Viscous Demand

Roy Radner
Leonard N. Stern School of Business
New York University

August 1999

Working Paper Series
Stern #IS-99-4

Viscous Demand

ROY RADNER
STERN SCHOOL OF BUSINESS
NEW YORK UNIVERSITY
44 W. FOURTH STREET
NEW YORK, NY 10012

May 18, 1999

ABSTRACT. In many markets, demand adjusts slowly to changes in prices, i.e., demand is "viscous." For such a market, the time path of a firm's prices acquires added significance, compared with the case of instantaneous demand response. In this paper I explore some problems in strategic dynamic pricing of a service, in the presence of viscous demand, for simple models of a monopoly and a duopoly.

1. INTRODUCTION

In many markets, demand adjusts slowly to changes in prices, i.e., demand is "viscous." For such a market, the time path of a firm's prices acquires added significance, compared with the case of instantaneous demand response. In this paper I explore some problems in strategic dynamic pricing of a service, in the presence of viscous demand, for a monopoly and a duopoly. In particular, the viscosity of demand confers on each firm a kind of monopoly power, since it can raise its price above that of its competitors without immediately losing all of its customers. As we shall see, this phenomenon can lead to equilibrium pricing behavior and market outcomes that differ significantly from what one would predict in the absence of demand viscosity. In particular, it provides a rationale for the importance of market share, and it provides an explanation for the "kinked demand curve." It also explains how apparently "competitive" pricing behavior can lead to outcomes that mimic those of collusion.

There are many reasons for the viscosity of demand. In the case of a service (which will be the focus of this paper), such as a subscription to a magazine, newspaper, or long-distance telephone carrier, the viscosity of demand is probably best explained by an "attention budget." The (potential) consumer cannot be thinking every hour, or even every week, about which long-distance carrier to use. Rather, the consumer rethinks such decisions from time to time, regularly or at some random intervals, perhaps triggered by some events. We may think of the consumer as a "server" for a queue of decision problems, which are served according to some system of priorities. The time it takes for a decision problem to be "served" will depend not only on the duration of the service time, but also on the pattern of arrivals of decision problems at the queue. Decisions about a service that belongs to a higher priority class will be served more quickly, and so the demand for that service will display less viscosity. For example, it has been verified empirically that, on the average, consumers who spend more on long-distance telephone service also exhibit less viscosity of demand. (For more details, see Section 2.)

In the case of a durable good, a person with a piece of equipment that has been recently acquired will typically not switch to another model as soon as a supplier lowers its price, but will wait until the equipment has suffered some wear and tear, or has otherwise depreciated. The analysis of a a market for durable goods requires a quite different model, and is beyond the scope of this paper.

Viscosity of demand is to be distinguished from "stickiness of demand" due to the cost of switching suppliers. In a pure switching-cost model, a consumer will switch suppliers as soon as the gross saving from switching exceeds the cost. By contrast, in the viscosity model, if two suppliers offer identical services at different prices, eventually all customers will switch to the lower-price supplier.

In the model explored in this paper, at every instant of time (time is continuous), each consumer purchases a service at a rate equal to 0 or 1. If there is more than one supplier, then a consumer who buys the service must also choose the supplier. I shall consider two cases: (1) monopoly, and (2) duopoly; in the latter case I shall assume that the services provided by the two suppliers are identical. With the exception of one result, I also assume that all consumers are identical in their two demand parameters : (1) the long-run willingness-to-pay (WTP) for the service, and (2) the "viscosity coefficient" (see below). Finally, for mathematical convenience, I assume that there is a continuum of consumers, which I normalize to have total mass one.

In the case of a monopoly (Section 3), at every time t, let $X(t)$ denote the mass of consumers who are actually purchasing the service at that time (the "customers"), and let $P(t)$ denote the price of the service, per unit time. One usually calls $X(t)$ the market penetration at time t. Let w denote the long-run willingness-to-pay for the service (the same for all consumers), and suppose that the price of the service is exogenously constrained not to exceed a value m. The market penetration evolves according to the differential equation,

$$X'(t) = \begin{cases} k[w - P(t)][1 - X(t)], & \text{if } 0 \leq P(t) \leq w, \\ -k[P(t) - w]X(t), & \text{if } w \leq P(t) \leq m. \end{cases}$$

Here the strictly positive parameter k is the reciprocal of the "viscosity coefficient;" *smaller values of* k correspond to *a higher viscosity.* According to the first line of the equation, if the price is less than the WTP, then noncustomers become customers, at a rate proportional to the mass of noncustomers. On the other hand, according to the second line, if the price exceeds the WTP, then customers stop buying the service, at a rate proportional to the mass of customers. If the price remains constant, say p, and p is less than w, then the market penetration will approach unity (in the limit), whereas if p exceeds w, then the market penetration will approach zero. The market penetration will remain unchanged from its initial value if p just equals w.

Since customers do not disappear immediately when the price is raised above the WTP, increasing the market penetration by lowering the price represents a kind of *investment*.

One interpretation of the upper bound, m, is that $X'(t)$ equals minus infinity (i.e., the market penetration falls immediately to zero) if the price $P(t)$ ever exceeds m. Other plausible laws of motion for the market penetration, which are qualitatively similar to the one described above, are discussed in Section 2.

For this simple model, I characterize the dynamic price policy that maximizes the monopolist's total discounted profit, for the special case in which the monopolist's cost is proportional to the market penetration, i.e., there is a constant marginal cost, say c and zero fixed cost. (We may assume that $c < w$; otherwise the monopolist would not be in business.) There are two cases to be distinguished. In Case 1, the discount rate, say r, is not too large relative to the other parameters of the model, and the maximum price, m, is sufficiently close to w (these conditions can be made precise). In this case, the optimal policy for the monopolist is a *target penetration policy*, namely there is an optimal target market penetration, such that the monopolist sets the price equal to zero if the current market penetration is strictly less than the target, and sets the price equal to w, the WTP, if the current penetration is at least as large as the target. The optimal target, s, is given by the simple formula,

$$s = k(w - c)/[r + k(w - c)]. \tag{1.1}$$

Note that in this formula the optimal target, s, is a decreasing function of the marginal cost, c, but that the target penetration is not complete even when the marginal cost is zero. Also, s is increasing in k, the inverse of the viscosity, and decreasing in r, the discount rate.

Note, too, that although the consumers are acting in a "myopic" manner at those times when they consider whether to buy the service or not, in fact, given the price strategy of the monopolists, their choices are optimal.

In Case 2 (when the conditions of Case 1 do not hold), the target penetration policy described above is dominated by a policy in which the price oscillates rapidly between 0 and m. In fact, *strictly speaking there is no optimal policy; we may say that in the "optimal" policy the price oscillates infinitely fast between 0 and m !* (One can provide a precise meaning of this statement.) Such a situation is hardly realistic, and provokes a reconsideration of the behavioral assumptions of the model when it has these parameter values. In particular, if prices are oscillating very quickly, one would not expect consumers to react so myopically as they do in the model described above. For example, one might expect (boundedly rational) consumers to forecast prices in some "adaptive" manner, e.g., with a moving average of past prices. Although a complete characterization of the monopolist's optimal policy in the face of consumers with adaptive expectations is unavailable, it is possible to show that, if it exists, it will typically lead (roughly speaking) to cyclical fluctuations of prices and market penetration. Such pricing could be interpreted as a policy of intermittent "sales."(The results for Case 2 of the monopolist are a product of joint work with Thomas Richardson; see Radner and Richardson, 1997).

