MARKETING SCIENCE
Vol. 13, No. 3, Summer 1994
Printed in U.S.A.

NONLINEAR PRICING TO PRODUCE INFORMATION

DAVID J. BRADEN AND SHMUEL S. OREN
University of Rochester
University of California, Berkeley

We investigate the firm's dynamic nonlinear pricing problem when facing consumers whose tastes vary according to a scalar index. We relax the standard assumption that the firm knows the distribution of this index. In general the firm should determine its marginal price schedule as if it were myopic, and produce information by lowering the price schedule; "bunching" consumers at positive purchase levels should be avoided.

As a special case we also consider a market characterized by homogeneous consumers with a static, but unknown, demand curve. We show that when there are repeat purchases the forward-looking firm should tend towards penetration pricing; otherwise its strategy should tend towards skimming. We extend our insights to more general settings and discuss implications for pricing product lines.

(Pricing, Segmentation)

1. Introduction

Nonlinear pricing appears across such a wide variety of goods and services that researchers have given it considerable attention since Oi's (1971) classic work on two-part pricing. Most of this literature considers markets where consumers' demand curves can be indexed by a scalar parameter that varies with tastes, incomes, or other characteristics. One can also view this parameter as indexing the same demand curve in different states of nature; in the words of Leland and Meyer (1976), "uncertainty is a special case of 'heterogeneous markets'." Research to date, however, has assumed the firm knows the index's distribution—a generally inappropriate assumption, particularly for firms entering new markets. Our principal contribution is to relax this assumption and provide insight into nonlinear pricing when a firm has incomplete information about the structure of demand for its product.

To better grasp the issues we are addressing, consider the Arena Shipper™, a collapsible container patented by A. R. Arena Products, Inc. It offers substantially lower costs of handling, transportation, storage and disposal than its competing technologies, which are primarily 55-gallon steel drums and wooden crates with plastic liners. Arena sells its Shippers to a few new customers each year in quantities ranging from several units to thousands of units. Each prospective customer has a demand curve that depends upon his degree of interest and sophistication in minimizing shipping costs. Arena's quantity discount schedule, in turn, depends upon its assumptions about the distribution of this "involvement index" among its potential customers.

0732-2399/94/1303/0310$01.25

Although Arena doesn't know the index's distribution, it can learn about the distribution through experience in the market. Each sale gives Arena additional information about the distribution of the index. A missed sale is informative as well, but the quality of the information is lower, since Arena's inferences about the index—and hence its distribution—aren't as sharp as when it makes a sale. By lowering its price schedule Arena decreases the chance of excluding a prospective customer from making a purchase, thereby increasing its chance of generating high-quality information about the demand structure it faces.

To get a handle on the interplay between nonlinear pricing and information production, we investigate two types of firms: myopic and experimenting. Each faces a market with stochastic demand, is uncertain about its distribution, and infers outcomes of stochastic variables by observing consumption decisions. The sole distinction between the firms is that the myopic one does not take the future benefits of learning into account when determining current price schedules, whereas the experimenting one does—it prices to produce information. For each of the firm types we characterize the optimal consumption patterns and pricing policies. We establish that it is the consumption pattern's monotonicity (with respect to the consumer index), and not its magnitude, that is relevant to learning, and that the firm loses information when it "bunches" purchases at the same consumption level for several values of the index. The experimenting firm should almost never bunch purchases at positive levels and should lower the entire price schedule to reduce the number of consumers priced out of the market. We also show that, in general, the experimenting firm can determine its marginal price schedule as if it were myopic.

We then consider the special case of a market where demand curves are static with an unknown distribution. For both types of firms, the price schedule decreases over time until the firm resolves its uncertainty. One of our insights from the general model—that the experimenting firm should price lower than the myopic one—of course continues to hold. This seems to contradict Lazear's (1986) result that a firm using *linear* pricing should price *higher* when it takes the future into account. What drives his result, however, is the assumption that there are no repeat purchases. We show that Lazear's insight when the firm is restricted to linear pricing also holds when the firm can price nonlinearly. That is, when there are repeat purchases the experimenting firm should tend towards penetration pricing; otherwise its strategy should tend towards skimming. The final section extends our insights to more general settings, and discusses implications for pricing product lines.

Related Research

Researchers have long recognized the potential for using prices to "produce information" in the sense that the resulting consumption pattern reduces the firm's uncertainty, enabling it to make more informed future decisions. Representative work includes Rothschild (1974), Grossman, Kihlstrom and Mirman (1977), Lazear (1986), McLennan (1984) and Aghion et al. (1991). The literature, however, has restricted the firm to linear pricing. When feasible, nonlinear pricing is not only more profitable but offers greater potential for eliciting information about the distribution of consumer preferences.

Leland and Meyer (1976), Spence (1977), and Goldman, Leland and Sibley (1984), among many others, discuss the single-period nonlinear pricing problem. Their work has been extended in two significant ways. Oren, Smith and Wilson (1983) generalized the analysis by considering a symmetric oligopoly in a Nash-Cournot framework. Littlechild (1975) and Oren, Smith and Wilson (1982) extend the single-period analysis to markets with positive demand externalities.

Robinson and Lakhani (1975), Dolan and Jeuland (1981), and Kalish (1983) consider dynamic pricing when there are word-of-mouth effects and cost reductions through learning. Dhebar and Oren (1985, 1986) consider price dynamics in markets with demand

externalities. In contrast to these papers, we focus on the dynamics that arise as a result of the firm's changing state of information about its demand structure employing concepts from the adaptive control literature.

Adaptive control models in marketing have been traditionally associated with advertising and promotions, dating back to the seminal work by Little (1966). Numerous subsequent papers published in this area address various issues arising in the context of simultaneous experimentation and control of an advertising or promotional policy. The experimentation is aimed at estimating the effectiveness of the policy, or more specifically the model parameters, while the control objective is to maximize, say, the net present value of profits. One of the most advanced treatments of such problems is due to Pekelman and Tse (1980), who proposed a general adaptive control scheme that combines experimentation in advertising levels, estimation of sales-advertising relationships, and the determination of advertising budgets when the unknown parameters of the underlying model may vary over time. Their paper also contains an excellent review of related work.

While from a theoretical perspective our paper employs concepts of adaptive and "dual" control similar to the above, the context and insights are totally different. We focus on situations in which the dominant feature is market heterogeneity and the goal of the nonlinear pricing policy is to exploit that aspect. Hence the experimentation is aimed at acquiring information about the distribution of customer types while the control amounts to designing the appropriate nonlinear price schedule.

2. Formulation

Consider a firm whose product can be bought in any nonnegative quantity. Technological, legal or economic circumstances preclude secondary resale. At the start of each period the firm chooses a price schedule $R(q; t)$ that specifies the total charge for supplying q units of its product to a consumer in period $t \in \{0, 1, \ldots, T\}$. Assume $R(0; t) = 0$, and that $R(\cdot; t)$ is nondecreasing and lower semicontinuous in its first argument. At any point q' where R isn't differentiable, define $R_q(q') = \lim_{x \uparrow q'} R_q(x)$.[1] Note there may be an upward discontinuity at $q = 0^+$: Although the firm cannot charge a positive amount for zero consumption, it can charge a "subscription fee" for the right to buy any positive amount.

