

THE IGNORANT MONOPOLIST: OPTIMAL LEARNING WITH ENDOGENOUS INFORMATION*

BY DANIEL TREFLER¹

Economics lacks a good theory of the pricing and output decisions of a monopolist which does not know its demand—we inevitably assume that the monopolist knows much more about demand conditions than is reasonable. I present a model in which demand information is generated endogenously. When information is endogenous the monopolist has an incentive to experiment with price and quantity. I derive the direction of experimentation, characterize an important value function arising from dynamic programming problems with learning, and relate the results to Blackwell's comparison of experiments.

1. INTRODUCTION

Economics lacks a good theory of the pricing and output decisions of a firm that does not know its demand. We always assume the firm has complete demand information or has exact knowledge of the stochastic process generating demand. Yet no explanation is given for how the firm comes by this information. In this paper demand information is generated endogenously. I examine an important, oft-cited, example involving a monopolist who wishes to learn about the demand process she faces. She may do so by experimenting with price and observing the resulting quantity. While such price experimentation reduces expected profits in the current period, the loss can possibly be recouped in subsequent periods through use of the improved information.

This pricing example is one of many settings in which actions taken today determine not only the reward today, but also the information available tomorrow. To capture this I consider an infinite-horizon control problem under uncertainty in which Bayesian learning provides the link between periods. In each period the agent takes an action (e.g., price), observes the outcome (e.g., demand), and updates her beliefs about the unknown parameters of the stochastic process generating outcomes. I address two questions. First, when does one action result in larger expected value of information than another? Second, what implications do endogenous information and learning have for the sequence of optimal actions i.e., what is the direction of experimentation?

One model of a price-setting monopolist that has received considerable attention is the case of a monopolist facing a linear demand with an unknown slope and an

* Manuscript received May 1991; first revision received March 1992; final revision received September 1992.

¹ I am indebted to John Riley for many helpful comments. The anonymous referees suggested several important improvements, especially to the proofs of Theorems 1 and 2. I am also grateful to Larry Epstein, Angelo Melino, and participants of the University of Toronto and UCLA Theory Workshops.

additive, normally distributed, error term.² In this paper I use a related model, though without the normality assumption. In particular, I entertain demand processes which are derived directly from readily observable features of the monopolist's market. Hence, the only distributional information I allow the monopolist is that which she can reasonably be expected to know. I show that in two types of markets, what I call the "reservation price" and "posted price" markets, strong conclusions emerge about the expected value of information and the direction of experimentation.

The simplest market in which the monopolist operates is the "reservation price" market. Imagine that the monopolist operates a shop in which one customer arrives in each period and decides whether or not to make a purchase. There is only one indivisible good for sale and each customer buys at most one unit. Due to unobservable customer heterogeneity the monopolist views customer purchase decisions as stochastic. At a given price some customers buy, others do not. Thus, customer purchases follow a binomial process with probability of purchase inversely related to price. It is this inverse relationship which the monopolist seeks to learn through price experimentation. More generally, the monopolist may offer several product varieties and the customer chooses only one of them so that customer purchases follow a multinomial process. Notice how the multinomial restriction on demand falls naturally out of the market in which the monopolist operates.³

An undesirable feature of the reservation price market is that customers arrive one per period regardless of the price charged. This is not realistic in a market with repeat customers or advertised prices. Where price information is readily available a more appropriate environment is what I call the "posted price" market. Here the monopolist posts a price. Customers observe this price and decide whether or not to enter the shop and make a purchase. If the arrival rate of customers depends only on the posted price then the demand process is Poisson. The monopolist seeks to learn the inverse relationship between the posted price and customer flow. More generally, the monopolist may post several prices, one for each product or one for each segmented market. Again, the Poisson restriction on demand falls naturally out of the market in which the monopolist operates.

The main conclusion of this paper is *a characterization of the expected value of information and the direction of experimentation in the reservation price and posted price markets*. The result is related to work by Prescott (1972) and Grossman, Kihlstrom, and Mirman (1977). Prescott was able to establish the direction of experimentation for the case of an additive, normally distributed, error term and a normal prior on the unknown slope. Grossman, Kihlstrom, and Mirman showed that the Prescott result holds when the normality restriction on the prior is relaxed. My work builds on these papers by examining the role of the additive normal error process. In this light, my conclusions are surprising since *normality is*

² For examples, see Prescott (1972), Grossman, Kihlstrom and Mirman (1977), and Easley and Kiefer (1987).

³ The binomial reservation price model was proposed by Rothschild (1974). Lazear (1986) and Stokey and Lucas (1989) considered simpler binomial models with identical customers.

critical in both these papers and even simple examples indicate that without strong distributional restrictions no conclusions are possible.

A second contribution of this paper is a generalization of the Prescott (1972) and Grossman, Kihlstrom, and Mirman (1977) results from a finite-horizon problem with a *single* control and a *single* unknown parameter to an infinite-horizon problem with *multiple* controls and unknown parameters. I do this for the reservation price and posted models as well as for a model with an additive normal error process. For example, unlike earlier results, my results apply to the case of a multiproduct monopolist who can price-discriminate between multiple markets with differing, unknown demands.

This paper is related to Fusselman and Mirman (1989) and Mirman, Samuelson, and Urbano (1989) who provided nonparametric results on the direction of experimentation for a two-period problem in which the unknown parameter takes on only two values. It is also related to Creane (1991) who analyzed a version of the Mirman, Samuelson, and Urbano (1989) model with heteroscedastic noise.⁴

The paper is organized as follows. Section 2 sets up the dynamic programming problem facing the monopolist. Since the state space is the space of prior beliefs, a space of probability measures, the setup is unavoidably technical and some readers may prefer only to browse through it. In Section 3 I define the expected value of information in accordance with the literature on Blackwell's (1953) "comparison of experiments." This section describes the most innovative result of the paper, that for the reservation price and posted price markets the expected value of information is monotonic in actions. Section 4 uses the monotonicity results to determine the direction of experimentation. It is this key result on the direction of experimentation that is applied in Section 5 to the problem of the ignorant monopolist. Proofs appear in the Appendix.

