Title
Dynamic Oligopoly: Estimation and Tests of Market Structure

Authors
Karp, Larry
Perloff, Jeffrey M

Publication Date
1988-10-01

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

Working Paper No. 468

DYNAMIC OLIGOPOLY: ESTIMATION AND TESTS OF MARKET STRUCTURE

by

Larry S. Karp and Jeffrey M. Perloff

California Agricultural Experiment Station
Giannini Foundation of Agricultural Economics
October 1988

DYNAMIC OLIGOPOLY:
ESTIMATION AND TESTS OF MARKET STRUCTURE

by

Larry S. Karp and Jeffrey M. Perloff*

Department of Agricultural and Resource Economics

University of California at Berkeley

*We would like to 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.

ABSTRACT

A linear-quadratic dynamic oligopoly model is developed and applied to the world coffee export market. The model nests various market structures using either open-loop or feedback strategies. The theoretical properties of this model are described. For given observed behavior, the assumption of feedback strategies implies a less competitive market structure than open-loop strategies.

DYNAMIC OLIGOPOLY:

ESTIMATION AND TESTS OF MARKET STRUCTURE

A dynamic model is used to identify, estimate, and test market structure and the types of strategies used by firms where adjustment is costly. The model is applied to the world coffee export market. Previous papers have used static oligopoly models to estimate market structure (Iwata, 1974; Gallop and Roberts, 1979; Sumner, 1981; Appelbaum, 1982). The static assumption of these papers is inappropriate where there are substantial adjustment costs in training or in capital accumulation, or where there is learning over time. By incorporating costs of adjustment in a dynamic model, we can estimate and test the competitiveness of a market under different assumptions about the rationality of firms.

The game-theoretic literature abounds with dynamic models of oligopoly that are too general to be usable in estimation. Our model is restricted to a specific family of equilibria to facilitate estimation. This family nests the collusive, price-taking, and Nash-Cournot-in-quantities markets among others. The justifications for collusive and price-taking behavior are obvious.

There are two reasons for including the Nash-Cournot model. First, including this model allows us to compare our results to earlier empirical studies. Second, the dynamic version of the Nash-Cournot model may be reasonably motivated. Suppose that there are discrete time periods (such as growing seasons) during which firms cannot vary their output levels. Thus, the firm is correct when it makes the Nash-Cournot assumption that its competitors cannot respond to changes in its output levels within a time period. Nonetheless, firms can respond over time.

We consider two types of equilibria. First, firms may choose an initial set of output levels and stick to that path. Second, firms may choose strategies, i.e., rules which give their output as a function of the state. These two equilibria are, respectively, open loop and feedback.

The open-loop and feedback models are identical where firms collude or act as price takers. In other oligopolistic models, such as where firms make the Nash-Cournot assumption within a period, the two models imply different adjustment paths and steady-state output levels. Under the open-loop model, firms do not expect to revise their strategies after an unexpected shock (such as bad weather) affects the output levels of various firms. This failure to anticipate revision is irrational.

The feedback equilibrium is subgame perfect, but for general functional forms is difficult to estimate. To be able to estimate practically a feedback model, we use a variation of the well-known solution to the open-loop and feedback linear-quadratic model (Starr and Ho, 1969).[1] The general open-loop model can be estimated,[2] but the linear quadratic specification makes it possible to compare the open-loop and feedback equilibria.

Our methodology allows us to distinguish between the four models described above: collusion, price-taking, Nash-Cournot open-loop, and Nash-Cournot feedback. A number of other dynamic oligopolistic games would produce output paths that lie between those of collusion and price taking. Rather than try to explicitly model each of these games, we generalize our specification so as to allow for intermediate paths and steady-state output.

We use an index of behavioral assumptions by firms within a single time period to approximate these other games. This index is analogous to a conjectural variation in a static game. Where firms may use feedback strategies, a conjectural variations interpretation is inappropriate. While this index is not the explicit outcome of a game, it allows us to easily approximate a range of games. In particular, the collusive, price-taking, and Nash-Cournot models are obtained as special cases of this more general model. Thus, estimation of the market structure only requires the estimation of this index. We treat this index as a single parameter but, more generally, it might be a function of exogenous variables.[3]

We start by describing the model. Next, we discuss the method of nesting and the leading cases of competition, collusion, and Nash-Cournot-in-quantities markets. The third section provides qualitative analysis of the model and the fourth gives the details necessary for econometric implementation. The following section provides an econometric application: the international coffee export market. We end with a summary and conclusions.

1. Definitions and the Model

Estimation relies on the discrete time model in which the length of a period is ε. The continuous model, obtained as $\varepsilon \to 0$, is used to obtain most of our analytic results (see Karp and Perloff, 1988).

The industry consists of $n + 1$ firms where $n \geq 1$. At time t, firm i decides how much to produce in the current period q_{it} or, equivalently, determines its change in output, $u_{it}\varepsilon = q_{it} - q_{i,t-\varepsilon}$, where ε is in units of time and u_{it} is a rate. Firm i incurs a quadratic cost of adjustment,

$$\left(\delta_{0i} + \delta \frac{u_{it}}{2}\right) u_{it}\,\varepsilon,$$

and a quadratic cost of production,

$$\left(\theta_{0i} + \theta \frac{q_{it}}{2}\right) q_{it}\,\varepsilon.$$

In period t, firm i faces the linear demand curve

$$p_{it} = a_i - b \sum_j q_{jt}. \tag{1}$$

Firm i's revenue in period t is $p_{it}\, q_{it}\, \varepsilon$. Given an instantaneous interest rate of r, the one-period discount rate is $e^{-r\varepsilon}$, and the objective of firm i is to maximize its discounted stream of profits,

$$\sum_{t=1}^{\infty} e^{-r(t-1)\varepsilon} \left[\left(p_{it} - \theta_{0i} - \frac{\theta}{2} q_{it} \right) q_{it} - \left(\delta_{0i} + \delta \frac{u_{it}}{2} \right) u_{it} \right] \varepsilon . \tag{2}$$

For simplicity in the theoretical analysis, we set $a_i = a$ and assume $\theta_{0i} = 0 = \delta_{0i}$. The last equality implies that adjustment costs are minimized when there is no adjustment. As a result, the steady-state levels of output in the open-loop, collusive, noncooperative Nash-Cournot, and price-taking equilibria are equal to their static analogs. This equality holds for general cost and revenue functions and not simply the quadratic ones assumed here.[4]

For estimation, it is convenient to allow the parameters δ_{0i}, θ_{0i}, and a_i to be firm specific and nonstationary. The resulting estimation model reflects quality differences, transportation costs, or other specific costs. This flexibility eliminates a set of restrictions, but the remaining restrictions implied by the model are sufficient to identify the characteristics of the market. The parameters r, b, θ, and δ are assumed common to all firms. This assumption is relaxed in the discussion of estimation.