The firm incurs a cost $C(q; t)$ for providing q units of its product to an individual in period t. We make the usual assumption that $C(\cdot; t)$ is nondecreasing and continuously differentiable for all $q > 0$, and that $C(0; t) = 0$. As with the price schedule, we allow the possibility of an upward discontinuity at $q = 0^+$; that is, the firm might incur a fixed cost in providing its good to an individual.

Consumers

Consumers can be indexed by a scalar parameter θ that varies with tastes, incomes, or other characteristics. The index θ is a continuous random variable with support $\Theta = [0, \bar{\theta}]$. Each consumer has a maximum willingness to pay $W(q, \theta; t)$ for quantity q in period t. Given θ, the consumer seeks to maximize his surplus

$$CS^*(\theta; t) = \max_{q \geq 0} W(q, \theta; t) - R(q; t),$$

choosing the highest level of consumption that does so:

$$q^*(\theta; t) = \max \{\arg\max_{q \geq 0} [W(q, \theta; t) - R(q; t)]\}, \qquad \theta \in \Theta = [0, \bar{\theta}], \quad t \in \{0, \ldots, T\}. \tag{1}$$

[1] Most functions in this paper have time t among their arguments; its absence, as in $R(q)$, means "this holds for all $t \in \{0, 1, \ldots, T\}$". Subscripts denote (partial) differentiation.

The purchase decision depends solely upon W, the index θ, the current price schedule R, and time t.

Another interpretation of this formulation is to regard consumers as homogeneous and finite in number, with a willingness to pay that is subject to shocks each period through the variate θ. Under this interpretation, take—without loss of generality—the number of consumers to be one.

Regarding the willingness to pay function W we adopt the usual assumptions, which are stated formally in the appendix: No one is willing to pay for zero consumption. Once the (inverse) demand curve W_q falls below the marginal cost curve, it stays below it; and for every possible demand curve, there is a finite point at which this occurs. Demand curves are non-crossing and parameterized to shift back towards the origin as θ increases ($W_{q\theta} < 0$); that is, the marginal willingness to pay for any q decreases as θ increases.

Goldman, Leland and Sibley (1984) show that the self-selection condition (1) and our assumptions about W imply that any price schedule induces a consumption pattern q^* that is nonincreasing and upper semicontinuous on Θ. Let $\mathcal{Q}$ denote the set of all such functions:

$$\mathcal{Q} \equiv \{q\colon \Theta \to \mathbb{R}_+ \mid q \text{ is nonincreasing and upper semicontinuous}\}.$$

Our assumptions also imply that CS^* is continuous, and strictly decreasing wherever CS^* is positive. Thus there is an *inclusion index* $M(t)$ such that

$$M(t) = \begin{cases} \bar{\theta} & \text{if} \quad CS^*(\bar{\theta}) > 0, \\ \min\{\theta \mid CS^*(\theta; t) = 0\} & \text{otherwise.} \end{cases}$$

If $\theta < M(t)$ the consumer buys the product; if $\theta > M(t)$ there is no purchase.

The Firm's Learning Process

Suppose the consumer index θ has a probability distribution with density $g(\theta)$. Then at the beginning of any period t the firm's expected profit is

$$\int_\Theta [R(q^*; t) - C(q^*; t)]g(\theta)d\theta,$$

where we suppress the arguments of q^* when the context is unambiguous. When the firm knows W, C and g, its optimal strategy is to choose price schedules that maximize current expected profits, so its dynamic optimization problem can be solved as a sequence of single-period optimization problems.

But suppose the firm doesn't know some (possibly vector-valued) parameter ξ of the density g, yet can represent its beliefs about the true value of ξ with a probability distribution $F(\cdot \mid I(t))$. Its state of information $I(t)$ about ξ evolves as it accumulates experience in the market. We envision learning as a process whereby $I(t+1)$ consists of $I(t)$, the price schedule $R(\cdot; t)$, and the actual consumption decision $\hat{q}(t)$ corresponding to the outcome $\hat{\theta}(t)$ of θ.

The firm uses Bayes' rule to revise its beliefs about ξ. Since the firm observes $\hat{q}$ and not $\hat{\theta}$, its inferences about ξ depend on $\mathcal{S}(q^*, \hat{q}) = \{\theta \mid q^*(\theta) = \hat{q}\}$. When q^* is strictly decreasing at $\hat{\theta}$ the firm can infer the precise value of $\hat{\theta}$ through observation of $\hat{q}$. But when $\hat{\theta}$ is greater than the inclusion index M or occurs at a point where q^* isn't strictly decreasing, the firm knows only the interval in which $\hat{\theta}$ occurred. Consequently, the prior density on ξ is updated according to

$$f(\xi \mid I(t+1)) = \Phi[f(\xi \mid I(t)), q^*(\cdot; t), \hat{q}(t)], \tag{2}$$

where Ξ is the support of ξ, and

$$\Phi[f(\xi \mid I), q^*(\cdot), \hat{q}] = \begin{cases} \dfrac{f(\xi \mid I) g(\theta \mid \xi)}{\int_{\Xi} g(\theta \mid \xi) dF(\xi \mid I)} & \mathcal{S}(q^*, \hat{q}) \text{ a singleton,} \\ \dfrac{f(\xi \mid I) \int_{\mathcal{S}} g(\theta \mid \xi) d\theta}{\int_{\Xi} \int_{\mathcal{S}} g(\theta \mid \xi) d\theta dF(\xi \mid I)} & \text{otherwise.} \end{cases}$$

Dynamic Nonlinear Pricing with Passive Learning

The myopic firm maximizes its expected current profit

$$\int_{\Theta} [R(q^*; t) - C(q^*; t)] h(\theta \mid I(t)) d\theta$$

period-by-period subject to consumer self-selection (1). When doing so it uses its predictive density

$$h(\theta \mid I(t)) = \int_{\Xi} g(\theta \mid \xi) dF(\xi \mid I(t)),$$

where Ξ is the support of ξ. Each period the firm revises its beliefs about ξ in the manner described above (2).

Since any feasible price schedule induces a $q^* \in \mathcal{Q}$, profit maximized over $\mathcal{Q}$ is at least as great as that maximized over the set of feasible R. It is straightforward to show that any $q^* \in \mathcal{Q}$ is supported by

$$R(q; t) = \int_0^q W_q(s, \theta^*(s); t) ds, \tag{3}$$

where

$$\theta^*(q) = \max \{\theta \mid q^*(\theta) \geq q\}. \tag{4}$$

Goldman, Leland and Sibley (1984) show in fact that (3) is the most profitable price schedule. In light of this, the myopic firm's problem is

$$\underset{q^* \in \mathcal{Q}}{\text{maximize}} \int_{\Theta} [R(q^*) - C(q^*)] h(\theta \mid I) d\theta \tag{5}$$

subject to the self-selection condition (1), as well as (3) and (4).