I next consider (Section 4) a model of a duopoly with a law of motion analogous to that of the monopoly model. In a duopoly there are three classes of consumers: (1) customers of firm 1, (2) customers of firm 2, and (3) noncustomers, i.e., consumers who are not customers of either firm. Again, all consumers have the same long-run willingness to pay, but each duopolist controls his own price dynamically. The *state of the system at time t* describes the number (mass) of consumers in each class at that time. Depending on the firms' prices, relative to each other and to the consumers' long-run willingness-to-pay, w, consumers will flow from one class to the other. More precisely, if the lowest price is less than w, then consumers will flow to the firm with that price, whereas if the lowest price exceeds w, then customers will flow from both firms into the class of noncustomers. When both firms charge the same price, and it is less than w, then noncustomers will flow to both firms in proportion to the firms' current stocks of customers. When the firms both charge a price equal to w , the masses of consumers in the two firms will remain constant. The total number of customers of the two firms will be called the *market penetration*, and the ratio of the number of customers of a firm to the market penetration will be called that firms's *market share.* Again, as in the case of a monopoly, if a firm raises its price above that of its competitor and above the WTP, it does not immediately lose all of its customers. For this reason, *lowering its price to increase its market share represents a kind of investment.*

In the context of such a model I shall describe a dynamic game in which the players are the two duopolists. I shall describe, and demonstrate the existence of, a family of (Nash) equilibria with (roughly) the following properties: (1) the strategies of the two players are *stationary*, i.e., at each time each firm's price depends only on the current state of the system (such an equilibrium is usually called *Markovian*); (2) each equilibrium in the family is characterized by two parameters, which may be interpreted as *the target market penetration* of the two firms and the *target market share of firm 1* (the target market share of firm 2 is, of course, one minus the target market share of firm 1); (3) if a firm's market share is strictly less than its target, then it charges a price equal to zero, and the other firm charges a price equal to m (the maximum price); (4) if both firms' market shares are equal to their targets, then they both charge a price equal to zero if the (total) market penetration is strictly less than the target, and a price equal to w if it is greater than or equal to the target. In order for a strategy-pair to form an equilibrium, the parameters of the model must satisfy certain conditions (similar to those in the monopoly case), and the target penetration and market shares must lie in a certain (nonempty) set. To simplify the analysis, I assume that the cost parameter, c, is zero, so that a firm's profit equals its revenue.

To describe the results for a duopoly more fully, I need some additional notation. Let S denote the target market share of Firm 1, and let $(1-Z)$ denote the target market penetration (in other words, Z is the target mass of noncustomers). I call the pair of strategies described above a (Z, S) *target strategy-pair.* Let $\zeta \equiv 1-s$ [cf. (1.1) above]. Under assumptions that correspond to Case 1 of the monopoly model, I demonstrate that there exists a number $\zeta' < \zeta$ such that, if m id sufficiently close to w, and if

$$\begin{aligned} \zeta' &\leq Z \leq \zeta, \\ 1-s &\leq S \leq s, \end{aligned}$$

then *the corresponding* (Z, S) *target strategy-pair is an equilibrium of the game* (Theorem 2).

Since the total mass of consumers is unity, we can characterize the system state at any time by the vector (x, z), where x is the mass of customers of Firm 1, and z is the mass of noncustomers. A comparison with the monopoly case shows that the equilibrium path is *efficient,* in the sense that the total profit of the two firms is maximized, if and only if (1) the initial state vector is on the line $x = S(1-z)$, and (2) $Z = \zeta$. Thus, if these conditions are satisfied, then the industry outcome as a whole mimics the monopoly outcome. On the other hand, if $Z < \zeta$, then the asymptotic market penetration will be greater than it would be in the corresponding monopoly, and the system spends more time in the regime in which one or both firms charge a zero price. In this sense, the equilibrium can be more "competitive" than the monopoly outcome. Note that as r/kw approaches zero, the minimum target market penetration, $1-\zeta$, approaches unity.

An implication of Theorem 2 is that a division of the market into shares S and $(1-S)$ is self-sustaining, so that no "explicit collusion" is required once the target S is determined. On the other hand, since thare is a nondegenerate interval of market shares that can be so sustained, some kind of "coordination" on a particular value of S is required. The same is true of the target market penetration, Z.

When the target market penetration and market shares have been reached, if one firm *lowers its price below* w, *the other will do so, too,* whereas *if a firm raises its price above* w, *the other firm will not respond.* The effect of this is that each firm's demand curve will not be differentiable at the point at which its price equals w. Thus

Theorem 2 provides a game-theoretic explanation of the so-called "kinked demand curve" in a duopoly. (I owe this observation to T. Groves.)

It is clear that the models analyzed in Sections 3 and 4 are quite special, and even for the special duopoly model the results are incomplete. In Section 5, I describe various extensions of the analysis, as well as some open problems. First, an obvious question is whether there are other equilibria of the duopoly game of Section 4. I cannot characterize the full set of equilibria, nor do I know whether there are other Markovian equilibria. However, I can show that a variant of Anderson's (1985) concept of quick-response equilibrium yields an equilibrium outcome that is identical to that of the "efficient" (Z, S) target strategy-pair, but with somewhat different strategies. In this equilibrium, (1) if the initial total market penetration of the two firms is less than s , then both firms charge a zero price until the market penetration reaches s, after which they both charge a price equal to w; (2) if the initial total market penetration is at least s, then both firms charge w; (3) once the total market penetration reaches or exceeds s, if either firm charges a price strictly less than s, then the other firm will "immediately retaliate" by charging a price equal to *zero* (in fact, both firms will switch to zero); (4) on the other hand, if in cases (1)-(2) either firm *raises* its price, the other firm will not change its own price. *Again, the equilibrium strategies have the effect that the industry as a whole imitates a monopolist's behavior, while the two firms maintain their initial relative market shares.*

In the second part of Section 5, I sketch a model of "adaptive expectations" and summarize the results of Radner and Richardson (1997) concerning the optimal monopoly pricing strategy corresponding to the model of Section 3.

In the third part of Section 5, I generalize the monopoly model to allow for the possibility that different consumers have different willingness-to-pay. I give a heuristic argument to suggest that *if there is a sufficient dispersion of WTP, w, in the population of consumers, then under the optimal pricing policy of the monopolist there is no steady state of the system.*

In the last part of Section 5, I discuss some of the issues that arise in the analysis of a general oligopoly and a "competitive" market.

In Section 6, I provide some bibliographic notes on the few previously published papers about somewhat related models of demand, notably by Selten, Phelps and Winter, Rosenthal, and Rosenthal and Chen.

2. A MODEL OF DEMAND VISCOSITY

As noted in the Introduction, I envisage the viscosity of demand as resulting from the fact that consumers typically have an "attention budget." For decisions about whether to start or stop a service (such as a subscription to a newspaper), switch suppliers (such as switching from one long-distance carrier to another), or switch brands of a commodity bought repeatedly (such as breakfast cereal), a consumer will devote only limited attention to the decision problem during any day or week. Rather, the consumer rethinks such decisions from time to time, regularly or at some random intervals. We may think of the consumer as a "server" for a queue of decision problems, which are considered according to some system of priorities. The time it takes for a decision problem to be "served" will depend not only on the duration of the service time, but also on the pattern of arrivals of decision problems at the queue. Decisions about a service that belongs to a higher priority class will be served more quickly, and so the demand for that service will display less viscosity. For example, it has been verified empirically that, on the average, consumers who spend more on long-distance telephone service also exhibit less viscosity of demand.