The ith firm's objective (2) is written in matrix form as

$$\sum_{t=1}^{\infty} e^{-r(t-1)\varepsilon} [a\, e_i'(q_{t-\varepsilon} + u_t\, \varepsilon) - \frac{1}{2} (q_{t-\varepsilon} + u_{t\varepsilon}\, \varepsilon)'\, K_i(q_{t-\varepsilon} + u_t\, \varepsilon) - \frac{1}{2} u_t'\, S_i\, u_t]\, \varepsilon, \tag{2a}$$

where e_i is the ith unit vector and e is a column vector of 1's, K_i = $b(e\, e_i' + e_i\, e') + \theta\, e_i\, e_i'$ (so K_i is a matrix with b's on the ith row and column except for the (i, i) element which is $2b + \theta$; all other elements are 0), and $S_i = e_i\, e_i'\, \delta$. As $\varepsilon \to 0$, this expression approaches

$$\int_0^{\infty} e^{-rt} [a\, e_i'\, q_t - \frac{1}{2} q_t'\, K_i\, q_t - \frac{1}{2} u_t'\, S_i\, u_t]\, dt. \tag{2b}$$

We assume that q_{it} is unconstrained so that negative sales are possible. Negative prices can be interpreted as very low prices. When prices fall below a certain level, firms would prefer to be buyers rather than sellers; they must bear the adjustment cost to make the transition.

Alternatively, the model can be interpreted as a standard investment problem in which q_{it} is firm i's capacity, and sales lie in the interval $[0, q_{it}]$. This interpretation requires additional assumptions. Provided that initial capacity lies within a certain range (an n + 1 dimensional set) that depends on the market structure, firms will produce at capacity (for the example of a Nash-Cournot market, see Reynolds, 1987). Given an initial condition in this range, the open-loop and feedback solutions are as shown below. The reader can either adopt the literal interpretation or regard the model as the standard investment problem in which the initial conditions are such that the capacity constraint is always binding.

2. Two Families of Equilibria

We consider two families of equilibria: open loop and feedback. Members of each family are indexed by a parameter v, which describes the behavioral assumption that determines the outcome. This parameter is defined by $v = \partial u_{jt}/\partial u_{it}$ for $i \neq j$ and all t. In a static model, v can be interpreted as a constant conjectural variation. Since the open-loop game is equivalent to a static problem, the same interpretation can be adopted in that case. We adopt the neutral description of v as a player's behavioral assumption. This assumption is taken as primitive and not explained by strategic considerations.

This procedure is justified on pragmatic, empirical grounds. The leading cases where $v = -1/n$, 0, or 1 (the latter for identical firms) result in the price-taking, Nash-Cournot, and collusive equilibria, respectively. The estimation of v provides a measure of the closeness of the observed market to a particular ideal market. If $v = -1/n$, each firm acts as if it believes its rivals will exactly offset its own deviation from equilibrium. Since the good is homogeneous, the firm acts as a price taker. If $v = 1$ and firms are identical, each firm acts as if its rivals will punish it for deviating from the equilibrium by making equal changes in their own output. This assumption is equivalent to a market-sharing agreement and leads to the collusive outcome.

In the open-loop equilibrium, each player chooses a sequence of changes in output, using a particular behavioral assumption, v. The equilibrium levels can be expressed in feedback form; in this case, strategies are open loop with revisions that are unanticipated. When players choose their current levels, they act as if they were also making unconditional choices regarding future levels.

In the feedback equilibrium, players recognize that their future choices will be conditioned on the future state; players select control rules rather than levels. The feedback equilibrium is obtained by the simultaneous solution of the n + 1 dynamic programming equation[5]

$$J_i(q_{t-\varepsilon}) = \max_{u_{i,t}} \left[\left(a\, e_i'\, q_t - \frac{1}{2} q_t'\, K_i\, q_t - \frac{1}{2} u_t'\, S_i\, u_t \right) \varepsilon + e^{-r\varepsilon} J_i(q_t) \right] \quad (3)$$

where $q_t = q_{t-\varepsilon} + u_t\varepsilon$. The particular behavioral assumption, v, determines the control rule for player i and his value function, $J_i(\cdot)$.

When v = 0, the result is the feedback Nash-Cournot game. It is well known that the open-loop and feedback equilibria differ in this case (e.g., Starr and Ho, 1969). When v = -1/n or v = 1, the open-loop and feedback equilibria are identical since, if players either take price as given or share the market in each period, it does not matter whether they choose levels or control rules.

3. Characteristics of the Model

In Karp and Perloff (1988), we derive a number of properties of these models analytically and through simulations. These properties are briefly summarized here:

1. For $v \in (-1/n, 1)$ and for given symmetric initial output level q_0, output at t is greater under the feedback equilibrium than under the open-loop equilibrium; convergence to steady state is faster in the latter case.

2. Industry profits are higher (and social surplus lower) under the open-loop equilibria. That is, feedback strategies are relatively procompetitive.[6]

3. For $v = -1/n$ or 1, the trajectories and control rules are identical in feedback and open-loop models.

4. Under both open-loop or feedback policies, output decreases in v.

The open-loop equilibrium can be obtained as the solution to a control problem for arbitrary v.[7] We use this result to illustrate the relationship between the feedback game and the control problem. Consider the standard control problem:

$$J(q) = \max_{u} \int_0^{\infty} e^{-rt} \left[ae' q - \frac{1}{2} q' Kq - \frac{\delta}{2} u' u\right] dt \tag{4}$$

subject to

$$\dot{q} = u, \qquad q_0 \text{ given,}$$

where $K = k_1 ee' + (k_0 - k_1) I$; that is, K is an $(n + 1) \times (n + 1)$ matrix with the parameter k_0 on the principal diagonal and the parameter k_1 elsewhere. In addition, K is positive-semidefinite and $\delta > 0$.

Define Q_t^i as aggregate output at t, where the superscripts ($i = c, 0, f$) indicate the paths given by the solutions to the control problem, open-loop, and feedback games, respectively. Assume that the intial output is the same for all firms, $q_0 = (e\, Q_0)/(n + 1)$, for each of the three cases. Then Q_t^i is the solution to

$$\dot{Q}_t^i = \gamma^i + \rho^i Q^i/\delta. \tag{5}$$

The proofs of the following four propositions, which are based on a comparison of the systems of equations that define γ^i and ρ^i, are in Karp and Perloff (1988):

Proposition 1. A sufficient condition for the open-loop and feedback equilibria to be identical is $v = 1$ or $v = -1/n$.

Simulation results (Karp and Perloff, 1988) indicate this condition is also necessary.

Proposition 2. Under the Nash-Cournot assumption, output is smaller in the open-loop equilibrium and converges to its steady state more rapidly than in the feedback equilibrium.

