COURNOT VERSUS BERTRAND: A DYNAMIC RESOLUTION

by

Kenneth L. Judd
Hoover Institution and
National Bureau of Economic Research

May 1989

Latest Revision April 1996

I thank Ray Deneckere, Avinash Dixit, Garth Salonar, Michael Riordan, and Michael Whinston, and participants of seminars at the Hoover Institution, Stanford University, Northwestern University, Princeton University, University of California-Berkeley, and Harvard University for helpful comments. I acknowledge the excellent research assistance of Kwang-Soo Cheong and the research support of NSF grants SES-8810935 and SES-9012128.

inventors.tex

Cournot Versus Bertrand: A Dynamic Resolution

1: Introduction

Formal analysis of oligopoly has focussed on two basic models: Cournot and Bertrand. Cournot analysis assumes that a firm determines its sales while price is determined by some unspecified agent so that market demand equals the total amount offered. Bertrand analyis assumes that a firm determines the price at which it sells its output with firms being absolutely obligated to immediately meet the resulting customer demand. Unfortunately, these models have very different predictions in many applications. This paper examines a dynamic game wherein the firm sets both its output and its price, imposing neither of the artificial and unrealistic strategic limitations of the Cournot and Bertrand models. In doing so, we find that noncooperative dynamic oligopoly looks like neither simple static model in general but, in a limited sense, incorporates both as special cases. Despite its hybrid nature, followup work has shown that this model is less ambiguous when addressing specific issues, such as merger incentives. We show that dynamic analysis incorporating both output and price strategies is tractable, and argue that it is more representative of actual markets.

It is well-known that the Cournot and Bertrand models are substantial simplifications of firms' strategies, but are used in a hope to approximate the "true" game. Since neither is a clearly superior approximation, both are used whenever a specific issue is examined. While this is great for fattening journals and vitas, the result has been a rich diversity of theories yielding no clear predictions about important questions since the answers generally depend on the choice of strategic instrument. For example, Gal-Or (1985) shows that the incentives of oligopolists to share information depends whether competition is in prices or quantities. Fershtman and Judd (1987) show that owners' incentives to strategically manipulate the incentives of their managers depends on whether competition is in price or quantity. Even the question of whether firms want to merge depends on the model used (see Salant et al.). Bulow, Geankoplos and Klemperer (1985) (BGK) show that this ambiguity is pervasive in oligopoly theory. They argue that a critical feature of any model is whether the strategic variables are strategic complements or strategic substitutes, that is, whether a marginal increase in my strategy will increase or decrease the marginal profitability of my opponents'

strategies. They show that many standard arguments concerning strategic investment, entry deterrence, product proliferation, and learning curves are reversed when reasonable alternative strategic formulations are used. While the BGK focus is often helpful, it is limited to special contexts where all variables are either strategic complements or complements.

This sensitivity to strategic formulation is a major weakness of oligopoly theory, and leads us to ask how the correct specification of oligopolistic interaction is related to the structural elements of an industry. This problem has been addressed in several inventive ways. Some have examined the issue by slightly extending the static game into a dynamic model. A good example of this is the analysis of Kreps and Scheinkman (1983). They argue that if firms first choose their capacity, and only later are allowed to commit to a price, the outcome will be the Cournot equilibrium. This is arguably a reasonable game to examine since it models the intuition that capacity choices are made prior to any binding price decisions. However, they make special assumptions concerning demand, costs, and rationing. In particular, Davidson and Deneckere (1986) have shown that this result is sensitive to reasonable but not compelling assumptions made by Kreps and Scheinkman concerning the off-equilibrium path disposition of residual demands.

Others have appended the usual static game with a preliminary game wherein the players register their preferences over the games to be played. Leading examples of this approach are Singh and Vives (1984), and Klemperer and Meyer (1986). While these analyses indicate which game the players want to play, they suffer from a lack of realism since we never observe such a game being played. They argue that such choices are implicit in investment decisions. While that may be somewhat true, few options force either extreme mode of behavior. A robust analysis of these issues would permit the continuum of possibilities which lie between the extremes of precommitment to price and precommitment to quantities.

The conjectural variations literature is another attempt which tries to represent dynamic considerations in a static framework. However, the lack of any extensive form game basis for conjectural variation analysis makes it difficult to evaluate. This is particularly unfortunate since it is often used as a basis for empirical analysis of strategic interdependence (e.g., see Gollop and Roberts). However, our model could be viewed as a foundation for the empirical

version of conjectural variations since we can compute the reaction of a firm to recent (but past) behavior of competitors. This view of conjectural variations is the same as examined in Riordan (1985), the difference being that Riordan concentrated on equilibrium reactions to information flows, whereas we are examining reactions to a different set of strategies. [1]

This paper begins with the premise that efforts to determine which static game is correct are futile since both are so clearly wrong. In the real world, firms leave neither output nor price for an automatic mechanism to determine. Firms choose both the price at which they sell their output and the amount which is produced; any formulation of firms' decisionmaking which eliminates either choice is implausible, particularly if answers to important questions depend on which strategy is arbitrarily thrown out of the analysis. This view is implicit in the earlier work cited above, but those models do not allow simultaneous and repeated choices of output and price, dynamic features of real life.

In static models, firms cannot be allowed to choose simultaneously both price and quantity [2] since the final price-quantity outcome must be consistent with consumer demand. To avoid these problems, this paper adds a generally neglected aspect of the real world: inventories. There has been some effort to include inventories in oligopoly analysis, with Kirman and Sobel(1974) being the earliest and almost only example of dynamic oligopoly analysis with inventories [3]. If a firm has the ability to draw on existing inventories when current demand exceeds current output, or augment them when current demand falls short of current output, then it can choose both the current price and the current level of output. The realism of such a specification is clear since firms do choose both the price at which they sell their output and the amount which they produce, and unsold output generally survives (albeit at a cost and with some decay) to be offered for sale the next day. This is true even

[1] In fact, future work will concentrate on integrating the strategic concerns of this paper and Riordan's informational concerns.

[2] Strictly speaking, it is not necessary to choose between the strategic variables if rationing rules are specified to handle the cases where demand and output are not equal. We do not pursue analysis of such games here because it would be contrary to this essay's point that any such static analysis is, at best, an unreliable approach to modelling a continuing stream of simultaneous output and price decisions by firms in a dynamic world.

[3] See also Arvan (1985), which focussed on general existence problems, assumed Cournot competition, and assumed only a two period horizon. Inventories have also been used by Rotemberg and Saloner(1989) to examine collusion incentives in otherwise standard oligopoly models; this paper focusses on non-supergame phenomena

of some "perishables", such as "fresh" milk. In most industries, unsold output does not decay immediately, and unexpectedly heavy demand can be met out of inventories as well as current output.

Adding inventories to our analysis is really not new, since they are *implicit* in all oligopoly models if those models are to be taken as snapshots of a dynamic reality. The assumptions implicit in oligopoly models are, however, never mentionned. One way to view the Cournot and Bertrand models is that they implicitly assume that any nonzero level of inventories is infinitely painful for firms; therefore, the Bertrand model forces each firm to produce realized output, and the Cournot forces each firm to sell all output. Inventories are more explicitly present and somewhat more reasonably handled in Kreps and Scheinkman, and Davidson and Deneckere, where the value of unsold output is zero. However, in these models, there is a cost to not having enough output to meet demand, implying that the marginal value of inventories in equilibrium has a discontinuity at zero. If we were at the end of time or if the good were infinitely perishable, the zero marginal valuation of inventories is reasonable, but outside of those cases a smoother treatment of the valuation of inventories is more reasonable. We are going to make less extreme assumptions concerning the costs of inventory holding.