For example, consider a consumer who is a potential or actual subscriber to home delivery of the New York Times. Let $P(t)$ denote the subscription rate (price) per unit of time (e.g., per week) at date t, and let $x(t)$ be 1 or 0 according as the consumer is or is not a subscriber at date t; call $x(t)$ her *state at date t*. At stochastic *decision times*, $T_1, T_2, \ldots$, the consumer considers whether to be a subscriber or not, i.e. whether or not to change her state. Suppose that she does so by comparing the current price with her *long-run willingness-to-pay*, say w.Assume further that if w is greater than the current price, then she remains or becomes a subscriber, whereas if w is less than the current price, then she remains or becomes a nonsubscriber. (If w and the price are equal, she leaves her state unchanged.) A decision time may be triggered by various events: receiving a bill in the mail, seeing a commercial on TV, talking to a friend, etc. But even the news of a price change may not be sufficient to engage the consumer's immediate attention; if the change is small enough, she may put off the decision problem until a less busy day. On the other hand, if she learns that the weekly subscription rate is about to go up to $1000 per week, she will no doubt cancel her subscription immediately.

Notice that there are two aspects of bounded rationality embodied in this model. First, the consumer is not continuously deciding whether or not to subscribe, but only *visits this decision problem from time to time.* Second, when the consumer does reconsider her decision whether or not to subscribe, she makes her decision *myopically*, comparing w with the current price, rather than attempting to forecast what the price will be until her next decision time. The first aspect is inevitable in almost all decision making, although some decisions are programmed to be made automatically by a computer. [Examples are (1) programs used by large businesses to select the "optimal" long-distance carrier for every long-distance call, and (2) "programmed trading" by large traders in securities markets.] Nevertheless, the mean rate at which decisions are reconsidered will be influenced by the "importance" of the decision, so that more important decisions may be revied more frequently. The second aspect of bounded rationality, the myopia of the decision criterion, is more plausible the less frequently prices and other relevant variables change in time. As we shall see in

Section 3, for some values of the model's parameters, the firms' optimal price strategy may require rapidly oscillating prices, in which case myopic decision behavior by the consumer may no longer be plausible, which will lead me to consider a so-called "adaptive expectations" model of consumer behavior (Section 4).

If there are very many consumers in the NYTimes market, behavior that was described (qualitatively) above will imply that consumers will be observed to "flow" in and out of the customer (subscriber) category, at rates that depend on the current price and on the joint distribution of the frequency of decision making and w in the population of consumers. In this paper, aside from some remarks in Section 5, I shall be content to analyze a model in which *all consumers have the same long-run willingness-to-pay*, w. I shall also (1) represent the the population of consumers by a continuum, which I can conventionally take to have mass one, and (2) represent time as a continuous variable, t. Denoting the mass of customers at time t by $X(t)$, I shall assume that the motion of the mass of customers is governed by a differential equation,

$$X'(t) = \mu[P(t), X(t)]. \tag{2.1}$$

Even so, the formulation in (2.1) is too general for my purposes. We can expect the *law of motion*, μ, to be nonlinear, with the following properties:

(1) When the price exceeds w, consumers will flow out of the customer pool, at a rate that is proportional to the mass of customers.

(2) When w exceeds the price, consumers will flow out of the noncustomer pool, at a rate that is proportional to the mass of noncustomers.

(3) The rates of flow will be monotone in the (absolute value of the) difference between w and the price.

(4) If the price exceeds w by too much, consumers will flow out of the customer pool at a very rapid rate.

A simple formula for the law of motion that has these properties is:

$$\mu(p, x) = \begin{cases} k(w-p)(1-x), & \text{for } p \leq w \\ -k(p-w)x, & \text{for } w \leq p \leq m, \\ -\infty, & \text{for } p > m. \end{cases} \tag{2.2}$$

where m is a parameter, typically $> w$. Also, I shall assume that the price must be nonnegative.

Other similar formulas might also be reasonable, depending on the application. In particular, the piecewise linearity of (2.2) and the jump down to $-\infty$ should be thought of as an approximation to a smoother function. Smoothness would also be introduced by "noise" in the flows, caused by movements of consumers that are not explained by price and willingness-to-pay alone. Such noise is observed in practice. For example, a nonlinear law of motion can be derived from the assumption that consumers' choices (when made) are generated by a discrete-choice logit model.

In the case of a duopoly, there are three groups of consumers: customers of firm 1, customers of firm 2, and noncustomers. If at least one of the firms has a price lower than w, then consumers flow to the lowest price firm; whereas if both firms' prices exceed w, then their customers flow to the noncustomer group. The precise model will be spelled out in Section 4; in particular, specific assumptions need to be made for the case of ties.

In Section 5, I shall introduce a monopoly model in which w has a (nondegenerate) distribution in the population of consumers, although the available results for that case are meager.

3. MONOPOLY

In this section I shall present the analysis of a model of a monopoly for a service, as described in Section 2. Recall that the total mass of consumers is taken to be unity, the mass of actual customers at time t is denoted by $X(t)$, and the price of the service (per unit time) is denoted by $P(t)$. All of the consumers have the same long-run willingness-to-pay (WTP) for the service, denoted by w. (See Section 5 for a discussion of generalizations of this assumption.) The *law of motion* for $X(t)$ is given by (2.1) and (2.2) in Section 2, which I reproduce here for the convenience of the reader:

$$X'(t) = \mu[P(t), X(t)], \tag{3.1}$$

$$\mu(p,x) \equiv \begin{cases} k(w-p)(1-x), & \text{for } p \leq w \\ -k(p-w)x, & \text{for } w \leq p \leq m, \\ -\infty, & \text{for } p > m. \end{cases} \tag{3.2}$$

I assume that the monopolist's cost per unit time is proportional to the mass of customers, i.e., is equal to $cX(t)$, where c is a nonnegative constant. [Note: One could add a fixed cost, but its magnitude would not affect the optimal pricing policy, although it would influence the net profitability of the service.] The monopolist's total discounted profit is therefore

$$V = \int_0^\infty e^{-rt}[P(t) - c]X(t)dt, \tag{3.3}$$

where $r > 0$ is the exogenously given rate of interest. Given the initial mass of customers, $X(0)$, the monopolist wants to choose a price path to maximize the profit V in (3.3). For reasons that will be explained below, I make the following assumptions:

$$0 < r < k(w-c); \tag{3.4a}$$

$$0 < w \leq m; \tag{3.4b}$$

$$0 \leq P(t) \leq m. \tag{3.4c}$$

In view of the third line of (3.2), the second inequality of (3.4c) is not really an assumption, but it is included there for completeness.

By Blackwell's Theorem, one can without loss of generality take the optimal price policy to be *stationary*, in the sense that, for some function Φ,

$$P(t) = \Phi[X(t)]. \tag{3.5}$$

Theorem 1. *If (3.4) is satisfied, and m is sufficiently close to w, then the optimal (stationary) policy is given by:*

$$\Phi(x) = \begin{cases} 0, & \text{if } 0 \leq x < s, \\ w, & \text{if } s \leq x \leq 1, \end{cases} \tag{3.6a}$$

where

$$s \equiv \frac{k(w-c)}{r + k(w-c)} > 1/2. \tag{3.6b}$$

The maximum profit is

$$\begin{aligned} V &= \frac{ws}{r}D(x) - c\left[\frac{rx + kw - (1-s)kwD(x)}{r(r+kw)}\right], \quad \text{for } x < s, \\ &= \frac{(w-c)x}{r}, \qquad \text{for } x \geq s, \qquad \text{where} \\ D(x) &\equiv \left(\frac{1-s}{1-x}\right)^a, \quad \text{where } a \equiv r/kw < 1. \end{aligned}$$

[For the proof of the theorem, see the Appendix.]