For the moment, consider maximized consumer surplus CS^* as a state variable. One can easily verify that

$$dCS^*(\theta) = W_\theta(q^*, \theta) d\theta \quad \text{for all } \theta \in [0, M). \tag{6}$$

Because the price schedule is single crossing, this is both necessary and sufficient for the self-selection condition (1). Substituting $W - CS^*$ for R in the objective (5) and using (6), the myopic firm's problem becomes

$$\underset{q^* \in \mathcal{Q}, M \in \Theta}{\text{maximize}} \int_0^M [W(q^*, \theta) - CS^*(\theta) - C(q^*)] h(\theta \mid I) d\theta \tag{7}$$

subject to (6) and $CS^*(M) = 0$, where we recall that for $\theta > M$, consumption, and hence profit, is zero. Integrating the CS^* term in (7) by parts and using the constraints in turn yields

PROBLEM (MP):

$$\underset{q^* \in \mathcal{Q}, M \in \Theta}{\text{maximize}} \int_0^M \pi(q^*(\theta; t), \theta) d\theta$$

where

$$\pi(q, \theta; t) = [W(q, \theta; t) - C(q; t)]h(\theta \mid I(t)) + W_\theta(q, \theta; t)H(\theta \mid I(t)),$$

H being the cumulative distribution associated with h. Denote by q^L and M^L the values of q^* and M that solve (MP).

Dynamic Nonlinear Pricing with Active Learning

The experimenting firm chooses the schedule $R^E(\cdot; t)$ that maximizes its expected sum of discounted profits. In doing so it recognizes that its choice of R^E affects both current profits and, through the learning effect, future expected profits. With the aid of Bellman's equation we can write the experimenting firm's objective as seeking q^E and M^E that solve

PROBLEM (EP):

$$J[f(\cdot \mid I(t)); t] = \max_{\substack{q^* \in \mathcal{Q}, \\ M \in \Theta}} \int_0^{M(t)} \pi(q^*(\theta; t), \theta; t)d\theta + \rho \int_\Theta J[f(\cdot \mid I(t+1)); t+1]h(\theta \mid I(t))d\theta$$

for all $t \in \{0, 1, \ldots, T\}$, where

$$\pi(q, \theta; t) = [W(q, \theta; t) - C(q; t)]h(\theta \mid I(t)) + W_\theta(q, \theta; t)H(\theta \mid I(t)),$$

$$f(\xi \mid I(t+1)) = \Phi[f(\xi \mid I(t)), q^*(\cdot; t), \hat{q}(t)],$$

and $J[f; T+1] = 0$.

For every period but the last, the experimenting firm's objective includes a term that measures the extent to which learning from current experience enables the firm to increase its profit by making more informed future decisions.

3. Analysis

We will first discuss how a consumption pattern affects learning, and then characterize the optimal consumption pattern and inclusion index for both firm types. This section concludes with a description of how to construct the optimal price schedule.

The Production of Information

After the firm has chosen $q^*(\cdot; t)$ it will observe consumption $\hat{q}(t) = q^*(\hat{\theta}; t)$. Recall that for a given $q^* \in \mathcal{Q}$ each possible consumption outcome $\hat{q}$ corresponds to a set $\mathcal{S}(q^*, \hat{q})$ of states θ for which $q^*(\theta) = \hat{q}$. The collection of these sets forms a partition $\mathcal{P}(q^*)$ of the set Θ of states of the environment.

We define q_1^* as *more informative* than q_2^* if the partition it induces is a refinement of that induced by q_2^*. Since updating requires knowledge only of the set $\mathcal{S}$ that corresponds to the outcome of $\theta(t)$ in the previous period, *only the partition induced by a consumption pattern is relevant to the expected stream of future profits; the* magnitude *of $q^*(\cdot)$ is irrelevant to subsequent decisions.*

The information content of a consumption pattern is tied directly to its monotonicity. In Figure 1, for example, q_1^* is more informative than both q_2^* and q_3^*, since its "flat spots" are over subintervals of those of q_2^* and q_3^*. Less bunching is more informative because it provides a richer sample observation space. On the other hand gaps in the consumption pattern—downward discontinuities in q^*—add no noise whatsoever to the observation.

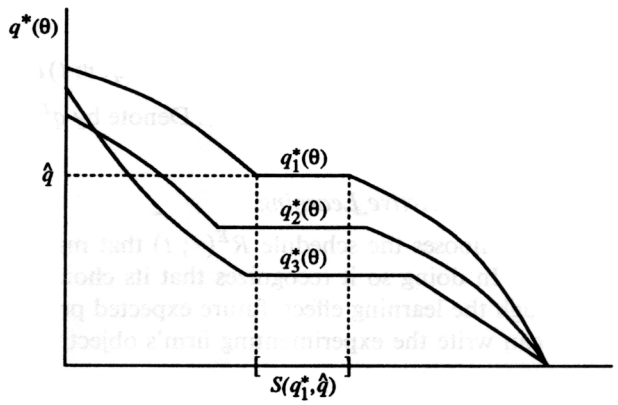

FIGURE 1. q_1^* is more informative than q_2^* and q_3^*.

The expected stream of discounted future profits is useful for comparing the *value* of information provided by different consumption patterns. We say that q_1^* provides *more valuable information* than q_2^* when it yields a higher discounted stream of expected future profits, which is when

$$\int_\Theta J(\Phi[f(\xi|I), q_1^*(\cdot), \hat{q}])h(\theta|I)d\theta > \int_\Theta J(\Phi[f(\xi|I), q_2^*(\cdot), \hat{q}])h(\theta|I)d\theta.$$

When the inequality is weak we say the information is *at least as valuable*.

If a policy is more informative than another then its information is at least as valuable. It follows that when q_1^* is at least as informative as q_2^* and yields current expected profit that is at least as high, the experimenting firm would never prefer to use q_2^*. In fact, the experimenting firm always chooses a consumption pattern that yields future expected profits that are at least as high as the myopic firm's (Proposition 1 in the Appendix).

In general it is difficult to compare the information content of two consumption patterns, since neither of the corresponding partitions may be a refinement of the other. In Figure 1, for example, we cannot say that q_2^* is more or less informative than q_3^*, since neither $\mathcal{P}(q_2^*)$ nor $\mathcal{P}(q_3^*)$ is a refinement of the other. There is, however, an important class of consumption patterns that are comparable: namely, those that strictly decrease wherever positive. For these consumption patterns all the bunching occurs at $q = 0$, so the degree of information depends solely upon the inclusion index $M = \sup\{\theta | q^*(\theta) > 0\}$.

A higher inclusion index is more informative and, by implication, at least as valuable. Thus in solving for an optimal consumption pattern, the experimenter could proceed as if myopic. Proposition 2 in the Appendix establishes that if the resulting solution is strictly decreasing, then it is optimal for the experimenter, given the inclusion index M. (Proposition 3 gives sufficient conditions for the myopic consumption pattern to be strictly decreasing.) Because a higher inclusion index is more informative, we expect the experimenting firm to choose M^E to be higher than M^L; Theorem 4 formalizes this intuition. In short, the firm that takes future profits into account when reaching current decisions should exclude less of the market than it would otherwise.

The Optimal Consumption Pattern

Our approach to determining price schedules for problems (MP) and (EP) is to solve for the optimal consumption pattern parametric on M, and to then characterize the optimal inclusion index M^*.