2. THE DECISION PROBLEM

In each period t the decision-maker chooses an action $x_t \in X$ and observes an outcome $y_t \in Y$. The probability of outcome y_t depends both on x_t and on an unknown parameter $\theta \in \Theta$. Assume that x_t and θ are conformable for matrix multiplication and let $\nu(\cdot|x_t\theta)$ be a probability measure over outcomes. The decision-maker begins with a prior $\mu_0 \in P(\Theta)$ about the unknown parameter θ . In period t her beliefs about θ are given by μ_t . At the end of period t , having chosen x_t and observed y_t , she uses Bayes rule to update her beliefs about θ to μ_{t+1} . After choosing x_t and observing y_t , the decision-maker receives a reward $r(x_t, y_t, \theta)$. Her objective is to maximize expected discounted rewards over an infinite horizon, $E[\sum_{t=0}^{\infty} \delta^t r(x_t, y_t, \theta)]$, where the expectation operator is described below and $0 < \delta < 1$ is the discount factor. The salient features of the setup are that x_t and θ enter $\nu(\cdot|x_t\theta)$ multiplicatively and that the only intertemporal link is Bayesian learning.

Before proceeding with a rigorous exposition, I offer an informal discussion of

⁴ This paper is also related to a growing body of literature on whether the monopolist's beliefs (posteriors) and optimal actions converge over time. For examples, see Rothschild (1974) and Easley and Kiefer (1988).

the Markovian structure of the dynamic programming problem. The state space is the space of prior beliefs, $P(\Theta)$. The transition from the state in period t (the prior μ_t) to the state in period $t + 1$ (the posterior μ_{t+1}) follows Bayes rule. In choosing x_t , the decision-maker must take into account the effect of x_t on future beliefs, μ_{t+1} . At the time x_t is chosen, however, y_t and hence μ_{t+1} are unknown so that the decision maker must calculate the distribution of future beliefs μ_{t+1} across the values that y_t may take. This forms the transition probability from μ_t to μ_{t+1} which is described by the measure $q(d\mu_{t+1}|\mu_t, x_t)$. To illustrate the use of q , consider the value function for this dynamic programming problem, $V(\mu_t)$. The expectation of $V(\mu_{t+1})$ given the information available at time t , $\mathbf{E}[V(\mu_{t+1})|\mu_t, x_t]$, is just $\int_{P(\Theta)} V(\mu_{t+1})q(d\mu_{t+1}|\mu_t, x_t)$ or more simply, $\int_{P(\Theta)} V(\bar{\mu}) dq(\mu_t, x_t)$. The remainder of this section is devoted to establishing the existence of a value function $V(\mu_0)$ and an optimal sequence of Markov actions $a^\infty(\mu_0)$. This is not the focus of the paper so that the reader uninterested in existence issues will prefer to browse through the rest of the section.

The following assumptions are made throughout.

(i) The parameter space Θ is a compact subset of \mathbb{R}^k . $\theta \in \Theta$ is taken to be a k -dimensional column vector. Let B_Θ be the Borel σ -algebra for Θ and let $P(\Theta)$ be the set of probability measures on (Θ, B_Θ) . μ_t is an element of $P(\Theta)$. Endow $P(\Theta)$ with the weak topology.⁵ Then $P(\Theta)$ is a compact, separable, metric space (Parthasarathy 1967, Section II.7).

(ii) The outcome space Y is a complete subset of \mathbb{R}^n . (It need not be compact.) $y \in Y$ is taken to be an n -dimensional column vector. Let B_Y be the Borel σ -algebra for Y and for each $(x_t, \theta) \in X \times \Theta$, let $\nu(\cdot|x_t, \theta)$ be a probability measure on (Y, B_Y) . Assume that for each $(x_t, \theta) \in X \times \Theta$, $\int_Y |y| \nu(dy_t|x_t, \theta) < \infty$ and that for each sequence $\{(x^n, \theta^n)\}$ of elements of $X \times \Theta$ converging to (x^0, θ^0) , $\nu(\cdot|x^n, \theta^n)$ converges weakly to $\nu(\cdot|x^0, \theta^0)$.

(iii) The action space X is a compact subset of \mathbb{R}^{nk} . $x \in X$ is taken to be an $n \times k$ -dimensional matrix.

(iv) The reward function $r: X \times Y \times \Theta \rightarrow \mathbb{R}$ is an element of C , the set of uniformly bounded, continuous, real-valued, functions on $X \times Y \times \Theta$.

The decision-maker's objective is to maximize the expected discounted reward, $\mathbf{E}[\sum_{t=0}^{\infty} \delta^t r(x_t, y_t, \theta)]$. To make clear the dependence of this expectation on μ_t define

$$u(x_t, \mu_t) = \int_{\Theta} \int_Y r(x_t, y_t, \theta) \nu(dy_t|x_t, \theta) \mu_t(d\theta).$$

Then the problem is to maximize $\mathbf{E}[\sum_{t=0}^{\infty} \delta^t u(x_t, \mu_t)]$, where μ_t evolves according to Bayes rule. Under the above assumptions $u: X \times P(\Theta) \rightarrow \mathbb{R}$ is bounded. Assume it is continuous. (A sufficient condition for the continuity of u is the uniform continuity of r .)

⁵ A sequence of measures $\{\mu^n\}$ in $P(\Theta)$ converges in the weak topology (converges weakly) to a measure μ^0 if and only if $\int_{\Theta} h(\theta) \mu^n(d\theta) \rightarrow \int_{\Theta} h(\theta) \mu^0(d\theta)$ for every bounded, continuous, real-valued function h on Θ .

Under the above assumptions, $q(d\mu_{t+1}|\mu_t, x_t)$ is well-defined and continuous. For further discussion of q under a somewhat different set of assumptions, see Easley and Kiefer (1988).

Using q and u , the problem may be reduced to a standard dynamic programming problem with action space X and state space $P(\Theta)$. For $\mu \in P(\Theta)$ define

$$(1) \quad V(\mu) = \max_{x \in X} \left\{ u(x, \mu) + \delta \int_{P(\Theta)} V(\bar{\mu}) dq(\mu, x) \right\}.$$

If there is a function V satisfying (1) then it gives the value of the problem to a decision-maker with prior beliefs μ who behaves optimally. Let $a : P(\Theta) \rightarrow X$ be an optimal Markov action. For each $\mu \in P(\Theta)$, $a(\cdot)$ selects an action $a(\mu) \in X$. A Markov plan $a^\infty(\mu)$ is a sequence of optimal Markov actions when the initial state is μ .

