

The Review of Economic Studies Ltd.

A Bayesian Approach to the Production of Information and Learning by Doing

Author(s): Sanford J. Grossman, Richard E. Kihlstrom, Leonard J. Mirman

Source: *The Review of Economic Studies*, Vol. 44, No. 3 (Oct., 1977), pp. 533-547

Published by: [The Review of Economic Studies Ltd.](#)

Stable URL: <http://www.jstor.org/stable/2296906>

Accessed: 09/12/2010 02:44

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=resl>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The Review of Economic Studies Ltd. is collaborating with JSTOR to digitize, preserve and extend access to *The Review of Economic Studies*.

<http://www.jstor.org>

A Bayesian Approach to the Production of Information and Learning By Doing

SANFORD J. GROSSMAN

Stanford University

RICHARD E. KIHLSSTROM and LEONARD J. MIRMAN

University of Illinois

1. INTRODUCTION

The desire to acquire economically valuable information provides a powerful explanation for many empirically observed economic phenomena. Two examples which have been extensively studied by economists are investment in "human capital" through education and expenditures on research and development. In these examples, information is explicitly purchased. In economic contexts where individuals and firms learn from experience, the demand for information may be manifested in other ways. One such example, arises when new and untested products such as new drugs are introduced to the market. It is plausible to hypothesize that consumers confronted with new products experiment with them to gain information. The demand for experimental consumption might be expected to increase total demand for new products over what it otherwise would be.

Firms which enter new and unfamiliar markets may also learn from experience. Indeed, firms facing unknown demand curves will often find it profitable to experiment with price in an effort to improve the information acquired through experience. The exact nature of the effect which firm experimentation has on observed prices and supplies would appear to be less predictable than the consumer response to new products, however.

This paper analyses the phenomena of learning and experimentation in the context of a dynamic economic model which incorporates a Bayesian expectation revision mechanism. In this model, the individual (or firm) responds to new information as it is received, but he is not passive about the information he obtains. Indeed, he recognizes that his future expectations, and therefore his future decisions, will depend on the information which is acquired by observing the consequences of his present actions. He is also aware that the quality of information acquired from experience may be affected by the specific course of action followed in the present. In this model, individuals find it profitable to modify their behaviour as a means of improving the information on which future decisions are based. In deciding exactly how much experimentation to engage in, individuals weigh the benefits from more informed future decisions against the costs incurred because present experimental actions differ from those which would be optimal if learning from experience did not occur.

The model is first interpreted as a description of the situation faced by a consumer who buys, in addition to other goods, a drug of unknown reliability. If the consumer's health is affected by random factors as well as by the drug, his experience with the drug provides less than completely dependable, i.e. "noisy", information about its reliability. In this model, the consumer's drug purchases reflect a desire to learn through experimentation. We

formulate a dynamic model to analyse the effects of experimentation on the amount of the drug consumed. The consumer is assumed to maximize the present value of the expected utility derived from consumption of the drug and other goods. The formal mechanism by which the information acquired through experimentation is assimilated is Bayes' Rule.

It is shown that, when drug consumption affects health through a linear regression equation, the possibility of learning from experience induces experimentation which in turn causes the consumer to buy more of the drug than he would if no learning took place, other things being equal. This result generalizes to a broader class of probability distributions a result obtained earlier by Prescott [13].

The final section of the paper reinterprets the model and our conclusions to analyse the effects of experimentation by monopolists who are attempting to learn their demand curve.

The idea that learning from experience affects economic behaviour has been investigated by Arrow [1]. Arrow assumes that over time firms accumulate experience which increases productivity. His paper is not, however, founded on a formal statistical model of the information generating process which results in "learning by doing". The present paper can be interpreted as providing a model of this process.

This paper can also be viewed as a complement to the papers of Kihlstrom [9], [10] which study the demand for information about product quality on the part of Bayesian consumers. In Kihlstrom's work information is actually purchased in markets which exist for the explicit purpose of selling information; consumers do not experiment to learn about product quality. In the model studied here, information demand arises in an implicit form and it is satisfied by the consumer himself when he experiments with his consumption choices.

2. THE MODEL

Consider the idealized problem of a consumer who receives a stationary income over time and uses it in each period to buy two goods, one of which is a drug of unknown quality. The other good can be interpreted as a composite good which provides fixed proportions of all other commodities.

We let

$$y_t = \text{drug consumption in period } t,$$

and

$$x_t = \text{consumption of the other good in period } t.$$

The consumer's periodic income is $I > 0$. Prices are normalized so that the price per unit of the composite good is 1. The per unit price of the drug is $p > 0$.

The fact that drug quality is unknown is assumed to imply that the consumer views the effect of drug consumption on health as a random, but non-cumulative relationship. Specifically, it is assumed that health in period t can be measured by a variable z_t which is related to drug consumption in period t by the linear equation

$$\tilde{z}_t = \alpha + \beta y_t + \tilde{\varepsilon}_t,$$

where $\{\tilde{\varepsilon}_t\}$ is a sequence of intertemporally independent and unobserved normal random variables each, with mean zero, and variance one. (In this equation, as in the remainder of the paper, random variables are denoted by a " \sim ".) This specification implies that random variations in health are unrelated to drug intake and occur even if the consumer abstains from drug use. Without drug use, health is measured by $\alpha + \tilde{\varepsilon}_t$ which is a normal random variable with mean α and variance 1. It is plausible to assume that the consumer knows the distribution of health when no drugs are used. Thus he is assumed to know his "average health" parameter α and the variance of ε_t which, for convenience, we assume to be 1.

The parameter β is the contribution of each unit of drug consumption to health. This parameter measures drug quality and is unknown. The consumer has beliefs about the

true value of β which change from period to period as experience with the drug accumulates. These beliefs are assumed to be represented by a probability distribution, which is revised in each period as information is received. We assume that there are n numbers which could possibly be the true value of β . That is, there exists a set $\{\beta_1, \beta_2, \dots, \beta_n\}$ of possible β values each of which has positive probability in the "prior" probability distribution that represents the consumer's initial beliefs. This assumption is made strictly for expositional convenience. The results obtained hold for a more general class of prior measures which are concentrated on a compact set. We will return to this point several times in the exposition which follows.