Thus, the feedback Nash-Cournot equilibrium is farther from the monopoly solution than is the open-loop Nash-Cournot equilibrium, so the feedback solution is relatively procompetitive. Fershtman and Kamien (1987) and Reynolds compare steady-state values in open-loop and feedback equilibria. Proposition 2 generalizes their results in that it compares the entire equilibrium path. The intuition is that, under the feedback assumption, capacity discourages rivals' investment. Therefore, firms have a greater incentive to invest today as a means of preempting their rivals' future investment. Thus, they develop larger capacities and hence larger output levels.

Since Proposition 1 provides a sufficient but not a necessary condition, we cannot prove that the comparison in Proposition 2 also holds for $v \neq 0$. Extensive simulations (Karp and Perloff, 1988), however, support the intuition that the result does hold for $v \neq 0$.

Two possible sources of error in estimating market structure in dynamic models are the assumption of a restrictive functional form and the assumption of open-loop rather than feedback strategies. If one maintains the linear-quadratic assumption, it is easy to determine the likely magnitude of the second source of error. This approach provides informal evidence of whether it is reasonable to tolerate the second source of error in an attempt to alleviate the first. Our simulation results suggest that the assumption of open-loop strategies is less serious when the market lies between Nash-Cournot and collusive than when it lies between Nash-Cournot and competitive (see Karp and Perloff, 1988).

We can formally show that the open-loop equilibrium can be obtained by solving a control problem:

Proposition 3. Aggregate output in the open-loop equilibrium and the control problem are the same if and only if $k_0 + nk_1 = [2 + n(v + 1)]\, b + \theta$. If $k_0 = (2 + nv)\, b + \theta$ and $k_1 = b$, the levels of output for the two problems are the same even if the n + 1 firms have different initial levels of output.

For the Nash-Cournot case ($v = 0$), Proposition 3 reproduces the result used by Hansen, Epple, and Roberds (1985). For the price-taking case, ($v = -1/n$), the integrand in (4) gives social surplus; this reproduces the well-known result that the competitive equilibrium can be obtained by solving the social planner's problem. For the collusive case ($v = 1$), the control problem that matches the open-loop game has $k_0 = (2 + n)\, b + \theta$, $k_1 = b$; whereas the control problem for the monopolist sets $k_0 = 2b + \theta$, $k_1 = 2b$. Therefore, the collusive game gives rise to the monopoly solution only if all firms produce equal quantities. Analogously, in static games, a conjectural variation of one produces the monopoly solution only in a symmetric equilibrium.

Proposition 3 shows that open-loop equilibria can be obtained by solving a control problem which depends on the parameters of the open-loop game. Therefore, estimation of the market structure under the assumptions of linear-quadratic functions and open-loop strategies is equivalent to estimating the parameters of a control problem. The methods described by Hansen and Sargent (1980) or Chow (1981) can be used.

Replacing the feedback game with a control problem requires knowing the solution to the game rather than merely the parameters of the game. Although of no computational assistance, the device is of interest for two reasons. First, it leads to the recognition that the feedback equilibrium with a homogeneous good is observationally equivalent to a particular open-loop game with heterogeneous goods.[8] This equivalence has obvious econometric implications.

Second, the device provides some intuition about why it is difficult to prove the feedback game is stable. Given the solution to the feedback game, it is possible to construct a K matrix such that the steady-state control rule of the resulting control problem is equal to the control rule of the game. However, the K matrix need not be positive definite, which violates one of the sufficiency conditions used to prove stability in the control problem. Details of this argument are provided in Karp and Perloff (1988).

For given $v \in (-1/n, 1)$, the equilibrium output depends on ε in the feedback game; output is independent of ε for the open-loop game. Consider, for example, the Nash-Cournot case where $v = 0$. Under the feedback model, a firm expects its rivals to react to its current decision only after an interval of ε. Its current decision, therefore, depends on ε. For the open-loop model, a firm expects no response on the part of its rivals (for $v = 0$), and the equilibrium is, therefore, independent of ε. Simulations show that the steady-state feedback output rises as ε falls (periods become shorter). Thus, the estimated index of market structure depends on the assumed period of adjustment in the empirical analysis. The importance of the period of commitment in dynamic Nash games has been noted by Reinganum and Stokey (1985).

Finally, as the number of firms $(n + 1)$ increases, the equilibrium trajectories change. By setting $\theta = 0$ and normalizing so that $\delta = (n + 1)c$ where $c > 0$ is constant, the price-taking and collusive equilibria are invariant to n. As n becomes large, the adjustment cost for each firm becomes infinite so each firm makes only infinitesimal adjustments and thus captures only an infinitesimal share of the market. Thus:

Proposition 4. Given $\theta = 0$ and the normalization $\delta = (n + 1)c$, the open-loop and feedback Nash-Cournot converge to the competitive equilibrium as $n \to \infty$.

From Proposition 2, the open-loop and feedback equilibria are identical for $v = 1/n$ which goes to 0 as $n \to \infty$. Since the open-loop model is a static game, it is well known that the Nash-Cournot equilibrium converges to the competitive equilibrium as $n \to \infty$. See Karp and Perloff (1988) for a rigorous proof.

4. Estimation Methods

Our objective is to obtain a consistent estimate of the index of market structure, v. In the process we can also estimate the adjustment parameter, δ. We estimate the adjustment equation,

$$\underline{q}_t = g(t) + G\,\underline{q}_{t-1}, \tag{6}$$

where $g(t)$ is an unrestricted function of exogenous variables and $\underline{q}_t$ is a vector of individual firms' output in time t. That is, 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 affine costs are constant over time and across firms. The elements of G are used to infer the parameter v.[9]

The most obvious reason for this estimation strategy is its simplicity. The type of market structure is logically distinct from the rational expectations hypothesis. If that hypothesis is true, there is a loss in efficiency from ignoring it, but our estimate is still consistent.

Throughout the estimation, we treat $\beta \equiv e^{-r\varepsilon}$ as known and common to all firms and assume $\theta = 0$. We designate as "symmetric" the case where the good is homogeneous ($p_i = a_i - b \Sigma q_j$), and all firms have the same δ. "Weak symmetry" allows demand to be of the form

$$p_i = a_i - b_1 q_i - b_0 \sum_{j \neq i} q_j. \tag{1'}$$

In either case it is natural to assume that the same value of v is common to all firms.