LEMMA 1. *There exists a unique solution $V : P(\Theta) \rightarrow \mathbb{R}$ to (1). This solution is continuous. Further, there exists a Markov plan a^∞ and a corresponding q -measurable optimal Markov action $a : P(\Theta) \rightarrow X$ such that $V(\mu) = u(a(\mu), \mu) + \delta \int_{P(\Theta)} V(\bar{\mu}) dq(\mu, a(\mu))$.*

PROOF. Apply the results of Blackwell (1965) and Maitra (1968). □

3. MONOTONICITY OF THE EXPECTED VALUE OF INFORMATION

An action taken today provides valuable information which may be used to make improved decisions tomorrow. Before observing y , the expected value of improved decisions due to observing y is called the expected value of information. Following DeGroot (1962), I define the expected value of information as

$$I(x; \mu, r) = \int_{P(\Theta)} V(\bar{\mu}) dq(\mu, x) - V(\mu).$$

In the terminology of Marschak and Miyasawa (1968), a result about $I(x; \mu, r)$ that holds for all initial priors $\mu \in P(\Theta)$ and all objective functions or rewards $r \in C$ is said to hold for all decision makers.⁶ $I(x; \mu, r)$ may be used to define a partial ordering over actions x .

DEFINITION. *An action x is said to be more informative than an action x' if $I(x; \mu, r) > I(x'; \mu, r)$ for all decision makers, i.e., for all $\mu \in P(\Theta)$ and $r \in C$.*

I will use Blackwell's (1953) well-known result that "sufficiency" implies "more informative." A proof at the level of generality needed here appears in Nyarko (1990). Let y and y' be the outcomes associated with actions x and x' , respectively.

⁶ Recall that r is the reward function and C is the space of uniformly bounded, continuous functions. r enters $I(x; \mu, r)$ through the value function V which depends on rewards.

Let $m(dy'|y)$ be a probability measure giving the probability of y' conditional on y .⁷

LEMMA 2. *If there exists a conditional probability measure m such that*

$$\nu(A|x'\theta) = \int_Y m(A|y)\nu(dy|x\theta) \quad \forall A \in B_Y \text{ and } \theta \in \Theta,$$

then $I(x; \mu, r) \geq I(x'; \mu, r)$ for all $\mu \in P(\Theta)$ and $r \in C$.

To interpret the Lemma, suppose y has a density function $f(y|x\theta)$. Then x is more informative than x' if there exists a conditional density $g(y'|y)$ such that

$$f(y'|x'\theta) = \int_Y g(y'|y)f(y|x\theta) dy \quad \forall y' \in Y \text{ and } \theta \in \Theta.$$

While g may depend on x and x' , the condition implies that g is independent of θ . See Kihlstrom (1984) for further discussion.

The remainder of this section is devoted to characterizing the dependence of $I(x; \mu, r)$ on x for the reservation price (multinomial), posted price (Poisson), and normal models. For these models it is notationally more familiar to work with the probability mass functions or density functions $f(\cdot|x\theta)$ and $g(\cdot|y)$ corresponding to $\nu(\cdot|x\theta)$ and $m(\cdot|y)$, respectively. In defining f let I_Y be the indicator function on Y . Since I have all but abandoned the time subscripts on x and y , it will create no confusion to denote elements of y and θ by subscripts: $y = (y_1, \dots, y_n)$ and $\theta = (\theta_1, \dots, \theta_k)$. Many of the economically interesting applications of the model occur when x is a diagonal matrix. In these cases write

$$x = \text{diag}(x_i) = \begin{bmatrix} x_1 & 0 & \cdots & 0 \\ 0 & x_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & x_k \end{bmatrix} \text{ and}$$

$$a(\mu) = \text{diag}(a_i(\mu)) = \begin{bmatrix} a_1(\mu) & 0 & \cdots & 0 \\ 0 & a_2(\mu) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \cdot & \cdots & a_k(\mu) \end{bmatrix}.$$

DEFINITION. (Y, X, Θ, f) is a multinomial model if

$$f(y|x\theta) = I_Y(y) \left(1 - \sum_{i=1}^k x_i \theta_i \right)^{1 - \sum_{i=1}^k y_i} \prod_{i=1}^k (x_i \theta_i)^{y_i},$$

⁷ Formally, $m: B_Y \times Y \rightarrow [0, 1]$ is a stochastic kernel: for each $y \in Y$, $m(\cdot|y)$ is a probability measure on (Y, B_Y) and for each $A \in B_Y$, $m(A|\cdot)$ is a B_Y -measurable function. Recall that B_Y is the Borel σ -algebra for Y .

$$Y = \left\{ y : y_i \in \{0, 1\} \forall i, \sum_{i=1}^k y_i \leq 1 \right\}, \text{ and}$$

$$X \times \Theta \subset \left\{ (x, \theta) : x = \text{diag}(x_i), x_i \theta_i \geq 0 \forall i, \sum_{i=1}^k x_i \theta_i \leq 1 \right\}.$$

A key result of this paper is that for the reservation price model, if $|x_i| > |x'_i|$ for $i = 1, \dots, k$ then x is more informative than x' . That is, $I(x; \mu, r)$ is monotonic in x for all decision makers.

THEOREM 1 (Multinomial). *If $|x_i| \geq |x'_i|$ for each i then $I(x; \mu, r) \geq I(x'; \mu, r)$ for all $\mu \in P(\Theta)$ and $r \in C$.*

The same result holds for the Poisson model.

DEFINITION. *(Y, X, Θ, f) is a Poisson model if*

$$f(y|x\theta) = I_Y(y) \prod_{i=1}^k e^{-x_i \theta_i} (x_i \theta_i)^{y_i} / y_i!$$

$$Y = \{y : y_i \in \{0, 1, 2, \dots\} \forall i\}, \text{ and } X \times \Theta \subset \{(x, \theta) : x = \text{diag}(x_i), x_i \theta_i > 0 \forall i\}.$$

THEOREM 2 (Poisson). *If $|x_i| \geq |x'_i|$ for each i then $I(x; \mu, r) \geq I(x'; \mu, r)$ for all $\mu \in P(\Theta)$ and $r \in C$.*

The definition of $X \times \Theta$ together with the fact that the action space X cannot depend on the unknown parameter θ implies that for each i , either $x_i \geq 0$ and $\theta_i \geq 0$ for all $(x, \theta) \in X \times \Theta$ or $x_i \leq 0$ and $\theta_i \leq 0$ for all $(x, \theta) \in X \times \Theta$, but not both. Hence, in Theorems 1 and 2 $|x_i| \geq |x'_i|$ only occurs if $x_i \geq x'_i \geq 0$ or $x_i \leq x'_i \leq 0$.

Under normality with x and θ taken as *scalars*, the monotonicity of a function related to $I(x; \mu, r)$ has been investigated by Grossman, Kihlstrom, and Mirman (1977). I generalize their result to the case where x is a *matrix* of controls (not necessarily diagonal) and θ is a *vector* of unknown parameters. Let a ‘‘T’’ superscript denote matrix transposition.