It is hypothesized that the consumer behaves as though he were a Bayesian statistician who uses Bayes' Rule to revise his expectations when new information is received. If the consumer buys y_i units of the drug in period i and observes a health level z_i in that period, then his experience in period i is summarized by the vector $w_i = (z_i, y_i)$. If we let f_i^t be the probability that $\beta = \beta_i$ given the information available at time t , and let

$$f^t = (f_1^t, f_2^t, \dots, f_n^t),$$

then f^t is the posterior mass function which represents the consumer's beliefs in period t , before the random vector \tilde{w}_t is observed. The posterior f^t reflects his previous experience as well as the a-priori mass function f^0 which describes his initial beliefs about drug quality. We let g_t be the function (implied by Bayes' Rule) which relates f^t to experience and to f^0 . The experience acquired up to and including period t is summarized by the vector $w^t \equiv (w_0, w_1, \dots, w_t)$. Let $z^t \equiv (z_0, z_1, \dots, z_t)$ and $y^t \equiv (y_0, y_1, \dots, y_t)$. We will sometimes write $w^t = (z^t, y^t)$. The set of possible w^t vectors is denoted by W^t . Thus we can express Bayes' Rule as $f^t = g_t(f^0; w^{t-1})$. Then

$$f^t = g_1(f^{t-1}; w_{t-1}).$$

In making his drug and goods purchase decision in each period the consumer is assumed to maximize expected utility. In period t , the utility of (x_t, z_t) is $u(x_t, z_t)$, where u is a strictly concave utility function with positive marginal utility of both goods. Since the utility function u is invariant over time, the consumer's preferences for *health* and other goods are stationary. In spite of this stationarity, the consumer's preferences for *drugs* and goods are not invariant over time. These preferences are represented by an expected utility function that varies with beliefs formed on the basis of experience. In period t , this expected utility function is

$$\begin{aligned} U(x_t, y_t | f^t) &= \sum_{i=1}^n \int u(x_t, \alpha + \beta_i y_t + \varepsilon_t) h(\varepsilon_t | 0) f_i^t d\varepsilon_t \\ &= \int u(x_t, z_t) m(z_t | f^t; y_t) dz_t, \end{aligned} \quad \dots(1)$$

where $h(\cdot | \mu)$ is the normal density function with mean μ and variance 1 and $m(\cdot | f^t; y_t)$ is the predictive density defined by

$$m(z_t | f^t; y_t) = \sum_{i=1}^n h(z_t | \alpha + \beta_i y_t) f_i^t.$$

At the beginning of period t the consumer has observed some realization of \tilde{w}^{t-1} , say w^{t-1} . Given this information he makes a decision about how much x_t and y_t to consume. Each period the consumer faces the budget constraint

$$p y_t + x_t = I, \quad \dots(2)$$

where p and I are respectively the (positive) price of drugs and the consumer's (positive) income.

In each period t , the consumer can, therefore, be viewed as choosing y_t , and setting $x_t = I - py_t$. The y_t level chosen must, of course, satisfy the restriction

$$0 \leq y_t \leq (I/p). \quad \dots(3)$$

In addition, the y_t choice will be influenced by w^{t-1} , the experience accumulated to time t .

Expected utility at time t , $U(I - py_t, y_t | f^t)$ is influenced by current beliefs as represented by f^t . But $f^t = g_t(f^0, w^{t-1})$, where $w^{t-1} = (z^{t-1}, y^{t-1})$. This means that the process of information assimilation, captured in the function g_t , introduces an intertemporal element to the problem of choosing an optimal drug purchase. Intuitively, the beliefs which provide the basis for preference formation in period t are arrived at by interpreting previously observed health levels z^{t-1} in the light of earlier drug consumption choices y^{t-1} . In particular, earlier drug consumption levels determine the extent to which observed health levels can be relied on as evidence about β . The consumer's problem, then, is to choose a level of drug demand which attains an optimal balance between informational gains that accrue later and current health gains.

The possibilities for experimentation in this model can be made explicit if we state the problem in the framework of dynamic programming. To do this let $V^T(f)$ be defined as

$$V^T(f) \equiv \max_{y_0, \{y_t\}_{t=1}^T} E[U(I - py_0, y_0 | f) + \sum_{t=1}^T \delta^t U(I - p\tilde{y}_t, \tilde{y}_t | g_t(f; \tilde{w}^{t-1}))], \quad \dots(4)$$

where $y_0 \in [0, I/p]$, $\tilde{y}_t: G_t \rightarrow [0, I/p]$, $G_t \equiv \{f_t: f_t = g_t(f; w^{t-1}); w^t \in W^t\} \equiv$ the set of possible posteriors. We will give a more explicit form for (4) just before the proof of Lemma 2. Lemma 3 proves that $V^T(f)$ is well defined. At this point, it should be noted that, in the definition (4), the consumer's strategy at t is a drug consumption choice $y_t(f_t)$ which depends on his posterior at that date. Thus in (4) y_t refers to a realization of

$$\tilde{y}_t = y_t(g_t(f; \tilde{w}^{t-1}));$$

i.e. $y_t = y_t(g_t(f; w^{t-1}))$, where w^{t-1} is a realization of \tilde{w}^{t-1} . It follows from (4) that, for $T \geq 0$,

$$V^T(f) = \max_{0 \leq y \leq I/p} \{U(I - py, y | f) + \delta E[V^{T-1}(g_1(f, w))]\}, \quad \dots(5)$$

and

$$V^T(f) \equiv 0 \quad \text{if} \quad T \leq 0.$$

In (5), we have omitted the subscript "zero" on y and z and the superscript "zero" on w . Thus in (5), $w = (z, y)$. We will continue to use the more convenient notation throughout the remainder of the paper. The expression $V^T(f)$ is the maximum future utility attainable when T periods remain and consumer beliefs are represented by f .

Note that

$$E\{V^{T-1}[g_1(f, w)]\} = \int_{-\infty}^{\infty} V^{T-1}[g_1(f; z, y)]m(z | f; y)dz. \quad \dots(6)$$

Equation (6) is an expression for the expected future utility when y units of the drug are consumed in period zero. In computing the expectation (6), the value $V^{T-1}[g_1(f; z, y)]$ associated with each posterior $g_1(f; z, y)$ is weighted by the probability $m(z | f; y)$. Define the function

$$H^T(y, f) \equiv E\{V^{T-1}[g_1(f; w)]\}, \quad \dots(7)$$

where it will be recalled that $w = (z, y)$. Using (7), equation (5) may be rewritten as

$$V^T(f) = \max_{0 \leq y \leq I/p} [U(I - py, y | f) + \delta H^T(y, f)]. \quad \dots(8)$$

An "experimenting" consumer will choose y to maximize

$$U(I - py, y | f) + \delta H^T(y, f). \quad \dots(9)$$

A maximizer, which may not be unique, is denoted by $y^T(f)$. Lemma 3 below will demonstrate that $y^T(f)$ exists by establishing that the function in (9) is continuous.

The experimental design aspects of the consumer's choice problem are apparent in (5) and its alternative expression (9). The first term in (5) and (9) measures the present utility of y units of drug consumption. The second term measures the expected future utility of the improvements in information made possible because y was chosen. In choosing $y^T(f)$ to maximize the sum of these utilities the consumer is, as stated above, arriving at an optimal balance between present utility and future information.