Define $\underline{v}_i$ as an n + 1 dimensional column vector with 1 in the ith position and v_{ij} elsewhere; $v_{ij} = \partial u_{jt}/\partial u_{it}$ for $i \neq j$. (This approach generalizes the previous section where $v_{ij} = v$ for $i \neq j$ was assumed.) Given an assumed value of β, an estimated matrix G and demand slope b (and hence K_i), $\underline{v}_i$ and δ for the open-loop equilibrium satisfy (see the Appendix)

$$K_i \underline{v}_i = [G^{-1}(I - G)(I - \beta G)]' e_i \delta. \tag{7}$$

The derivation of (7) does not depend on either symmetry assumption, but the solution of (7) requires either symmetry or a similar assumption. The matrix K_i is of rank 2 so, in general, there are either infinitely many solutions to (7) or no solutions. Given either symmetry assumption, equilibrium requires that the diagonal elements of G are equal as are off-diagonal elements. That is, G is of the form

$$G = \begin{pmatrix} g_1 & g_2 & \cdots & g_2 \\ g_2 & g_1 & & \vdots \\ \vdots & & \ddots & g_2 \\ g_2 & \cdots & g_2 & g_1 \end{pmatrix}. \tag{8}$$

If G is estimated subject to this restriction, all elements except the ith in the column vector on the right-hand side of (7) are equal and there exists a unique solution to (7); this assumes that $v_{ij} = v$ for all $i \neq j$.

More generally, suppose that the symmetry assumption is not used. For example, set i = 1 and let

$$p_1 = a_1(t) - \sum_{j=1}^{n+1} b_{1j}\, q_{jt}.$$

Then (7) can be rewritten as

$$\begin{pmatrix} 2b_{11} & b_{12} & \cdots & b_{1,n+1} \\ b_{12} & 0 & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ b_{1,n+1} & 0 & \cdots & 0 \end{pmatrix} \begin{pmatrix} 1 \\ v_{11} \\ v_{12} \\ \vdots \\ v_{1n} \end{pmatrix} = \begin{pmatrix} y_{11} \\ y_{12} \\ \vdots \\ y_{1,n+1} \end{pmatrix} \delta_1$$

where y_{ij} are obtained from the right side of (7) and depend only on β and the elements of G. Consistency requires

$$\frac{b_{ij}}{b_{ik}} = \frac{y_{ij}}{y_{ik}} \tag{9}$$

for all i and all j, $k \neq i$ and gives $n^2 - 1$ restrictions involving the demand system, b_{ij}, and the feedback system, G. Under the symmetry assumption and G as in (8), these restrictions are satisfied. This approach is the simplest but not the most general way to satisfy (9).

If (9) is satisfied, δ_i can be uniquely estimated as $\delta_i = b_{ij}/y_{ij}$, $j \neq i$. That is, given the estimated G matrix, δ_i is linear in the estimated demand slope coefficient(s). Using the previous equation to eliminate δ_i gives

$$\sum_{j \neq i} \frac{b_{ij}}{b_{ii}} v_{ij} = \frac{y_{ii}}{y_{ik}} \frac{b_{ik}}{b_{ii}} - 2 \tag{10}$$

for all i, $k \neq i$. There are n + 1 equations in (n + 1) n unknowns. An additional assumption, such as $v_{ij} = v_i \; \forall \; j \neq i$, is required. That is, for each firm, we can obtain an (possibly firm specific) aggregate index (or aggregate conjectural variation), but we cannot distribute this index over a firm's rivals.

If we assume that $b_{ij} = b_{ik}$, $\forall$ j, $k \neq i$ (a weaker condition than either symmetry assumption) and also that $v_{ij} = v_i$, then (9) and (10) simplify to

$$1 = \frac{y_{ij}}{y_{ik}} \qquad \forall \qquad j, k \neq i \tag{9'}$$

and

$$v_i = \frac{(y_{ii} - 2y_{ik})}{ny_{ik}}. \tag{10'}$$

Thus, it is possible to estimate the $(n + 1)^2$ elements of G subject to the $n^2 - 1$ restrictions of (9') and use (10') to infer v_i. This approach does not require estimation of the demand parameters, b_{ij}. Those parameters are necessary only to recover δ_i and, of course, to test whether the hypothesis $b_{ij} = b_{ik}$ ($j, k \neq i$) is reasonable.

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

$$w_i = [I - \beta(G' \otimes G')]^{-1} [(G' \otimes G') (\text{vec } K_i)]$$

$$x_i = [I - \beta(G' \otimes G')]^{-1} [(G' \otimes G') - (I \otimes G') - (G' \otimes I) + I] [\text{vec}(e_i e_i')]$$

where vec (Z) stacks the columns of the matrix Z. "Rematricize" (inverse vec operation) w_i and x_i to obtain the $(n + 1) \times (n + 1)$ matrices W_i and X_i. Notice that W_i is linear in firm i's demand coefficient(s) and that X_i depends only on β and G. If agents use feedback strategies, v and δ must satisfy (see the Appendix).

$$[K_i + \beta W_i + (e_i e_i' + \beta X_i) \delta_i] \underline{v}_i = G'^{-1} e_i \delta_i \equiv y_i^* \delta_i. \qquad (11)$$

Given the complexity of W_i and X_i, it is difficult to analyze (11) for the general case. However, for the symmetric case, the left side of (11) is of rank 2. Under the assumption of symmetry, the estimate of v is independent of b; recall that a stronger result holds in the open-loop model. Define the matrix A^i such that $b A^i \equiv K_i + \beta W_i$ and define $B^i \equiv e_i e_i' + \beta X_i$, so A^i and B^i depend only on β and G. To recover v and δ, rewrite the <u>i</u>th and the <u>k</u>th $(k \neq i)$ equation of (11) as

$$b(A_{ii} + v \sum_{j \neq i} A_{ij}) + (B_{ii} + v \sum_{j \neq i} B_{ij}) \delta = y_{ii}^* \delta \qquad (12a)$$

and

$$b(A_{ki} + v \sum_{j \neq i} A_{kj}) + (B_{ki} + v \sum_{j \neq i} B_{kj}) \delta = y_{ik}^* \delta, \qquad (12b)$$

where subscripts designate the element of A^i, B^i, and y_i, and the superscript i is suppressed. Solving (12b) for δ gives δ as a linear function of b and a nonlinear function of v. Substituting this function into (12a) gives a quadratic in v which is independent of b. Hence, under the symmetry assumption, v can be estimated with knowledge of only β and G.

Although there are two solutions to (12a), extensive simulation experiments show that one value is close to the open-loop value and that the other is implausible ($v < -1/n$, $v > 1$, or $\delta < 0$); therefore, in practice it is easy to choose the correct root. Using simulations, a comparison of the estimated v and δ under the open-loop and feedback models, maintaining the assumption of symmetry [and G as in (8)], shows that v and δ are larger under feedback (Proposition 2).

It is possible for a value (or values in the asymmetric case) of v and of δ to satisfy (9) or (11) without the implied game being meaningful. Given the solution to (9) or (11), it is necessary to check that each player's second-order conditions are satisfied and that the underlying Ricatti difference equations are stable. These tests are straightforward in the open-loop game. It is necessary only to check that $\delta > 0$ and that the K matrix defined in Proposition 3 is positive semidefinite. For the feedback game, it appears necessary to solve the game using the estimated δ and v; but this computation is straightforward.

