NOTES ON EXPECTATIONS EQUILIBRIA IN BAYESIAN SETTINGS

by

Kenneth J. Arrow and Jerry R. Green

Working Paper No. 33
August 1973

Prepared Under
NATIONAL SCIENCE FOUNDATION GRANT GS-3269-A2

THE ECONOMICS SERIES
INSTITUTE FOR MATHEMATICAL STUDIES IN THE SOCIAL SCIENCES
Fourth Floor, Encina Hall
Stanford University
Stanford, California
94305

NOTES ON EXPECTATIONS EQUILIBRIA IN BAYESIAN SETTINGS*

by

Kenneth J. Arrow and Jerry R. Green

1. Introduction

In removing assumptions of perfect and complete information, economists have begun to confront problems of expectations formation. Although these have been acknowledged implicitly for a long time in econometrics and economic theory, the specific incorporation of seeking and using information into optimizing models has proven to be a difficult task. Still more troublesome has been the problem of formulating equilibrium models in which learning processes as well as the ordinary economic activities are being optimized simultaneously by all participants.

The first problem to be faced is that of defining the concept of an equilibrium in expectations. Several possibilities have been used in various contexts. Muth [1961] defined an expectation (i.e., a probability distribution describing belief about a real-valued random variable) to be fulfilled if it has the same mean as the observations on this variable. An equilibrium will then be a set of actions, optimal with respect to the expectations of all agents, such that everyone's expectations are fulfilled in this sense.

---

*We are indebted to David Starrett for pointing out an error in an earlier version. This work was supported in part by National Science Foundation Grant GS-3269-A2 to the Institute for Mathematical Studies in the Social Sciences, Stanford University and NSF Grant GS-31683 to Harvard University.

A second possibility is to think of expectations as subjective certainties at every future date. In this sense, an expectation pattern is a function which gives the value of the variable in question at every time in the future. This removes the static property of the previous definition, but also makes it impossible to consider beliefs that recognize subjective uncertainty at a given point in time. An equilibrium of this system, as studied by Brock [1972], is an expectations pattern which, if held by everyone, would generate a realization in the economy with which it is identical.

The third possibility, and the one we shall follow, is similar to the first but stronger. For equilibrium, it is required that the realized distribution be identical with the anticipated distribution, not only the same in mean. This definition was proposed by Lucas and Prescott [1971]. However, they assumed that expectations have this characteristic, whereas our concern will be to study the circumstances under which such an equilibrium is possible, its characteristics, and how it may be attained through a learning process by simultaneously maximizing agents whose initial beliefs are not in equilibrium.

This approach has been followed by Cyert and de Groot [1971] in an oligopoly model in which firms behave as Bayesian statisticians in search of the kink in their kinked demand curves. Grossman [1972] pursues a similar model of a perfectly competitive industry with a continuum of producers whose production functions are stochastic. Rothschild [1973] explores the question of consumers learning about prices for products sold at different locations, when the statistical distribution of prices is unknown. Rothschild [1971] treats the problem of setting prices when there is uncertainty about a stochastic demand function.

In any learning situation in a competitive model there are two conflicting forces at work. Each agent may be faced with a tradeoff between learning more quickly and operating in more advantageous ways during the learning process. But at the same time the learning processes themselves are interconnected because the actions of one agent influence the experiences, and hence the learning, of all others. In this paper we concentrate exclusively on the second problem.

As in the works cited above, we will proceed by using specific examples. But unlike the others (except Rothschild [1971]) our results are largely negative. To say what "negative" means in this context, one must say what types of results might be expected, or at least desired. This is discussed in the next section. Section 3 sets forth the general theoretical structure of the paper. Section 4 presents and discusses the examples which form our main conclusions.