The advantage of adding inventories to our analysis is also clear when we want to investigate particular issues of how investment and output are distorted for strategic purposes. Since we do not have to make a choice between price and quantity as the strategic variable, results concerning such issues cannot be sensitive to that choice. An appeal to strategic complementarities will not suffice since we will have a mixture of strategic substitutes and complements. We also argue that there is no tractibility excuse for focussing on the simpler Bertrand and Cournot models.

One approach may be to just add a value function for inventories. That, however, is not an appropriate way to proceed since then results will depend critically on the arbitrary assumptions made concerning the valuation of inventories. In reality, the value of inventories is not exogenous to the game since it is determined by the nature of future competition. Furthermore, changes in the nature of competition, such as when a merger occurs, will affect the valuation of inventories. Therefore, if we want to examine issues related to changes in

market structure, it is not legitimate to treat the valuation of inventories as being exogenous. Therefore, we specify the cost of inventories, a natural technical parameter, but then examine games of arbitrary duration. In this way we do not let arbitrary specifications of the terminal valuation of inventories affect the results. In this way (as well as others) this approach differs from that of Friedman (1988) who also discusses price and quantity competition, but with a two-stage game.

We examine the basic nature of competition, determining when the dynamic equilibrium looks Cournot, Bertrand, or of some intermediate nature. In order to focus clearly on the main ideas we examine linear-quadratic models, a particularly tractable class. These models assume quadratic utility functions, quadratic inventory holding cost functions, and quadratic production costs. If we stopped here, some may argue that the contest was rigged because our model does not allow for the precommitment aspect of output which is part of the appeal of the Cournot model, as formalized in Kreps and Scheinkman. To examine these arguments, we next also allow adjustment costs. With these models, we will examine not only the actual equilibrium prices, but also the strategic effects of investment, learning curves, and strategic incentive manipulation, issues which are all sensitive to strategic formulation.

In these models we generally find that neither model is consistently a better approximation of our dynamic game. If short-run marginal costs are invariant to the rate of current and past output and investment, then the Bertrand model appears to do better, where by "better" we mean that prices tended to be closer to the static Bertrand level. However, if there are substantial decreasing returns to scale or adjustment costs, then the Cournot model does better in predicting the long-run nature of equilibrium. Our computational approach is able to make some judgments about the likelihood of these cases. For example, examination of the structural parameter values in a linear-quadratic version of our model and their implication for the critical elasticities indicates that the adjustment costs must be relatively high for the Cournot model's predictions to dominate.

Some aspects of linear-quadratic structures are unappealing; for example, they assume that inventory holding costs are symmetric around some cost-minimizing level, and negative inventories are permissible. To address concerns of robustness along those dimensions, ??

discusses numerical partial differential equation techniques which can be used to examine general models. A discussion of those results would take us far afield from the basic story here, however initial results indicate that linear-quadratic model results are robust to more reasonable specifications of inventory holding costs. However, much more work is needed.

2: General Model

We first describe a general duopoly model of differentiated products with inventories, abstracting from adjustment costs, learning, and investment. Firm i faces a demand for its product in period t equal to $d_t^i = D^i\left(p_t^1,\, p_t^2\right),\ i = 1,\ 2$. To assure the continuity of profits in prices, we assume that each firm's demand is a continuous function of both firms' prices. While this rules out the case of perfect substitutes, it is an inessential restriction since we can come arbitrarily close to the perfect substitutes case. Note that the presence of some imperfect substitution implies that the Bertrand equilibrium price will not be marginal cost.

We assume that production costs for firm i in period t equal a convex function of output, $C_i(q_t^i)$. If firm i's inventory at the beginning of period t is I_t^i, the cost to firm i of holding inventory is assumed to be a convex function of I_t^i, $H_i(I_t^i),\ i = 1,\ 2, t,\ i = 1,\ 2$. At this point we allow negative inventories. They can be thought of as unfilled orders, with the holding cost representing the cost of keeping customers happy waiting for delivery. The possibility of negative inventories is not pleasing, but we will make specifications which should reduce any undesirable properties which might arise. For example, we may assume that a positive level of inventories is desirable for standard operations management reasons. For example, grocery stores carry inventories of milk, not for competitive reasons but to minimize handling costs. In fact, the only way to avoid having inventories of milk held by someone is to have cows behind the store, certainly not a low cost alternative. Since we will usually make assumptions about H which model these considerations, our equilibria will generally involve positive inventories. In that case, negative inventories are contingencies which are away from the equilibrium paths which we will compute, hopefully making assumptions about these contingencies less important. While there may be some weaknesses in our formulation of inventories, these assumptions allow us to tractably examine price and quantity competition in a truly dynamic model, and are more realistic than the assumptions implicit in

conventional oligopoly analysis.

In each period, each firm independently chooses both the price of its good and its level of output for that period. If output exceeds demand, the unsold output is put into inventory, and if demand exceeds output then inventory is drawn down to cover the excess. Firm i's inventory evolves according to

$$I^i_{t+1} = (1-\delta)(I^i_t + q^i_t - d^i_t)$$

indicating that next period's inventories equals current inventories plus net additions less depreciation, at rate of $\delta \in [0,1)$, between today and tomorrow.

Note that this formulation of inventories assumes that the good is somewhat durable, but that our formulation of demand assumed that only current prices affect demand, implying that there is no durability from the consumers' perspective. We justify this on several grounds. First, there may be substantial economies of scale in storage, making it impractical for individuals to store the good. For example, would you buy milk from your neighbor, even if it had the same expiration date as that in the store? Second, the point of introducing inventories into our analysis is to allow firms to choose price and quantity, and as long as the goods are not significantly durable, ignoring consumer storage should not affect our results concerning price and output substantially.

Within any period, the firm incurs production and inventory maintenance costs, and earns income from sales. The net income in period t equals

$$\pi^i(p_t,\, q_t,\, I_t) = p^i_t\ D^i(p^1_t,\, p^2_t) - C_i(q^i_t) - H_i(I^i_t)$$

where $p_t = (p^1_t,\, p^2_t)$ and $q_t = (q^1_t,\, q^2_t)$.

We assume that inventories are perfectly observed by both firms at the beginning of each period. This implies that we have a complete information dynamic game. We focus on the closed-loop equilibrium [4] of our long-but-finite-horizon games. In our case, closed-loop equilibria are those subgame perfect equilibria in which each player's strategy is only a

[4] Some game theorists have recently begun using the phrase "Markov perfect equilibria" to describe what appears to be the same concept as was earlier defined to be "closed-loop equilibria." Since this paper uses the old technology and extensions of this model will naturally use the literature which still uses the original terminology, I will also use the original terminology.

function of the current level of inventories. These restrictions are important. For example, inventories have been studied as a tool of implicit collusion (see Rotemberg and Salonar), but only in infinite-horizon models. Our focus on closed-loop equilibria in finite horizon games will not find such equilibria since our finite-horizon equilibria will be unique. We choose to examine the purely noncooperative equilibria since we want to make comparisons with the predictions of the static models, which are always purely noncooperative.

It is clear, however, that our results will be important for considerations of noncooperatively supported cooperation. The possibility of tacit collusion depends on the strategic formulation, with collusion being easier with price competition. The closed-loop equilibria computed here can serve as the building blocks for constructing trigger strategy equilibria which implement tacit collusion, and evaluating whether the Bertrand or Cournot model is better at predicting the viability of tacit collusion.