Remark 1. Call $X(t)$ the *market penetration*, and s the *target penetration*. If the initial market penetration is strictly less than the target, then, under the policy Φ, the penetration will increase monotonically to the target, reaching it in finite time. On the other hand, any penetration greater than or equal to the target is a steady state. (These conclusions hold even if the target, s, does not satisfy (3.6b), i.e., even if it is not optimal.)

Remark 2. Under the optimal policy, the market penetration never reaches, or even approaches, unity (unless it starts there), so that a strictly positive fraction of the consumers never become customers. These results are intuitively plausible in the light of the first line of the law of motion, (3.2); as the market penetration increases, the remaining mass of noncustomers, $[1 - X(t)]$, decreases towards zero, so that eventually the incremental discounted value of adding to the current customer base (market penetration) is unable to make up for the corresponding incremental loss of revenue from the current customer base

Remark 3. The optimal target penetration, s, is *decreasing* in the marginal cost, c, the viscosity, $(1/k)$, and the discount rate, r, approaching unity as $k \to \infty$ and/or $r \to 0$. However, $s < 1$ even when the marginal cost is zero.

Remark 4. The value function V is increasing and differentiable, *strictly convex* for $x < s$, and *linear* for $x > s$.

Remark 5. We have assumed that our boundedly-rational consumers are "myopic," in the sense that whenever they make a decision they do so on the basis of the currently prevailing price, not a projection of future prices. However, one can easily verify that, if the monopolist uses the policy Φ of the theorem, *then the consumers' postulated behavior is an optimal response at those instants of time when they make a decision.* (In fact, this last statement will remain true in the face of any "target market penetration" price policy, not necessarily an optimal one.)

I shall now briefly discuss the nature of the optimal price policy if the assumptions of Theorem 1 are not satisfied. The parameters of the model are: w, k, r, c, and m, all of which are assumed to be nonnegative. In fact, the parameter space can be partitioned into two parts, say R and R', such that (1) the conclusion of the theorem holds in R, whereas (2) in the set R' there is no exactly optimal policy, but the supremum of the profit will be approached as the price oscillates faster and faster between 0 and m. The assumptions stated in the hypothesis of Theorem 1 determine a strict subset of R, so that the conclusion is actually valid for a somewhat larger set. [For a full treatment of both cases, and a precise characterization of the two sets R and R', see (Radner and Richardson, 1997).]

In control theory, the kind of policy that is "optimal" in the set R' is sometimes called a *measure-valued* or *generalized* control. It should be clear that the behavior of even boundedly-rational consumers facing a very rapidly oscillating price is unlikely to conform to the kind of model of viscous demand described in Section 2 (and in the hypothesis of Theorem 1). When faced with such a price policy, even "myopic" consumers are more likely to react to some (possibly weighted) *average of past prices,* rather than to the current price at the instant of decision. Such a model of consumer behavior, and its implications, will be described briefly in Section 5.

4. DUOPOLY

In a duopoly there are three classes of consumers: (1) customers of firm 1, (2) customers of firm 2, and (3) noncustomers, i.e., consumers who are not customers of either firm. The *state of the system at time t* describes the number (mass) of consumers in each class at that time. Depending on the firms' prices, relative to each other and the consumers' long-run willingness-to-pay, w, consumers will flow from one class to the other. More precisely, if the lowest price is less than w, then consumers will flow to the firm with that price, whereas if the lowest price exceeds w, then customers will flow from both firms into the class of noncustomers. When both firms charge the same price, and it is less than w, then noncustomers will flow to both firms in proportion to the firms' current stocks of customers. When the firms both charge a price equal to w , the masses of consumers in the two firms will remain constant. The total number of customers of the two firms will be called the *market penetration*, and the ratio of the number of customers of a firm to the market penetration will be called that firms's *market share.*

In the context of such a model I shall describe a dynamic game in which the players are the two duopolists (see below for a precise mathematical formulation). I shall describe, and demonstrate the existence of, a family of equilibria with (roughly) the following properties: (1) the strategies of the two players are *stationary*, i.e., at each time each firm's price depends only on the current state of the system (such an equilibrium is usually called *Markovian*); (2) each equilibrium in the family is characterized by two parameters, which may be interpreted as *the target market penetration* of the two firms and the *target market share of firm 1* (the target market share of firm 2 is, of course, one minus the target market share of firm 1); (3) if a firm's market share is strictly less than its target, then it charges a price equal to zero, and the other firm charges a price equal to m (the maximum price); (4) if both firms' market shares are equal to their targets, then they both charge a price equal to zero if the (total) market penetration is strictly less than the target, and a price equal to w if it is greater than or equal to the target.In order for a strategy-pair to form an equilibrium, the paramters of the model must satisfy certain conditions (similar to those in the monopoly case), and the the target penetration and market shares must lie in a certain (nonempty) set. To simplify the analysis, I assume that the cost parameter, c, is zero, so that a firm's profit equals its revenue. (See remarks in the next section.)

I now turn to a precise description of the model and results. Let $X(t)$ and $Y(t)$ denote, respectively, the masses of customers of firms 1 and 2. Then, adopting the convention that the total mass of consumers is unity, the mass of noncustomers is $Z(t) = 1 - X(t) - Y(t)$. Hence we can take the state of the system at time t to be $[X(t), Z(t)]$.

Let $P(t)$ and $Q(t)$ denote the prices at time t of firms 1 and 2, respectively. Suppose that at a given time t,

$$\begin{aligned} P(t) &= p, \quad Q(t) = q, \quad M = \min\{w, q\}, \\ X(t) &= x, \quad Y(t) = y, \quad Z(t) = z. \end{aligned}$$

To describe the law of motion of the system, let $X'(t)$ denote the time-derivative of $X(t)$; then the following table shows the values of $X'(t)/k$ for the various cases of the relative magnitudes of $p, q, w,$ and M:

Case	$X'(t)/k$
$p < M$	$(w-p)z + (q-p)y,$
$p > M$	$-(p-M)x,$
$p = q < w$	$[x/(x+y)](w-p)z,$
$p = w < q$	$[x/(x+z)](q-p)y,$
$p = q = w$	0

Tablel 4.1. $X'(t)/k$ as a function of (p, q)

(Recall that k is the inverse of the viscosity coefficient.) The law of motion for $Y(t)$ is determined symmetrically and, since the total mass of consumers is unity, $X'(t)+Y'(t)+Z'(t) = 0$. In particular, if $p = q = w$, then $X'(t) = Y'(t) = Z'(t) = 0$.

Note that $X(0) > 0$ implies that $X(t) > 0$ for all t, and similarly for Y and Z. Unless I explicitly mention otherwise, I shall assume that

$$X(0), Y(0), Z(0) \text{ are all } > 0. \tag{4.1}$$

I shall also assume, as in the monopoly model, that each firm's prices are confined to the closed interval $[0, m]$, where $m \geq w$ is an exogenously given parameter.

A *history, H(t)*, of the system at time t describes the time-path of the state of the system up to and including time t, and the time path of prices up to but not including time t, i.e.,

$$H(t) = [\{X(s), Z(s); 0 \leq s \leq t\}, \{P(s), Q(s); 0 \leq s < t\}]. \tag{4.2}$$

A *strategy* for a firm is a mapping that determines, for each time t, its price at time t as a function of the history $H(t)$. A pair of strategies is called *feasible* if it determines a time path, $[X(t), Z(t), P(t), Q(t), t \geq 0]$, such that the payoffs of the two firms are well defined. The *payoff* (total discounted profit) of firm 1 is given by

$$V = \int_0^\infty \exp(-rt)P(t)X(t)dt, \tag{4.3}$$

where $r > 0$ is an exogenously given rate of interest. (Recall that costs are zero, so that profit equals revenue.) Firm 2's payoff is defined analogously.