Consider first the myopic firm's problem, and define $\Pi^L(M)$ as the maximal value of the objective in problem (MP), given the inclusion index M:

$$\Pi^L(M) = \max_{q^* \in \mathcal{Q}} \int_0^M \pi(q^*(\theta), \theta) d\theta. \tag{8}$$

We know that

$$\int_0^M \max_{q \geq 0} \pi(q, \theta) d\theta \geq \Pi^L(M),$$

so any $q^*(\theta) \in \mathcal{Q}$ that maximizes the functional in (8) pointwise is an optimal solution for (8). With this in mind, we approach the firm's problem by determining the point-to-set correspondence

$$\Psi(\theta) = \arg\max_{q \geq 0} \pi(q, \theta)$$

and setting $q^*(\theta) = \max \{q \mid q \in \Psi(\theta)\}$. If $q^* \in \mathcal{Q}$ then it is optimal for the myopic firm; assume for now that this monotonicity constraint is nonbinding.

By the Maximum Theorem (Harris 1987, pp. 10–12), if $\Psi(\theta)$ is a singleton for each $\theta \in [0, M]$ then the optimal consumption pattern is continuous on $[0, M]$. This implies that if an optimal solution q^L is discontinuous at some point θ', then there are multiple maxima of $\pi(q, \theta')$. So the first-order conditions for problem (MP) are $\pi_q(q^L, \theta) = 0$ for $\theta \in [0, M]$ at which q^L is continuous, and at points of discontinuity

$$\pi_q(q^L(\theta^-), \theta^-) = \pi_q(q^L(\theta^+), \theta^+) = 0 \quad \text{and} \quad \pi(q^L(\theta^-), \theta^-) = \pi(q^L(\theta^+), \theta^+),$$

where

$$\pi_q(q, \theta) \equiv [W_q(q, \theta) - C_q(q)]h(\theta \mid I) + W_{q\theta}(q, \theta)H(\theta \mid I).$$

Note that at $\theta = 0$, the optimal quantity q^L is precisely the value at which marginal cost equals marginal willingness to pay, the socially efficient level of consumption.

Let's turn now to the experimenting firm's choice of an optimal consumption pattern, parametric on the inclusion index. Define

$$\Pi^E(M) = \max_{q^* \in \mathcal{Q}} \int_0^M \pi(q^*, \theta; t) d\theta + \rho \int_\Theta J[f(\cdot \mid I(t+1)); t+1] h(\theta \mid I(t)) d\theta,$$

where J is the stream of maximized expected future profits. Under our assumption that the monotonicity constraint is nonbinding, the myopic firm's optimal pattern q^L is also optimal for the experimenting firm, given the inclusion index (Proposition 2 in the Appendix). Note that q^L need not be continuous for this to hold: Unless they serve to break up intervals of bunching, gaps in the consumption pattern don't affect the quality of information produced by q^L.

The Optimal Inclusion Index

After determining the optimal consumption pattern parametric on M, the firm finds the optimal inclusion index by maximizing $\Pi(M)$. Consider first the experimenting firm's choice of M. Once the optimal consumption pattern has been determined, the experimenting firm's problem (EP) reduces to a stochastic control problem with $f(\cdot \mid I(t))$ as the state variable and $M(t)$ as the control variable. Differentiating $\Pi^E(M)$ with respect to M (taking into account that $M \leq \bar{\theta}$) yields the first-order condition

$$\pi^E(q^E, M) + \rho \frac{d}{dM} \int_\Theta J[\Phi(f(\xi \mid I), q^E, \hat{q})] h(\theta \mid I) d\theta \geq 0,$$

where the inequality becomes an equality if $M < \bar{\theta}$. Assuming M is an interior point, we can rewrite this as

$$W(q^E, M) - C(q^E) + W_\theta(q^E, M)\frac{H(M|I)}{h(M|I)}$$
$$= \frac{\rho}{h(M|I)}\frac{d}{dM}\int_\Theta J[\Phi(f(\xi|I), q^E, \hat{q})]h(\theta|I)d\theta$$

at $M = M^E$ and $q^E = q^E(M^E)$.

To find the myopic firm's optimal inclusion index, recall that its problem (MP) is just the experimenting firm's problem (EP) with $\rho = 0$. Consequently, the myopic firm chooses its inclusion index to satisfy

$$W(q^L, M) - C(q^L) + W_\theta(q^L, M)\frac{H(M|I)}{h(M|I)} = 0$$

at $M = M^L$ and $q^L = q^L(M^L)$.

Binding Monotonicity Constraint

The firm's problem becomes somewhat trickier when the monotonicity constraint is binding. For the myopic firm the solution is simple. Using a variational argument, one can show that if the optimal consumption pattern equals some positive constant q' over an interval $(\theta^1, \theta^2]$, then it satisfies the first-order conditions $\pi_q(q', \theta^1) = \pi_q(q', \theta^2) = 0$, and

$$\int_{\theta^1}^{\theta^2} \pi_q(q', \theta)d\theta = [(W_q(q', \theta) - C_q(q'))H(\theta|I)]_{\theta^1}^{\theta^2} = 0.$$

For an example of a binding monotonicity constraint on q^L, let

$$W(q, \theta) = (6 - \theta)q, \qquad C(q) = q^2/2,$$

and

$$H(\theta) = (4\theta^2 - .8\theta^3 + 0.25\theta^4)/12.8, \qquad \theta \in \Theta = [0, 4].$$

Then

$$\pi(q, \theta) = [(6 - \theta)q - q^2/2]h(\theta) - qH(\theta).$$

Setting $\pi_q = 0$ and rearranging terms yields

$$q(\theta) = 6 - \theta - \frac{4\theta - 1.8\theta^2 + 0.25\theta^3}{8 - 5.4\theta + \theta^2}.$$

As shown in Figure 2, this violates both the monotonicity constraint and the requirement that $q(\theta)$ be nonnegative. The optimal solution is to use the unconstrained consumption pattern for $\theta \leq 2.36$, and to bunch all $\theta > 2.36$ at $q^L(\theta) = 0.373$.

If it is indeed optimal for the myopic firm to bunch different customer types at the same consumption level, then in general no solution will exist for the experimenting firm. This is not because the functional (EP) is unbounded; rather, it is because the quality—and hence the value—of information changes discontinuously when a flat region in q^* is perturbed to have a slight slope. The firm can always increase expected profits by adjusting q^* ever closer to the optimal myopic solution q^L, but it decreases expected profits discontinuously if it actually adopts q^L. The practical implication for the experimenting firm is to deviate from the optimal myopic solution only to the extent necessary to detect different outcomes of θ. In addition, a bunching strategy can only be supported

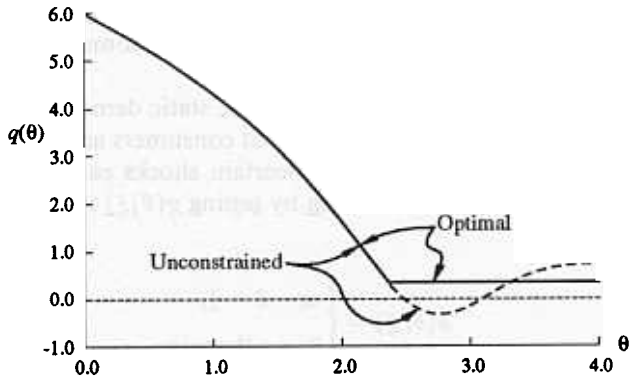

FIGURE 2. Example of Constrained Consumption Pattern.

by a nonconcave price schedule; if the firm decides *a priori* to only use concave price schedules then the bunching issue becomes moot.