2. Bayesian Inference in Competitive Situations

In standard statistical theory it is shown that a statistician who is uncertain about a parameter of a distribution will learn its true value through repeated applications of Bayes' Theorem. This basic result is true in the limit independent of the prior beliefs of the individual--as long as a neighborhood of the true parameter value had positive weight in the prior. Similarly, since the convergence of beliefs is almost-sure, it is (almost-surely) independent of the actual sample received. We shall call these properties <u>prior independence and sample independence</u> respectively.

In the model we study, as in those of others mentioned in the previous section, each agent acts as a Bayesian statistician. We might hope that results analogous to prior independence or sample independence could be obtained. That is, regardless of initial beliefs, posterior distributions would converge to an equilibrium for almost-every evolution of the economy, and therefore in the limit all agents would agree--and be correct--concerning the true economic environment. Further, if their criteria were identical and the environment symmetric, they would ultimately take identical actions and would be achieving identical results.

The examples we present in Section 4 show that these properties do not hold in general. In particular, equilibria may fail to exist; there may be great asymmetries in equilibrium situations; and these may be attained as equilibria of learning processes when the agents have different priors. Further, and not surprising in light of these possibilities, the prior of one agent may influence the learning of others through its effect on the actions of this agent which in turn affect the observations of everyone in the system.

We have no example, thus far, of sample dependence. However, it seems clear that this should also be expected in general.

3. Basic Model

In this section we present the structure of a general model designed to capture the problems mentioned above. The following section contains some special cases of this model in which we explore the equilibria and the learning process that leads to them.

We assume that there are n economic agents indexed by $i = 1,\ldots,n$. The agents choose an action at each stage of the process. The set of all possible actions for the i-th agent is denoted X_i, and a typical action is an element x_i of X_i. These actions jointly determine an outcome which may itself be a random variable. The outcome is a complete description of the state of the economy after the results of the actions taken and after the influence of whatever random parameters are relevant to that stage of the process are felt. Let Z be a space in which the random influences on the economy can be measured. Let Y_i be the set of all possible observations on the state of the economy for individual i. A typical observation for this individual is denoted by y_i.

We suppose that the welfare of the individual is affected by his action and his observation at the date in question, but not by his actions or observations at other dates.[1] Thus the utility received by i at any date is a function

$$u_i: X_i \times Y_i \to R .$$

[1] Formally of course this would be generalized in a straightforward fashion by considering each stage of the process to be a T-tuple of dates, where T is the longest lag in the process. The action $x_i(t)$ would then be a vector of the last T actions $x_i(t),\ldots,x_i(t-T+1)$.

Let $x = (x_1,\ldots,x_n)$, $y = (y_1,\ldots,y_n)$, $X = X_1 \times \ldots \times X_n$, and $Y = Y_1 \times \ldots \times Y_n$. The process can be described by a function

$$G: \; X \times Z \rightarrow Y \quad .$$

The interpretation of this is that the actions x and the random variable with value z produce a new state of the economy which is observed by the individuals as $y = G(x,z)$. Some of the observations may include knowledge about the actions taken by other individuals. In extreme cases it will be either the entire x vector, complete information, or none of it, complete ignorance. In this paper we will be primarily concerned with the latter case. Undoubtedly, intermediate cases are of great importance and we introduce this formalism with the hope that it will be explored at a future date.

We assume that the random parameter z is independent of x, and distributed independently and identically in each iteration. The distribution of z will be denoted μ, and the density of μ will be denoted by f. By the true model, we will mean the pair (G,μ).

The essential feature of the model is that individuals are not assumed to know the true model (G,μ). Each individual knows that he is one participant among others in the economy, and he has his own model of the system as it is relevant to him. The model is assumed to have as variables only those which he either observes or chooses directly. Unexplained variations in the observation are embodied in a personal random variable $z_i \in Z_i$. Thus the individual uses the information that he receives during the process to estimate the parameters of a <u>personal model</u>

$$y_i = g_i(x_i, z_i) \quad .$$

The random variable z_i is assumed by the individual to have a distribution μ_i parameterized by $\theta_i \in \Theta_i$. The learning that takes place concerns these parameters. Each individual has a prior at each date on θ_i, which may be described by a probability distribution with density ν_i on Θ_i. After receiving the observation y_i as a consequence of having taken the action x_i, he computes $g_i^{-1}(y_i;x_i) = z_i$ and applies Bayes theorem to this observation, which gives rise to a posterior distribution over Θ_i.