The set of feasible strategy-pairs is not a product space, so it is not possible to define a game and its associated equilibria in the normal way. Instead, I shall use the concept of a *generalized game* (REF?). For any strategy ψ of Firm 2, let $\Phi(\psi)$ denote the set of strategies ϕ of Firm 1 such that the strategy-pair (ϕ, ψ) is feasible; such a strategy ϕ will be called *a feasible response to* ψ. (Note that the set of feasible responses may be empty.) The set $\Psi(\phi)$ of feasible responses by Firm 2 to a strategy ϕ of Firm 1 is defined analogously. A feasible strategy-pair (ϕ, ψ) is called a *(Nash) equilibrium* if neither firm can increase its payoff by unilaterally switching to another feasible response. (Subgame-perfection can be defined analogously.) A firm's strategy is called *stationary* if its current price is a function of the current state of the system only (not the full history at the current date). An equilibrium strategy-pair is called *Markovian* if the strategies are stationary. (Markovian equilibria are automatically subgame-perfect.)

I shall demonstrate the existence (under certain assumptions) of a family of particularly simple Markovian equilibria, indexed by two parameters, a target total market penetration, and a target division of the market between the two firms. Formally, let Z and S be numbers between zero and one, where $(1 - Z)$ is interpreted as the *target market penetration*, and S and $(1 - S)$ are interpreted as the *target market shares* of Firms 1 and 2, respectively. The pair (ϕ, ψ) of stationary strategies will be called a *(Z,S) target strategy-pair* if the prices $p = \phi(x, x)$ and $q = \psi(x, z)$ are given by the following table. The table divides the (x, z) state space into four regions, and shows the corresponding prices and laws of motion in each region.

Case	Region	Prices	Law of Motion
Case 1A	$x = S(1 - z)$, $z \leq Z$	$p = q = w$	$X'(t) = Z'(t) = 0$
Case 1B	$x = S(1 - z)$, $z > Z$	$p = q = 0$	$X'(t) = [x/(x + y)]kwz, \quad Z'(t) = -kwz$
Case 2	$x < S(1 - z)$	$p = 0, q = m$	$X'(t) = kw(1 - x), \quad Z'(t) = -kwz$
Case 3	$x > S(1 - z)$	$p = m, q = 0$	$X'(t) = -kwx, \quad Z'(t) = -kwz$

Table 4.2. A (Z, S) Target Strategy-Pair

Figure 4.1 indicates the motion of the state vector in the triangle

$$\Delta \equiv \{(x, z) : x \geq 0, z \geq 0, x + z \leq 1\}. \tag{4.4}$$

From any point in Δ the state vector, $[X(t), Z(t)]$, moves to the line $x = S(1 - z)$, with $Z(t)$ decreasing. Once on this line, say at (x, z), if $z > Z$ then the state vector moves down the line until $Z(t) = Z$, and stays there; if $z \leq Z$ then the state vector stays at (x, z). Thus any point (x, z) on the line segment

$$\begin{aligned} x &= S(1 - z), \\ 0 &\leq z \leq Z, \end{aligned}$$

is a steady state of the system. Figure 4.2 indicates the corresponding motion of the state in the (X, Y) plane.

Define

$$\begin{aligned} a &\equiv \frac{r}{kw}, \\ s &\equiv \frac{1}{a+1} \\ \varsigma &\equiv \frac{a}{a+1}. \end{aligned} \tag{4.5}$$

Make the following assumptions:

$$\begin{aligned} 0 &< a < 1, \\ 0 &< w \leq m, \\ 0 &\leq P(t) \leq m, \\ 0 &\leq Q(t) \leq m. \end{aligned} \tag{4.6}$$

Theorem 2. *If the assumptions (4.6) are satisfied, and if m is sufficiently close to w, then there exists $\zeta' < \varsigma$ such that, if*

$$\begin{aligned} \zeta' &\leq Z \leq \varsigma, \\ 1 - s &\leq S \leq s, \end{aligned}$$

then the (Z, S) target strategy-pair is an equilibrium.

[See the Appendix to this section for the proof of the theorem.]

Remark 1. A comparison with Theorem 1 shows that the equilibrium path is *efficient,* in the sense that the total profit of the two firms is maximized, if and only if (1) the initial state vector is on the line $x = S(1 - z)$, and (2) $Z = \zeta$. Thus, if these conditions are satisfied, then the industry outcome as a whole mimics the monopoly outcome. On the other hand, if $Z < \zeta$, then the asymptotic market penetration will be greater than it would be in a corresponding monopoly, and the system spends more time in the regime in which one or both firms charge a zero price. In this sense, the equilibrium can be more "competitive" than the monopoly outcome.

Remark 2. An implication of the theorem is that a division of the market into shares S and $(1 - S)$ is self-sustaining, so that no "explicit collusion" is required once the target S is determined. On the other hand, since thare is a nondegenerate interval of market shares that can be so sustained, some kind of "coordination" on a particular value of S is required. The same is true of the target market penetration, Z.

Remark 3. As $a \equiv (r/kw)$ approaches zero, the minimum target market penetration, $1 - \zeta$, approaches unity.

5. EXTENSIONS AND PROBLEMS

In this section I describe three extensions of the previous analysis, and also discuss some open problems. (See Sec. 1 for a summary.)

5.1. Quick-Response Equilibrium. In this subsection I describe a non-Markovian equilibrium equilibrium of the duopoly model of Section 4 that formalizes behavior in which each firm retaliates against a price cut by the other firm with a price cut of its own. Although the strategies are different, the equilibrium outcome is the same as that for the particular (Z, S) target equilibrium described in Section 4 in which $Z = \zeta$ and S is the initial market share of Firm 1. The equilibrium strategies have the following properties: (1) if the initial total market penetration of the two firms is less than $s \equiv (1 - \varsigma)$, then both firms charge a zero price until the market penetration reaches s, after which they both charge a price equal to w; (2) if the initial total market penetration is at least s, then both firms charge w; (3) once the total market penetration reaches or exceeds s, if either firm charges a price strictly less than s, then the other firm will "immediately retaliate" by charging a price equal to *zero* (in fact, both firms will switch to zero); (4) on the other hand, if either firm *raises* its price above athe equilibrium price, then the other firm will not change its own price. (Once again, each firm faces a "kinked demand curve.") Note that *on the equilibrium path the industry as a whole imitates a monopolist's behavior, while the two firms maintain their initial relative market shares.* This property of the equilibrium path constrasts with the seemingly "competitive" behavior of the firms (a point made by Anderson (1985) in a different context).

Technically, the game-theoretic approach differs somewhat from that of Section 4. I use a concept of equilibrium derived from the approaches of T. A. Marschak and R. Selten (1978) and R. M. Anderson (1985). Following Anderson, I shall call this a "quick-response equilibrium." This concept formalizes the intuitive notion that if time is continuous then one firm can respond "immediately" to changes in the other firm's price. In order to sidestep some of the difficulties of doing game theory with continuous time (see, e.g., Stinchcombe, 1992), this approach deals with a family of discrete-time approximations to the continuous-time model. The framework is consequently notationally more complicated than that of Section 4. I shall present here only the model and the results. (For proofs, see [Radner, 1997.]

