k > 1$	1.7	1.2	11.0	11.0
Total rejections (1 - p)	26.5	26.5	29.7	29.7
Asymptotic standard error				
of p ($\sqrt{p(1-p)/T}$)	0.0073	0.0073	0.0109	0.0109

Table 3. Distribution of v^k Based on Bayesian Estimates

Proportion of weight between[a]		Importance Sampling		Bootstrap	
		v^o	v^f	v^o	v^f
-1	0	93.5%	92.8%	90.3%	89.2%
0	1/2	4.2	4.8	5.8	6.3
1/2	1	2.3	2.4	3.8	4.4
-1	v_c^k	34.0	35.1	53.7	51.7
v_c^k	0	59.5	57.7	36.6	37.6
-1	v_b^k	67.2	65.5	72.9	71.1
v_b^k	0	26.3	27.3	17.5	18.2

[a] The classic estimate is v_c^k and v_b^k is the Bayesian estimate based on a quadratic loss function (k = o or f).

Table 4. Simulated Exports under Various Models

Model	Brazilian and Colombian Exports (Static or Steady-State)	Percentage of Price-Taker Output
Static and Dynamic Price-Taking (v = -1)	25789	100
Feedback (v_c^o)	24665	96
Feedback (v_c^f)	24308	94
Open Loop and Static Models (v_c^o)	23868	93
Open Loop and Static Models (v_c^f)	23469	91
Feeback (v_s^f)	22601	88
Feedback (v_b^f)	22504	87
Open Loop and Static Models (v_s^o)	22232	86
Open Loop and Static Models (v_b^o)	21948	85
Feedback Cournot (v = 0)	17557	68
Open-Loop and Static Cournot (v = 0)	17192	67
Static and Dynamic Collusion (v = 1)	12894	50

Note: Exports are in thousands of 60 kilogram bags. The estimated dynamic models are shown in bold. The v_b^k are the Bayesian quadratic loss estimates and the v_s^k are the bootstrap estimates (k = o or f).

Figure 1
Exports under Various Models

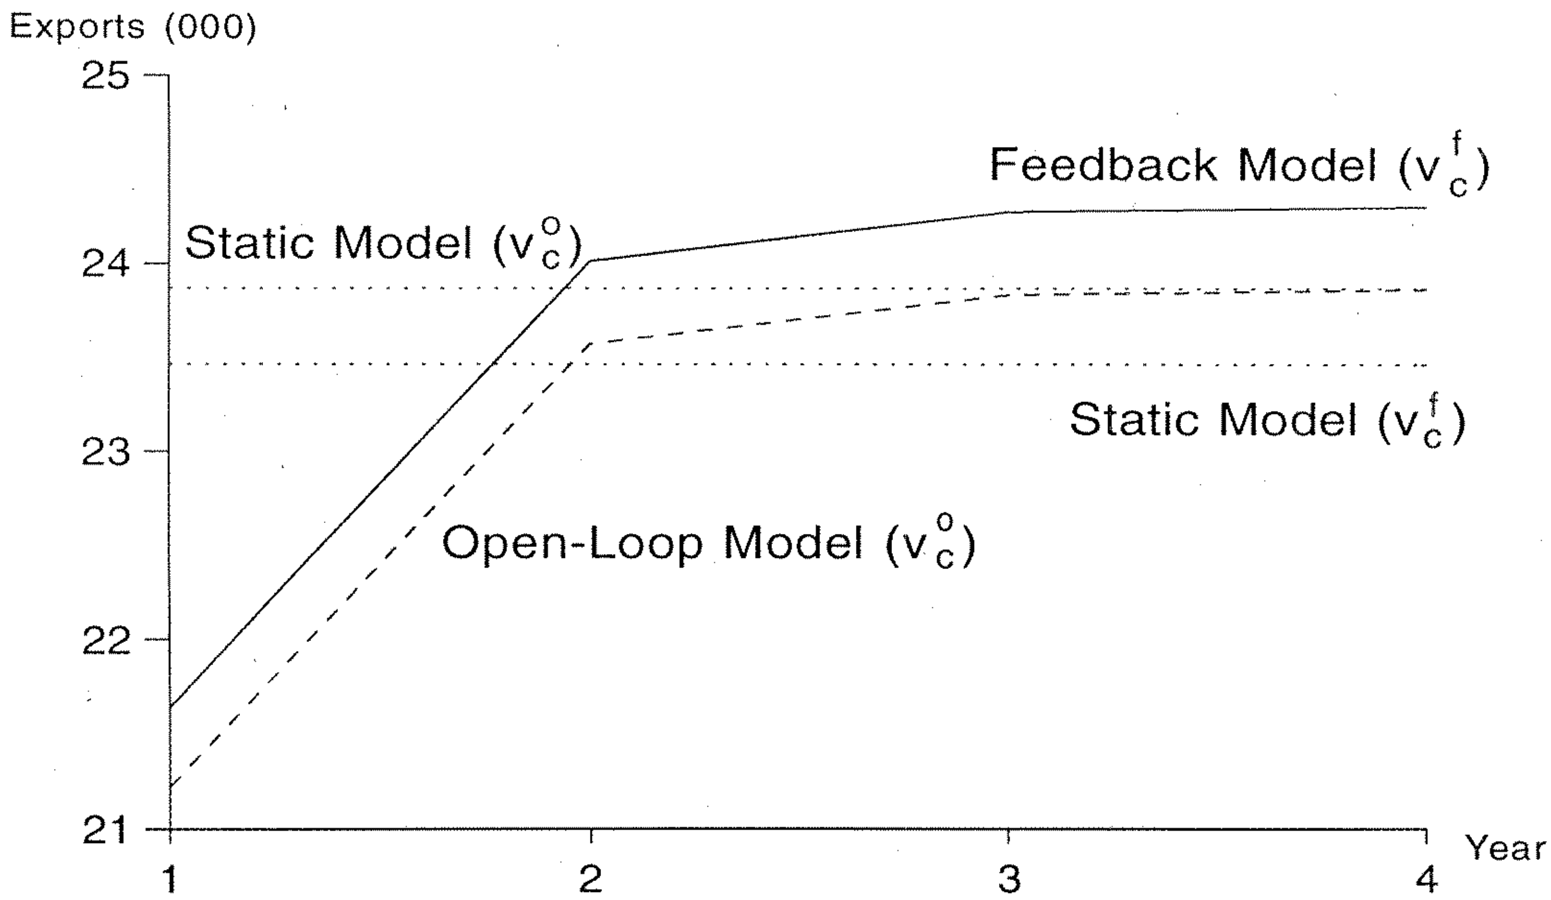

Figure 2
Actual and Simulated Exports