The state variable, I, of the game is the vector of the firms' beginning-of-period inventory holdings, $(I^1,\, I^2)$. Closed-loop equilibria can be expressed as solutions to a pair of dynamic programming equations since each firm is solving a dynamic programming problem given the strategy of its opponent. We define the value function of each firm to be a function of I:

$$V^i(I) = \max_{p^i, q^i} \left\{ \pi^i(p,\, q,\, I) + \beta V^i\Big((1-\delta)(I+q-d)\Big) \right\}$$

We need examine only pure strategy equilibria. If I is the current state of inventories then $P^1(I)$ will denote its equilibrium choice for p^1 and $Q^1(I)$ will denote its choice for q^1. $P^2(I)$ and $Q^2(I)$ will denote firm two's equilibrium strategies. Finally, if $Q \equiv (Q^1,\, Q^2)$ and $P \equiv (P^1,\, P^2)$, then

$$I^+(I) = (1-\delta)\left(I + Q(I) - D\left(P(I)\right)\right.$$

expresses the equilibrium value of next period's inventories if the current inventory holdings are I.

The first-order condition for firm one's choice of p^1 implies that $P^1(I)$ satisfies

$$P^1 d_1^1(P^1,\, P^2) + d^1(P^1,\, P^2) = \beta(1-\delta)\left(V_1^1\left(I^+(I)\right) d_1^1 + V_2^1\left(I^+(I)\right) d_1^2\right)$$

and firm two's equilibrium choice for q^1 must satisfy

$$C_1'(q^1) = \beta V_1^1\left(I^+(I)\right)(1-\delta)$$

Existence and uniqueness are nontrivial problems in general, but are easily handled in the case of the linear-quadratic specification which we study below. We shall first discuss the qualitative features indicated by the first-order conditions.

The price condition states that the current marginal revenue of lowering price is equated with the marginal impact on the equilibrium value to firm one of the inventory levels and the output condition equates the marginal cost of current production with its contribution to the next period's value of the game to firm one. If the demands were independent, that is, the demand for each good is unaffected by the price of the other, then one solution is that each firm ignores the other and chooses a monopoly strategy. In this case, the first-order conditions for price and quantity imply

$$p^1 d_1^1 + d^1 = C'(q^1)\, d_1^1$$

which is the standard marginal revenue equals marginal cost condition. In general, we have

$$p^1 d_1^1 + d^1 = C'(q^1)\, d_1^1 + \beta(1-\delta)\, V_2^1\, d_1^2$$

implying that marginal revenue equals marginal cost plus the marginal value to firm one of an increase in firm two's inventories. If η_{ij} is the elasticity of demand for good i with respect to p_j, we can rewrite this as

$$p_1\,(1 + 1/\eta_{11}) = C'(q_1) + \beta(1-\delta)\, V_2^1\,(\eta_{21}/\eta_{11})\,(d^2/d^1)$$

In this case of substitutes, $\eta_{21} > 0$.

It is natural to examine the case $V_2^1 < 0$ since this says that firm one's profits decline as firm two's inventories rise. This is natural since inventories are a low cost source of output for the rival; hence, large inventories likely reduce the marginal cost of fulfilling demand making a rival more aggressive. If $V_2^1 < 0$ the first-order condition shows that the standard monopoly pricing formula is augmented by a positive term representing the ratio of elasticities and demands, and V_2^1/p_1. Under these conditions, the inventory game acts as if we took the static game and increased marginal cost, indicating a less competitive outcome than the static Bertrand case. While this observation is suggestive, we must of course wait for a complete analysis to see if $V_2^1 < 0$ in general.

Theorem 1 is immediate from the first-order conditions.

Theorem 1: If the marginal cost of production, $C'(q)$, is constant over $q \in [-\infty, \infty]$, then the set of equilibrium outcomes include the collection of static Bertrand equilibrium prices. In general, if $V_2^1 < 0$ then price exceeds the static Bertrand equilibrium prices.

Proof: Let I^* be the level of inventories which minimizes the present value of inventory holding costs. Consider the strategy of charging a price consistent with some Bertrand equilibrium, and producing enough output so as to begin the next period with inventory equal to I^*. The value functions for these strategies satisfy the equilibrium conditions stated above.

Theorem 1 indicates that the Bertrand model is better under conditions where inventory holding has no impact on the equilibrium marginal costs since marginal costs are constant. Note that any static Bertrand equilibrium may result. Their may be others, but Theorem 1 shows that the Bertrand equilibrium are always present in these cases. In cases where equilibrium is unique, as in the linear-quadratic models examined below, the Bertrand equilibrium is the only equilibrium.

3: Linear-Quadratic Model

The general model is difficult to examine. While existence is not a serious problem[5] it is unlikely that we would be able to determine any qualitative properties about equilibrium. In order to obtain precise results, we will examine a simple linear duopoly model of differentiated products with inventories.

We assume that firm i faces a demand for its product equal to

$$d^i = a_i p_i - b_i p_j - c_i, \ i, \ j = 1, \ 2, \ i \neq j$$

We assume that production costs for firm i equal

$$C_i(q_i) = m_i q_i + n_i q_i^2 \ , \ i = 1, \ 2.$$

[5] For example, if the model were formulated in continuous time with a Brownian noise additive shock to the state variables, then existence of equilibria with smooth value functions for the firms could be proved using standard existence theorems for stochastic differential games.

The inventory technology is also given a linear-quadratic specification in order to maintain tractability. The cost to firm i of holding an inventory I_i equals

$$H_i(I^i) = h_i I^i + g_i (I^i)^2 \ , \ i = 1, \ 2.$$

with inventories following $I^i_{t+1} = (1 - \delta) \ (I^i_t + q_{it} - d_{it})$.

While the linear-quadratic specification has some unappealing properties when viewed globally (e.g., negative inventories are possible), as long as h_i is sufficiently negative the efficient level of inventories is positive, a reasonable assumption. Hopefully, our linear-quadratic specification operates essentially as only a local restriction on the nature of the inventory holding costs. Below we discuss ways to check this presumption.

The resulting game between the two firms is a simple linear-quadratic game with well-understood properties. We will assume that the game has an infinite horizon, but examine only the equilibria which are limits of finite horizon games. Given the linear-quadratic specification of the game, there is only one such equilibrium. This equilibrium is also known as the closed-loop equilibria since it is a subgame perfect equilibria in which each firm's strategy depends only on the current state of the game, i.e., the inventory holdings.

We will base our analysis on the variable y_t,

$$y_t \equiv [1, \ I^1_t, \ q_{1t}, \ p_{1t}, \ I^2_t, \ q_{2t}, \ p_{2t},]'$$

and the vector of control variables is

$$x_t \equiv [q_{1t}, \ p_{1t}, \ q_{2t}, \ p_{2t}]'$$

Since profits equal

$$\pi_i(y) = a_1 p_1^2 - b_1 p_1 p_2 - c_1 p_1 - m_1 q_1 - n_1 q_1^2/2 - h_1 I^1 - g_1 (I^1)^2$$

we can express the game's payoffs as a quadratic form,

$$\pi_i(y_t) = (1/2) \, y_t' \, R_i y_t, \ i = 1, \ 2$$

where R_1 can be

$$R_1 = \begin{pmatrix} 0 & -h_1 & -m_1 & -c_1 & 0 & 0 & 0 \\ -h_1 & -2g_1 & 0 & 0 & 0 & 0 & 0 \\ -m_1 & 0 & -n_1 & 0 & 0 & 0 & 0 \\ -c_1 & 0 & 0 & 2a_1 & 0 & 0 & -b_1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -b_1 & 0 & 0 & 0 \end{pmatrix}$$

and R_2 is similary defined.