The underlying model used here is that described at the beginning of Section 4, and thus has time varying continuously. In particular, the law of motion is the one given in Table 4.1. However, in a *quick-response equilibrium (QRE)* one does not define a game directly for the situation of continuous time, but rather approximates that situation with a family of discrete-time games. Accordingly, for each number $h > 0$, define a game $G(h)$ as follows: for every nonnegative integer multiple nh of h, the two firms simultaneously choose respective prices that will be operative during the half-open interval, $[nh, (n+1)h)$, and in that interval the masses of customers of the two firms, $X(t)$ and $Y(t)$, evolve according to the law of motion described in Table 4.1. Thus a *strategy* in the game $G(h)$ is defined in the usual way for a discrete-time game, and each strategy determines a time path, $[P(t), Q(t), t \geq 0]$ of the prices of the two firms. The *payoff* (total discounted profit) of firm 1 is given by

$$V = \int_0^\infty \exp(-rt)P(t)X(t)dt,$$

where $r > 0$ is the exogenously given rate of interest. (Recall that costs are zero, so that profit equals revenue.) Firm 2's payoff is defined analogously. Note that the integral (4.2) is well-defined, since P is a simple function, which we may take to be right-continuous.

Suppose that $\sum = \{\Sigma(h)\}$ is a family of strategy pairs such that $\Sigma(h)$ is a strategy-pair in the game $G(h)$, and let $[V(h), W(h)]$ denote the corresponding payoffs of the two firms, respectively. The family $\sum$ is a *quick-response equilibrium (QRE)* if the following two conditions hold: (1) for every initial state $(x, y) >> 0$ there exists a number $h(x, y) > 0$ such that, for every strictly positive number $h \leq h(x, y)$, the strategy pair $\Sigma(h)$ is a Nash equilibrium of the game $G(h)$; (2) the limit payoffs exist, namely,

$$V = \lim_{h \to 0} V(h), \text{ and } W = \lim_{h \to 0} W(h).$$

The numbers (V, W) will be called the *QRE payoffs*.

For each h, let $[X(t; h), Y(t; h)]$ denote the state of the system at time t determined by the (QRE) strategy-pair $\Sigma(h)$ in game $G(h)$, and let the corresponding prices be $[P(t; h), Q(t; h)]$. If, in addition, the limit trajectory exists, namely

$$[X(t), Y(t), P(t), Q(t)] = \lim_{h \to 0}[X(t; h), Y(t; h), P(t; h), Q(t; h)],$$

then I shall call the limit trajectory the *QRE path*. In the QRE of the duopoly model that I shall describe, the QRE path will exist, and furthermore the QRE payoff for each firm will be its discounted profit along the QRE path.

Note that the definition of QRE given thus far does not include any notion of subgame-perfecton, i.e, it is not required that the "threats" of retaliation against price cuts be "credible." It is therefore desirable to define a stronger version version of QRE that responds to this need. Accordingly, for every $\theta > 0$, let $\Delta(\theta)$ denote the open triangle,

$$x > \theta, y > \theta, x + y < 1 - \theta.$$

I shall say that the QRE is $\sum$ is *quasi-subgame-perfect (QSP)* if, for every $\theta > 0$, there is an $H(\theta) > 0$ such that, for every positive $h < H(\theta)$, every time t, and every history of the game $G(h)$ through time t for which

$$[X(s), Y(s)] \text{ is in } \Delta(\theta), \quad 0 \leq s \leq t,$$

the continuation of the strategy-pair $\Sigma(h)$ from time t on is a QRE of the continuation game.

Here is a heuristic description of the QRE strategies.Fix $h > 0$, and let I be a positive integer. It suffices to describe firm 1's strategy, firm 2's being defined symmetrically. Firm 1 charges a price equal to zero until the first date nh at which $Z(nh) \leq \zeta$ (i.e., total market penetration reaches or exceeds the target). Thereafter, firm 1 charges a price equal to w, with the following exception: if at some date nh firm 2 *undercuts* firm 1 by charging a price strictly less than w, then firm 1 will retaliate by charging a price equal to zero for the next I periods, and then return to the price w at date $(n + I + 1)h$; by symmetry, firm 2 will do likewise. The sequence of retaliation periods will also be started anew after any failure of either firm to carry out the prescribed retaliation. I shall show that by taking I large enough, firm 1 can deter firm 2 from any deviation from the QRE path, since whatever value firm 1 can gain in period nh will be offset by a sufficiently large loss in the subsequent I periods. It will be important to show that the number I can be taken to be independent of h, although it will depend on the state, $[X(nh), Y(nh)]$, in which the deviation occurs. Note that firm 1 does *not* respond if firm 2 *raises* its price above firm 1's price (no matter what the value of $Z(t)$ at the time). (For a precise description of the QRE strategies, see Radner, 1997.)

Theorem 3. *There exists a choice of the function I such that the family $\sum$ of strategy-pairs is a quasi-subgame-perfect quick-response equilibrium of the family of games $\{G(h)\}$. Furthermore, the QRE path exists, and the QRE payoff for each firm is its total discounted profit along the QRE path. Firm 1's QRE payoff is given by*

$$\begin{aligned} V &= \left(\frac{wx}{r}\right)\frac{(1-\zeta)\zeta^a}{(1-z)z^a}, \quad \textit{for } z > \zeta, \\ &= \left(\frac{wx}{r}\right), \quad \textit{for } z \leq \zeta. \end{aligned}$$

A corresponding equation holds for Firm 2.

We see that, for fixed z, Firm 1's QRE payoff is linear in its initial market penetration, x. On the other hand, by the Lemma of Section 4, the function f defined by

$$f(z) \equiv (1-z)z^a$$

is decreasing if $z > \zeta$ (and, incidentally, increasing if $z < \zeta$); use the fact that

$$\zeta = a/(1+a).$$

Hence, for fixed x, firm 1's QRE payoff is independent of z for $z < \zeta$, and is increasing in z for $z > \zeta$.

5.2. Adaptive Expectations. As noted in Section 3, for some parameter values the monopolist's "optimal" price oscillates "infinitely fast" between zero and the maximum value. As noted in Section 1, purely myopic choices by consumers would implausible under such circumstances. Suppose, therefore, that - when considering whether or not to purchase the service - a consumer forecasts the future price to be some moving average of past prices, and makes the purchase decision on the basis of that forecast. Accordingly, let $\tilde{P}(t)$denote the price "forecast" at time t, and suppose that $\tilde{P}(t)$ is determined by

$$\tilde{P}(t) = \theta \int_0^\infty e^{-\theta s} P(t-s) ds, \tag{5.2.1}$$

where $\theta > 0$ is a given parameter of the model. Assume that the law of motion (3.1) is modified to read

$$X'(t) = \mu[\tilde{P}(t), X(t)]. \tag{5.2.2}$$

Following a precedent in the literature on expectations (Arrow and Nerlove, 1958), I shall call this the *adaptive expectations model*, or more precisely, the θ-AR model (since the model is parametrized by θ). Note that

$$\lim_{\theta \to \infty} \tilde{P}(t) = P(t),$$

so that the model of Section 3 may be considered a limiting case of the adaptive expectations model, which one might denote the ∞-AR model. In fact, although an explicit solution for the monopolist's optimal pricing strategy is not know for the θ-AR model (with $\theta < \infty$), one can show that (1) an optimal policy exists, and (2) for large finite θ the corresponding optimal policy is approximately optimal for $\theta = \infty$. In particuar, when the optimal price for the ∞-AR model oscillates infinitely fast, the optimal price for large finite θ also oscillates, but at a finite rate (see Radner and Richardson, 1997). *Such behavior by the monopolist might be interpreted as a policy of "intermittent sales."*

5.3. A Monopolist Facing Consumers with a Distribution of Willingness-to-Pay. Up to this point, the entire analysis for both monopoly and duopoly has been carried out under the assumption that all the consumers have the same willngness-to-pay (WTP) for the service. Extending the analysis to the case in which which the consumers are heterogeneous with respect to WTP would be desirable, but thus far I have been unable to do this in any generality in the context of the present model. In this subsection I give a heuristic argument that suggests that oscillatory pricing may be common when the WTP is sufficiently dispersed in the population of consumers.