The Optimal Price Schedule

After determining its optimal inclusion index, the firm can construct the total charge R as follows. The first-order conditions implicitly define the optimal consumption pattern. Given this, the firm can construct the optimal marginal price schedule through the consumer's self-selection condition, setting $R_q(q) = W_q(q, \theta^*(q))$, where $\theta^*(q) = \max\{\theta \mid q^*(\theta) \geq q\}$. By definition, $R(q(M)) = W(q(M), M)$. This plus knowledge of R_q for $q > q(M)$—corresponding to the consumption decisions for all $\theta < M$—enable the firm to specify $R(q)$ over all $q > q(M)$. For $0 \leq q < q(M)$ the firm could set $R(q)$ equal to $W(q, M^E)$. Note that when $q(M) > 0$, the optimal price schedule is not unique: As long as R isn't less than $W(q, M)$ on the interval $[0, q(M)]$, its value within that interval is irrelevant to both the firm and consumers.

From Theorem 4 we know that $M^E \geq M^L$. If in fact the inequality is strict, then *the experimenting firm produces information about the market by lowering the price schedule*. This is because $dW(q^*(\theta), \theta)/d\theta < 0$ when $q^*(\theta)$ is positive. So $M^E > M^L$ implies that $W(q^E(M^E), M^E)$ is less than $W(q^L(M^L), M^L)$, and hence that the total charge R is lower.

We have developed useful rules of thumb for dynamic nonlinear pricing and gained insight into the relationship between pricing and learning. We have shown conditions under which the experimenter chooses the marginal price schedule as if he were myopic and that taking the potential benefits of learning into consideration leads one to lower the price schedule. Because he has a more complete sample observation space, the experimenter will be able to resolve his uncertainty about the unknown parameter ξ more quickly than if he were myopic. However, but for one exception, the stochastic nature of the market generally precludes prediction of the time profile of the price schedules. We turn now to that exception.

4. Special Case: Static Unknown Demand

We have assumed that demand depends on a variate θ generated each period by a density g, and that the firm doesn't know some fixed (possibly vector-valued) parameter ξ of g. This implies the firm's predictive density

$$h(\theta \mid I) = \int_{\Xi} g(\theta \mid \xi) dF(\xi \mid I).$$

We could view θ as either an index for heterogeneous consumers whom the firm contacts one at a time, or as a parameter in the demand curve of homogeneous consumers that varies period by period.

Now suppose instead that the firm faces a single static demand curve with some unknown parameter. The interpretation here is that consumers are homogeneous and that their willingness to pay is not subject to uncertain shocks each period. That this is a special case of our basic model can be seen by setting $g(\theta \mid \xi)$ to the Kronecker impulse function

$$\delta(\theta, \xi) = \begin{cases} \infty & \theta = \xi, \\ 0 & \text{otherwise}, \end{cases}$$

which implies that $h(\theta \mid I) = f(\xi \mid I)$ at $\xi = \theta$. Use $\hat{\theta}$ to denote the true—but initially unknown—value of θ, and $h(\theta \mid I(t))$ to denote the firm's predictive density at the start of period t.

When there is no consumption in period t, both the myopic and experimenting firm choose $M(t) \geq M(t-1)$. Assume each period's optimal consumption pattern strictly decreases wherever positive. Hence the firm either observes an outcome $\hat{q}(t) = q^*(\hat{\theta}; t) > 0$ from which it infers the true value $\hat{\theta}$, or it observes no consumption and can infer only that $\theta \in (M(t), \bar{\theta}]$. We can view $M(\cdot)$ as the firm's information state variable, and write the firm's prior density in period t as $h(\theta \mid M(t-1))$. Until it learns $\hat{\theta}$, the firm updates its prior according to

$$h(\theta \mid M(t)) = \frac{h(\theta)}{1 - H(M(t-1))}, \qquad \theta \in (M(t-1), \bar{\theta}),$$

where $h(\theta)$ denotes $h(\theta \mid I(0))$.

Now consider how the firm's choice of $M(t-1)$ affects its choice of $q^*(\cdot; t)$ the following period. Theorem 5 (in the Appendix) shows that a higher value of $M(t-1)$—which corresponds to lower potential demand in period t—leads the firm to assign a *higher* consumption level to each $\theta > M(t-1)$, and to choose a higher inclusion index $M(t)$ for the current period. To support the higher consumption levels and inclusion index, the firm must lower the marginal price schedule and the total charge at $q^*(M)$. This is the case for both the myopic and experimenting firms. Along with our observation that the inclusion index is nondecreasing, this also implies that if W and C are stationary and the conditions of Theorem 5 hold, then both the marginal price and the total price schedules decrease over time until the firm learns the true value of $\hat{\theta}$.

So how *does* an experimenting firm's pricing strategy differ from the myopic firm's? We know that for a given state of information $M(t-1)$, the experimenting firm should choose the same consumption pattern—and hence marginal price schedule—as the myopic firm, since by assumption q^* is strictly decreasing (Proposition 2 in the Appendix). Thus any difference in the pricing strategies lies in the choice of the inclusion index M. It is instructive to consider two cases: The first retains the assumption that the consumer makes repeat purchases, whereas the second supposes the firm is interested in making only one sale.

Case 1. Once the firm learns $\hat{\theta}$ it should set a price schedule that extracts all social surplus $SS(\hat{\theta}; t) = W(q, \hat{\theta}; t) - C(q; t)$. The optimal quantity to sell is the $q^c(\hat{\theta}; t)$ that solves $C_q(q^c; t) = W_q(q^c, \hat{\theta}; t)$. The firm can support this through a two-part tariff with slope $r = R_q(q; t) = W_q(q^c, \hat{\theta}; t)$ and fixed fee $R(0; t) = W(q^c, \hat{\theta}; t) - rq^c(\hat{\theta}; t)$. Substituting and rearranging terms, the experimenting firm's problem is

$$\max_{q_0^*, M(0)} \left\{ \int_0^{M(0)} \left[\pi_0(q_0^*, \theta_0) + \sum_{j=1}^{T} \rho^j SS(\theta_0; j) h(\theta_0) \right] d\theta_0 \right.$$

$$+ \cdots + \max_{q_t^*, M(t)} \left[\int_{M(t-1)}^{M(t)} \left[\rho^t \pi_t(q_t^*, \theta_t) + \sum_{j=t+1}^{T} \rho^j SS(\theta_t; j) h(\theta_t) \right] d\theta_t \right.$$

$$\left. \left. + \cdots + \max_{q_T^*, M(T)} \left\{ \int_{M(T-1)}^{M(T)} \rho^T \pi_T(q_T^*, \theta_T) d\theta_T \right\} \cdots \right] \cdots \right\}$$

where $\pi_t(q, \theta)$ is

$$[W(q^*, \theta; t) - C(q^*; t)]h(\theta) + W_\theta(q^*, \theta; t)[H(\theta) - H(M(t-1))].$$

As in our earlier analysis, the firm could, for each period t, determine the optimal q^E parametric on M, and then solve for the optimal M. The optimal inclusion index satisfies

$$\begin{aligned} \pi(q^E, M; t) &= -\rho \int_M^{M(t+1)} \frac{d\pi_{t+1}(q_{t+1}^*, \theta)}{dM} d\theta \\ &= \rho \int_M^{M(t+1)} W_\theta(q_{t+1}^*, \theta; t+1) h(M) d\theta \\ &= \rho h(M) \int_M^{M(t+1)} dCS(\theta; t+1) \\ &= -\rho h(M) CS(M; t+1). \end{aligned}$$

The third equality comes from (6). Thus the experimenting firm's first-order condition for the optimal inclusion index M is

$$W(q^E, M; t) - C(q^E; t) + W_\theta(q^E, M; t) \frac{H(M|M(t-1))}{h(M|M(t-1))} = -\rho CS(M; t+1).$$

That is, when choosing the inclusion index for the current period, the experimenting firm should take into account the *additional* profit $CS(M; t+1)$ it would earn if it knew that $\hat{\theta} = M$, instead of knowing only that $\hat{\theta} > M$. This difference is shown, for some hypothetical demand curves and price schedules, as region A in Figure 3.