In order that this procedure be feasible, it is necessary that $g_i^{-1}(y_i;x_i) = \{z_i \in Z_i | g_i(x_i,y_i) = z_i\}$ be a one-element set for all observations y_i and actions x_i such that y_i is a possible observation given x_i. That is, we assume more formally the following <u>uniqueness criterion</u>. If $x_i \in X_i$, $y_i \in Y_i$, and $z \in Z$ satisfy

$$y_i = G_i(x,z)$$

for some $(x_1,\ldots,x_{i-1},x_{i+1},\ldots,x_n) \in X_1 \times \ldots \times X_{i-1} \times X_{i+1} \times \ldots \times X_n$ and $f(z) > 0$ or z is an atom of μ, then

$$\{z_i \in Z_i | g_i(x_i,z_i) = y_i\}$$

consists of exactly one element.

This assumption is a joint condition on G, μ, and $\{g_i\}$. The single-valuedness condition implies that the learning process of the individual is essentially independent of his action, and therefore a separation exists between the formation of expectations and the choice of an optimal action at each date. This removes the relevance, for the present purposes, of experimental design problems (see, e.g., Rothschild [1971]), which are undoubtedly of considerable importance and should be treated simultaneously with this problem in a more general analysis.

Thus the individual is assumed to choose x_i at each date to maximize

$$E\ u(x_i, y_i)$$

where the expectation is taken with respect to $y_i = g_i(x_i, \cdot)$ over all possible values of θ_i, weighted by the prior.

An <u>equilibrium</u> of this system will be a state in which no further learning takes place and in which expectations are correct in the sense that the anticipated distribution of z_i is realized. More formally, an equilibrium is a collection of values of the parameter $\bar{\theta}_i \in \Theta_i$ $i = 1, \ldots, n$ which give rise to optimal actions $\bar{x}_i \in X_i$ such that the random variable

$$G_i(\bar{x}, z)$$

has the same distribution as the random variable

$$g_i(\bar{x}_i, z_i)$$

where z_i is distributed with parameters $\bar{\theta}_i$.[2/]

The interpretation of this is that $\bar{\theta}_i$ is the point at which individual i's prior is degenerate, and he is correct in the sense that the random variable he observes relative to his model has the distribution that he believes it will have. Of course his model itself may be wrong in some sense--but the problem considered herein is only an equilibrium of a particular learning process and does not consider the choice of such a process.

In order for the above definition to be a meaningful criterion, we must have the possibility that there exists a point in the personal parameter space of any agent such that the distribution he observes given constant actions by everyone else is the same as that which he anticipates. That is, we assume that the following consistency criterion is satisfied by the models $\{g_i, \theta_i\}$.

For all $x \in X$, there exists for each $i = 1, \ldots, n$, a point $\theta_i \in \Theta_i$ such that $G_i(x,z)$ is distributed the same as $g_i(x_i, z_i)$ where z_i has distribution parameterized by θ_i.

This condition is a joint restriction on G, μ, g_i, and θ_i. The examples of the next section will show that it is compatible with the uniqueness criterion stated above.

[2/] See Hahn, F. [1973].

4. Examples

In this section, which forms the central part of this paper, we present a set of examples which are highly specialized cases of the above rather general structure. The point of these examples is to display some seemingly pathological, or at least mildly unexpected, properties of their equilibria.

All the examples are most easily thought of as quantity setting monopolistic competitors, who do not receive information on the quantities produced by their rivals, but only observe the market price. This is determined by the interaction of aggregate supply with a random demand curve.