As we are interested in the effect of experimentation on consumption we propose to investigate the relationship between the (possibly non-unique) strategy $y^T(f)$ and the consumption strategy which would be optimal if the possibilities for learning from experience are non-existent or ignored. To facilitate this comparison we first study the optimal consumption decisions made by a consumer who assumes that his future beliefs will be unchanged by his current experience. Such a consumer will choose a sequence of drug consumption levels to solve the problem

$$\max_{\{y_t\}_{t=1}^T} \sum_{t=1}^T \delta^t U(I - py_t, y_t | f). \quad \dots(10)$$

The solution to this problem is obtained by choosing, in each period t , the consumption level which satisfies (3) and maximizes

$$U(I - py, y | f). \quad \dots(11)$$

The value y which maximizes (11), is denoted by $y^0(f)$, and is called the optimal non-experimental consumption policy. Under our assumptions $y^0(f)$ is unique.

To emphasize the difference between $y^0(f)$ and $y^T(f)$ recall that, from (9), the function which $y^T(f)$ maximizes is a sum of two terms. The first term represents expected current utility, while the second term measures the extent to which learning from current experience enables the consumer to increase his utility by making more informed future decisions. If we compare (9) with (11), it is seen that (11) is the first term in (9), i.e. (11) is the current expected utility of consumption. Thus $y^0(f)$, unlike $y^T(f)$, is chosen without regard to the effects of learning from experience.

Since z_t observations provide information about β when y_t is positive, the consumer can, in essence, produce information by consuming drugs. The "amount" of information he produces will depend on the amount of drugs he consumes. This relationship between information and drug consumption can be interpreted as a technology for information production in which the "input" is y_t . The "output" (measured in expected utility terms) can be interpreted as $\delta H^T(y, f)$. Of course an experimenting consumer who avails himself of this technology pays a price. Indeed, one can think of the cost of the information provided by y as $C(y, f)$, where

$$C(y, f) \equiv U[I - py^0(f), y^0(f) | f] - U(I - py, y | f).$$

$C(y, f) \geq 0$ because $y^0(f)$ maximizes (11). $C(y, f)$ gives the one-period cost of choosing a drug consumption which is designed to give more future information about β than the drug consumption which is non-experimentally optimal provides.

We can now prove that the possibility of experimentation causes consumers to buy more of the drug than they would otherwise; i.e. $y^T(f) \geq y^0(f)$. (Since $y^T(f)$ may not be unique, we must show that this inequality holds for all $y^T(f)$ which maximize (9).) We prove this by showing that the inequality holds precisely because larger drug consumptions lead to more "informative" experiments in the sense of Blackwell [2]. (In Blackwell's terminology, observation of the consumer's health level when a large amount of the drug has been consumed is an experiment which is sufficient for the observation of health when small amounts of the drug have been used.) The most complicated part of the proof involves showing that this fact implies that $H^T(y, f)$ is an increasing function of y —Theorem 1. Theorem 2 uses Theorem 1 to show that $y^T(f) \geq y^0(f)$.

These theorems and their proofs use the notation appropriate to the case in which the prior measure is finite-discrete, i.e. concentrated on a finite number of possible values. As mentioned earlier the proofs of these theorems can be extended to a more general class of cases in which the support of the prior is contained in a compact set.

Theorem 1. $H^T(y, f)$ is a non-decreasing function of y for each f . We defer the proof of Theorem 1 until later.

Theorem 2. Assume that $\int_{-\infty}^{\infty} u(I, \alpha + \beta_i I/p + \varepsilon) h(\varepsilon | 0) d\varepsilon < \infty$, for $i = 1, 2, \dots, n$. Also suppose that a $y^T(f)$ exists. Then $y^T(f) \geq y^0(f)$.

Proof. The Lebesgue dominated convergence theorem implies that $U(x, y | f)$, defined in (1), is a continuous function of (x, y) for all (x, y) such that $py + x \leq I$, since

$$u(x, \alpha + \beta_i y + \varepsilon) \leq u(I, \alpha + \beta_i I/p + \varepsilon) \quad \text{for all } i = 1, 2, \dots, n.$$

Therefore $U(I - py, y | f)$ is a continuous function of y for $0 \leq y \leq I/p$. Thus a maximizer of (11) exists. The maximizer is unique for each f because $u(x, y)$ is assumed strictly concave in (x, y) . Hence $y^0(f)$ is well defined.

Suppose a $y^T(f)$ exists such that $y^T(f) < y^0(f)$. Since $y^0(f)$ is uniquely maximal

$$U[I - py^T(f), y^T(f) | f] < U[I - py^0(f), y^0(f) | f]. \quad \dots(12)$$

By Theorem 1

$$H^T[y^0(f), f] \geq H^T[y^T(f), f]. \quad \dots(13)$$

Adding (12) and (13) implies

$$U[I - py^T(f), y^T(f) | f] + \delta H^T[y^T(f), f] < U[I - py^0(f), y^0(f) | f] + \delta H^T[y^0(f), f]. \quad \dots(14)$$

But (14) contradicts the assumption that $y^T(f)$ is a maximizer of

$$U(I - py, y | f) + H^T(y, f). \quad \parallel$$

Theorem 1 is the key to Theorem 2. Theorem 1 states that more valuable information is provided by larger drug consumptions. The proof of this result is based on Blackwell's approach to the comparison of experiments, which we now digress to discuss. In the course of this discussion, we will make clear the formal meaning of the term "more informative" which has been used informally up to here.

As in the previous discussion, $h(z | \alpha + \beta y)$ is a normal density function of z with mean $\alpha + \beta y$ and variance 1. Define

$$k(z | \beta, y) \equiv h(z | \alpha + \beta y). \quad \dots(15)$$

Let $F \equiv [k(\cdot | \cdot, y) | y \in R]$. F is called a family of experiments. (If $k \in F$, $k' \in F$, then $k = k'$ if and only if $k = (z | \beta, y) = h(z | \alpha + \beta y)$, $k' = (z | \beta, y') = h(z | \alpha + \beta y')$ and $y = y'$.) An experiment k is *sufficient* for an experiment k' , if there exists a function $v(z' | z) \geq 0$ such that for all z' and all β

$$k(z' | \beta, y') = \int_{-\infty}^{\infty} v(z' | z) k(z | \beta, y) dz, \quad \dots(16)$$

and for all z

$$\int_{-\infty}^{\infty} v(z' | z) dz' = 1. \quad \dots(17)$$

Note that $v(z' | z)$ must not depend on β . To interpret this definition one might envisage an experimental apparatus k which yields observations z . Since the distributions of z and z' depends on β , k or k' can be used to learn about β . When (16) and (17) hold an observer who only has access to the apparatus k can reproduce the apparatus k' using $v(\cdot | \cdot)$. This can be done as follows. If the apparatus k yields an observation z , then draw

an observation z' from an urn for which the density of z' is given by $v(z' | z)$. If this procedure is followed, then z' will have a density given by $h(z' | \beta, y')$ which is exactly what the apparatus k' would have yielded. Note that $v(z' | z)$ must not depend on β . If it does, the observer who does not know β cannot construct the urn. We will sometimes refer to k as *more informative* than k' if k is sufficient for k' . See De Groot [5, p. 433] and Kihlstrom [10] for more details and references to the literature on Blackwell's sufficiency theory.