The law of motion can be expressed as

$$y_t = Ay_{t-1} + Bx_t$$

where the coefficients are:

$$A = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ (1-\delta)c_1 & 1-\delta & 1-\delta & -(1-\delta)a_1 & 0 & 0 & (1-\delta)b_1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ (1-\delta)c_2 & 0 & 0 & (1-\delta)b_2 & 1-\delta & 1-\delta & -(1-\delta)a_2 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

and

$$B = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

The solution to V^i (see any reference on linear-quadratic games, Kydland (1975) being a good one) takes the form of coupled Ricatti equations and can be computed recursively given a terminal valuation for inventories. Suppose that the solution to V^i is a quadratic form:

$$v_{it}(y_{t-1}) = (1/2)\, y'_{t-1} S_{it}\, y_{t-1}$$

Define B to be

$$B = [b_1,\ b_2]$$

Let H_t be

$$H_t = \begin{bmatrix} b'_1 \sum_{1t} \\ b'_2 \sum_{2t} \end{bmatrix}$$

where Σ_{it} is

$$\Sigma_{it} = R_t + \beta S_{i,\,t+1}, i = 1,\ 2$$

The first-order conditions for firm 1 when choosing $x_{1t} \equiv (p_{1t}, q_{1t})$ imply that

$$0 = b_1'\Sigma_{1t}b_1x_{1t} + b_1'\Sigma_{1t}b_2x_{2t} + b_1'\Sigma_{1t}Ay_{t-1}$$

First-order conditions for firm 2 imply a symmetric condition. If we combine these two systems of linear equations, we find that the equilibrium rule is

$$x_t = G_t y_{t-1}$$

where

$$G_t = (H_tB)^{-1}H_tA$$

Furthermore, the quadratic form for the current value function is

$$S_{it} = (A + BG_t)'\Sigma_{it}(A + BG_t)$$

The foregoing equations form a recursive set of matrix conditions. If period T were the last period and S_{iT} were firm i's final payoff, then we can compute the payoff for periods $t < T$ by iterating the above process. Theorem 2 is standard; see, for example, Kydland.

Theorem 2: There exists a unique closed-loop sequence of equilibria given any concave terminal valuation function.

One of the special, but useful, features of a linear-quadratic framework is the certainty equivalence property. More specifically, if the laws of motion for the state variables are subject to additive shocks, then the equilibrium strategies for the stochastic game are the same as for the deterministic game. Such additive shocks could come from noise in either production or demand. Therefore, we will be able to interpret our results as applying to such stochastic environments.

4: Adjustment Costs, Investment, and Learning

The model described in the previous section is too simple to be of much interest. At this point we add details which make the model more realistic, albeit still confined to a linear-quadratic framework.

One argument for the Cournot model is that it corresponds to the observation that it is much easier to alter one's price than one's output. Firms are locked into a scale of production

by previous investment choices, whereas prices are set at the last moment before a sale. This intuition lies at the heart of the Kreps and Scheinkman analysis.

This consideration is largely absent in our model so far since, other than through inventories, current decisions have no impact on future conditions. In fact, inventories cannot represent any adjustment cost since an unexpectedly high inventory, possibly arising from unexpectedly high output in the previous period, will often lead to reduced output today, just the opposite from what an adjustment costs story would indicate. In order to consider the Kreps-Scheinkman argument we add adjustment costs to our model. More precisely, we now assume that total production costs for firm i in period t equal

$$C_i(q_{it}) = m_i q_{it} + n_i q_{it}^2 + \gamma_i (q_{it} - q_{i,t-1})^2/2.$$

This, of course, means adding the outputs of each firm to the list of state variables.

This structure allows us to distinguish between short-run and long-run costs. The adjustment costs are part of short-run marginal costs. However, in a steady state they contribute nothing to total costs. Hence, m_i and n_i represent the long-run marginal cost function. We will often focus on the case $n_i = 0$ since that is the case of linear long-run costs, the assumption in Kreps and Scheinkman.

As we discussed in the introduction, investment and learning have received attention in the literature on preemption, with the results often relying on a Cournot specification. To examine this sensitivity, we will add both learning and investment to our model. To do so, we now make m_i, firm i's marginal cost at zero output, a state variable obeying the law

$$m_{i,\,t+1} = \; -\lambda_i q_{it} - \mu_i f_{it} + \psi_i m_{it} + (1 - \psi_i)\bar{m}$$

where f_{it} is firm i's investment effort at t, and $\bar{m}$ is the long-run marginal cost if there is neither output nor investment. We assume that the cost of investment effort is

$$\xi_i f_i + \theta_i f_i^2$$

We assume λ_i, μ_i, ξ_i, θ_i, and ψ_i, and are all positive, and $\psi_i \leq 1$. This specification models both learning and investment since previous output and investment reduce current costs. The

assumptions concerning ψ help to keep the equation stable since they generate a tendency for m to return to $\bar{m}$.

Adding adjustment costs, learning, and investment to our model allows us to address a much richer set of issues. We should note, however, that the strictures of the LQ form limit the robustness of the analysis. For example, investment here causes the marginal cost curve to shift out in a parallel fashion. A CRTS specification for technology in terms of capital and labor, such as Cobb-Douglas, would instead have investment in capital cause a clockwise rotation of the marginal cost curve. While we cannot model this in the LQ model, we will focus our discussion on aspects of the results which are most likely to be robust.

5: Steady-State Prices and Equilibrium Reactions

We shall now begin comparing the results of our dynamic model with the Cournot and Bertrand models. The first index we use for this is the steady state price of our dynamic model. A focus on the steady state is proper since the system will spend almost all of its life there. Examination of the price is also important since that will be strongly related to the allocative efficiency of equilibrium. Then we will examine the reactions in the equilibrium strategies. This will tell us about biases in resource allocation due to strategic effects.

We cannot solve for the equilibrium in a reduced form fashion. However, numerical computation of equilibrium is easy. We shall examine a range of examples, focussing on their steady-state prices and the equilibrium strategies when the horizon is long. Table 1 summarizes the results. In Table 1, we assume a utility function over the two goods equal to

$$U(x, y) = -(x + y)^2 + (x + y) - B(x - y)^2 + M$$

where M is money expended on other goods. This utility function is symmetric in the two goods, and by varying B we can represent any symmetric quadratic utility function. When $B = 1$ the products have independent demand, and when $B = 0$ they are perfect substitutes. We will allow $B = .5$, .15, and .02; the value of B is displayed in the first column of each panel. In the next five columns, we vary the adjustment cost parameter, allowing $\gamma = 0.0$, 0.5, 2.0, 10.0, and 50.0 as we go left to right. We assume that $\bar{m}_i = 30$. The last column in the first two panels give the equilibrium price and quantity for the static Cournot case.

Finally, the numeraire is chosen so that demand for each good is zero if the price for each good is 100.