We denote by π, the profit per unit of production which is made by every producer. This will be the only observed variable and it will be observed by every producer. Any non-negative real number is a possible production level. That is, $X_i = R_+$ and $Y_i = R$, for all i.

Example 1:

This example is characterized by the model,

$$(\pi,\ldots,\pi) = y = G(x,z) = b - z\sum_{k=1}^{n} x_k$$

where z is exponentially distributed with parameter $\bar{\theta}$ and b is a constant, and the personal models,

$$\pi = y_i = g_i(x_i,z_i) = b - z_i x_i$$

where z_i is exponentially distributed with parameter θ_i and b is known (with certainty) to be the same constant as in the specification of the "true" process above.

To check that this example satisfies the consistency criterion, it is only necessary to observe that for any $x = (x_1,\ldots,x_n)$, the observations

$$z_i = \frac{b-\pi}{x_i} = \frac{z\sum_k x_k}{x_i}$$

will be exponentially distributed with mean $\sum_k x_k/\bar{\theta}x_i$ and hence with parameter $x_i\bar{\theta}/\sum_k x_k$.

The uniqueness criterion is satisfied as long as $z_i \geq 0$, which follows immediately from the fact that z is exponentially distributed and hence, almost surely non-negative.

The optimal output for firm i with a prior degenerate at θ_i is clearly $b\theta_i/2$, since profit is $x_i\pi_i = bx_i - z_i x_i^2$ and expected profit is $bx_i - x_i^2 Ez_i$.

We can therefore write the perceived parameter values $\hat{\theta}_i$ as a function of those believed *ex ante* as

$$\hat{\theta}_i = \frac{x_i\bar{\theta}}{\sum_k x_k} = \frac{\theta_i\bar{\theta}}{\sum_k \theta_k} .$$

A collection of θ_i such that $\hat{\theta}_i = \theta_i$ for all i will be an equilibrium in the sense defined above. This is satisfied whenever

$\sum_k \theta_k = \bar{\theta}$, which can obviously hold for infinitely many combinations of θ_i.

Intuitively, the multiplicity and variety of equilibria are possible because a firm which has quite optimistic expectations, that is, large θ_i implying inelastic demand, will produce large outputs; hence the other firms will receive pessimistic observations (high elasticity of demand) which will justify the pessimistic expectations that they hold in equilibrium.

The asymmetry of the possible equilibria are the main point of this rather simplistic example. We will return to it as example 6 to examine the dynamics of the learning process, and in particular to show how the asymmetries of the equilibria may be produced from differences in initial priors.

Example 2:

Example 1 is quite unrealistic because each monopolistic competitor is behaving as if he were a monopolist in the sense that he believes that the demand curve he faces has the same intercept, b, as the true demand curve for the entire market.

Let us therefore modify example 1 to the extent that

$$(\pi_1,\ldots,\pi_n) = y = G(x,z) = (b_i - z\sum_k x_k) \qquad i = 1,\ldots,n$$

and

$$\pi = b_i - z_i x_i \quad , \quad i = 1,\ldots,n$$

where the b_i may differ from each other and the stochastic specifications of z and $\{z_i\}$ are the same as above.

The consistency and uniqueness criteria may be verified as in example 1. In equilibrium, the mappings

$$\hat{\theta}_i = \frac{b_i \theta_i \bar{\theta}}{\sum_k b_k \theta_k}, \quad i = 1,\ldots,n$$

have a fixed point at $\bar{\theta}_i$, $i = 1,\ldots,n$. Substituting $\hat{\theta}_i = \bar{\theta}_i$ and $\theta_i = \bar{\theta}_i$ above, this is shown to be impossible unless all the b_i are equal. That is, no equilibrium exists in this example--except in the borderline case in which it degenerates to example 1. The technical reason for this is that the functions above are defined on the non-negative reals which are not compact and hence no fixed point need exist in general. This shows that the principles discovered in example 1 should not be overstated, because they may be spurious as they are only possible in a hairline specification of the model.