We can test whether the equilibrium is open loop or feedback, in addition to estimating the degree of competitiveness (v). If symmetry is assumed so that G is estimated as in (8), exactly the same restrictions are imposed in estimating the parameters of demand and the control rule under both open loop and feedback. In order to distinguish the two, an overidentifying restriction is needed. For example, given information on cost, it would be possible to estimate jointly a cost function involving δ and the demand function and control rule subject to (9) or (11). In principle, one could apply methods of nonnested hypothesis testing or, less formally, compare the values of the likelihood functions under the two sets of restrictions. Unfortunately, reliable cost data are rarely available. Most firm-specific cost data are constructed by allocating total cost to a set of categories which does not include "adjustment." It would be surprising if, using this data, one could obtain a reliable estimate of δ.

In the absence of cost data it is, in principle, possible to test open-loop vs. feedback behavior by dropping the symmetry assumption. The estimation of b_{ij}, G, v_i^k, and δ^k (k = 0, f) subject to (7) or (11) will, in general, result not only in different estimates but also different values of the likelihood function. That is, in the absence of symmetry, the two sets of restrictions, (7) and (11), are not equivalent.

We demonstrated this nonequivalence by means of an example. For n = 2, we chose an arbitrary G, b_{11}, and b_{12} and then chose b_{13} to satisfy (9); by construction, a unique estimate of v_1 and δ_1 satisfies the open-loop restrictions (7). However, for these values of G and b_{1j}, the feedback restrictions constitute three independent equations in δ_1 and v_1, and no solution exists. In order to make these equations consistent, we could, for example, change b_{13}. In that case the open-loop restrictions would cease to be consistent. Therefore, the value of the likelihood function may be either greater or less under open loop: The open-loop and feedback models are observationally distinct even without cost data. Unfortunately, imposing the restrictions in (7) and (11), other than by using symmetry, is computationally difficult.

The previous example also demonstrates that the feedback model is capable of providing more information than the open loop. As mentioned above, it is possible to estimate only a single value v_i in the open-loop model. However, for the example we constructed, the feedback restrictions constitute three independent equations so it would have been possible to estimate v_{12} and v_{13}.

5. An Example: Coffee

We have applied our estimation technique to the coffee export market. As we are primarily interested in illustrating our basic technique, we use (and test) the symmetry assumption. Both classical and Bayesian methods are used to estimate the market structure parameter.

Background

Since 1959, most exporting and importing countries have participated in a series of International Coffee Agreements (ICAs). These agreements set quotas for exporting countries, but there is substantial evidence that many countries violate the agreement.

Gilbert (1986) surveys the various ICAs. He quotes critics who claim that the agreements "are no more than an internationally sanctioned producer's cartel" (p. 602); his judgment is that the agreements have resulted in higher prices rather than simply more stable prices (p. 604) but suggests that this view is controversial. Greenstone (1981) also asserts that the large coffee producers have behaved as a cartel. During 1974, in the absence of an ICA, Brazil and Colombia attempted to form an explicit producers' cartel; they were later joined by other (smaller) producers. Static econometric models including de Bries (1975), Akiyama and Duncan (1982), Palm and Vogelvang (1986), and Herrmann (1986) argue that the ICAs have resulted in higher prices for member importing countries but lower prices for nonmember importing countries.

The institutional arrangements in the coffee market constitute circumstantial evidence of an intent, by large producers, to exert market power.[10] The difficulties of negotiating these agreements, and the failure to comply fully with them, indicate that it is very unlikely that producers have behaved as monopolists. Thus, the hypothesis that the market structure lies "between" monopoly and competition is plausible. The complexity and inconstancy of the institutions and the paucity of data make it unreasonable to attempt to estimate the explicit game that producers play. The index v, described above, provides a measure of the market structure.

There are two reasons why it is important to use a dynamic model in attempting to estimate the market structure of coffee. First, changes in levels of production almost certainly involve nonlinear costs. There is a lag of 2 to 5 years between planting and the 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 costs of adjustment increase with the size of the change. Second, the two largest producers, Brazil and Colombia, maintain substantial stockpiles. Standard inventory models assume that there are nonlinear costs of adjusting the level of inventory (see, e.g., Blinder, 1982) and often approximate these costs using a quadratic function.

The costs of adjusting exports, therefore, consist of costly adjustment of production and/or inventories. A more fully specified model would include both stocks and lagged production level as state variables. Although in principle additional state variables may be introduced, there is a substantial increase in the complexity of the calculations, and much heavier demands are placed on the data. We, therefore, treat lagged exports as the state variable and interpret the costly adjustment of this variable as consisting of the adjustment costs of production and of inventories. This approach is defended as an approximation; its chief merit is that it leads to an estimable model.

We examine the period 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 rest of South America's share was 4 percent, the rest of Latin America's was 17 percent, Asia's was 6 percent, and Africa's was 29 percent. The two largest African exporters' shares were 7 percent (Ivory Coast) and 5 percent (Uganda).

For simplicity, we assume that Brazil and Colombia are engaged in a dynamic game in which they treat the rest of the world as a fringe with exogenous output. Most Latin American and some African countries produce Arabica (mild) coffee, which is 70 percent of all coffee produced; most African and Asian countries produce Robustas coffee (a harsher coffee), which is mainly used in instant coffees and is an imperfect substitute for Arabica coffee.

Many countries are thought to violate the ICA quotas regularly, especially the African countries. At various times, especially during the extended negotiations on ICA from 1978-1980, Brazil and Colombia were thought to have intervened in markets to maintain stability.

It is reasonable to treat Brazil and Colombia as "firms" since 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.

Quantity data for this study come from Coffee: World Coffee Situation (various years) published by the U. S. Department of Agriculture, Foreign

Agriculture Service. The price of coffee is an average of the prices of all coffee traded in the New York market, the major market for coffee. Price data, 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).

Estimation

A linear demand curve, equation (1), that treats Arabica coffee and Robustas coffee from various countries as imperfect substitutes was estimated using instrumental variables. Other right-hand side variables included world gross domestic product and a beverage price index. All prices were deflated by the world commodity wholesale price index. The coefficient b on Arabica coffee was 0.00002874 with a t-statistic (for $b = 0$) of 2.42. The equation had an $R^2 = 0.97$ and a Durbin-Watson statistic of 1.85. We do not discuss the demand results further, since b only affects the estimated value of δ and not v.

The adjustment equations (6) were estimated 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 major freeze in Brazil in 1977-78. We imposed the cross-equation symmetry constraints that the coefficient on the own lagged exports was equal across equations as was the coefficient on the other country's lagged exports. That is $g_{11} = g_{22} \equiv g_1$, and $g_{12} = g_{21} \equiv g_2$ where g_{ij} is the (i,j) element of G. The F-statistic on these restrictions was 0.64 with 2 and 34 degrees of freedom, so we cannot reject them. That the coefficients on the lagged exports are zero independently or collectively can be rejected on the basis of t-tests, F-tests, and likelihood ratio tests.