Table 1: Linear-Quadratic Model, No Investment, No Learning

mbar=30, n=0, g=100, h=0, $\delta = 0$, cinvl=0, cinvq=2, ainv=0, lrn = 0

$\gamma =$	0	.5	2.0	10.0	50.0	Cournot
Equilibrium Price						
B = 0.50	58.00	58.20	58.56	59.10	59.44	60.00
0.15	44.49	46.32	48.89	51.66	53.04	55.56
0.02	32.64	37.54	43.02	47.49	49.47	53.64
Equilibrium Output						
0.50	10.50	10.45	10.36	10.22	10.13	10.00
0.15	13.88	13.42	12.77	12.08	11.74	11.11
0.02	16.84	15.61	14.24	13.12	12.63	11.59
Price-to-Marginal Cost Ratio						
0.50	1.93	1.94	1.95	1.97	1.98	
0.15	1.48	1.54	1.62	1.72	1.76	
0.02	1.08	1.25	1.43	1.58	1.64	
Bertrand vs. Cournot						
0.50	0	0.10	0.28	0.55	0.72	
0.15	0	0.17	0.39	0.64	0.77	
0.02	0	0.23	0.49	0.71	0.80	
Own Price Elasticity						
0.50	-2.07	-2.09	-2.12	-2.17	-2.19	
0.15	-3.07	-3.31	-3.67	-4.09	-4.33	
0.02	-12.36	-15.32	-19.25	-23.06	-24.96	
Cross Price Elasticity						
0.50	0.69	0.69	0.70	0.72	0.73	
0.15	2.27	2.44	2.71	3.02	3.20	
0.02	11.87	14.72	18.49	22.16	23.98	
Elasticity of Marginal Cost						
0.50	0.00	0.17	0.69	3.41	16.90	
0.15	0.00	0.22	0.85	4.03	19.56	
0.02	0.00	0.26	0.95	4.38	21.05	
Output Reaction to Output						
0.50	0	-0.03	-0.06	-0.06	-0.03	
0.15	0	-0.15	-0.21	-0.15	-0.07	
0.02	0	-0.47	-0.38	-0.20	-0.09	
Price Reaction to Output						
0.50	0	-0.12	-0.40	-0.99	-1.50	
0.15	0	-0.31	-0.95	-1.98	-2.73	
0.02	0	-0.52	-1.36	-2.48	-3.24	

We make assumptions concerning the inventories which hopefully minimize any idiosyncratic role they play. We assume that the cost-minimizing level of inventories is zero and that the inventory holding cost function has very high curvature. This is in an approximation implicit in conventional Bertrand and Cournot analysis that inventories are prohibitively costly. Centering the cost function at zero implies that the dynamically efficient level of inventories is nearly zero; if this were not true then some output in each period would be

devoted to replacing depreciated inventory, forcing the firm to operate at a higher point on the marginal cost curve, and making comparisons with the static model difficult. The high curvature assumption is made to eliminate other perversities, such as negative inventories, from playing a role. We also assume that $\delta = 0$. These restrictions are sensible at this point since the purpose of this exercise is not to focus on inventories as a strategic tool, but rather to compare the dynamic model with the static models.

In Tables 1 we focus on one basic case; the text below will discuss the impact of perturbations involving investment and learning. The basic case, described in Table 1, assumes a CRTS technology, $n_i = 0$, no investment, $\mu_i = 0$, and no learning, $\lambda_i = 0$. Hence the marginal and average cost are both 30. The various panels of Table 1 display the results.

The "Equilibrium Price" panel gives the steady-state price for each pair of B and γ. The final column is the Cournot price. Of course, Theorems 1 and 2 tell us that the first column of prices is the static Bertrand price. To get a handle on the relation between our equilibria and the static Cournot and Bertrand predictions, we look at the "Bertrand vs. Cournot" panel. The number measures the closeness to the Cournot price relative to the Bertrand price. Specifically, if the number is .1, then the steady state price in that case equals the Bertrand price plus ten per cent of the Cournot-Bertrand difference.

The results are informative and conform to some intuitive arguments. Generally speaking, as long as adjustment costs are small, implying that the slope of the marginal cost function is not large, the steady state price is close to the Bertrand level, as indicated by Theorem 1. However, as adjustment costs and the slope in marginal cost become significant, the steady state price is more like the Cournot price, a result similar to that of Kreps and Scheinkman.

The critical role of adjustment costs is intuitively clear. This can be explained in our model by considering the firms' equilibrium reaction functions. Suppose both firms were charging the Bertrand price. If one firm unexpectedly raised its price, then it would lose sales and enjoy an increase in its inventory, whereas its competitor would have to deplete its inventory to satisfy unanticipated demand. This would, in the short-run, raise the the competitor's marginal cost since it would now have to increase output above its long-run

level to get back to the desired level of inventories. The first firm could could continue to sell at the old rate and produce at a lower rate and lower marginal cost, but charge a higher price since its competitor's marginal cost has increased, for its output. Since we began with the Bertrand equilibrium, the static effects on the first firm's profits of this deviation are nil, whereas the dynamic effects are beneficial. Therefore, firms will want to raise price above the Bertrand level. As the costs of adjustment rise, this ability to increase one's opponent's short-run costs also increases, leading to greater incentives to raise price.

However, a close examination of the quantitative side of our equilibria shows that the results are not too supportive of the Cournot model as a model of long-run equilibrium. This is indicated by the panel labeled "Elasticity of Marginal Cost." The numbers in that panel tell the percentage increase in marginal cost which would accompany a one per cent increase in output. This is a measure of slope for the marginal cost curve. For example, in the Cobb-Douglas case, if capital share is one-third, then this elasticity is .5. We see that the cases where the steady-state price is more than halfway to the Cournot level correspond to high marginal cost elasticities. We say high because the interest rate is chosen to correspond to a time period of a year. The results here are even less supportive of the welfare implications of the Cournot model, since, if goods are nearly perfect substitutes, even when the price is eighty per cent of the way to Cournot, the efficiency cost is only two-thirds of the Cournot model.

These results tend to support the basic intuition of Kreps and Scheinkman, that if past actions create a very steep marginal cost function, then the result will be similar to the Cournot price. The advantage of our approach is that we were not distracted by the details of rationing in off-equilibrium nodes of the game. Also, we can examine these arguments in the context of a more realistic game.

The next several panels of Table 1 gives descriptive indices for the steady states of each case. We see that the price-marginal cost ratios are within a broad, but sensible, as are the own- and cross-price elasticities. These are displayed to show that the examples are not perverse in the underlying structural paramaterizations.

The rest of the panels examine the equilibrium reactions to one's opponents past behav-

ior, as well as to one's own deviations. Since Table 1 assumes that investment expenditures do not affect anything, all reactions to investment decisions are zero. Of more interest are the price and quantity reactions. First note that when adjustment costs are zero, firms do not react to their competitor's output decisions, nor does a firm respond to a competitor's price deviation. If an opponent raises his price, his opponent's current sales increases, implying an increase his output in the next period to bring inventories back to the target level. All these facts were implied by Theorem 1.

In general, with positive adjustment costs, we find that an increase yesterday in a competitor's output causes a firm to reduce its price and output today. This appears to be a Cournot-like reaction. On the other hand, firms react to price increases by increasing price and output, a Bertrand-like reaction. In our model, we really appear to have an intermediate situation. In the static Cournot model, firms cannot alter output, and in the static Bertrand model, firms are not allowed to alter their prices. In our dynamic model, the descriptions of the true reactions have both Cournot and Bertrand features.

I will argue, however, that the reactions are more Bertrand than Cournot in substance. This is seen by considering whether a firm benefits form his competitor's reactions. Suppose that we had just a Cournot model with adjustment costs. Then extra output by a firm today would cause his competitor to reduce his output in the future, a beneficial reaction. On the other hand, if we had a dynamic Bertrand model, a price reduction (resulting in higher output) will cause a competitor to reduce future prices, a harmful reaction. In this model, a choice to increase output or to reduce price will cause a competitor to react by, among other things, reducing future prices. Since the equilibria are stable with positive roots, the competitor's reaction dampens over time, but never changes sign. Therefore, procompetitive deviations prompt harmful reactions from competitors, an important feature of the static Bertrand model.