Example 3:

Fortunately, the pathology of example 2 need not be too worrysome because with a minor modification in the direction of reality and generalization, the possibility of an equilibrium is restored. The cost of this, however, is a violation of the uniqueness criterion, which makes a dynamic analysis using the specification we have set forth above impossible. It should be possible to study the learning processes involved in a dynamic formulation of this example using the methods of Bayesian econometrics. At present the dynamics of this example remain open, and we present only an analysis of the equilibria.

It is characterized by the model

$$(\pi_1,\ldots,\pi_n) = y = G(x,z) = (b_i - z\sum_k x_k) \qquad i = 1,\ldots,n$$

where z is exponentially distributed with parameter $\bar{\theta}$ and b_i are constants, and the personal models

$$\pi_i = g_i(x_i,\alpha_i,z_i) = b_i - z_i\alpha_i - z_i x_i$$

where b_i is the (correctly known) intercept of the demand curve, α_i is an unknown constant and z_i is exponentially distributed with parameter θ_i.

The consistency condition is satisfied since, for any $\{x_k\}$, the parameter values $\alpha_i = \sum_{k\neq i} x_k$ and $\theta_i = \bar{\theta}$ will give rise to anticipated distributions for (α_i,z_i) that correspond to the perceived frequencies.

Profit maximization implies that x_i will be set so as to satisfy

$$x_i = \frac{b_i}{2Ez_i} - \frac{\alpha_i}{2}$$

which can be written

$$x_i = \frac{b_i\theta_i - \alpha_i}{2}$$

since z_i is exponentially distributed. The observed z_i's are

$$z_i = \frac{z\sum_k x_k}{\alpha_i + x_i} .$$

Taking the mean of the random variables on both sides of this equation,

$$Ez_i = \frac{\sum_k x_k}{\alpha_i + x_i} Ez .$$

If expectations are being fulfilled when each agent chooses his optimal output,

$$\frac{1}{\theta_i} = \frac{\sum_k \left(\frac{b_k}{2Ez_k} - \frac{\alpha_k}{2}\right)}{\alpha_i + \frac{b_i}{2Ez_i} - \frac{\alpha_i}{2}} Ez$$

for all i, or,

$$\frac{\alpha_i}{\theta_i} + b_i = \frac{1}{\theta} \sum_k (b_k \theta_k - \alpha_k)$$

for all i.

Let $y = (\alpha_i/\theta_i) + b_i$, which is independent of i in any equilibrium because it is equal to the right-hand side of the above equation. One constraint on the parameter space is that $\alpha_i \geq 0$ for all i. Hence,

$$y \geq b_i \quad \text{for all } i .$$

Substituting the definition of y in the expression for x_i,

$$x_i = b_i\theta_i - \frac{\theta_i y}{2} .$$

Since $x_i \geqq 0$ is a further constraint on choice, we have that $y \leqq 2b_i$ for all i in any equilibrium in which all firms are operating--that is, when the first-order condition for the optimal output is satisfied with equality.

Therefore, in all such equilibria (and we shall concentrate only on those equilibria in which all firms are active),

$$\max_k b_k \leqq y \leqq 2\min_k b_k .$$

Further, from the definition of y, we have

$$y = \frac{1}{\bar\theta}\sum_k (b_k\theta_k - \alpha_k) = \frac{1}{\bar\theta}\sum_k (2b_k\theta_k - \alpha_k y)$$

and therefore,

$$\sum b_k\theta_k = \frac{y}{2}(\bar\theta + \sum\theta_k) .$$

Since $\sum x_k = \sum b_k\theta_k - y(\sum\theta_k)/2$, the last result yields

$$\sum_k x_k = \frac{y\bar\theta}{2} .$$

Thus, in any equilibrium, aggregate output is determined by the parameter y and is constrained by the inequalities above.