DEFINITION. *(Y, X, Θ, f) is a normal model if $Y = \mathbb{R}^n$, $X \subset \mathbb{R}^{nk}$, $\Theta \subset \mathbb{R}^k$, and $f(y|x\theta) = I_Y(y) (2\pi)^{-n/2} |\Sigma|^{-1/2} \exp\{-(y - x\theta)^T \Sigma^{-1} (y - x\theta)/2\}$ where Σ is a known $n \times n$ symmetric positive definite matrix which may depend on x .*

When Σ depends on x there is heteroscedasticity in the sense of Creane (1991). For the normal model Blackwell’s ‘‘sufficiency’’ condition in Lemma 2 can be enormously simplified. The following definition and its use in Theorem 3 shows how. Let Σ and Σ' be the covariances of y and y' , respectively.

DEFINITION. An $n \times n$ matrix S is a sufficient correlation matrix for (x, x') if

- (i) $x' = S\Sigma^{-1}x$
- (ii) $M = \begin{bmatrix} \Sigma & \\ & \Sigma' \end{bmatrix}$ is positive definite, and
- (iii) S is independent of θ .

THEOREM 3 (Normal). If there exists a sufficient correlation matrix S for (x, x') then $I(x; \mu, r) \geq I(x'; \mu, r)$ for all $\mu \in P(\Theta)$ and $r \in C$.

When x is a diagonal matrix, $\Sigma = \Sigma'$, and the y_i are independent so that Σ is a diagonal matrix, then Theorem 3 is the multivariate generalization of Grossman, Kihlstrom, and Mirman (1977). Their insightful and constructive proof forms the basis for the proof of Theorem 3.

COROLLARY 1 (Normal). Let $\Sigma = \Sigma'$ and let Σ and x be diagonal matrices. If $|x_i| \geq |x'_i|$ for each i then $I(x; \mu, r) \geq I(x'; \mu, r)$ for all $\mu \in P(\Theta)$ and $r \in C$.

PROOF. Write $\Sigma = \text{diag}(\sigma_i^2)$. Let $S = R\Sigma$ for some diagonal matrix R with typical element ρ_i . $S = R\Sigma$ is convenient since then $x' = S\Sigma^{-1}x$ reduces to $x' = Rx$ or $\rho_i = x'_i/x_i$ for $i = 1, \dots, k$. Verify that S is a sufficient correlation matrix by verifying that M is positive definite. Positive definiteness holds if for all vectors $a = (a_1, \dots, a_n)$ and $b = (b_1, \dots, b_n)$ with $(a \ b)$ nonzero, $0 < (a \ b)M(a \ b)^T = a\Sigma a^T + 2bR\Sigma a^T + b\Sigma b^T = \sum_{i=1}^k \sigma_i^2 ((a_i + \rho_i^2 b_i)^2 + (1 - \rho_i^2) b_i^2)$. This is satisfied if and only if $|\rho_i| = |x'_i/x_i| < 1$. Hence, there exists a sufficient correlation matrix if $|x'_i/x_i| < 1$ for all i . The conclusion follows from Theorem 3. \square

This proof was presented in order to show how easy it is to verify the existence of a sufficient correlation matrix. In contrast, the condition in Lemma 2 is often difficult to verify. This completes the discussion of the monotonicity in x of the expected value of information, $I(x; \mu, r)$.

4. DIRECTION OF EXPERIMENTATION

In the nonlearning problem the decision maker ignores the impact of current actions on future information and myopically chooses $x = a^{NL}(\mu)$ to maximize one period rewards $u(x, \mu)$. NL stands for "nonlearning." Experimentation may be defined as $a(\mu) - a^{NL}(\mu)$ i.e., the deviation from the nonlearning optimal action which is motivated by the desire for improved information. Theorem 4 states the main result on experimentation.

THEOREM 4 (Multinomial, Poisson, and Normal ($\Sigma = \Sigma'$)). Assume that $a^{NL}(\mu)$ is unique and let X contain only diagonal matrices. Then $|a_i(\mu)| \geq |a_i^{NL}(\mu)|$ $i = 1, \dots, k$ for all measurable selections $a(\cdot)$, $\mu \in P(\Theta)$, and $r \in C$.

Theorem 4 applies in each period. That is, $|a(\mu_t)| \geq |a^{NL}(\mu_t)|$ for all t . Further, there need not be a unique optimal plan $a(\mu)$ for the infinite horizon problem—Theorem 4 applies to any measurable selection $a(\mu)$ from the set of optimal actions.

Theorem 4 shows that for demand processes which stem naturally from the market in which the monopolist operates, namely, the reservation price and posted price markets, the direction of experimentation can be determined.

5. THE IGNORANT MONOPOLIST

Economics lacks a good theory of the pricing and output decisions of a monopolist who does not know her demand. The preceding Theorems can be used to explain how the monopolist can generate demand information endogenously through optimal price experimentation.

The Reservation Price Market.

Example 1. Consider a market in which one customer arrives per period and decides whether or not to buy one unit of a good. Let $y = 0$ indicate a decision to purchase. The monopolist chooses a price $x \in X = [0, 1]$ which determines the probability of purchase $1 - x\theta$. Let $\Theta = [0, 1]$ so that the probability of purchase lies between 0 and 1. As in the Lazear (1986) model, at very low prices there is little uncertainty about the probability of purchase $1 - x\theta$ and in the extreme when $x = 0$ the customer always purchases regardless of the "true" value of θ . On the other hand, at the highest price ($x = 1$) uncertainty is greatest in the sense that the probability of purchase $1 - \theta$ can lie anywhere between 0 and 1.⁸ Thus, the direction of experimentation is expected to be upwards. Indeed, Theorem 4 states that $a(\mu) \geq a^{NL}(\mu)$: the learning monopolist always sets a price at least as high as her nonlearning counterpart who ignores the informational content of prices and observed quantities.