With this framework, we can make some distinctions which are not possible in a stylized, two-period analysis. The adjustment costs and the rising marginal cost curves correspond to two different contributions to marginal costs. If $n = 0$, but adjustment costs are high, we have a situation where the short-run marginal cost curve is nearly vertical at any point

in time, but the long-run marginal cost curve is flat. Our initial results indicate one needs to make strong assumptions about adjustment costs in order to come close to the Cournot outcome. Other computations show that the Cournot model becomes a better approximation when $n > 0$, i.e., an increasing long-run marginal cost curve, even without adjustment costs. The Cournot outcome is approximated by moderate choices for n and adjustment costs.

The third case of $n < 0$ is also of interest, particularly since our focus is on oligopolistic industries where one possible reason for concentration is declining marginal costs. In this case, calculations not displayed here show that the steady-state price may actually fall below the static Bertrand price (remember, the static Bertrand price is above marginal cost). Therefore, the static Bertrand and Cournot equilibria do not even bracket the long-run possibilities in our dynamic model. The intuition for this result is straightforward. Suppose that both firms are charging the Bertrand price. If one firm cuts its price a little and increases its output to satisfy the extra demand, then its competitor will lose sales and have higher inventories than anticipated. This increase in inventories will imply that the competitor need not produce as much tomorrow, causing it to operate at a higher marginal cost in the presence of falling marginal costs. The opposite argument also explains why prices move towards the Cournot level as the marginal cost curve steepens.

The relation between the slope of the marginal cost and the equilibrium price-cost margin raises some interesting possibilities for price-cost margins over time. Recently, there has been substantial interest in the question of how price-cost margins move over the business cycles (see Domowitz, Hubbard, and Petersen (1986), and Hall). Suppose that there are persistent shocks to demand, as in a business cycle. If the short-run cost function displays decreasing returns to scale, then the firms operate at a steeper region of their marginal cost functions when demand is high than when demand is low. Since price-cost margins in our model are greater when the marginal cost curve is steeper, this suggests that price-cost margins are greater when demand is high. While there have been other arguments predicting the movement of price-cost margins, this version is purely noncooperative and makes general and unbiased assumptions about demand and costs. Unfortunately, this conjecture can be investigated only in nonlinear versions of our model.

While comparing steady-state prices may provide some useful information, it is unlikely that the closeness of the steady-state price to either the Bertrand or Cournot price will tell us anything about issues other than allocative efficiency. Any specific question should be reexamined in this model to determine which static model is a better approximant. The nature of the reactions will be particularly important in addressing these questions. We next reexamine some specific problems in strategic interaction where the results have been sensitive to strategic formulation.

The cases of learning and investment were investigated by allowing various choices for μ and λ. The case of learning only was modelling in Table 2 by choosing $\lambda_1 = \lambda_2 = .1$ This resulted in lower prices and greater output, of course. The reactions were very similar, except for the price reaction to a competitor's price. When adjustment costs were low, an increase in a competitor's price caused a firm to react by reducing price. This is because a firm has unanticipated sales when a competitor raises his price, resulting in greater output in the future and, because of the learning, lower costs. However, in the presence of nontrivial adjustment costs, the reaction returned to being positive as in Table 1.

Table 2: Linear-Quadratic Model, Investment, No Learning

mbar=30, n=0, g=100, h=0, $\delta = 0$, cinvl=0, cinvq=2, ainv=1, lrn = 0

$\gamma =$	0	.5	2.0	10.0	50.0	Cournot
Equilibrium Price						
0.50	56.5763	56.7803	57.1406	57.7004	58.0646	58.6525
0.15	42.3165	44.0834	46.6685	49.5817	51.0854	53.8081
0.02	29.9916	34.0788	39.8715	44.7756	47.0277	51.6642
Equilibrium Output						
0.50	10.8559	10.8049	10.7149	10.5749	10.4839	10.3369
0.15	14.4209	13.9791	13.3329	12.6046	12.2286	11.5480
0.02	17.5021	16.4803	15.0321	13.8061	13.2431	12.0839
Marginal Cost						
0.50	27.6272	27.6243	27.6222	27.6290	27.6419	
0.15	27.2686	27.1812	27.1362	27.1814	27.2477	
0.02	27.2462	26.5945	26.6720	26.8792	27.0130	
Price-to-Marginal Cost Ratio						
0.50	2.0478	2.0555	2.0686	2.0884	2.1006	
0.15	1.5518	1.6218	1.7198	1.8241	1.8748	
0.02	1.1008	1.2814	1.4949	1.6658	1.7409	
Output Reaction to Investment						
0.50	-0.0321	-0.0235	-0.0126	-0.0031	-0.0004	
0.15	-0.2051	-0.1114	-0.0424	-0.0076	-0.0009	
0.02	-1.9907	-0.3286	-0.0767	-0.0108	-0.0013	
Price Reaction to Investment						
0.50	-0.0857	-0.0837	-0.0717	-0.0392	-0.0133	
0.15	-0.2140	-0.2155	-0.1702	-0.0790	-0.0244	

0.02	-0.3123	-0.3550	-0.2431	-0.0994	-0.0291

In all cases, however, the decision to increase output will cause a competitor to reduce his price, a harmful reaction. Therefore, at the margin, the strategic effects reduce learning, as in the Bertrand model, not the Cournot model.

Table 3 next examines the case of investment (but no learning) by taking the basic case studied in Table 1, but setting $\mu_1 = \mu_2 = 1.0$. Since the investment reactions are nontrivial in this case, we include these in Table 3. The important results are that a firm reduces its price in response to a competitor's investment. This is a harmful reaction, implying that, at the margin, investment will be reduced due to strategic effects, as in the Bertrand model.

Table 3: Linear-Quadratic Model with Learning, No Investment

mbar=30, n=0, g=100, h=0, $\delta = 0$, cinvl=0, cinvq=2, ainv=0, lrn = 0.1

$\gamma =$	0	.5	2.0	10.0	50.0	Cournot
Equilibrium Price						
0.50	56.75	56.97	57.34	57.91	58.28	59.40
0.15	42.15	44.11	46.88	49.89	51.39	54.78
0.02	28.57	33.94	40.14	45.18	47.42	52.77
Equilibrium Output						
0.50	10.81	10.75	10.66	10.52	10.42	10.14
0.15	14.46	13.97	13.28	12.52	12.15	11.30
0.02	17.85	16.51	14.96	13.70	13.14	11.80
Marginal Cost						
0.50	28.91	28.92	28.93	28.94	28.95	
0.15	28.55	28.60	28.67	28.74	28.78	
0.02	28.21	28.34	28.50	28.63	28.68	
Output Reaction to Output						
0.50	-0.003	-0.03	-0.06	-0.06	-0.03	
0.15	-0.019	-0.16	-0.22	-0.15	-0.07	
Price Reaction to Output						
0.50	-0.01	-0.12	-0.40	-0.99	-1.50	
0.15	-0.02	-0.33	-0.96	-1.98	-2.73	
0.02	-0.04	-0.54	-1.38	-2.49	-3.25	

6: Comparisons with Dynamic Cournot and Bertrand Equilibrium

In the previous analysis, we compared our dynamic equilibrium with the static Cournot and static Bertrand equilibria. It is more appropriate to compare our dynamic equilibria with the dynamic equilibria which would result if we had kept the usual Cournot or Bertrand specification. In this section we do that, showing how our model differs from both Cournot and Bertrand in interesting ways.