We now ask if there exist any constraints on $\{x_k\}$ other than

$$x_i \geqq 0 \quad \text{for all } i$$

$$\frac{\bar{\theta}}{2} \max_k b_k \leqq \sum x_k \leqq \bar{\theta} \min_k b_k \quad .$$

Note that if $\max b_k \geqq 2\min b_k$ then equilibria exist, but if the b_k are too different, then there are no equilibria, at least with all firms active. In this sense the present example generalizes the earlier ones and demonstrates that the properties exhibited there may not be as spurious as might otherwise have been expected.

It can be seen that, in fact, the constraints above completely characterize all the equilibria. Let $\{x_k\}$ satisfy them, then

$$\theta_i = \frac{2x_i}{2b_i - y} \quad , \text{ for all } i \quad .$$

The θ_i are non-negative by virtue of condition $y \leq 2\min_k b_k$. Also, by construction,

$$\alpha_i = \theta_i(y - b_i)$$

is non-negative by the other part of the same condition.

This example is of interest because, unlike the previous ones, in equilibrium every individual has a personal model which is literally exact. That is, he perceives the constant term in the unit profit function to be $b_i - z_i \sum_{k \neq i} x_k$ and the slope of this function to be $-z_i$, with exactly the correct stochastic specification. We have shown that if the firms are sufficiently similar, equilibria exist and, in such cases, a continuum of equilibria are possible. As in example 1, the multiplicity of equilibria can be explained because the perceptions of one agent indirectly affect the perceptions of all others as they are linked through the true economic model which they observe only partially.

We cannot at this time provide an answer to the dynamics of this model, since the failure of the uniqueness condition gives rise to issues of experimental design and identification problems in their estimation. These interesting questions should be pursued in later research.

Example 4:

In this example we move towards a perfectly competitive system and away from the monopolistically competitive models of the earlier cases. It is characterized by the model

$$(\pi_1,\ldots,\pi_n) = y = G(x,z) = \langle p(\sum_k x_k, z) - \frac{1}{x_i}c(x_i)\rangle \qquad i = 1,\ldots,n$$

where $p(\sum_k x_k, z)$ is a decreasing function of its first argument for every value of the random variable z, and $c(\cdot)$ is a convex function. Naturally, we think of p as the market price and c as the total cost function.

In the spirit of perfect competition, all firms will perceive that price is a random variable, unaffected by their own production level. That is, their personal models are

$$\pi_i = g_i(x_i, z_i) = z_i - \frac{1}{x_i} c(x_i) \quad .$$

Note that we assume that they correctly perceive their total cost function. We assume that G, z, and Θ_i are such that the consistency condition is satisfied. The uniqueness condition clearly holds in this example.

In any stationary state, $p(\sum_k x_k, \cdot)$ is a random variable to which there corresponds $\theta_i \in \Theta_i$ for every i, such that z_i has the same distribution when parameterized by θ_i.

Optimal behavior implies that

$$Ez_i = c'(x_i) \quad , \quad \text{for all } i \quad ,$$

since the second-order conditions are insured by the convexity of c. Since the observed z_i are equal to p and hence identical across individuals, it is clear that in any equilibrium the expectations of all individuals must be the same. The monotonicity assumption on G implies therefore that only one equilibrium is possible.

Thus, when individuals believe that they do not affect market variables by their actions, the link between the expectations of one agent and the equilibrium optimal actions of others is broken and the possibility of multiple asymmetric equilibria is removed.

<u>Example 5</u>:

In this example we return to the first one to explore the dynamics of the learning process. In particular we try to show that the multiplicity and asymmetry of the equilibria in this example are actually attained as the limit of the learning-optimizing process of participants in this market. Further, unlike in the typical results of Bayesian statistical inference, we find that the equilibrium attained in the limit is not independent of the prior beliefs. This settles, in the negative, one of the questions raised in Section 3. However, the problem of sample dependence remains open, for in this example at least, the equilibrium depends only on the priors, with probability one. We present a conjecture concerning sample dependence in the next section.