TABLE 1

Coffee Equations
Adjustment Equations (Export Quantity on Lagged Export Quantities)

	Brazil	Colombia
Constant	12,986.0 (4.99)[a]	6,967.9 (4.28)
Brazilian freeze (1977-78)	-9,980.7 (-4.66)	843.7 (0.74)
Time (1, 2, ...)	22.4 (0.30)	124.8 (2.59)
Lagged Brazil exports	0.302 (2.27)	-0.192 (-2.42)
Lagged Colombia exports	-0.192 (-2.42)	0.302 (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.

Based on these estimated coefficients, and assuming $\beta = .95$, the open-loop model implies a v^o of -0.84, while the feedback model implies a v^f of -0.80 as shown in Table 2. Using a Taylor expansion, the standard errors on these two estimates of v are 0.27 and 0.31, respectively. Thus, we cannot reject the $v = -1$ hypothesis (price taking) but can reject the $v = 0$ and $v = 1$ hypotheses at the 0.05 level.[11]

We can use a Bayesian approach to obtain alternative estimates of v and of the hypothesis tests. Geweke (1986) has developed a method using Monte Carlo numerical integration to impose inequality constraints in a normal linear regression model.[12] Importance sampling (using a multinomial normal) is used. We need to impose three sets of restrictions (Bayesian priors). First, the system is stable ($-1 < g_1 + n\, g_2 < 1$ and $-1 < g_1 - g_2 < 1$). Second, the market structure lies between monopoly and price taking, $1 > v^k = \phi_i(G) > -1/n$ (k = o or f) where the v^k are nonlinear functions, ϕ_i, of the terms in G as shown in (7) and (11). Third, $\delta^k = \Psi(G) \geq 0$ (k = o or f), where the δ^k are the adjustment parameters for the two models. Our classical point estimates of the elements of G and our estimates of v^k meet these restrictions.

Using Geweke's Bayesian approach, we can test these restrictions as shown in Table 2 (based on 5,000 importance-sampling replications). The stability conditions are almost always met. In a quarter of the replications, $v^k < -1$, $v^k > 1$ or $\delta^k < 0$. Thus, since in roughly three-quarters of the cases all the restrictions hold, imposing these restrictions seems reasonable.

Imposing the restrictions, we obtain distributions of v^o and v^f. The values that minimize the loss function are the mean if we use a quadratic loss function and the median if we use an absolute difference loss function (Zellner, 1971, pp. 24-25).

TABLE 2

Classical and Bayesian Inequality Constrained Estimates

	v^o	v^f		
Classical estimates				
v^k (unrestricted)	-0.84	-0.80		
Standard deviation (Taylor approximation)	0.27	0.31		
Bayesian inequality constrained estimates	Importance sampling[a]		Bootstrap[b]	
	v^o	v^f	v^o	v^f
Quadratic loss				
v^k (mean)	-0.65	-0.62	-0.68	-0.63
Standard deviation (sd)	0.35	0.36	0.43	0.44
Precision of the mean of v^k ($sd/\sqrt{t}$)	0.0059	0.0060	0.0096	0.0097
Absolute loss				
v^k (median)	-0.76	-0.73	-0.86	-0.81
Standard deviation	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	1.19	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

[a]5,000 replications (t).

[b]2,000 replications (t).

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

The importance-sampling approach assumed that the likelihood function was normal. We can relax the assumption of normality by using a bootstrapping approach, as reported in Table 2.[13] The bootstrap estimates show slightly higher standard deviations corresponding to the mean v^k estimates (0.43 and 0.44), a lower probability of rejecting due to $v^k \leq -1$ (16 percent), and a higher probability of rejecting due to $v^k < 1$ (10 percent).

Using the Bayesian estimates, we can calculate the probability that the market structure or behavior parameter v^k lies within a certain range since we have an entire distribution for v^k. Some of the interesting ranges are summarized in Table 3. Assuming normality, the probability that v^k lies between -1 (price taking) and 0 (Nash-Cournot) is greater than 90 percent. There is a slightly higher probability that v^k lies between the classical estimate and 0 (nearly 50 percent) than that it lies between -1 and the classical estimate (over 40 percent). Two-thirds of the distribution lies below the mean (quadratic loss) estimates of v^k.

The bootstrap distribution has thicker tails, as shown in Table 3. For example, the probability that v^f lies in the right tail between 0 and 1 is 10.7 percent using bootstrapping and 7.2 percent with normality. The probability that v^f lies in the left tail between price taking (-1) and the classical estimate (-0.80) is 51.7 percent in the bootstrap distribution and 35.1 percent in the distribution based on normality.

TABLE 3

The Distribution of v^k Based on Bayesian Estimates

Proportion of weight between[a]		Normality v^o	Normality v^f	Bootstrap v^o	Bootstrap v^f
		percent			
-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^{kc}	34.0	35.1	53.7	51.7
v^{kc}	0	59.5	57.7	36.6	37.6
-1	v^{kb}	67.2	65.5	72.9	71.1
v^{kb}	0	26.3	27.3	17.5	18.2

[a]The classical estimate is v^{kc} (k = o or f) and v^{kb} is the Bayesian estimate based on a quadratic loss function.

We have assumed that the Brazilians and Colombians are dynamic oligopolists who face an exogenous fringe. Based on either the classical or Bayesian approaches, their behavior appears close to price taking. The probability that they are at least as noncompetitive as Nash-Cournot is no more than 11 percent.

Simulations

These estimates have implications for steady-state outputs. We normalize so that, if the two countries were price takers, each country's steady-state output would be 100. If the countries played open-loop Nash-Cournot, the steady-state output would be 66.67; and if they were a perfect cartel, their output would be 50.

Using the classical point estimates, the open-loop ($v = -0.84$) steady-state output is 93. The feedback ($v = -0.83$) steady-state output is 95. Using the Bayesian quadratic loss estimates, the open-loop ($v = -0.65$) output is 86 and the feedback ($v = -0.62$) output is 87.

Thus, while the estimated market structure is "close" to price taking in the sense that v is close to -1, the steady-state outputs are below the price-taking levels by between 7 and 14 percent, depending on the model and the market structure estimate used.

The cost of adjustment affects the steady state in the feedback model. In a static model with a $v = -0.83$ (classical estimate), each firm produces 92, but in the dynamic model the steady-state output is 95. Similarly, with a $v = -0.67$ (Bayesian quadratic loss estimate), the static output is 84 and the dynamic model's steady-state output is 87. Thus, where costs of adjustment are positive, more is produced in the steady state.

6. Summary and Conclusions

This paper develops a method to determine the degree of competition among dynamic oligopolists and to test whether they use open-loop or feedback strategies. New families of open-loop and feedback models that nest behavioral assumptions are presented.