We will consider two alternatives, assuming in both cases linear demand and quadratic short-run costs. First, we will call the Cournot game the dynamic game where firms choose

output in each period, prices are set so as to clear the market, no inventories are allowed, and there are quadratic adjustment costs in output. Second, we will call the Bertrand game the dynamic game where firms choose price in each period, with output being automatically set so as to satisfy demand, no inventories are allowed, and with quadratic adjustment costs in output. Again, the equilibria can be calculated as before since the games are both still LQ games. We will focus on the effects of increasing the adjustment cost parameter since it played an important role in parameterizing the Kreps-Scheinkman intuition.

The results were interesting for both the Bertrand and Cournot games. In the Cournot game, when adjustment costs are absent the steady state price is the static Cournot price. As adjustment costs rise, the steady state price falls. The intuition is simple. In the presence of adjustment costs, an increase in output today causes a fall in tomorrow's short-run marginal cost and an increase in output, which in turn causes a competitor to reduce output, a favorable response. If the firms were at the static Cournot equilibrium, the static effect of raising output would be zero but the dynamic effect would be positive. Hence, we expect the firms to raise output above the Cournot level.

The Bertrand result is initially surprising. In the absence of adjustment costs the steady state price is the static Bertrand price. The knee-jerk conjecture is that as adjustment costs rise, steady-state price should rise because the Bertrand model usually gives the opposite result. One effect would argue for that conjecture: if I raise my price today, I reduce current output, precommitting myself to a higher cost curve tomorrow, which should cause a higher price from my competitor, a beneficial response. However, note that our Bertrand model is not a complete complement of the Cournot model. The complement of the Cournot model would have adjustment costs only on prices, not output. In that case, adjustment costs would cause prices to increase. However, since this is not a macroeconomics paper, we put adjustment costs on output only. In our Bertrand model, there is a second effect. If I raise my price today, I will also cause demand to rise for my competitor, automatically raising his output and lowering his future marginal cost curve. Since own costs are more important than a competitor's costs, the net result is that if I raise my price today, my competitor lowers his price tomorrow, a damaging response. Therefore, in a Bertrand model prices also

fall as adjustment costs increase. The dominant effect is that by lowering my price today, I raise my competitor's costs and future prices, and lower mine.

A comparison of the three models shows that our hybrid model differs from both the dynamic Bertrand and Cournot models since it yields the static Bertrand price in the absence of adjustment costs, but then yields higher prices as adjustment costs increase. Given the close relationship between the hybrid model and the Bertrand model seen above, the differences here are particularly revealing. In the hybrid model, if I raise my price my opponent meets the extra demand initially out of inventory, thereby raising his future marginal cost curve. My competitor is allowed to respond in the short-run by either raising price or output, a reasonable description of real-life flexibility. In the Bertrand, however, my competitor is forced to immediately raise output. That specification limits my competitor's flexibility of response in an unrealistic fashion and turns out to be substantively important.

7: Comparisons with Stackelberg Equilibria

Another common game form examined in oligopoly theory is that of Stackelberg leadership. In these games, one firm is allowed to move first, followed by other firms who know the leader's (irreversible) move. These games are used to model asymmetric market power. Sometimes it is advantageous to be a leader, but not always (see Bulow et al.). For example, in Cournot oligopoly it is advantageous to be the leader since one can precommit to a high output, and this commitment will cause the followers to reduce their output, a favorable response for the leader.

In a dynamic world, each firm is both a leader and a follower. There is no natural asymmetry in timing when one takes a dynamic game perspective. However, there are various asymmetries which are possible. For example, firms may differ in their adjustment costs, their inventory costs, or their production costs. The structural asymmetries may lead to asymmetries in solutions which differ from the asymmetries implied by static Nash-Cournot game forms.

In fact, one such asymmetry does lead to Stackelberg-like results. Suppose that the firms differ in their adjustment costs, with firm 1 having higher adjustment costs. Then we find that in the steady state of the dynamic game, firm 1 has a higher market share, charges a

higher price, and makes more profits. The intuition for this is clear. If firm 1 increases output today, then his future costs are reduced. Output today acts as a precommitment device for output tomorrow. If firm 1 has higher adjustment costs, then it is a stronger precommitment device for him then firm 2. On seeing an increase in output by firm 1, firm 2 will know of the precommitment, and will be able to adjust accordingly since his adjustment costs are lower. Such downward adjustments in output are good for firm 1. Therefore, knowing of firm 2's propensity to adjust and react favorably, firm 1 will be more likely to increase his output. The result is that firm 1 has a larger market share, as if he were a Stackelberg leader.

This is just one example of where a structural asymmetry results in an asymmetry which is missing in any static analysis. Recall that one of the reasons for using Cournot analysis is the difficulty in adjusting output. We see here that asymmetries in that adjustment lead to an outcome different from the static Cournot outcome. Other asymmetries are also likely to result in important difference, and will be pursued in further work.

8: General Specifications

While somewhat informative, linear-quadratic models bring many special features to the analysis, some of which are unappealing. For example, the inventory holding cost function is symmetric and even permits negative inventories. It is even possible to have price above the choke price and negative output. One would never examine cases where the steady state or any equilibrium path from a sensible initial condition possessed any of these perverse properties.

However, that is not enough for a subgame perfect analysis since it is possible that our treatment of these off-equilibrium path states will affect the equilibrium. Some considerations indicate that it is unlikely that this is a severe problem. For example, in optimal growth models, which possess structures similar to this model, behavior in any neighborhood of the steady state is invariant to substantial changes in production and utility specifications outside of the path to the steady state. In this model, while a firm may be allowed to run negative levels of inventories, it is unclear why there would be any advantage as long as the holding costs were not trivial. However, even if we don't believe that these problems are likely to be important, it would be desirable to make some check.

Another problem is, of course, the fact that linear demand and cost functions are special. For example, monopolists facing linear demand will increase price by less than a dollar if they experience a one dollar increase in (constant) average costs, whereas a monopolist facing a constant elasticity of demand will increase price by more than a dollar. Also, there are many aspects of industrial strategy which we would like to examine which do not fit easily into a linear-quadratic framework, such as investment and learning curves. These considerations all demand a way to check the robustness of our results.

In this section, we will outline an approach which can be used to address these concerns. We will take a continuous time approach, appealing to the theory of differential games. Let $u^i \equiv (p_i,\, q_i),\ u \equiv (u^1,\, u^2),\ U^i = (P^i,\, Q^i)$, and $U \equiv (U^1,\, U^2)$. The state variable is again the level of inventories, now following the differential equation

$$\dot{I} = f^i\,(u,\, I) \equiv q_i - d^i\,(p_1,\, p_2) - \delta I_i$$

If $V^i(I)$ is the current value of future profits of firm i if the current state is I, then the Bellman equation for firm i is given by the partial differential equation

$$0 = \max_{u^i} \{\pi^i(u^i,\, U^{-i}(I),\, I) + V^i_j(I)\, f^j(u^i,\, U^{-i}(I),\, I)\} - \rho V(I) \tag{1.i}$$

These equations imply the system

$$0 = \pi^i_{u_\ell}\,(U(I),\, I) + V^i_j(I)\, f^j_{u_\ell} \tag{2.$i.\ell$}$$

$$0 = \pi^i\,(U(I),\, I) + V^i_j(I) f^j(U,\, I) - \rho V^i(I) \tag{3.i}$$

Therefore, equations (2) and (3) express 6 conditions that any solutions to $V^i(I)$ and $U^i(I)$ must satisfy.