Example 2. Consider a multi-product monopoly selling many varieties of a good. The monopoly might be a public utility facing a peak-load pricing problem in which electricity is sold at k different rates depending on the time of day and the customer decides at which time to use electricity. Alternatively, the monopoly might be a car dealership and the customer decides which of k car models to purchase or if to purchase at all. In the dealership example let $y_i = 1$ indicate a decision to purchase car model i and let $\sum_{i=1}^k y_i = 0$ indicate a decision not to purchase any model. As before, let $X = \Theta = [0, 1]$, but let the specification of uncertainty be opposite that of Example 1 by having uncertainty being most pronounced at low prices. For concreteness, let $p_i = 1/x_i$ be the price of model i , $p_i(\mu) = 1/a_i(\mu)$ be the optimal price, and $p_i^{NL}(\mu) = 1/a_i^{NL}(\mu)$ be the optimal nonlearning price. As p_i rises, the model i purchase probability θ_i/p_i approaches zero with certainty whereas at very low prices the purchase probability is very uncertain. Thus, the expected direction of price experimentation is downwards and Theorem 4 states that $p_i(\mu) \leq p_i^{NL}(\mu)$

⁸ At very high prices it is reasonable to expect demand to be zero so that no learning takes place. However, the requirement of nonnegative purchase probabilities means demand is never zero even at the highest price. More precisely, if $\text{prob}\{\theta = 1\} < 1$ and $x = 1$, then $1 - \theta > 0$ with positive prior probability. The reason that a biting upper bound on price does not appear in other studies is that other studies typically allow quantity demanded to be negative. I do not. (Note that in Example 2 below prices are not bounded from above.)

for $i = 1, \dots, k$: the learning monopolist sets *each* price at least as low as her nonlearning counterpart who ignores the informational content of prices and observed quantities.⁹

The Posted Price Market. In the posted price market the monopolist posts a price which determines the flow of customers wishing to purchase the good. When $k > 1$ the model can also be interpreted as a price-discriminating monopolist operating in k segmented markets or as a multi-product monopolist selling k products.

Example 3. Consider a multi-product firm producing k different products. For example, the firm might be a pharmaceutical company for which R&D externalities link products on the cost side. Of course, the specification of cost externalities is unimportant since the results on experimentation will hold for all cost functions. Let y_i be the quantity demanded of the i th product. Mean demand is given by $E[y|x\theta] = (x_1\theta_1, \dots, x_k\theta_k)$. Let $p_i = 1/x_i$ be the price of the i th product so that its mean demand, θ_i/p_i , is downward sloping. Since demand uncertainty is most pronounced at lower prices the direction of experimentation is expected to be downwards and Theorem 4 states that $p_i(\mu) \leq p_i^{NL}(\mu)$ for *each* product i .

Another interpretation of the posted price market is that of a monopolist producing a single homogeneous good which is sold in k segmented markets. Then in Example 3, p_i is reinterpreted as the price charged in market i .

Price Setting versus Quantity Setting in the Normal Model. Grossman, Kihlstrom, and Mirman (1977) considered a special case of this paper where $y = \alpha + x\theta + \varepsilon$, ε is normal with zero mean and unit variance, and α is a known intercept. They arrive at the following “apparently contradictory results.” For the price-setting monopolist, x is price and y is quantity demanded. Theorem 4 then states that $a(\mu) \geq a^{NL}(\mu)$: the monopolist always experiments upwards. One then incorrectly expects the quantity-setting monopolist to experiment downwards. For the quantity-setting monopolist facing an inverse demand function $y = \alpha + x\theta + \varepsilon$, x is quantity supplied and y is price. Theorem 4 still states that $a(\mu) \geq a^{NL}(\mu)$: the monopolist experiments upwards. Both the price and quantity setters experiment in the same direction! These “apparently contradictory results” were observed by Grossman, Kihlstrom, and Mirman (1977) who offered little explanation, merely pointing out that the results were not logically inconsistent. The explanation is not complicated. Consider Figure 1. Write $f(y - \alpha - x\theta_i)$ in place

⁹ Example 2 features the independence of irrelevant alternatives (IIA) in that the probability of purchasing variety j is independent of the price of variety i . While this may be desirable in the electric utility example, it is less so in the car dealership example where a higher price of the luxury model can be expected to increase the probability of purchasing mid-line models. The IIA feature is well-known from the literature on multinomial logits and there are many ways of relaxing it (McFadden 1984). We consider only one. In Example 2 let p_i be price and let $P_i = 1/x_i$ be price normalized by some price index. For example, let $P_i = p_i/\sum_{i=1}^k p_i$ or use a CES index to obtain $P_i = p_i/(\sum_{i=1}^k (p_i)^{\sigma-1})^{1/(\sigma-1)}$ where σ is the cross-elasticity of demand. For reasonable indexation, including the CES index, raising p_i raises P_i and lowers the model i purchase probability θ_i/P_i while raising p_i lowers P_j for $j \neq i$ and raises the probability of purchasing all other models, θ_j/P_j . That is, the monopolist recognizes the dependence of the model i purchase probability on the prices of all other models.




FIGURE 1

PRICE-SETTING MONOPOLIST

of $f(y|x\theta_i)$. The expected demand functions, $Ey = \alpha + x\theta$, for the price-setting monopolist are drawn for two values of θ . The price-setter knows her expected demand at low prices, but is less certain at high prices. Thus, the larger is x the less $f(y - \alpha - x\theta_1)$ overlaps with $f(y - \alpha - x\theta_2)$ and the more clearly a realization of y reflects on θ . That is, a higher price is more informative. For the quantity-setting monopolist y is price and x is quantity so that the expected inverse demand functions hinge on the price axis. See Figure 2. Hence, larger quantities are more informative. This point is taken up in Mirman, Samuelson, and Urbano (1989).

This resolves the apparently contradictory results. It also highlights that in using the results of this paper one must be sure that there are economic reasons for the particular form of the stochastic process. Lazear (1986) provides a justification for assuming that a price-setting retailer faces a demand which “hinges” on the quantity axis. At high prices the retailer is unsure if he can move stock, but at low prices he is sure of sales. On the other hand, Tonks’ (1983) result is turned on its head by an economically unimportant change in his specification.

Dichotomies. There has been considerable attention in the literature to the case where θ can take on only two values i.e., $\Theta = \{\theta_1, \theta_2\}$. Fusselman and Mirman (1989), Mirman, Samuelson, and Urbano (1989), and Creane (1991) examined the two period problem with x , y , and θ as scalars and $y = x\theta + \epsilon$. For the infinite-horizon problem, Theorem 5 reports a similar result.¹⁰

¹⁰ In the context of the literature on experiments an unusual and very useful feature of these studies is that they allow for particular forms of demand nonlinearity. These nonlinearities can be incorporated into Theorem 5. Theorem 5 differs from the cited results in that f need not be a density and $a(\mu)$ and $V(\mu)$ need not be differentiable. (Kiefer 1989 showed that for even simple infinite-horizon problems differentiability may not hold.)