I shall say that *a consumer is of type w* if his WTP for the service is w. Although the consumers are heterogeneous with respect to WTP, I assume that the monopolist can charge only a single price at any given time. In other words, the monopolist is unable to discriminate among the consumers according to their type. Of course, if the monopolist could do so, then the problem would reduce to that of Section 3, for each type of consumer.

Suppose that the type w lies in the interval $[0, W]$ and has an absolutely continuous distribution in the population of consumers, and let g denote the density function of this distribution. The absolute continuity of the distribution expresses the assumption that the WTP is "dispersed" in the population of consumers. Let $X(w,t)$ denote the "fraction of consumers of type w" who are customers of the monopolist at time t, or to be more precise, the total mass of customers of type not exceeding w at time t is given by the integral,

$$\int_0^w X(u,t)g(u)du.$$

Let $X(t)$ denote the function $X(.,t)$; then $X(t)$ is the *state of the system at time t.* Thus a state of the system is a function, say x, from $[0, W]$ to $[0,1]$ such that the integral

$$\int_0^w x(u)g(u)du$$

exists for each w. Assume that, for each w, the state variable $X(w,t)$ obeys the law of motion (3.1)-(3.2).

A *steady state* for a particular pricing policy would be a state-price pair, say (ξ, φ), such that

$$[X(0), P(0)] = (\xi, \varphi) \Rightarrow [X(t), P(t)] = (\xi, \varphi) \text{ for all } t > 0.$$

(For example, in the case of a single WTP (Section 3), if $s \leq x \leq 1$, then (x, w) is a steady state. The law of motion implies that a steady state (ξ, φ) satisfies

$$\xi(w) = 0 \text{ or } 1 \text{ according as } w < \text{ or } > \varphi. \tag{5.3.1}$$

(Note that the law of motion implies nothing about $\xi(p)$.) I shall give a heuristic argument that suggests that *there is no steady state for an optimal policy.*

The heuristic argument uses the so-called "Bellman Optimality Conditions" (see the Appendix to Section 3). Suppose that an optimal pricing policy exists. For any state x , let $V(x)$ denote the monopolist's maximum profit, starting from the state x; V is the monopolist's *value functional.* Note that V is a mapping from the infinite-dimensional space of states to the real numbers. Although the state-space is infinite-dimensional, one can still (under suitable conditions) formulate the notion of *the partial derivative of V with respect to x(w)*, which I shall denote by $V'(w,x)$. [REF?] Correspondingly, the Bellmanian Functional for the monopolist's optimization problem (cf. the appendix to Section 3) is

$$\begin{aligned} B_V(p,x) &\equiv p\int_0^W x(w)g(w)dw - rV(x) - \int_{w<p} k(p-w)x(w)V'(w,x)dw \\ &\quad + \int_{w>p} k(w-p)[1-x(w)]V'(w,x)dw. \end{aligned}$$

If (ξ, φ) is a steady state for the optimal policy, then by (5.3.1), if $p > \varphi$,

$$B_V(p,\xi) \equiv p\int_\varphi^W g(w)dw - rV(\xi) - \int_\varphi^p k(p-w)V'(w,x)dw,$$

and hence

$$\begin{aligned} \frac{\partial B_V(p,\xi)}{\partial p} &= \int_\varphi^W g(w)dw - \int_\varphi^p kV'(w,x)dw \\ &\rightarrow \int_\varphi^W g(w)dw \text{ as } p \searrow \varphi. \end{aligned}$$

Therefore, if $\varphi < W$, then the Bellmanian is strictly increasing in the price p for $p > \varphi$ and sufficiently close to φ. Hence (under suitable regularity conditions) $\varphi < W$ cannot be optimal at the state ξ.

On the other hand, if (ξ, W) were a steady state for the optimal policy, then by (5.3.1), $\xi(w)$ would be zero for all $w < W$, and hence $V(\xi)$ would be zero. However, starting from such a state ξ it is possible for the monopolist to make a strictly positive profit, e.g., by setting the price equal to zero for a positive amount of time, and then setting the price equal to any positive value thereafter. Hence $\varphi = W$ cannot be optimal, either.

To make this heuristic argument rigorous, I would have to do two things. First, I would have to postulate conditions on the model such that a theory of the Bellmanian Functional for an infinite-dimensional state space is valid. (Alternatively, one could formulate a model with a large but finite number of consumer types, whose masses are uniformly small.) Second, I would have to provide conditions such that the corresponding "Bellman Conditons" for optimality are necessary rather than sufficient. Such an analysis is beyond the scope of this paper.

5.4. Oligopoly and Competition. In this subsection I sketch some of the problems to be faced in generalizing the analysis to the cases of oligopoly (with more than 2 firms) and "competition."

The formulation of a model with more than 2 identical firms is straightforward, if one assumes that the consumers always flow to the firm(s) with the lowest price, or flow into the noncustomer category if all firms charge a price greater than the WTP (assuming for the moment that all consumers have the same WTP). One might even conjecture that there are Markovian equilibria analogous to that of the duopoly of Section 4, in which each firm has a target market share, and sets its price low or high according as its current market share is less than or greater than its target.

The assumption that the firms are identical is of course problematic, even in the case of a duopoly. Differentiation of the firms could have two consequences: (1) each consumer would have different WTPs for the services of different firms; (2) consumers could differ in their WTP profiles, some consumers preferring the services of one firms, and other preferring the services of another firm (different firms have different "clienteles"). Given the difficulty suggested in Section 5.3 above, I shall not even venture a conjecture about the nature of equilibria in such cases.

If by "competition" (perfect or imperfect) one means a "large" number of "small" firms, then one might want to reconsider the law of motion. First, with a law of motion similar to that described above, if a small firm's price were the lowest among a large number of firms, then that firm would face a relatively enormous rate of increase in demand, which it might not be able to meet in the short run. Second, a consumer facing a large number of firms might not immediately be able to identify the most preferred one (e.g., the one with the lowest price), and hence might migrate in stages from less preferred firms to more preferred ones. (Models studied by Selten and by Phelps and Winter have the flavor of the latter phenomenon; see Section 6.)

The analysis of models that incorporate all these considerations would appear to present daunting difficulties. In fact, my limited exposure to practical problems of this kind suggests to me that managers are not (knowingly) following optimal policies in the pricing situations that they face. If this impression is correct, then a satisfactory theory would have to model the bounded rationality of the managers, as well as that of the consumers.