Recall that this model is specified by

$$(\pi,\ldots,\pi) = y = G(x,z) = b - z\sum_{k=1}^{n} p_k$$

and

$$\pi = y_i = g_i(x_i,z_i) = b - z_i x_i \quad , \quad \text{for all } i \ ,$$

where b is a constant and z is distributed exponentially with parameter $\bar{\theta}$. In the learning process, each individual believes that his personal random variable z_i is distributed exponentially with parameter θ_i and is uncertain about the value of θ_i. Expectations about θ_i are assumed to be summarized by a prior which has the Gamma distribution with parameters α and β. Learning involves revising α and β in the light of new evidence received.

The Gamma and exponential distributions are called a conjugate pair[3/] because, if priors concerning the parameter of an exponential distribution are distributed according to a Gamma distribution, then regardless of the evidence obtained, Bayes Theorem will lead to a posterior which is also Gamma. In fact, Gamma is a two-parameter family and it can be shown that the posterior is related to the prior in the following simple way.

If α_t and β_t are the parameters of the prior at time t, then

$$\alpha_{t+1} = \alpha_t + 1$$

and

$$\beta_{t+1} = \beta_t + z$$

are the parameters of the posterior after the observation z has been received. This property makes these distributions particularly well suited to calculation for the purposes of this example.

Let us suppose that, for all individuals, i, $\alpha_{it} = t$ and let $t = 2, \ldots$ for technical convenience. Let β_{i2}, $i = 1, \ldots, n$ be arbitrary positive numbers. We calculate the optimal output at date t, $x_{it} = b/2E_{it}z_i$ where the expectation $E_{it}z_i$ depends on i's prior at date t, and hence on α_{it} and β_{it}, as follows:

[3/] See de Groot [1970].

$$E_{it}z_i = \int \frac{1}{\theta}\, d\Gamma(\alpha_{it},\beta_{it})(\theta)$$

$$= \frac{\beta_{it}^{\alpha_{it}}}{(\alpha_{it}-1)!}\int_0^\infty \theta^{\alpha_{it}-2} e^{-\beta_{it}\theta} d\theta \ .$$

This integral can be evaluated using the standard formula,[/]

$$\int x^n e^{ax} dx = e^{ax}\left[\frac{x^n}{a} - \frac{nx^{n-1}}{a^2} + \ldots + \frac{(-1)^{n-1} n!x}{a^n} + (-1)^n \frac{n!}{a^{n+1}}\right] \ .$$

Noting that the definite integral in the expression for $E_t z_i$ goes from 0 to ∞, we have

$$E_{it}z_i = \frac{\beta_{it}}{t-1} \ ,$$

and hence

$$x_{it} = \frac{b(t-1)}{2\beta_{it}} \quad \text{for all } i, t \ .$$

Thus the actual unit profits received are

$$\pi_t = b - z_t \frac{b(t-1)}{2}\left(\sum_k \frac{1}{\beta_{kt}}\right)$$

where z_t is the realized value of the random variable z.

The perceived personal random variables, as they depend on the $\{\beta_{kt}\}_{k=1,\ldots,n}$ are given by

---

[/] See

$$z_{it} = z_t \beta_{it} \sum_k \frac{1}{\beta_{kt}} \quad , \quad \text{for all } i \quad .$$

Note that $z_{it} \geqq 0$ since z_t is exponential and hence non-negative and β_{kt} is non-negative, for all k and all t by virtue of the revision rule. Using this revision formula,

$$\beta_{i,t+1} = \beta_{it}(1 + z_t \sum_k \frac{1}{\beta_{kt}}) \quad , \quad \text{for all } i \quad ,$$

and hence

$$\frac{\beta_{i,t+1}}{\beta_{j,t+1}}$$

is a constant for any particular pair of agents i and j, and is independent of the realization z_t. Therefore we can let $\beta_{it}/\beta_{jt} = c_{ij}$ for any pair i,j, $i \neq j$. Then we can write

$$\beta_{i,t+1} = \beta_{it} + z_t \sum_k c_{ik}$$

and

$$\beta_{i,t+s} = \beta_{it} + (\sum_{v=0}^{s-1} z_{t+v})(\sum_k c_{ik}) \quad .$$

Thus

$$\frac{\beta_{i,t+s}}{s} \text{ converges to } (Ez)(\sum_k c_{ik}) \text{ for all } i$$

by the law of large numbers.