FIGURE 2

QUANTITY-SETTING MONOPOLIST

THEOREM 5 (Dichotomy). *Let $x, y,$ and θ be scalars. Assume (i) Θ has only two elements, (ii) Y is bounded, (iii) $x\theta$ is a location parameter of y i.e., $f(y|x\theta) = h(y - x\theta)$ for some density or probability mass function $h,$ and (iv) f has a monotone likelihood ratio. Then $x > x' \geq 0$ or $x < x' \leq 0$ implies $I(x; \mu, r) \geq I(x'; \mu, r)$ for all $\mu \in P(\Theta)$ and $r \in C.$*

As the proof of Theorem 5 makes clear, the additive error assumption $y = x\theta + \varepsilon$ (condition (iii)) can be replaced by the much weaker condition $x\theta \geq x'\theta \Rightarrow F(\cdot|x\theta) \leq F(\cdot|x'\theta)$ where F is the cumulative probability function. Further, if f is symmetric then Theorem 5 holds for $|x| > |x'|.$ The question of whether Theorem 5 and the Figure 1 and 2 argument generalize to more than two values of θ is still open.¹¹

Heteroscedasticity. In Figures 1 and 2, x affects the location of the density, but not its shape; in particular, the variance of y is independent of $x.$ Creane (1991) calls this the homoscedastic case. Unlike the distribution represented in Figures 1 and 2, the multinomial and Poisson distributions display heteroscedasticity since the variance of y depends on x (as well as $\theta).$ For the Poisson distribution both the mean and variance of y equal $x\theta,$ so that the tails of $f(y|x\theta)$ become fatter as x increases. Thus, while a larger x reduces the overlap between $f(y|x\theta_1)$ and $f(y|x\theta_2)$ by spreading out the means, a larger x also increases the overlap by fattening the tails. Similarly, for the multinomial distribution the variance is $x\theta(1 - x\theta),$ which increases in x for $x\theta < 1/2.$ *Due to heteroscedasticity, the monotonicity of the*

¹¹ The argument of Figures 1 and 2 and the proof of Theorem 5 are based on a theorem by Blackwell which has yet to be generalized to the case where θ has more than two elements. See Blackwell's (1953) discussion preceding his Theorem 10.

expected value of information is not obvious for the multinomial and Poisson models.

Creane introduced heteroscedasticity into the two-period dichotomy problem by considering $y = \alpha + x\theta + \phi(x)\varepsilon$ for some known function ϕ . He showed that $I(x; \mu, r)$ is increasing (decreasing) in x if $\phi - x(\partial\phi/\partial x) > 0$ (< 0). For the infinite-horizon problem with ε distributed normally, Theorem 3 can be used to show that Creane's result remains valid *without any restrictions* on the number of values that θ can take. Further, the simplicity of the proof evidences the usefulness of Theorem 3.¹²

6. CONCLUSIONS

Economics lacks a good theory of the pricing and output decisions of a firm which does not know its demand. We always assume the firm has complete demand information or exact knowledge of the stochastic process generating demand. Yet no explanation is given for how the firm comes by this information. I presented a simple model in which demand information is generated endogenously. Initially, the monopolist only has the limited information which she can reasonably be expected to know from the readily observable features of the reservation price and posted price markets. All other demand information is generated endogenously through price and quantity experimentation. I characterized the direction of experimentation and showed how it relates to the nature of the demand uncertainty. Unlike the previous literature which usually assumes only a single control x this characterization applied to a price-discriminating monopolist operating in several segmented markets and to a multi-product monopolist producing either many distinct products or a line of related products. I also explained the "apparently contradictory results" of Grossman, Kihlstrom, and Mirman (1977) in a way that highlighted differences between the homoscedastic normal and dichotomy models and the heteroscedastic reservation price and posted price models. Finally, I related my results to the literature on the expected value of information and Blackwell's comparison of experiments. In so doing, I characterized a value function used in a wide class of dynamic programming problems for which actions taken today determine not only the reward today, but also the information available tomorrow.

University of Toronto, Canada

APPENDIX

PROOF OF THEOREM 1 (*Multinomial*). $y = (y_1, \dots, y_k)$ where $y_i \in \{0, 1\}$ so that without confusion write $y = i$ if $y_i = 1$ and $y = 0$ if $y_i = 0$ for all i .

¹² Without loss of generality let x, x' , and ϕ be positive. Define Σ, Σ', S , and M as in the definition of a sufficient correlation matrix. Then $\Sigma = \phi^2(x)$ and $\Sigma' = \phi^2(x')$. Let the candidate for S be $(x'/x)\phi^2(x)$ so that $x' = S \Sigma^{-1}x$. Then S is a sufficient correlation matrix if M is positive definite. Positive definiteness holds for $|S| < \phi(x)\phi(x')$ or $(x'/x)\phi^2(x) < \phi(x)\phi(x')$ or $\phi(x')(x - x') - x'(\phi(x) - \phi(x')) > 0$. Dividing by $x - x'$ and letting $x \rightarrow x'$ yields Creane's result: $I(x; \mu, r) \geq I(x'; \mu, r)$ if $x > x'$ and $\phi - x\partial\phi/\partial x > 0$ or if $x < x'$ and $\phi - x\partial\phi/\partial x < 0$.

Accordingly, replace the set Y with the set $\{0, 1, \dots, k\}$. By Lemma 2, it suffices to find a conditional probability mass function (p.m.f.) $g(y'|y)$ satisfying

$$(2) \quad f(j|x'\theta) = \sum_{i=0}^k g(j|i)f(i|x\theta) \quad \text{for } j = 0, 1, \dots, k \text{ and for all } \theta,$$

where i indexes y and j indexes y' . Consider

$$g(0|i) = \begin{cases} 1 & i = 0 \\ 1 - x'_i/x_i & i > 0 \end{cases} \quad \text{and for } j > 0 \quad g(j|i) = \begin{cases} x'_j/x_i & i = j \\ 0 & i \neq j \end{cases}.$$

Then equation (2) holds. Further, g is a p.m.f. since $\sum_{j=0}^k g(j|i) = 1$ for all i and $g \geq 0$ if for each i , $x_i \geq x'_i \geq 0$ or $x_i \leq x'_i \leq 0$. \square

PROOF OF THEOREM 2 (Poisson). Find a conditional p.m.f. $g(y'|y)$ satisfying the condition in Lemma 2. Denote realizations of y and y' by (i_1, \dots, i_k) and (j_1, \dots, j_k) , respectively. Consider

$$g(j_1, \dots, j_k | i_1, \dots, i_k) = \prod_{m=1}^k h(m, i_m, j_m)$$

where

$$h(m, i_m, j_m) = \begin{cases} \binom{i_m}{j_m} \frac{(x'_m)^{j_m} (x_m - x'_m)^{i_m - j_m}}{(x_m)^{i_m}} & i_m \geq j_m \\ 0 & i_m < j_m \end{cases} \quad m = 1, \dots, k.$$