In the remainder of the paper, comparisons of experiments based on sufficiency are used to determine the relative value of alternative experiments. In Theorem 3, which follows, it is shown that if one experiment k is more informative than another k' , then k is more valuable to observers than k' .

As a preliminary to this theorem let η be a real valued function with domain $\{\beta_1, \dots, \beta_n\} \times \Delta$. The set Δ is interpreted as a set of possible decisions. A generic element of Δ is denoted by ξ .

Now let

$$V^*(f) \equiv \sup_{\xi \in \Delta} \sum_{i=1}^n \eta(\beta_i, \xi) f_i. \quad \dots(18)$$

As above,

$$g_1^i[f; (z, y)] \equiv \frac{f_i k(z | \beta_i, y)}{\sum_{j=1}^n f_j k(z | \beta_j, y)} \quad \dots(19)$$

and

$$g_1^i[f; (z', y')] \equiv \frac{f_i k(z' | \beta_i, y')}{\sum_{j=1}^n f_j k(z' | \beta_j, y')}, \quad \dots(20)$$

where f_i is the prior probability that β_i is the true value of β . The vector

$$g_1[f; (z, y)] = \{g_1^1[f; (z, y)], \dots, g_1^n[f; (z, y)]\}$$

is the posterior mass function that results when experiment k is run and z is observed. The mass function $g_1[f; (z', y')]$ is similarly defined.

Theorem 3. *If k is sufficient for k' , then*

$$EV^*\{g_1[f; (z, y)]\} \geq EV^*\{g_1[f; (z', y')]\},$$

where

$$EV^*\{g_1[f; (z, y)]\} = \int_{-\infty}^{\infty} V^*\{g_1[f; (z, y)]\} \sum_{i=1}^n h(z | \beta_i, y) f_i dz$$

and

$$EV^*\{g_1[f; (z', y')]\} = \int_{-\infty}^{\infty} V^*\{g_1[f; (z', y')]\} \sum_{i=1}^n h(z' | \beta_i, y') f_i dz'.$$

Theorem 3 shows that a "more informative" experiment is more valuable. That is, if a decision maker's choices depend on his beliefs about β , he will achieve higher expected utility if his beliefs are formed on the basis of a more informative experiment. Marschak and Miyasawa [12] and Blackwell [3] also prove this result.

The following Lemma, which states a result about sufficiency for the family F , is used to prove Theorem 1.

Lemma 1. *If $y > y' > 0$, then k is sufficient for k' .*

Proof. From (15), $k(z | \beta, y) = h(z | \alpha + \beta y)$ and $k(z' | \beta, y') = h(z' | \alpha + \beta y')$. Consider two random variables \tilde{z} and \tilde{z}' which are jointly normally distributed such that

$$E\tilde{z} \equiv \alpha + y, E\tilde{z}' \equiv \alpha + y', \text{var}(\tilde{z}) \equiv \text{var}(\tilde{z}') = 1,$$

and covariance $(\tilde{z}, \tilde{z}') \equiv y'/y$. From elementary normal distribution theory, the conditional density of \tilde{z}' given z , $v^*(z' | z)$, is of the normal form with mean

$$E[\tilde{z}' | z] = E\tilde{z}' + y'(z - E\tilde{z})/y = (1 - (y'/y))\alpha + y'z/y \quad \dots(21a)$$

and variance

$$\text{var} [\tilde{z}' | z] = 1 - (y'/y)^2. \quad \dots(21b)$$

By assumption $0 < y' < y$, so $\text{var} [\tilde{z}' | z] > 0$. Thus the joint distribution of (\tilde{z}, \tilde{z}') is well defined, and the conditional density of \tilde{z}' given z is independent of β by (21). By definition of \tilde{z} and \tilde{z}' , k is the marginal density of \tilde{z} and k' is the marginal density of \tilde{z}' . Since the marginal density of \tilde{z}' is equal to the conditional density of \tilde{z}' given z averaged by the density of z we have

$$h(z' | \alpha + \beta y') = \int_{-\infty}^{\infty} v^*(z' | z) h(z | \alpha + \beta y) dz.$$

Since v^* is a density of \tilde{z}' given z , $v^* \geq 0$ and $\int_{-\infty}^{\infty} v^*(z' | z) dz' = 1$ for all z . Therefore $v^*(\cdot | \cdot)$ is a function that satisfies (16) and (17).

Remark. The following outline of an alternative proof is suggestive of the motivation for Lemma 1. Note that, when α is known, observing

$$z = \alpha + \beta y + \varepsilon \quad (z' = \alpha + \beta y' + \varepsilon)$$

is equivalent to observing

$$r = (z - \alpha)/y = \beta + (\varepsilon/y) \quad [r' = (z' - \alpha)/y' = \beta + (\varepsilon/y')]$$

which is normal with mean β and variance $\sigma^2 = y^{-2} [(\sigma')^2 = (y')^{-2}]$. Now define \tilde{r}'' to be a normal random variable (independent of \tilde{r} and \tilde{r}') with mean β and variance $(\sigma'')^2 = \{[1/\sigma^2] - [1/(\sigma')^2]\}^{-1}$. It is obvious and easy to prove that joint observation of both independent variables \tilde{r}' and \tilde{r}'' is sufficient for observations of the single variable \tilde{r} . It is also relatively easy to demonstrate that the random variable

$$\tilde{r}^* = \left[\frac{1}{(\sigma')^2} \tilde{r}' + \frac{1}{(\sigma'')^2} \tilde{r}'' \right] \left[\frac{1}{(\sigma')^2} + \frac{1}{(\sigma'')^2} \right]^{-1}$$

is sufficient for joint observation of \tilde{r}' and \tilde{r}'' . But \tilde{r}^* is a normal random variable with mean β and variance

$$\left[\frac{1}{(\sigma')^2} + \frac{1}{(\sigma'')^2} \right]^{-1} = \left[\frac{1}{(\sigma')^2} + \left[\frac{1}{\sigma^2} - \frac{1}{(\sigma')^2} \right] \right]^{-1} = \sigma^2.$$

Thus \tilde{r}^* has the same distribution as \tilde{r} . As a consequence, observations of \tilde{r} are sufficient for observations of \tilde{r}^* . Since Blackwell has demonstrated that sufficiency is a transitive relation, the experiment "observe \tilde{r} " must therefore be sufficient for the experiment "observe \tilde{r}^* " and k is sufficient for k' .