The linear-quadratic estimation method is easily implemented, as shown in an application to the coffee export market. Both classical and a Bayesian technique imposing inequality constraints are demonstrated. The Bayesian results are generally close to the classical ones. An empirical Bayesian approach that drops the standard normality assumption also produces similar results. Moreover, the bootstrap standard errors are similar to the classical Taylor expansions. One advantage of the Bayesian approach is that an empirical approximation to the distribution of the behavior strategy variable is obtained, so the probability that market structure lies in a particular range can be easily calculated. In our example the steady-state outputs may be between 7 and 14 percent lower than if Brazil and Colombia were pure price takers.

Appendix

Derivation of Restrictions

Since we are only interested in imposing restrictions on the demand slopes and the coefficients on q in the control rule and not on the intercepts of the demand and control systems, we can restrict our attention to the quadratic part of the problems.

The Open-Loop Restrictions

The Lagrangian for the ith player is

$$Q_i = \sum_{\tau=t}^{T} \beta^{\tau-t} \left[-\frac{1}{2} q_\tau' K_i q_\tau - \frac{1}{2} \mu_\tau' S_i \mu_\tau + \lambda_{i\tau}'(q_{\tau-1} + \mu_\tau - q_\tau) \right] .$$

The first-order conditions for q_τ and $\mu_{i,\tau}$ are

$$-K_i q_\tau - \lambda_{i\tau} + \beta \lambda_{i,\tau+1} = 0 \tag{A.1a}$$

and

$$-\underline{v}_i' S_i \mu_\tau + \underline{v}_i' \lambda_{i,\tau} = 0. \tag{A.1b}$$

Use $\lambda_{i,\tau} = H_{i,\tau} q_\tau$, which can be shown by induction, and let $T \to \infty$ so that $H_{i\tau} \to H_i$; (A.1b) becomes

$$\underline{v}_i' H_i q_\tau = \delta_i \mu_\tau, \qquad i = 1, n + 1.$$

Stack up these conditions to obtain $E q_\tau = S \mu_\tau$ where the ith row of E is $\underline{v}_i' H_i$ and the ith row of S is $\delta_i e_i'$. This gives $E q_\tau = S(q_\tau - q_{\tau-1})$ or $q_\tau = G q_{\tau-1}$ where $G \equiv (S - E)^{-1} S$.

Use the previous definitions to rewrite (A.1a) as

$$0 = -K_i q_\tau - H_i q_\tau + \beta H_i q_{\tau+1} = (-K_i - H_i + \beta H_i G) q_\tau = 0$$

so

$$H_i(I - \beta G) = -K_i$$

and

$$H_i = -K_i(I - \beta G)^{-1}.$$

From the definition of G, we have

$$E = S(I - G^{-1}).$$

Premultiply both sides by e_i' and use the definition of E and the previous expression for H_i to obtain

$$e_i' E_i = \underline{v}_i' H_i = -\underline{v}_i' K_i(I - \beta G)^{-1} = e_i' S(I - G^{-1}) = \delta_i e_i'(I - G^{-1})$$

so that

$$-\underline{v}_i' K_i = \delta_i e_i'(I - G^{-1}) (I - \beta G)$$

and

$$K_i \underline{v}_i = -[(I - G^{-1}) (I - \beta G)]' e_i \delta_i.$$

Factoring out $G^{-1'}$ gives (7).

The Feedback Restrictions

The stationary dynamic programming equation is

$$- \frac{1}{2} q_{t-1}' H_i q_{t-1} = \max_{q_{i,t}} - \frac{1}{2} q_t'(K_i + S_i + \beta H_i) q_t + q_t' S_i q_{t-1} - \frac{1}{2} q_{t-1}' S_i q_{t-1}. \tag{A.2}$$

The first-order condition is

$$-\underline{v}_i'(K_i + S_i + H_i) q_t + \underline{v}_i' S_i q_{t-1}.$$

Stack the n + 1 first-order conditions to obtain

$$E q_t = S q_{t-1}$$

where the _i_th row of E is $\underline{v}_i'(K_i + S_i + H_i)$ and the _i_th row of S is $\delta_i e_i'$. Rewrite this as $q_t = G q_{t-1}$ where $G \equiv E^{-1} S$. Substitute this into the maximized value of (A.2) to obtain

$$H_i = G'(K_i + S_i + \beta H_i) G - G' S_i - S_i G + S_i.$$

Apply the vec operation and simplify to obtain vec $H_i = w_i + x_i \delta_i$ where w_i and x_i are defined in the text. Rematricize to obtain $H_i = W_i + X_i \delta$. Take the _i_th row of E G = S and use the definition of E to obtain

$$\underline{v}_i'[K_i + e_i e_i' \delta_i + \beta(W_i + X_i \delta_i)] G = e_i' \delta_i.$$

Rearranging this equation gives (11).

Footnotes

[1]The linear-quadratic cost-of-adjustment model has been used extensively (Sargent, 1978; Hansen and Sargent, 1980; and Blanchard, 1983). Our model is a generalizaton of those studies that assume a competitive structure. The paper by Hansen, Epple, and Roberds (1985) uses the dynamic linear quadratic model to study different open-loop markets as well as the open-loop and feedback Stackelberg models. It does not compare the open-loop and feedback symmetric firms markets which is the focus of this paper. Fershtman and Kamien (1987) and Reynolds (1987) compare the open-loop and feedback linear-quadratic Nash-Cournot models in theoretical models only.

[2]If open-loop strategies are assumed, there are at least two alternatives to the linear-quadratic model. One uses instrumental variables to estimate the game analog of the stochastic Euler equations (as in Hansen and Singleton, 1982; and Pindyck and Rotenberg, 1983). Similar methods could be used to estimate noncompetitive markets; but since the Euler equations restrict the equilibria to be open loop, this approach is not pursued here. The second method uses dynamic duality (McLaren and Cooper, 1980; and Epstein, 1981). Although, in principle, this method could be used to estimate both open-loop and feedback noncompetitive equilibria, it implies very complicated restrictions for the feedback case and may be of limited practical use.

[3]Gallop and Roberts (1979) implicitly made their conjectural variations a function of another variable in a static model.

[4]Treadway (1970) showed that the comparative statics of the steady state of cost-of-adjustment models differ from those of the "corresponding" static model. In a similar vein, Reynolds finds that the output under static Nash-Cournot and at the steady state of the open-loop dynamic Nash-Cournot models

are different. However, under the assumption that adjustment costs are minimized when adjustment is 0 (i.e., at the steady state), these results no longer hold. This assumption seems reasonable if the objective is to compare the various dynamic models with their static analogs.

[5]It is well known that, for infinite horizon games, there typically exist many equilibria even when these are required to be subgame perfect. We avoid the problem of nonuniqueness by considering the equilibrium strategies that result from the game with finite horizon T and letting $T \to \infty$.