6. BIBLIOGRAPHIC NOTES

The earliest theoretical paper that I am aware of that is related to the present one in spirit is (Selten, 1965). In Selten's model, finitely many firms repeatedly face a market. with an exogenously given total "demand potential" (see below). Each firm has a linear cost function (they may be different). Time is discrete. If we specialize the Selten model to the case of a duopoly, with total "demand potential" constant in time, then in the notation of Section 4 of the present paper, the demand (quantity sold) for firm 1 obeys the law of motion,

$$X(t+1) - X(t) = KP(t) + (1-K)Q(t) - P(t+1),$$

where K is a constant, the same for both firms. (Recall that firm 1's demand in period t is $X(t)$, and the prices of firms 1 and 2, resp., are $P(t)$ and $Q(t)$.) The law of motion for firm 2's demand is determined symmetrically. The initial conditions are

$$\begin{aligned} X(0) &= m - P(0), \\ Y(0) &= n - Q(0), \end{aligned}$$

where m and n are given constants, or equivalently,

$$\begin{aligned} X(0) + P(0) &= m, \\ Y(0) + Q(0) &= n. \end{aligned}$$

One can verify that, for all periods t,

$$X(t) + P(t) + Y(t) + Q(t) = m + n.$$

Selten calls

$$\begin{aligned} M(t) &\equiv X(t) + P(t) \text{ firm 1's ``demand potential,'' and,} \\ N(t) &\equiv Y(t) + Q(t) \text{ firm 2's ``demand potential.''} \end{aligned}$$

We may interpret a single firm's demand potential in a given period as what its demand would be if its price were zero in that period. The model implies that the *total* demand potential remains constant, i.e., equal to the initial value, $(m + n)$.

For a game with a fixed (and known) number of periods, Selten demonstrates the existence of a subgame-perfect equilibrium. (An infinite-horizon game is studied in Part II of the paper.) The model is related to the present one in that customers do not instantaneously react fully to changes in prices. However, it is not clear to me how to reconcile this law of motion with a model of consumer behavior like that sketched in Section 2 of the present paper. [Some readers will recognize Selten's paper as the one in which he introduced the concept of "subgame perfection."]

In a series of three papers, Rosenthal and Chen have studied related models of duopoly with "customer loyalties." Each of the models is a discrete-time, infinite-horizon, symmetric, non-zero-sum stochastic game, in which the players are the two firms, and there are finitely many identical customers who act in accordance with a fixed "rule of thumb" (different in each paper). In (Rosenthal, 1982), in each period "each buyer purchases from the same seller from whom he purchased in the last period unless that seller has raised his price, in which case the buyer purchases from the current-period low-pricesetter." In this game, Rosenthal demonstrates the existence of a Markov equilibrium in mixed strategies, and studies its properties. In (Rosenthal, 1982), customer loyalties are "weaker." In each period, "after the sellers have set their current-priod prices, a random device ... determines whether (with probability α) each buyer will remain loyal whenever his previous-period seller has not raised price, or whether (with probability $(1-\alpha)$) all buyers abandon their loyalties and purchase from the current-period low-price seller." In this game, Rosenthal demonstrates the existence of an epsilon-equilibrium in stationary mixed strategies. (There is an exact equilibrium from an initial state in which one seller has all the customers.) In (Chen and Rosenthal, 1996), in each period, if one seller's price is strictlly less than the other's, then one customer shifts from high-price seller to the other one (unless, of course, the high-price seller has no customers). Again, there is a Markov equilibrium in mixed strategies, and the authors study the effect on this equilibrium of changing the parameters of the model. [Note: I have omitted a description of the detailed assumptions in each paper.]

Phelps and Winter (1970) studied a moodel at the other end of the spectrum from duopoly, namely, one in which there is a very large number of small sellers. Each seller i "subjectively assumes" that the (continuous-time) law of motion for his market share, $X_i(t)$, is

$$X_i'(t) = \delta[p_i(t), P_i(t)]X_i(t),$$

where $p_i(t)$ is his current price, $P_i(t)$ is the customer-weighted mean of the other firms' prices, and δ is a skew-symmetric function with plausible properties. In fact, *the "true" law of motion in the model is different from, and more complicated than, the one "subjectively perceived" by the firms;* I omit the details. At time $t = 0$, each firm chooses a price path that maximizes the present value of of its discounted profits, *under the assumption that the average price $P_i(t)$ will remain equal to its initial value, $P_i(0)$, for all time*, in other words, assuming that the law of motion for its market share is

$$X_i'(t) = \delta[p_i(t), P_i(0)]X_i(t).$$

This represents a further simplification of the firm's subjective perception of the law of motion of its demand. (Alternatively, it may represent an approximation that may be reasonable for a short enough time interval.) In particular, each firm assumes that, if its price remains constant during some interval of time, then its demand will grow or decrease exponentially in that interval. This model is sufficiently far from the one in the present paper that I shall not attempt to summarize the results of the authors' analysis, except to report that any steady state of the system will depart from the standard picture of the equilibrium of a "neoclassical" model.

I have already noted the difference between the model of viscous demand presented here, and models intended to capture consumers' "switching costs." For a recent treatment, and references to earlier literature, see (Padilla, 1995).

7. Acknowledgements

I thank Peter B. Linhart for encouragement and helpful discussions throughout the research that led to this paper; he also read several previous drafts. Thomas J.

Richardson collaborated on the analysis of the case of monopoly, and in particular demonstrated the importance of considering generalized controls (see Radner and Richardson, 1997). Aldo Rustichini and Prajit Dutta participated in early stages of the research, and contributed helpful insights. Robert M. Anderson introduced me to the concept of quick-response equilibrium, and kindly allowed me to read and cite his unpublished paper (1985). Theodore Groves pointed out the connection with the kinked demand curve, and made other helpful suggestions, on the occasion of a presentation of an early draft at a seminar at the University of California, San Diego. I was also grateful for the opportunites to present earlier reports on this project at AT&T Bell Lalboratories, New York University, the Calif. Inst. of Technology, the Stockholm School of Economics, Boston University, and other institututions I may have forgotten. Last but not least, this research was supported in part by the Leonard N. Stern School of Business, New York University. As usual, any errors are attributable solely to me.

8. References

Anderson, R. M., "Quick-Response Equilibrium," Department of Economics, U. of Calif., Berkeley, Feb. 1985 (unpublished).

Arrow, K. J., and M. Nerlove, "A Note on Expectations and Stability," *Econometrica*, 26 (1958), 297-305.

Chen, Y., and R. W. Rosenthal, "Dynamic Duopoly with Slowly Changing Customer Loyalties," *International J. of Industrial Organization*, 14 (1996), 269-296.

Marschak, T. A., and R. Selten, "Restabilizing Responses, Inertia Supergames, and Oligopolistic Equilibria," *Quarterly J. of Econ.*, Feb. 1978, 71-93.

Padilla, A. J., "Revisiting Dynamic Duopoly with Consumer Switching costs," *J. of Econ. Theory*, 67 (1995), 520-530.

Phelps, E. S., and S. G. Winter, Jr., "Optimal Price Policy under Atomistic Competition," in E. S. Phelps, et al, *Microeconomic Foundations of Employment and Inflation Theory*, W. W. Norton, New York, 1970.

Radner, R., and T. J. Richardson, "Monopolists and Viscous Demand," AT&T Bell Laboratories, Murray Hill, NJ, 1997 (unpublished).

Rosenthal, R. W., "A Dynamic Model of Duopoly with Customer Loyalties," *J. of Econ. Theory*, 27 (June 1982), 69-76.

Rosenthal, R. W., "Dynamic Duopoly with Incomplete Customer Loyalties," *International Econ. Rev.*, 27 (June 1986), 399-406.

Selten, R., "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragetragheit," *Zeitschrift fur die Gesamte Staatwissenschaft*, 121 (1965), 301-324, 667-689.

Simon, L., and M. B. Stinchcombe, "Extensive Form Games in Continuous Time, Part I: Pure Strategies, *Econometrica*, 57 (1989), 233-253.

M. B. Stinchcombe, "Maximal Strategy Sets for Continuous-Time Game Theory," *J. of Econ. Theory*, 56 (April 1992),235-265.