Lemma 1 states that larger drug consumptions lead to more informative experiments. More informative experiments are more valuable by Theorem 3. So we would expect that $H^T(y, f)$ is increasing in y because H^T gives the informational value of changes in drug consumption. Lemma 2, below, is the key to the proof of this fact. However, it is necessary to introduce some notation in order to specify the consumer's maximum problem (5) in more detail. This we now do.

As above, we denote the range of the random variable \tilde{w}^{t-1} , $t \geq 1$, by W^{t-1} . For $t \geq 1$, let γ_t denote a function with domain W^{t-1} and range $[0, I/p]$; i.e. $\gamma_t: W^{t-1} \rightarrow [0, I/p]$. Denote the set of possible γ_t 's by

$$\Gamma_t = \{\gamma_t: W^{t-1} \rightarrow [0, I/p]\}.$$

Now let T be a finite integer, to be interpreted as the consumer's horizon. Denote any sequence $\{\gamma_t\}_{t=1}^T$ by γ^T and let Γ^T represent the set of possible γ^T 's; i.e.

$$\Gamma^T = \prod_{t=1}^T \Gamma_t.$$

For most of the following discussion we will suppress the dependence of Γ^T and γ^T on T , and write Γ and γ instead.

Once γ has been chosen, we can define new functions ζ_t ; $t = 0, \dots, T$; such that

$$\zeta_0: \Gamma \times [0, I/p] \rightarrow [0, I/p],$$

and

$$\zeta_t: \Gamma \times R_t \times [0, I/p] \rightarrow [0, I/p],$$

where R_t is t -dimensional Euclidean space.

We do this by letting

$$\zeta_0(\gamma; y_0) = y_0,$$

$$\zeta_1(\gamma; z_0, y_0) = \gamma_1(z_0, y_0),$$

and

$$\zeta_t(\gamma; z^{t-1}, y_0) = \gamma_t[z^{t-1}, \zeta^{t-1}(\gamma; z^{t-2}, y_0)]$$

where

$$\zeta^{t-1}(\gamma; z^{t-2}, y_0) = [\zeta_0(\gamma; y_0), \dots, \zeta_{t-1}(\gamma; z^{t-2}, y_0)] \quad \text{for } t = 2, \dots, T.$$

Intuitively, γ_t is a strategy or contingency plan for period t that specifies the consumer's choice of y_t for each "experience vector" w^{t-1} which might possibly be observed. Note that the vector of previous decisions is an argument of the function γ_t . But each previous decision, except the first, is chosen as a function of previous experience. Thus the decision y_t is ultimately dependent only on the previously observed health levels z^{t-1} and the first decision y_0 . Of course, the way it depends on these variables is determined by γ . The functions ζ_t express the functional relationship between y_t and the variables z^{t-1} , y_0 and γ .

If we now let $\lambda_t = (z^{t-1}, y_0)$ for $t = 1, \dots, T$ and $\lambda_0 = y_0$, then the consumer's problem is choose $y_0 \in [0, I/p]$ and a sequence $\gamma \in \Gamma$ to maximize the expected utility

$$\begin{aligned} & \int_{R_{T+1}} \sum_{i=1}^n \sum_{t=0}^T \delta^t u[I - p\zeta_t(\gamma; \lambda_t), z_t] f_i^0 \prod_{\tau=0}^t h[z_\tau | \alpha + \beta_i \zeta_\tau(\gamma; \lambda_\tau)] dz^T \\ & = \sum_{t=0}^T \delta^t \int_{R_{t+1}} \sum_{i=1}^n u[I - p\zeta_t(\gamma; \lambda_t), z_t] f_i^0 \prod_{\tau=0}^t h[z_\tau | \alpha + \beta_i \zeta_\tau(\gamma; \lambda_\tau)] dz^t. \quad \dots(22) \end{aligned}$$

This expected utility can be simplified in a useful way by substituting the expression for the posterior f^t , when $t \geq 1$. Specifically,

$$f_i^t = g_i^t(f^0; w^{t-1}) = \frac{f_i^0 \prod_{\tau=0}^{t-1} h(z_\tau | \alpha + \beta_i y_\tau)}{\sum_{i=1}^n f_i^0 \prod_{\tau=0}^{t-1} h(z_\tau | \alpha + \beta_i y_\tau)}.$$

If $y_\tau = \gamma_\tau(w^{\tau-1})$ for $\tau = 1, \dots, t$; then $w^{t-1} = [z^{t-1}, \zeta^{t-1}(\gamma; z^{t-2}, y_0)]$, and

$$\begin{aligned} f_i^t &= g_i^t\{f^0; [z^{t-1}, \zeta^{t-1}(\gamma; z^{t-2}, y_0)]\} \\ &= \frac{f_i^0 \prod_{\tau=0}^{t-1} h(z_\tau | \alpha + \beta_i \zeta_\tau(\gamma; \lambda_\tau))}{\mu(z^{t-1} | \gamma, y_0)}, \quad \dots(23) \end{aligned}$$

where

$$\mu(z^{t-1} | \gamma, y_0) = \sum_{i=1}^n f_i^0 \prod_{\tau=0}^{t-1} h[z_\tau | \alpha + \beta_i \zeta_\tau(\gamma; \lambda_\tau)].$$

(Of course, μ should be subscripted by $t-1$. This is omitted to simplify the notation which is already somewhat cumbersome.)

Substituting the expression (23) in the right side of (22) and using (1), we obtain

$$\begin{aligned} U(I - py_0, y_0 | f^0) + \sum_{t=1}^T \delta^t \int_{R_t} & U[I - p\zeta_t(\gamma; \lambda_t), \zeta_t(\gamma; \lambda_t) | g_t\{f^0; [z^{t-1}, \zeta^{t-1}(\gamma; \lambda_{t-1})]\}] \\ & \times \mu(z^{t-1} | \gamma, y_0) dz^{t-1} \end{aligned}$$

as the expression which the consumer wishes to maximize by his choice of y_0 and γ . To simplify notation, the above expression will often be written simply as

$$U(I - py_0, y_0 | f^0) + \sum_{t=1}^T \delta^t \int_{R_t} U(I - py_t, y_t | f^t) \mu(z^{t-1} | \gamma, y_0) dz^{t-1}. \quad \dots(22')$$

This is clearly equivalent to the expression being maximized in (5). The reader should be aware, however, that when this simplified notation is used y_t is always equal to $\zeta_t(\gamma; \lambda_t)$ and $f^t = g_t\{f^0; [z^{t-1}, \zeta^{t-1}(\gamma; \lambda_{t-1})]\}$.