Case 2. Lazear (1986) shows that an experimenting firm that prices *linearly* should price *higher* than when myopic, in contrast to our result that it should price *lower* when using *nonlinear* pricing. We now show that this difference is due not to the form of price schedule, but rather to Lazear's assumption that the firm makes but one sale of its product.

Lazear's model can be written in our notation as follows. The firm has one unit of its product, and faces a single consumer whose willingness to pay is $W(\hat{\theta}) = \bar{\theta} - \hat{\theta}$. The index $\hat{\theta} \in [0, \bar{\theta}]$ is fixed but unknown to the firm. The consumer buys the product at price P when $W(\hat{\theta}) \geq P$, whence the inclusion index M is such that $W(M) = \bar{\theta} - M = P$. In the initial period, the probability of a sale is $H(M_0)$, and the expected profit, in terms of the inclusion index, is $(W(M_0) - C)H(M_0)$, where C is the firm's cost of providing its product. The discounted sum of expected profits is

$$\sum_{j=0}^{T} \rho^j (W(M_j) - C)(H(M_j) - H(M_{j-1})),$$

where M_{-1} is defined to be 0. The myopic firm solves for M in period t that satisfies

$$W(M) - C + W_\theta(M) \frac{H(M) - H(M_{t-1})}{h(M)} = 0,$$

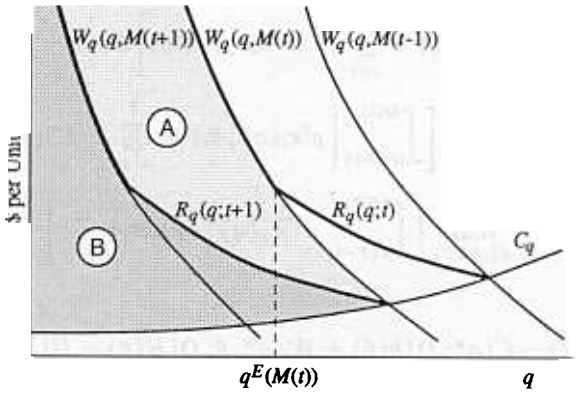

FIGURE 3. Marginal Sum of Discounted Profits When Purchases Are (A) Repeat and (B) Nonrepeat.

or, noting that $h(\theta \mid M) = h(\theta)/(1 - H(M))$,

$$W(M) - C + W_\theta(M)\frac{H(M|M_{t-1})}{h(M|M_{t-1})} = 0,$$

The experimenting firm, on the other hand, seeks M to solve

$$W(M) - C + W_\theta(M)\frac{H(M|M_{t-1})}{h(M|M_{t-1})} = \rho(W(M_{t+1}) - C).$$

The term on the right is the profit from the consumer with index M were he to buy the product the following period. Since this is positive, the experimenter chooses a lower index, and hence higher price, than the myopic firm: Recognizing it will have another opportunity to make a sale if it misses its sale in the current period leads the experimenter to price higher than he would otherwise.

To extend this to the nonlinear case, assume the firm has a product it will sell in any quantity to a single consumer; it will then close up shop. The sum of expected discounted profits is as in Case 1, except that the discounted social surplus terms (SS) vanish since there are no repeat sales. As a result, the inclusion index satisfies

$$\begin{aligned}\pi(q^E(M,t), M; t) &= \rho\pi_{t+1}(q^*_{t+1}, M) - \rho\int_M^{M(t+1)} \frac{d\pi_{t+1}((q^*_{t+1}, \theta)}{dM}\, d\theta \\ &= \rho h(M)[SS(M; t+1) - CS(M; t+1)],\end{aligned}$$

or

$$\begin{aligned}&W(q^E, M; t) - C(q^E; t) + W_\theta(q^E, M; t)\frac{H(M|M(t-1))}{h(M|M(t-1))} \\ &\quad = \rho[SS(M; t+1) - CS(M; t+1)].\end{aligned}$$

As in Lazear's model, the term on the right is the profit from consumer type M, shown as region B in Figure 3, were he to buy in the following period. Since it is positive, the optimal inclusion index should be *lower* than in the myopic case, and the entire schedule should be *higher*. We conclude then that when there are repeat purchases the forward-looking firm should tend towards penetration pricing; otherwise its strategy should tend towards skimming.

5. Discussion

In his review of the research on nonlinear pricing, Dolan (1987) writes that

> Additional research is required . . . to study the benefits of schedule complexity in various demand situations. Uncertainty in demand is a key factor to be studied. Research on what type of schedule yields maximal knowledge of buyers would be useful to impacting practice.

This paper is the first to address the above agenda rigorously in the context of nonlinear pricing. Certainly a firm can attain "maximal knowledge of buyers" most quickly by inducing all potential buyers to purchase, and can do so more profitably with a nonlinear schedule than with a linear one. Yet information production comes at a cost, since a pricing policy that includes all potential buyers will usually not be as profitable as one which excludes at least some segments in order to get higher prices from others.

We have presented techniques with which a firm can determine the implicit costs of resolving uncertainty about demand through its pricing policy and have shown that schedule complexity per se is not the key to information production; how quickly uncertainty is resolved depends solely on the set of buyers excluded from the market or bunched at positive purchase quantities. This insight, which is one of the key results of our analysis, can be interpreted as a "separation principle" (using the terminology of the stochastic control literature): The "estimation problem" of learning the parameters characterizing customer heterogeneity can be separated from the "control problem" of determining appropriate volume discounts to exploit that heterogeneity.

Like many theoretical analyses, this paper provides more insights than numbers. Some of these insights might seem obvious in hindsight, yet they serve a valuable purpose in enhancing managers' intuition and articulating the underlying assumptions that support that intuition. A good example of such insights is how an assumption about repeat purchases affects the price schedule over time. We have extended the common wisdom regarding "skimming" versus "penetration" pricing policies. Our analysis demonstrates that the justification for these respective policies in the appropriate market settings (saturated versus repeat purchase) follows not only from word-of-mouth considerations but also from demand uncertainty and learning considerations.

Managers concerned with test marketing aimed at learning about demand heterogeneity should be relieved to learn that they need not concern themselves with learning effects when designing volume discount schedules. Experimentation can be carried out simply by lowering the entire price schedule through, for instance, a temporary rebate that doesn't depend on purchase quantity. Of course a manager may wish to refine the volume discounts after observing the distribution of purchase quantities. However, such ex post adjustments constitute passive learning, which is considerably simpler to implement than active learning where the prospect of learning has to be taken into consideration in the original schedule design.