Show that $f(y'|x'\theta) = \sum_{y \in Y} g(y'|y)f(y|x\theta)$. Define $f_m(i_m | x_m \theta_m) = e^{-x_m \theta_m} (x_m \theta_m)^{i_m} / i_m!$ and note that

$$\begin{aligned} (3) \quad & \sum_{i_m=0}^{\infty} h(m, i_m, j_m) f_m(i_m | x_m \theta_m) \\ &= \sum_{i_m=j_m}^{\infty} \frac{i_m!}{j_m!(i_m - j_m)!} \frac{(x'_m)^{j_m} (x_m - x'_m)^{i_m - j_m}}{(x_m)^{i_m}} \frac{e^{-x_m \theta_m} (x_m \theta_m)^{i_m}}{i_m!} \\ &= e^{-x_m \theta_m} \frac{(x'_m \theta_m)^{j_m}}{j_m!} \sum_{i_m=j_m}^{\infty} \frac{(x_m \theta_m - x'_m \theta_m)^{i_m - j_m}}{(i_m - j_m)!} = f_m(j_m | x'_m \theta_m) \end{aligned}$$

where the third equality follows from $e^z = \sum_{n=0}^{\infty} z^n/n!$. Now

$$\begin{aligned}
 & \sum_{y \in Y} g(y'|y)f(y|x\theta) \\
 &= \sum_{i_k=0}^{\infty} \cdots \sum_{i_1=0}^{\infty} h(k, i_k, j_k)f_k(i_k|x_k\theta_k) \cdots h(1, i_1, j_1)f_1(i_1|x_1\theta_1) \\
 &= \sum_{i_k=0}^{\infty} \cdots \sum_{i_2=0}^{\infty} h(k, i_k, j_k)f_k(i_k|x_k\theta_k) \cdots h(2, i_2, j_2)f_2(i_2|x_2\theta_2) \\
 &\quad \times \sum_{i_1=0}^{\infty} h(1, i_1, j_1)f_1(i_1|x_1\theta_1) \\
 &= f_1(j_1|x'_1\theta_1) \sum_{i_k=0}^{\infty} \cdots \sum_{i_2=0}^{\infty} h(k, i_k, j_k) \\
 &\quad \times f_k(i_k|x_k\theta_k) \cdots h(2, i_2, j_2)f_2(i_2|x_2\theta_2)
 \end{aligned}$$

where appeal is made to equation (3). Repeating this factorization yields $\sum_{y \in Y} g(y'|y)f(y|x\theta) = \prod_{m=1}^k f_m(j_m|x'_m\theta_m) = f(y'|x'\theta)$. Hence, the condition in Lemma 2 is satisfied.

It remains to show that $g(y'|y)$ is a conditional density. The binomial theorem states

$$\sum_{j_m=0}^{i_m} \binom{i_m}{j_m} (x'_m)^{j_m} (x_m - x'_m)^{i_m - j_m} = (x'_m + (x_m - x'_m))^{i_m} = (x_m)^{i_m}$$

so that

$$(4) \quad \sum_{j_m=0}^{\infty} h(m, i_m, j_m) = 1 \quad \text{for } m = 1, 2, \dots, k.$$

Hence,

$$\begin{aligned}
 \sum_{y \in Y} g(y'|y) &= \sum_{j_k=0}^{\infty} \cdots \sum_{j_1=0}^{\infty} h(k, i_k, j_k) \cdots h(1, i_1, j_1) \\
 &= \sum_{j_k=0}^{\infty} \cdots \sum_{j_2=0}^{\infty} h(k, i_k, j_k) \cdots h(2, i_2, j_2) \sum_{j_1=0}^{\infty} h(1, i_1, j_1).
 \end{aligned}$$

Using equation (4) to simplify and repeating the factorization yields $\sum_{y \in Y} g(y'|y) = 1$. Finally, $g \geq 0$ if $x_i \geq x'_i \geq 0$ or $x_i \leq x'_i \leq 0$. \square

PROOF OF THEOREM 3 (Normal). I show that if a sufficient correlation matrix S for (x, x') exists then there exists a conditional distribution of y' given y which is independent of θ . It follows that this conditional distribution satisfies the conditions of Lemma 2. Write Z is $N(\alpha, \beta)$ if Z is normally distributed with mean α and variance β . y is $N(x\theta, \Sigma)$ and y' is $N(x'\theta, \Sigma')$. I am free to choose the joint distribution of y and y' provided the implied marginal distributions of y and y' are $N(x\theta, \Sigma)$ and $N(x'\theta, \Sigma')$, respectively. This is accomplished by letting $\begin{bmatrix} y \\ y' \end{bmatrix}$ be $N(\begin{bmatrix} x \\ x' \end{bmatrix}\theta, M)$, where M is given in the definition of a sufficient correlation matrix. For each M (and S) the conditional distribution of y' given y is $N(x'\theta + S \Sigma^{-1}$

$(y - x\theta), \Sigma' - S\Sigma^{-1}S^T$). Since $x' = S\Sigma^{-1}x$, the conditional distribution simplifies to $N(S\Sigma^{-1}y, \Sigma' - S\Sigma^{-1}S^T)$ which is independent of θ . \square

PROOF OF THEOREM 4. If $a^{NL}(\mu) \leq x$ for all $x \in X$ the conclusion follows immediately. Consider $x < a^{NL}(\mu)$ where “ $<$ ” means $x_i \leq a_i^{NL}(\mu)$ for all i with strict inequality for some i . From equation (1), $V(\mu)$ may be written as $V(\mu) = \max_{x \in X} \phi(x, \mu)$, where

$$\phi(x, \mu) = u(x, \mu) + \delta \int_{P(\Theta)} V(\bar{\mu}) dq(\mu, x).$$

Since $a^{NL}(\mu)$ is the unique maximizer of $u(x, \mu)$,

$$(5) \quad u(x, \mu) < u(a^{NL}(\mu), \mu) \quad x < a^{NL}(\mu).$$

By Theorems 1 and 2 and Corollary 1, $I(x; \mu, r) = \int_{P(\Theta)} V(\bar{\mu}) dq(\mu, x) - V(\mu)$ is nondecreasing in x . Hence,

$$(6) \quad \delta \int_{P(\Theta)} V(\bar{\mu}) dq(\mu, x) \leq \delta \int_{P(\Theta)} V(\bar{\mu}) dq(\mu, a^{NL}(\mu)) \quad \text{for } x < a^{NL}(\mu).$$

Summing (5) and (6) yields

$$\phi(x, \mu) < \phi(a^{NL}(\mu), \mu) \quad \text{for } x < a^{NL}(\mu).$$

Since $\phi(a(\mu), \mu) = \max_{x \in X} \phi(x, \mu) \geq \phi(a^{NL}(\mu), \mu)$, it follows that $a(\mu) \geq a^{NL}(\mu)$. \square

PROOF OF THEOREM 5 (*Dichotomy*). Let $F(y|x\theta)$ and $H(y - x\theta)$ be the cumulative distribution functions associated with y and $y - x\theta$, respectively. Let $\Theta = \{\theta_0, \theta_1\}$, let $\text{mlr}(y|x) = f(y|x\theta_1)/f(y|x\theta_0)$, and without loss of generality set $\theta_0 = 0$.