Define

$$\phi(y, \beta) \equiv \int_{-\infty}^{\infty} u(I - py, \alpha + \beta y + \varepsilon) h(\varepsilon | 0) d\varepsilon.$$

Then, from (1)

$$U(I - py, y | f) = \sum_{i=1}^n \phi(y, \beta_i) f_i.$$

Thus by (3), (4) and the definition of V^T

$$V^T(f^0) = \max_{(y_0, \gamma) \in [0, I/p] \times \Gamma} \left[\sum_{i=1}^n \phi(y_0, \beta_i) f_i^0 + \sum_{t=1}^T \delta^t \int_{R_t} \sum_{i=1}^n \phi(y_t, \beta_i) f_i^t \mu(z^{t-1} | \gamma, y_0) dz^{t-1} \right]. \quad \dots(24)$$

It should be recalled that in (24) $y_t = \zeta_t(\gamma; z^{t-1}, y_0)$ and

$$f^t = g_t\{f^0; [z^{t-1}, \zeta^{t-1}(\gamma; z^{t-2}, y_0)]\}, \quad \text{for } t \geq 1.$$

The above notation can now be used to prove the following lemma, from which the main theorem can be proved immediately. Again it should be emphasized that the proof does not depend crucially on the assumption of a discrete prior.

Lemma 2. *If $V^T(f)$ is well defined, then there exists a decision set Δ and a function $\eta_T: \{\beta_1, \dots, \beta_n\} \times \Delta \rightarrow R_1$ such that*

$$V^T(f) = \max_{\xi \in \Delta} \sum_{i=1}^n \eta_T(\beta_i, \xi) f_i. \quad \dots(25)$$

Proof. Substitute equation (23) in (24) to obtain

$$V^T(f^0) = \max_{(y_0, \gamma) \in [0, I/p] \times \Gamma} \left[\sum_{i=1}^n \phi(y_0, \beta_i) f_i^0 + \sum_{t=1}^T \delta^t \int_{R_t} \sum_{i=1}^n \phi(y_t, \beta_i) f_i^t \prod_{\tau=0}^{t-1} h(z_\tau | \alpha + \beta_i y_\tau) dz^\tau \right], \quad \dots(26)$$

where $y_t = \zeta_t(\gamma; \lambda_t)$ and $y_\tau = \zeta_\tau(\gamma; \lambda_\tau)$. If we now let

$$\Delta = [0, (I/p)] \times \Gamma \quad \text{and} \quad \xi = (y_0, \gamma)$$

we can define η_T by

$$\eta_T(\xi, \beta_i) = \left[\phi(y_0, \beta_i) + \sum_{t=1}^T \delta^t \int_{R_t} \phi(y_t, \beta_i) \prod_{\tau=0}^{t-1} h(z_\tau | \alpha + \beta_i y_\tau) dz^\tau \right],$$

where, as above, $y_t = \zeta_t(\gamma; \lambda_t)$ and $y_\tau = \zeta_\tau(\gamma; \lambda_\tau)$. Interchanging the order of summation and integration in (26) then yields

$$V^T(f^0) = \max_{\xi \in \Delta} \sum_{i=1}^n \eta_T(\xi, \beta_i) f_i^0. \quad \parallel$$

Theorem 1 is an immediate consequence of Theorem 3 and Lemmas 1 and 2.

Proof of Theorem 1. From (6) and (7)

$$H^T(y, f) \equiv \int_{-\infty}^{\infty} V^{T-1}[g_1(f; z, y)]m(z | f; y)dz. \quad \dots(27)$$

Suppose $y > y'$, let $k = k(z | \beta_i, y)$ and $k' = k(z' | \beta_i, y')$ then by Lemma 1, k is sufficient for k' . Using Lemma 2, we may set $V^*(f) \equiv V^{T-1}(f)$. Then by Theorem 3, if g_1 and g'_1 are defined as in (19) and (20) then $EV^*\{g_1[f; (z, y)]\} \geq EV^*\{g'_1(f; (z', y'))\}$. Recalling that $m(z | f; y) = \sum_{i=1}^n f_i h(z | \beta_i, y)$,

$$H^T(y, f) = EV^{T-1}\{g_1[f; (z, y)]\} \geq EV^{T-1}\{g_1[f; (z', y')]\} = H^T(y', f). \quad \parallel$$

Prescott [13] proved a similar theorem under the assumption of a normal prior on β . All of the results proved in this paper can also be shown to hold for any prior on β as long as $V^T(f)$ is well defined. Lemma 3, which follows, uses the assumption that the prior on β is finite-discrete to prove continuity of $H^T(y, f)$, existence of $y^T(f)$ and well definition $V^T(f)$. This is the only step in any of the arguments of this paper in which the discreteness assumption is actually used. The argument used to prove Lemma 3 can be employed to prove the same results when the prior probability distribution is concentrated on a compact set of possible β values and when the prior is either continuous or discrete with mass concentrated on a countable set. The adaptations required in the proof are discussed in the appendix.

Lemma 3. Let $\beta^* = \max \{\beta_1, \beta_2, \dots, \beta_n\}$. Assume that

$$M_0 \equiv \int_{-\infty}^{\infty} u(I, \alpha + (\beta^* I/p) + \varepsilon) h(\varepsilon | 0) d\varepsilon < \infty.$$

Then $H^T(y, f)$ is a continuous function of y for all (y, f) such that $0 \leq y \leq I/p$ and $f = (f_1, f_2, \dots, f_n)$, with $0 < f_i < 1$ for all i . Further $y^T(f)$ exists and $V^T(f)$ is well defined.

Proof. We prove by induction that, when T is finite, $H^T(y, f)$ is continuous in y , $y^T(f)$ exists and $V^T(f)$ is well defined.

In the proof of Theorem 2, it was shown that $U(I - py, y | f)$ is a continuous function of y . Thus $y^0(f)$ exists and $V^0(f) \equiv \max_{0 \leq y \leq (I/p)} U(I - py, y | f)$ is well defined. By (5)-(7),

$$H^1(y, f) \equiv \int_{-\infty}^{\infty} V^0[g_1(f; z, y)]m(z | f; y)dz. \quad \dots(28)$$

Since $\phi(y, \beta) \leq M_0$, for all (y, β) , $U(I - py, y | f) \leq M_0$, for all $y \in [0, I/p]$, and

$$V^0[g_1(f; z, y)] \leq M_0.$$

Thus

$$V^0[g_1(f; z, y)]m(z | f; y) \leq M_0 \psi(z), \quad \dots(29)$$

where

$$\psi(z) = \begin{cases} h(z | \alpha), & \text{if } z \leq \alpha, \\ \frac{1}{2\pi}, & \text{if } \alpha \leq z \leq \alpha + (\beta^* I/p), \\ h(z | \alpha + (\beta^* I/p)), & \text{if } \alpha + (\beta^* I/p) \leq z. \end{cases}$$

The bounding function $\psi(z)$ is shown in Figure 1. The right-hand side of (29) integrates to

$$M_0(1 + (1/2\pi)(\beta^* I/p)) < \infty.$$

Therefore (28) and the Lebesgue dominated convergence theorem imply that $H^1(y, f)$ is continuous in y , for $y \in [0, I/p]$.