This paper's techniques are appropriate for markets where buyers' tastes vary according to a parameter whose distribution is unknown. They are also applicable to markets with homogeneous buyers whose demand curve is subject to random shocks each period. Most of the analysis holds in even more general settings. For example, suppose buyers can be grouped into several segments that are internally homogeneous, and that each has a stochastic demand curve about which the firm has parametric uncertainty. While in practice this problem is difficult to solve, several key insights still hold: Solve for the marginal price path as if the firm were myopic, reduce bunching at positive quantities, and lower the total price schedule to produce information.

There is an obvious analogy between nonlinear pricing, and pricing a line of products that vary in a single dimension such as "quality" (Mussa and Rosen 1978, Moorthy 1984, and Smith 1986). From the perspective of product line pricing, our results indicate

that the experimenting firm should use a more extensive product line than the myopic firm. For example, if q is a scalar measure of a photocopier's "quality", then the experimenting firm would include in its line low-quality copiers that would not be in the myopic firm's line. Furthermore, since a product line necessarily bunches buyers at discrete quality levels, the experimenting firm would tend to choose a richer product line than the myopic firm, particularly when uncertainty about demand is greatest (say, just when the line is introduced).

Our results complement the findings of others in the dynamic pricing literature, who have shown that a forward-looking firm chooses a lower price than its myopic counterpart in order to stimulate word-of-mouth effects (Kalish 1983), lower its production costs (Robinson and Lakhani 1975), or establish a sufficient subscriber set when there are positive demand externalities (Dhebar and Oren 1986). This indicates, for example, that a firm facing both positive demand externalities and demand uncertainty will price even lower than it would were it to ignore the benefits of information production, thereby establishing its critical subscription set even sooner.

There are, of course, reasons other than uncertainty resolution for a firm to change price over time. Unfortunately, one can't tell from observations of price schedules and sales if a firm is doing what we suggest (priors are private) except perhaps in one respect. We stated that because a firm foregoes information when it bunches buyers, the experimenter wouldn't do so. Our prediction then is that we should see no bunching in situations where the firm can sell its product in any quantity.

Needless to say, inferring the distribution of customer type parameters from the purchase data is not a trivial matter and may still require some fairly strong assumptions regarding the demand functions and the form of the consumer type distribution. A good example of such inference is presented by Mitchell (1978) for pricing local telephone service. He assumed the demand function to be linear in price and proportional to the type parameter and the type parameter to have a Log-Normal distribution.

Finally, competitive issues may well loom large as a firm considers its choice of schedules. Our methodology is most valuable when uncertainty about demand is greatest, which tends to be early on in the game before competitive entry. Furthermore, our suggestion to lower price to produce information is consistent with a strategy that seeks to delay competitive entry by building up share early on. Unfortunately, modelling dynamic nonlinear pricing with learning in a competitive market is extremely difficult and certainly should be tailored to a specific industry. For the reader up to the modelling challenge, we suggest a review of Oren, Smith and Wilson (1983), who demonstrate the technical difficulties inherent in even the simplest setting—a static, symmetric Nash-Cournot framework.[2]

Acknowledgments. Many thanks to Tom Cooley, John Long, Glenn MacDonald, Sridhar Moorthy and Sherwin Rosen for their comments, and to Tony Arena of A. R. Arena Products, Inc. for his insights. We also thank John Hauser, Brian Ratchford and the anonymous reviewers for their thoughtful comments on earlier drafts. This work was based on Braden's unpublished PhD thesis, "Nonlinear Pricing to Reduce Demand Uncertainty", Stanford University, 1989, and was supported in part by a grant from the Center for Economic Policy Research, Stanford University.

[2] This paper was received February 11, 1992, and has been with the authors 3 months for 2 revisions. Processed by Brian T. Ratchford, Area Editor.

Appendix

Herein are technical details of our paper. Throughout we assume that the willingness-to-pay function $W: \mathbb{R}_+ \times \Theta \times \{0, 1, \ldots, T\} \to \mathbb{R}_+$ satisfies each of the following for all θ and t:

(A) $W_{qq\theta}(q, \theta; t)$ exists and is continuous on $\mathbb{R}_+ \times \Theta$.

(B) $W(0, \theta; t) = 0$.

(C) There is a finite quantity $q^c(\theta; t) \geq 0$ such that $W_q(q, \theta; t) > C_q(q; t)$ if and only if $q < q^c(\theta; t)$.

(D) $W_{qq}(q, \theta; t) < 0$ for all $q \in [0, q^c(\theta; t)]$.

(E) $W_{q\theta}(q, \theta; t) \leq 0$ for all $q \in [0, q^c(\theta; t)]$, with strict inequality for positive q; furthermore, $W_{q\theta}$ is bounded from below.

Denote by q^L and M^L the optimal consumption pattern and inclusion index for the myopic firm, and let q^E and M^E denote the experimenting firm's optimal choices.

PROPOSITION 1. *q^E is at least as valuable as q^L.*

PROOF. By definition of q^L and q^E, $\int_0^{M^L} \pi(q^L(\theta), \theta)d\theta \geq \int_0^{M^E} \pi(q^E(\theta), \theta)d\theta$, and

$$\int_0^{M^E} \pi(q^E(\theta), \theta)d\theta + \int_\Theta J[\Phi(f(\xi|I), q^E, \hat{q})]h(\theta|I)d\theta$$

$$\geq \int_0^{M^L} \pi(q^L(\theta), \theta)d\theta + \int_\Theta J[\Phi(f(\xi|I), q^L, \hat{q})]h(\theta|I)d\theta.$$

Subtracting the first equation from the second completes the proof. □

PROPOSITION 2. *If q^L is strictly decreasing wherever positive then it maximizes the experimenting firm's profits, subject to $M^E = M^L$.*

PROOF. Immediate from the fact that such a q^L is at least as informative as any q^*, given M. □

We'll now give conditions sufficient for q^L to be strictly decreasing wherever positive. Recall that if $f(x)$ is a real-valued function that is twice-differentiable, then it is strongly quasiconcave iff $f_x = 0 \Rightarrow f_{xx} < 0$ (Diewert, Avriel and Zang 1981).

PROPOSITION 3. *Suppose $q^*(\theta) \in \arg\max_{q \geq 0} \pi(q, \theta)$. If π is strongly quasiconcave in q and $\pi_{q\theta}(q^*(\theta), \theta) < 0$, then q^* is smooth and strictly decreasing wherever it is positive. Furthermore, q^* is optimal for problems (MP) and (EP), given M.*

PROOF. Since $q^*(\theta)$ maximizes $\pi(q, \theta)$, it is either 0 or it satisfies the first-order condition $\pi_q = 0$. Suppose it is positive. Since π is strongly quasiconcave and twice differentiable, $\pi_q = 0 \Rightarrow \pi_{qq} < 0$. So by the Implicit Function Theorem q^* is differentiable, with derivative $dq^*/d\theta = -\pi_{q\theta}/\pi_{qq} < 0$. Since q^* is either zero or strictly decreasing, $q^* \in \mathcal{Q}$; and because it is determined pointwise, Jensen's inequality implies that $q^L = q^*$, given M. Its strict monotonicity implies, by Proposition 2, that $q^E = q^*$, given M. □

We will need the following lemma for the subsequent theorem. It states that a consumption pattern that is strictly decreasing wherever positive is more informative than *any* consumption pattern with a lower inclusion index. Its proof is immediate from the definition of "informative".