The Neyman-Pearson test of $H_0: \theta = \theta_0$ against $H_1: \theta = \theta_1$ is $T(y|x, t) = 1$ if $y \in Y_1(x, t) = \{y: \text{mlr}(y|x) > c(t)\}$ and $T(y|x, t) = 0$ otherwise. $Y_1(x, t)$ is interpreted as the region of rejection for H_0 when t is the type I error or size of the test.

Consider first the case $x > 0$ and $\theta_1 > 0$. By the monotone likelihood ratio property $Y_1(x, t)$ simplifies to $Y_1(x, t) = \{y: y > y(t)\}$ where $y(t)$ satisfies $t = 1 - F(y(t)|0)$. Type II error is given by $F(y(t)|x\theta_1)$.¹³ Blackwell’s (1953) Corollary to his Theorem 10 implies that x is at least as informative as x' if for all test sizes $T(\cdot|x, t)$ has at least as small type II error as $T(\cdot|x', t)$. This is immediate from the discussion following the Corollary and the fact that for any size test the Neyman-Pearson test attains the infimum of type II error (Ferguson 1967, Section 5.1). Thus, x is at least as informative as x' if for all $t \in [0, 1]$, $F(y(t)|x\theta_1) \leq F(y(t)|x'\theta_1)$ or $H(y(t) - x\theta_1) \leq H(y(t) - x'\theta_1)$. This is satisfied for $x > x'$.

¹³ If f is a p.m.f. then the Neyman-Pearson test presented needs a slight amendment since there may not exist a $y(\cdot)$ satisfying $t = 1 - F(y(t)|0)$ for all $t \in [0, 1]$. To attain all test sizes randomize by letting $T(y|x, t) = \gamma$ when $\text{mlr}(y|x) = c$ where $\gamma \in [0, 1]$ is the probability of rejecting H_0 . See Ferguson (1967). This wrinkle plays no role.

Hence when $x > 0$ and $\theta_1 > 0$, $I(x; \mu, r)$ is nondecreasing for all $\mu \in P(\Theta)$ and $r \in C$.

The conclusion follows by finding $Y_1(x, t)$ and type II error for the cases $(x > 0, \theta_1 < 0)$, $(x < 0, \theta_1 > 0)$, and $(x < 0, \theta_1 < 0)$. \square

REFERENCES

- BLACKWELL, D., "Equivalent Comparison of Experiments," *Annals of Mathematical Statistics* 24 (1953), 265-272.
- , "Discounted Dynamic Programming," *Annals of Mathematical Statistics* 36 (1965), 226-235.
- CREANE, A., "Experimentation with Heteroscedastic Noise," mimeo, Department of Economics, Michigan State University, 1991.
- DEGROOT, M. H., "Uncertainty, Information and Sequential Experiments," *Annals of Mathematical Statistics* 33 (1962), 404-419.
- EASLEY, D. AND N. M. KIEFER, "Infinite-Horizon Bayesian Control of the Normal-Normal Regression Process," Working Paper No. 386, Cornell University, 1987.
- AND ———, "Controlling a Stochastic Process with Unknown Parameters," *Econometrica* 56 (1988), 1045-1066.
- FERGUSON, T. S., *Mathematical Statistics: A Decision Theoretic Approach* (New York: Academic Press, 1967).
- FUSSELLMAN, J. M. AND L. J. MIRMAN, "Experimental Consumption for a General Class of Disturbance Densities," working paper, University of Virginia, 1989.
- GROSSMAN, S. S., R. E. KIHLSSTROM AND L. J. MIRMAN, "A Bayesian Approach to the Production of Information and Learning by Doing," *Review of Economic Studies* 44 (1977), 533-547.
- KIEFER, N. M., "A Value Function Arising in the Economics of Information," *Journal of Economic Dynamics and Control* 13 (1989), 201-223.
- AND Y. NYARKO, "Optimal Control of an Unknown Linear Process with Learning," *International Economic Review* 30 (1989), 571-585.
- KIHLSSTROM, R. E., "A 'Bayesian' Exposition of Blackwell's Theorem on the Comparison of Experiments," in M. Boyer and R. E. Kihlstrom, eds, *Bayesian Models in Economic Theory* (Amsterdam: Elsevier Science Publishers B. V., 1984), 13-31.
- LAZEAR, E. P., "Retail Pricing and Clearance Sales," *American Economic Review* 76 (1986), 14-32.
- McFADDEN, D. L., "Econometric Analysis of Qualitative Response Models," in Z. Griliches and M. D. Intriligator, eds., *Handbook of Econometrics*, Vol. 2 (Amsterdam: Elsevier Science Publishers B. V., 1984), 1395-1457.
- MAITRA, A., "Discounted Dynamic Programming on Compact Metric Spaces," *Sankhya, Series A* 30 (1968), 211-216.
- MARSCHAK, J. AND K. MIYASAWA, "Economic Comparability of Information Systems," *International Economic Review* 9 (1968), 137-174.
- MIRMAN, L. J., L. SAMUELSON AND A. URBANO, "Monopoly Experimentation," working paper, University of Virginia, 1989.
- NYARKO, Y., "On Bayesian Learning in Optimal Control Problems," working paper, New York University, 1990.
- PARTHASARATHY, K., *Probability Measures on Metric Spaces* (New York: Academic Press, 1967).
- PRESCOTT, E., "The Multi-period Control Problem under Uncertainty," *Econometrica* 40 (1972), 1043-1058.
- ROTHSCHILD, M., "A Two-Armed Bandit Theory of Market Pricing," *Journal of Economic Theory* 9 (1974), 185-202.
- STOKEY, N. L. AND R. E. LUCAS, JR., *Recursive Methods in Economic Dynamics* (Cambridge: Harvard University Press, 1989).
- TONKS, I., "Bayesian Learning and the Optimal Investment Decision of the Firm," *Economic Journal* 93 (1983), 87-98.