FIGURE 1

Suppose now that $H^T(t, f)$ is continuous in y . We will now show that $y^T(f)$ exists, that $V^T(f)$ is well defined, and that H^{T+1} is continuous in y . From (6) and (7)

$$H^{T+1}(y, f) \equiv EV^T[g_1(f; z, y)]. \quad \dots(30)$$

As $H^T(y, f)$ is continuous in y , $y^T(f)$ exists and $V^T(f)$ is well defined by (8). By Lemma 2, there exists an $\eta_T(\beta, \xi)$, and Δ such that

$$V^T(f) = \max_{\xi \in \Delta} \sum_{i=1}^n \eta_T(\beta_i, \xi) f_i. \quad \dots(31)$$

It is immediate that $V^T(f)$ is a convex function of f . Hence it is continuous on open convex sets. Without loss of generality, it can be assumed that f is in the open convex set $\bar{S} = [(g_1, \dots, g_n): g_i > 0, \sum_{i=1}^n g_i = 1]$. Our assumptions guarantee that the probability of $g_1(f; z, y) \in \bar{S}$ is one. So $V^T(\cdot)$ is continuous at every possible posterior. Since $g_1(f; z, y)$ is a continuous function of y and z , $V^T[g_1(f; z, y)]$ is also continuous in y . Now by (31)

$$V^T(f) \leq \sum_{\tau=0}^T \delta^\tau M_0 \equiv M_T.$$

Therefore

$$V^T[g_1(f; z, y)]m(z | f; y) \leq M_T \psi(z). \quad \dots(32)$$

The right-hand side of (32) integrates to $M_T[1 + (1/2\pi)\beta^*(I/p)] < \infty$. The left-hand side is a continuous function of y and integrates to $H^{T+1}(y, f)$. Therefore by the Lebesgue dominated convergence theorem $H^{T+1}(y, f)$ is a continuous function of y for all (y, f) such that $0 \leq y \leq (I/p)$ and f in \bar{S} . \parallel

3. MONOPOLISTS WHO EXPERIMENT TO LEARN THEIR DEMAND CURVES: A REINTERPRETATION

The statistical model described and the theorems proved above can be reinterpreted to analyse experimental behaviour on the part of a monopolist who does not know the slope of his demand curve. In carrying out this reinterpretation, the monopolist can be viewed as choosing either price or quantity. If price is chosen, a stochastic demand curve determines the demand which that price calls forth. When quantity is chosen, the price at which that quantity can be sold is determined by the stochastic demand curve. Interestingly, the conclusions obtained in these two alternative interpretations appear to be contradictory. But as we shall show these interpretations are based on different statistical models, so the paradox is indeed only apparent.

First consider the case where quantity is the decision variable. Then we let

$$z_t = \text{price in period } t,$$

and

$$y_t = \text{quantity in period } t.$$

The demand curve in period t is

$$z_t = \alpha + \beta y_t + \varepsilon_t, \quad \dots(33)$$

where, as above, $h(\varepsilon_t | 0)$ is the density of ε_t . Also as above, α is known but the slope of the demand curve, β , is unknown. The firm's beliefs about β at time t are described by the mass function f^t which assigns positive probability to elements of a finite set $\{\beta_1, \dots, \beta_n\}$. In this case, all β_i are negative. (It should be noted that this model has one unfortunate feature; the demand function (33) permits negative prices to occur with positive probability.)

The firm's profit is, of course

$$z_t y_t - c(y_t),$$

where c is a cost function, with $c' > 0$ and $c'' > 0$. The firm is assumed to maximize the expected utility of profit. The utility function, \hat{u} , is assumed to have $\hat{u}' > 0$ and $\hat{u}'' < 0$.

A firm with horizon T can then be viewed as choosing (y_0, γ) to maximize (4) with $\hat{u}[z_t y_t - c(y_t)]$ replacing $u(I - py_t, z_t)$ in the computation of (4). In this interpretation, y_0 is the quantity supplied in period 0, and γ_t is a strategy for choosing the supply level in period t contingent on the demand experience prior to t .

The non-experimenting monopolist chooses y_t to maximize

$$\hat{U}(y_t | f^t) = \int_{-\infty}^{+\infty} \hat{u}[y_t z_t - c(y_t)] m(z_t | f^t; y_t) dz_t. \quad \dots(34)$$

This function is strictly concave under the assumptions made about \hat{u} and c . Thus if a solution $y^0(f^t)$ exists, it is unique. Since the set of possible y levels is $[0, \infty)$, it is not compact and the continuity of \hat{U} does not guarantee the existence of a maximum, however. The same difficulty frustrates attempts to prove Lemma 3 for this example. If, however, $y^T(f)$ and $y^0(f)$ exist, we can define

$$\phi(y, \beta) = \int_{-\infty}^{+\infty} \hat{u}[y(\alpha + \beta y + \varepsilon) - c(y)] h(\varepsilon | 0) d\varepsilon$$

and apply Lemma 2 to prove Theorems 1 and 2. Theorem 2 asserts that the experimenting monopolist never chooses to supply less than the non-experimenting monopolist. As a result, the average price, $\alpha + \beta y^T(f)$, paid to an experimenting monopolist will be lower than $\alpha + \beta y^0(f)$, the average price paid in markets supplied by a non-experimenting monopolist.

If price is the decision variable, then

$$z_t = \text{the demand at time } t,$$

and

$$y_t = \text{the price at time } t.$$

With this reinterpretation, the demand curve is (33). Again the set of β 's which occur with positive probability is assumed to contain only negative numbers. In this case profit at time t is

$$y_t z_t - c(z_t).$$

As above, the utility function, \hat{u} , is assumed to be strictly concave and exhibit a positive marginal utility of income. The utility of profit used in computing (4) is $\hat{u}[y_t z_t - c(z_t)]$, i.e. in (4) $U(I - py, y | g)$ is replaced by $\int_{-\infty}^{\infty} \hat{u}(yz - c(z)) m(z | g; y) dz$. The decision variable y_0 is the price chosen in period zero; γ_t is the pricing strategy for period t , which is again contingent on previous demand experience.

The expected utility function (34) maximized by the non-experimenting monopolist is now computed with $\hat{u}[y_t z_t - c(z_t)]$ replacing $\hat{u}[z_t y_t - c(y_t)]$. Again the strict concavity and convexity assumptions made about \hat{u} and c respectively, guarantee that (34) is strictly concave and that $y^0(f)$ is unique if it exists. The existence question remains open. If $y^0(f)$ and $y^T(f)$ exist, Lemma 2 and Theorems 1 and 2 imply $y^0(f) \leq y^T(f)$. Thus when price is the decision variable, the price charged by an experimenting monopolist is higher than that charged by a non-experimenter with the same initial beliefs. This contrasts with the situation that results when quantity is chosen. In that case, the average price of an experimenter is lower than the average price charged when experimentation is absent.