LEMMA. If q_1^* is strictly decreasing wherever positive, and $\sup\{\theta | q_1^*(\theta) > 0\} > \sup\{\theta | q_2^*(\theta) > 0\}$, then q_1^* is more informative than q_2^*.

THEOREM 4. *Assume solutions exist to problems (MP) and (EP), and that q^L is strictly decreasing wherever positive. If M^L is unique or M^E is the highest inclusion index that maximizes problem (EP), then $M^E \geq M^L$.*

PROOF. If $M^E < M^L$, then by the lemma q^L is more informative than q^E and hence at least as valuable. But by Proposition 1 this can hold only if q^E and q^L are equally valuable, implying $\int_0^{M^E} \pi(q^E(\theta), \theta)d\theta \geq \int_0^{M^L} \pi(q^L(\theta), \theta)d\theta$. Since q^L and M^L are optimal for problem (MP), it must be that the above inequality is in fact an equality. This means that M^L is not unique and that M^E isn't the largest point that maximizes problem (EP), a contradiction. □

The following pertains to the case of a static unknown demand curve, considered in Section 4.

THEOREM 5. *If, for $\theta \in [M^u, \bar{\theta}]$, $q^*(\theta)$ satisfies the second-order sufficiency conditions for maximizing π then it is an increasing function of M^u; furthermore, if M is such that $q^*(M) > 0$ and $\pi_\theta(q^*(M), M) < 0$, then M is also an increasing function of M^u.*

PROOF. The second-order conditions are that $\pi_q(q, \theta) = 0$ and $\pi_{qq}(q, \theta) < 0$ at $q = q^*(\theta)$, where

$$\pi(q^*, \theta) = [W(q^*, \theta) - C(q^*)]h(\theta|M^u) + W_\theta(q^*, \theta)H(\theta|M^u)$$

$$= [W(q^*, \theta) - C(q^*)]\frac{h(\theta|0)}{1 - H(M^u|0)} + W_\theta(q^*, \theta)\frac{H(\theta|0) - H(M^u|0)}{1 - H(M^u|0)}$$

Differentiate $\pi_q(q^*, \theta)$ with respect to M^u:

$$\frac{dq^*}{dM^u} = \frac{W_{q\theta}(q^*, \theta)h(M^u|0)}{\pi_{qq}(q^*, \theta)(1 - H(M^u|0))}.$$

The numerator is negative by Assumption E.

To show that M is an increasing function of M^u, we use the same approach with the myopic firm's transversality condition, $\pi(q^*(M), M) = 0$. Differentiating with respect to M^u implies that

$$\frac{dM}{dM^u} = \frac{W_\theta(q^*, M)h(M^u|0)}{\pi_\theta(q^*, M)(1 \quad H(M^u|0))}$$

Assumptions B and E together imply the numerator is negative, so $dM/dM^u > 0$. This holds for the experimenting firm as well, since it chooses an inclusion index at least as high as the myopic firm's. □

References

Aghion, Philippe, Patrick Bolton, Christopher Harris and Bruno Jullien (1991), "Optimal Learning by Experimentation," *Review of Economic Studies*, 58 (July), 621–654.

Dhebar, Anirudh S. and Shmuel S. Oren (1985), "Optimal Dynamic Pricing for Expanding Networks," *Marketing Science*, 4 (Fall), 336–351.

——— and ——— (1986), "Dynamic Nonlinear Pricing in Networks with Interdependent Demand," *Operations Research*, 34 (May–June), 384–394.

Diewert, W. Erwin, Mordecai Avriel, and Israel Zang (1981), "Nine Kinds of Quasiconcavity and Concavity," *Journal of Economic Theory*, 25 (December), 397–420.

Dolan, Robert J. (1987), "Quantity Discounts: Managerial Issues and Research Opportunities," *Marketing Science*, 6 (Winter), 1–22.

——— and Abel P. Jeuland (1981), "Experience Curves and Dynamic Demand Models: Implications for Optimal Pricing Strategies," *Journal of Marketing*, 45 (Winter), 52–62.

Goldman, M. Barry, Hayne E. Leland, and David S. Sibley (1984), "Optimal Nonuniform Prices," *Review of Economic Studies*, 51 (April), 305–319.

Grossman, Sanford J., Richard E. Kihlstrom, and Leonard J. Mirman (1977), "A Bayesian Approach to the Production of Information and Learning By Doing," *Review of Economic Studies*, 44 (October), 533–547.

Harris, Milton (1987), *Dynamic Economic Analysis*, New York: Oxford University Press.

Kalish, Shlomo (1983), "Monopolist Pricing with Dynamic Demand and Production Cost," *Marketing Science*, 2 (Spring), 135–159.

Lazear, Edward P. (1986), "Retail Pricing and Clearance Sales," *American Economic Review*, 76 (March), 14–32.

Leland, Hayne E. and Robert A. Meyer (1976), "Monopoly Pricing Structures with Imperfect Discrimination," *Bell Journal of Economics*, 7 (Autumn), 449–462.

Little, John D. C. (1966), "A Model of Adaptive Control of Promotional Spending," *Operations Research*, 14 (November–December), 1075.

Littlechild, Stephen (1975), "Two-Part Tariffs and Consumption Externalities," *Bell Journal of Economics*, 6 (Autumn), 661–670.

McLennan, Andrew (1984), "Price Dispersion and Incomplete Learning in the Long Run," *Journal of Economic Dynamics and Control*, 7 (September), 331–347.

Mitchell, Bridger M. (1978), "Optimal Pricing of Local Telephone Service," *American Economic Review*, 68 (June), 517–537.

Moorthy, K. Sridhar (1984), "Market Segmentation, Self-Selection, and Product Line Design," *Marketing Science*, 3 (Fall), 288–307.

Mussa, Michael and Sherwin Rosen (1978), "Monopoly and Product Quality," *Journal of Economic Theory*, 18 (August), 301–317.

Oi, Walter Y. (1971), "A Disneyland Dilemma: Two-Part Tariffs for a Mickey Mouse Monopoly," *Quarterly Journal of Economics*, 85 (February), 77–96.

Oren, Shmuel S., Stephen A. Smith and Robert B. Wilson (1982), "Nonlinear Pricing in Markets with Interdependent Demand," *Marketing Science*, 1 (Summer), 287–313.

———, ——— and ——— (1983), "Competitive Nonlinear Tariffs," *Journal of Economic Theory*, 29 (February), 49–71.

Pekelman, Dov and Edison Tse (1980), "Experimentation and Budgeting in Advertising: An Adaptive Control Approach," *Operations Research*, 28 (March–April), 321–347.

Robinson, Bruce and Chet Lakhani (1975), "Dynamic Price Models for New-Product Planning," *Management Science*, 21 (June), 1113–1122.

Rothschild, Michael (1974), "A Two-Armed Bandit Theory of Market Pricing," *Journal of Economic Theory*, 9 (October), 185–202.

Smith, Stephen A. (1986), "New Product Pricing in Quality Sensitive Markets," *Marketing Science*, 5 (Winter), 70–87.

Spence, Michael (1977), "Nonlinear Prices and Welfare," *Journal of Public Economics*, 8 (August), 1–18.