There are numerous techniques in the numerical partial differential equation literature for handling equations such as these. We chose a straightforward implementation of minimum weighted residual techniques. For a general discussion of these techniques and their applications to economic problems, see Judd (1989). While judiciously chosen finite-difference techniques may be more efficient, these techniques do not rely on subtle choices of differencing schemes. Instead they exploit smoothness properties which we expect to hold in equilibrium. They also are familiar to economists since they correspond to commonly used

ideas in (nonlinear) econometrics, and can be simply programmed. The basic idea is simple. The true solutions, $V^1(I)$ and $V^2(I)$, lie in an infinite dimensional space, e.g., the space of differentiable functions. It is impossible to search over that entire space for functions which satisfy (1.i). Instead we shall specify a finite-dimensional space of functions and find that element which "best fits" the system (1.i). The crucial elements of a minimum weighted residual approach are therefore, first, defining the space of functions over which we search, and, second, defining a measure of fit.

One specific approach is to take our candidate V^i from the space of functions of the form

$$V^i(I_1,\, I_2) = \Sigma_{k,\,\ell=1}^{n}\; a_{k,\,\ell}^{i}\, \varphi_k(I_1)\; \varphi_\ell\,(I_2) \tag{4}$$

where $\{\varphi_j\}$ is the sequence of Chebyshev polynomials over $[0,\, \bar{I}]$ where $\bar{I}$ is chosen sufficiently large so as to be sure that the steady state inventory levels lie well below $\bar{I}$. (Specifically, this is equivalent to assuming that any inventory in excess of $\bar{I}$ shrinks instantly to $\bar{I}$.) For each guess of the coefficients in (4), we substitute each V^i into the right hand side of $(1.i)$. The result will generally be a nonzero function, $e_i(I_1, I_2)$, representing the error of the approximation.

To measure the goodness of fit of an approximation, one can use one of several approaches. A least-squares approach defines a scalar index of the overall fitness,

$$F = \int_0^{\bar{I}} \int_0^{\bar{I}} (e_1(I_1, I_2)^2 + e_2(I_1, I_2)^2)\, dI_1\, dI_2$$

The least-squares method chooses the approximation which [6] solves the finite-dimensional nonlinear optimization problem

$$\min_{a_{k,\,\ell}^{i}} \quad F$$

The $a_{k,\,\ell}^{i}$ coefficients then give us a V^i which we take to be an approximate solution to $(1.i)$.

Alternative choices for the finite-dimensional space of candidate approximations and the measure of fit may be made. This choice was made because the Chebyshev polynomials are orthogonal, and are used in approximation theory to produce approximations which nearly minimize the maximum absolute error. The least squares measure of fit is easy

[6] More precisely, we used a weighted collocation method since a Gauss-Lagrange quadrature method was used to calculate the L^2 norm of the error.

and tractable. In fact, this method strongly resembles recently developed techniques in nonparametric econometrics. Many of the same considerations which govern the quality of fit apply here.

In order to test the robustness of our results, we added a (small) term proportional to I^{-1} to the inventory holding cost function. This makes it costly to hold very small inventories and impossible to hold negative inventories. The initial results are limited. However, in those cases where the steady-state inventories were nonnegative in the original linear-quadratic model, the new steady state was only slightly different in terms of equilibrium output and price. It appears that the option of negative inventories which exists in linear-quadratic models is unimportant if the steady state has positive inventories.

The general value of this approach is obvious. Many exercises become possible. For example, the issue of price-cost margins becomes tractable with these techniques. Also, investment is another firm decision which can be distorted for strategic reasons, with the direction of the distortion depending on the nature of strategic interaction. A natural specification for investment is that it pushes the short-run marginal cost curve to the right in an equiproportionate fashion. Since this cannot be modelled in a linear-quadratic model, a nonlinear model is necessary. We anticipate many applications of these techniques in future work.

9: Conclusion

We have explored basic issues of market conduct and performance in a dynamic model which allows firms to choose both output and price. In some cases, equilibrium looks like Cournot or Bertrand, but in general outcomes are spread out between these extremes. Constant marginal costs and no adjustment costs implies Bertrand outcomes, but both steeply rising marginal costs and high adjustment costs imply Cournot outcomes. This dynamic model also allows us to examine firm behavior without imposing untenable restrictions concerning behavior. For example, identification of these structural relationships will assist in determining whether observations correspond to noncooperative behavior, the only behavior assumption made here.

While only a few substantive questions were examined here, the linear-quadratic frame-

work developed is capable of incorporating other aspects of reality, such as investment, learning curves, cost-cutting R&D, and private information about demand and technology. Numerical analysis of nonlinear models will allow us to incorporate many of these elements. In summary, there is no reason to limit strategies to either prices or quantities when examining questions of market conduct and performance.

References

Arvan, Lanny. "Some Examples of Dynamic Cournot Duopoly with Inventory," *Rand Journal of Economics*, 16, 4, pp. 569-580.

Bulow, Jeremy, John D. Geanokoplos, and Paul D. Klemperer. "Multimarket Oligopoly: Strategic Substitutes and Complements," *JPE*, 93, 3, 1985, pp. 488–511.

Bresnahan, Timothy. "Duopoly Models with Consistent Conjectures," *American Economic Review*, 71, 1981, pp. 934–45.

Bresnahan, Timothy. "Identification of Market Power," Economics Letters.

Cournot, A. *Recherches sur Les Principes de la Theorie des Richesses*. Paris, 1938.

Davidson, Carl and Deneckere, Raymond. " Long-Run Competition in Capacity, Short-Run Competion in Price, and the Cournot Model," *Rand Journal of Economics*,17, 3, 1986, pp. 404-415.

Domowitz, Ian, R. Glenn Hubbard, and Bruce C. Petersen, "Business Cycles and the Relationship Between Concentration and Price-Cost Margins,"*Rand Journal of Economics* 17 (Spring, 1986), pp. 1-17.

Gal-Or, Esther. "Information Sharing in Olgopoly," *Econometrica*, 53, 1985, 329–43.

Kirman, Alan P., and Matthew J. Sobel. "Dynamic Oligopoly with Inventories," *Econometrica* 42, No. 2 (March, 1979): 279–287.

Klemperer, Paul D., and Margaret Meyer. "Price Competition vs. Quantity Competition: The Role of Uncertainty," *Rand Journal of Economics*, 17, 4, 1986, pp. 404–415.

Kreps, David and Scheinkman, Jose. "Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes," *Bell Journal of Economics*, 14, 1983, pp. 326–337.

Kydland, Fynn. "Noncooperative and Dominant Player Solutions In Discrete Dynamic Games, *International Economic Review*, 16, 2, (June, 1975), 321-335.

Riordan, Michael, "Imperfect Competition and Dynamic Conjectural Variations" *Rand Journal of Economics*, 16, 1, (Spring 1985), pp. 41-50.

Rotemberg, Julio, and Garth Saloner. "The Cyclical Behavior of Strategic Inventories," *Quarterly Journal of Economics* 104 (February, 1989): 73 –97.

Salant, Stephen W., Sheldon Switzer, and Robert J. Reynolds. "Losses from Horizon-

tal Merger: The Effects of an Exogenous Change in Industry Structure on Cournot-Nash Equilibrium," *Quarterly Journal of Economics* 98 (No. 2, 1983): 185-199.

Singh, N. and Vives, Xavier. "Price and Quantity Competition in a Differentiated Duopoly," *Rand Journal of Economics*, 15, 1984, pp. 546–54.

Vives, Xavier. "On the Efficiency of Bertrand and Cournot Equilibria with Product Differentiation," *Journal of Economic Theory*, 36, 1, 1985, pp. 166–175.

Vives, Xavier. "Duopoly Information Equilibria: Cournot and Bertrand," *Journal of Economic Theory*, 34, 1, 1989, pp. 71–94.