These apparently contradictory results are not logically inconsistent. To see why, consider the case in which z_t is price and y_t is output. Note that the demand curve (33) can be transformed to obtain a demand curve

$$y_t = \frac{\alpha}{\beta} + \frac{1}{\beta} z_t + \frac{\varepsilon_t}{\beta}. \quad \dots(35)$$

We have just shown that (33) implies a higher average price with experimentation than without. However, since the variance of $(\varepsilon_t/\beta) = (1/\beta^2)$, and β is unknown, the demand curve (35) fails to satisfy the assumptions which led us to the conclusion that the experimental price exceeds the non-experimental price.

4. CONCLUSIONS

Uncertainty is pervasive, but it can be reduced at a cost. Rather than assuming some ad-hoc cost of information we have modelled a process of endogenous information generation. The cost of information turns out to be the utility that must be foregone by the choice of a larger control variable than would be optimal given current information and ignoring experimentation. This model provides a statistical foundation for the ad-hoc process which Arrow assumed and called "learning by doing". The model is also consistent with the empirical observations which motivated his paper: that productivity increases with experience. This happens in our model of the consumer because more informed consumers choose better combinations of risky drugs and non-risky goods for consumption.

The model in this paper is one of intertemporal optimization for a consumer who takes prices as given. Grossman [6] and [7] analysed a competitive equilibrium model of a market where firms are uncertain about the productivity of an input. In the context of a rational expectations model he derived an algorithm which could be used to characterize the path of equilibrium price random variables as well as the optimal input policy for firms. An equilibrium version of our model could similarly be analysed using Grossman's approach. This is left for future work.

APPENDIX

The proof of Lemma 3 uses the assumption of a finite-discrete prior in two ways. First, this assumption guarantees that the posterior distribution is an element of a finite dimensional simplex. Thus the convex function V^T has its domain in a finite dimensional space. Because its domain is finite dimensional, V^T is easily shown to be a continuous function of β on open convex sets. But continuity on open convex sets is a property which convex functions may not possess when they are defined on infinite dimensional sets. If the prior on β is not discrete, then the domain of V^T is indeed infinite dimensional and a more difficult proof is required to establish continuity. This is done below under the assumption that the prior has compact support.

The finite discreteness of the prior was also used to obtain equation (19), Bayes' Law, which implies that $g_1(f; z, y)$ is a continuous function of y and z . When the prior is either continuous or concentrated on a countable set, Bayes' Law can be expressed by an equation analogous to (19) which again implies continuity in y and z .

We now establish the continuity of V^T under the assumption that the prior has its support on the compact interval $[a, b]$.

Let $M([a, b])$ be the set of finite measures over the measure space defined by $[a, b]$ and its Borel sets. (A measure ν on $[a, b]$ is finite, if $\nu([a, b]) < \infty$.) This set includes but is larger than the set of probability measures over $[a, b]$ and its Borel sets. Using natural definitions of addition and vector multiplication and using the Prohorov-metric topology, i.e. the topology of weak convergence, $M([a, b])$ can be shown to be a convex topological vector space. (The Prohorov metric is discussed in Hildenbrand [8]. A topological vector space is defined in Choquet [4].)

The function V^T can be defined on $M[a, b]$ by an expression analogous to (25). Specifically, if $\nu \in M[a, b]$,

$$V^T(\nu) = \sup_{\xi \in \Delta} \int \eta_T(\beta, \xi) \nu(d\beta). \quad \dots(A.1)$$

The function V^T is easily seen to be convex on $M[a, b]$. The continuity of V^T follows from the following proposition.

Proposition. *If*

$$M_0 = \int_{-\infty}^{+\infty} u(I, \alpha + (bI/p) + \varepsilon) h(\varepsilon | 0) d\varepsilon < \infty, \quad \dots(A.2)$$

then V^T is continuous on $M[a, b]$.

Proof. Assumption (A.2) and the definition of V^T imply that for any $K < \infty$, $V^T(\nu)$ is bounded on the set of ν 's for which $\nu([a, b]) < K$. Thus condition (iii) of Proposition 19.9, p. 341 of Choquet [4], is satisfied, and this proposition implies that V^T is continuous on $M[a, b]$. ||

First version submitted July 1975; final version accepted October 1976 (Eds.).

This work was supported by National Science Foundation Grants GS-05317 and SOC74-11446 at the Institute for Mathematical Studies in the Social Sciences, Stanford University. The authors are grateful for these supports.

REFERENCES

- [1] Arrow, Kenneth J. "The Economic Implications of Learning by Doing", *Review of Economic Studies* (June 1962), 155-173.
- [2] Blackwell, David. "The Comparison of Experiments", *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability* (Berkeley: University of California Press, 1951, pp. 93-102).
- [3] Blackwell, David. "Equivalent Comparisons of Experiments", *Annals of Mathematical Statistics* (June 1953), 265-272.
- [4] Choquet, Gustave. *Lectures on Analysis*, Vol. I (New York: W. A. Benjamin, 1969).
- [5] De Groot, Morris H. *Optimal Statistical Decisions* (New York: McGraw-Hill, 1970).
- [6] Grossman, Sanford. "Rational Expectations and the Econometric Modeling of Markets Subject to Uncertainty: A Bayesian Approach", *Journal of Econometrics*, 3 (1975), 255-272.
- [7] Grossman, Sanford. "Equilibrium under Uncertainty, and Bayesian Adaptive Control Theory", in Day, R. and Groves, T. (Eds), *Adaptive Economic Models* (Academic Press, 1975).
- [8] Hildenbrand, Werner. *Core and Equilibria of a Large Economy* (Princeton: Princeton University Press, 1974).
- [9] Kihlstrom, Richard E. "A Bayesian Model of Demand for Information about Product Quality", *International Economic Review* (February 1974), 99-118.
- [10] Kihlstrom, Richard E. "A Bayesian Exposition of Blackwell's Theorem on the Comparison of Experiments" (Unpublished paper, State University of New York at Stony Brook, March 1974).
- [11] Kihlstrom, Richard E. "A General Theory of Demand for Information about Product Quality", *Journal of Economic Theory* (August 1974), 413-439.
- [12] Marschak, Jacob and Miyasawa, Koichi. "Economic Comparability of Information Systems", *International Economic Review* (June 1968), 137-174.
- [13] Prescott, Edward C. "The Multi-Period Control Problem Under Uncertainty", *Econometrica* (November 1972), 1043-1058.