For the Gamma distribution with parameters α and β the mean is α/β, and the variance is α/β^2. Therefore the priors converge to degenerate distributions at $\theta_i = \bar{\theta}/\sum_k c_{ik}$ for all i, since $Ez = 1/\bar{\theta}$. Note that

$$\sum_i \theta_i = \bar{\theta}\sum_i \left(\frac{1}{\sum_k c_{ik}}\right) = \bar{\theta}\sum_i \frac{1}{\sum_k \frac{\beta_{i2}}{\beta_{k2}}} = \bar{\theta}\sum_i \frac{1}{\beta_{i2}\sum_k \frac{1}{\beta_{k2}}}$$

and letting $\sum_k \frac{1}{\beta_{k2}} = c$, we have that

$$\sum_i \theta_i = \bar{\theta} \ .$$

Hence the limiting parameter values of the learning process are equilibria as example 1 has demonstrated, and further any equilibrium is a possible limit of the process which will be achieved if the initial priors are parameterized by β's which stand in the same ratio.

5. Conclusion

This paper has been a preliminary effort at studying equilibria in expectations. Many open problems exist. In a negative vein, a possible conjecture is that the property of sample independence, which was possessed by example 5, would fail whenever the optimal actions are non-linear in the mean of the relevent random variable.

The most interesting class of problems surrounding this approach is to decide how an economic agent should choose and reject his personal model. We had been assuming that he retained his model throughout the estimation procedure, continually optimizing with respect to it. But if it were fundamentally incorrect in that the true functional forms did not correspond to any point in his personal parameter space, we should allow for the possibility that the individual will optimize on his choice of model as well. The hierarchy of rationalities that would result from this approach would be likely to suffer from the same basic problems as our simple models encounter. This leads to the hope that non-maximizing (satisficing) models might be fruitfully united with this structure.

There are also a variety of statistical problems which remain to be settled even in the simplest cases. For example, it would be highly desirable to study the learning dynamics of example 3, except that, to our knowledge, the corresponding solution to the Bayesian statistical problem that would have to be solved by each firm at every stage, is itself unknown.

Finally, it would be useful to synthesize the learning-optimizing interactions of this paper with the models concerning optimal experimental design in the pricing process.

References

Brock, W. A. [1972], "On Models of Expectations That Arise from Maximizing Behavior of Economic Agents over Time," Journal of Economic Theory, 5:3, pp. 348-376.

Cyert, R. M. and M. H. de Groot [1971], "Interfirm Learning and the Kinked Demand Curve," Journal of Economic Theory, 3:3, pp. 272-287.

de Groot, M. H. [1970], Optimal Statistical Decisions, McGraw-Hill.

Grossman, S. [1972], "A Bayesian Approach to Static, Stochastic Equilibria," unpublished manuscript, August, University of Chicago.

Hahn, F. [1973], On the Notion of Equilibrium in Economics, Cambridge University Press.

Lucas, R. E. and E. C. Prescott [1971], "Investment under Uncertainty," Econometrica, 39:5, pp. 659-681.

Muth, J. F. [1961], "Rational Expectations and the Theory of Price Movements," Econometrica, 29, pp. 315-335.

Rothschild, M. [1971], "Economics of Imperfect Information: A Survey," paper presented to the annual meetings of the Econometric Society, New Orleans, December.

Rothschild, M. [1973], "Searching for the Lowest Price When the Distribution of Prices is Unknown," Working Paper No. 1, Mathematical Social Sciences Board Workshop, July, University of California, Berkeley.