[6]Feedback policies require knowledge of the current state (output of all firms in the previous period), so a possible policy conclusion is that this information should be made available. However, this conclusion ignores the likelihood that the degree of collusion, measured by v, may increase as information is shared. Riordan (1985) models a dynamic oligopoly with stochastic demand where firms are unable to observe their rivals' output. He concludes that aggregate output is greater in this case than in the case where firms are able to observe their rivals' output. Riordan's model is quite different from the current one; nevertheless, the conflicting conclusions illustrate the difficulty of a general comparison of social welfare when firms do or do not know their rivals' output.

[7]Hansen, Epple, and Roberds (1985) show this result for the Nash-Cournot assumption ($v = 0$).

[8]This statement is actually too strong. Suppose that the game were completely stationary and firms completely symmetric so that it were practical to impose the restrictions implied by the constant part of the control rules. In that case the slope coefficients of the control rules of a homogeneous firm game with feedback strategies would be the same as the slope coefficients of

the rules of a heterogeneous (but symmetric) firm game with open-loop strategies; the intercepts would be different so the two could still be distinguished. However, for the econometric work, we do not wish to impose the restrictions on the intercepts of the control rules, so that nonstationarity or firm-specific features may be included in the parameters a, θ_0, and δ_0.

[9]We could impose or test the rational expectations hypothesis by including an exogenous state vector of current information in the feedback game (see Chow, 1981). A large literature (see the cites in footnote 1) shows how to do the same in an open-loop model.

[10]Marshall (1983) provides a detailed description of the world coffee market.

[11]Since Brazil and Colombia are not equal in size, $v = 1$ leads to the collusive solution only in the long run where their exports become equal. During the sample period, Colombia's share did rise relative to Brazil's.

[12]See Geweke (1986) and Chalfant and White (1988) for applications. Geweke has a brief discussion on lagged endogenous variables. Chalfant and White present a multi-equation generalization.

[13]See Freedman and Peters (1984) on bootstrapping with lagged endogenous variables. We assume (and classical tests support) the assumption that there is no autocorrelation. Thus, we can bootstrap by choosing rows of the original data (left- and right-hand-side variables). Experiments with the alternative method of choosing estimated standard errors produced similar results. This approach of using bootstrapping to replace the assumption of normality in Geweke's approach is suggested, but not employed, by Chalfant and White.

References

Appelbaum, E. "The Estimation of the Degree of Oligopoly Power," Journal of Econometrics, Vol. 19 (1982), pp. 287-299.

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

Blanchard, Olivier J. "The Production and Inventory Behavior of the American Automobile Industry," Journal of Political Economy, Vol. 9 (1983), pp. 365-400.

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

Chalfant, James A., and Kenneth J. White. "Estimation and Testing in Demand Systems with Concavity Constraints," University of California, Department of Agricultural and Resource Economics, Working Paper No. 454, Berkeley, 1988.

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

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

Epstein, Larry G. "Duality Theory and Functional Forms for Dynamic Factor Demands," Review of Economic Studies, Vol. XLVIII (1981), pp. 81-95.

Fershtman, Chaim, and Morton R. Kamien. "Dynamic Duopolistic Competition with Sticky Prices," Econometrica, Vol. 55 (1987), pp. 1151-1164.

Freedman, David A., and Stephen C. Peters. "Bootstrapping an Econometric Model: Some Empirical Results," Journal of Business and Economic Statistics, Vol. 10, No. 2 (April, 1984), pp. 150-158.

Gallop, F. M., and J. M. Roberts. "Firm Interdependence in Oligopoly Markets," Journal of Econometrics, Vol. 10 (1979), pp. 313-331.

Geweke, J. "Exact Inference in the Inequality Constrained Normal Linear Regression Model," Journal of Applied Econometrics, Vol. 26 (1986), pp. 295-321.

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

Greenstone, W. D. "The Coffee Cartel: Manipulation in the Public Interest," Journal of Futures Markets, Vol. 1 (1981), pp. 1-18.

Hansen, Lars Peter, D. Epple, and W. Roberds. "Linear-Quadratic Duopoly Models of Resource Depletion," in Energy Foresight and Strategy, ed. Thomas J. Sargent, Resources for the Future, Washington, D. C., 1985.

Hansen, Lars Peter, and Thomas J. Sargent. "Formulating and Estimating Dynamic Linear Rational Expectations Models," Journal of Economic Dynamics and Control, Vol. 2 (1980), pp. 7-46.

Hansen, Lars Peter, and Kenneth J. Singleton. "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," Econometrica, Vol. 50 (1982), pp. 1269-1286.

Herrmann, Roland. "Free Riders and the Redistributive Effects of International Commodity Agreements: The Case of Coffee," Journal of Policy Modeling, Vol. 8, No. 4 (Winter, 1986), pp. 597-621.

Iwata, G. "Measurement of Conjectural Variations in Oligopoly," Econometrica, Vol. 42 (1974), pp. 947-966.

Karp, Larry S., and Jeffrey M. Perloff. "Open-Loop and Feedback Models in Dynamic Oligopoly," unpublished manuscript (1988).

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

McLaren, Keith R., and Russel J. Cooper. "Intertemporal Duality: Application to the Theory of the Firm," Econometrica, Vol. 48 (1980), pp. 1755-1763.

Palm, F. C., and E. Vogelvang. "A Short-Run Econometric Analysis of the International Coffee Market," European Review of Agricultural Economics, Vol. 13 (1986), pp. 451-476.

Pindyck, Robert S., and Julio J. Rotenberg. "Dynamic Factor Demands and the Effects of Energy Price Shocks," American Economic Review, Vol. 73 (1983), pp. 1066-1079.

Reinganum, Jennifer F., and Nancy L. Stokey. "Oligopoly Extraction of a Common Property Natural Resource: The Importance of the Period of Commitment in Dynamic Games," International Economic Review, Vol. 26 (1985), pp. 161-173.

Reynolds, Stanley S. "Capacity Investment, Preemption and Commitment in an Infinite Horizon Model," International Economic Review, Vol. 28 (February, 1987), pp. 69-88.

Riordan, Michael H. "Imperfect Information and Dynamic Conjectural Variations," Rand Journal of Economics, Vol. 16 (1985), pp. 41-50.

Sargent, Thomas J. "Estimation of Dynamic Labor Demand Schedules under Rational Expectations," Journal of Political Economy, Vol. 86 (1978), pp. 1009-1044.

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

Sumner, D. "Measurement of Monopoly Behavior: An Application to the Cigarette Industry," Journal of Political Economy, Vol. 89 (1981), pp. 1010-1019.

Treadway, Arthur. "Adjustment Costs and Variable Inputs in the Theory of the Competitive Firm," Journal of Economic Theory, Vol. 2 (1970), pp. 329-347.

Zellner, Arnold. An Introduction to Bayesian Inference in Econometrics. New York: John Wiley and Sons, Inc., 1971.