Equilibrium under Uncertainty:

Multi-Agent Statistical Decision Theory

Edward C. Prescott

and

Robert M. Townsend

revised June 1978

This essay is to appear in Studies in Bayesian Econometrics and Statistics in Honor of Harold Jeffreys, edited by Arnold Zellner. We wish to acknowledge helpful comments from the editor as well as Boyan Jovanovic, Finn Kydland, and Michael Visscher.

1. Introduction

As a positive science, the purpose of economics is to explain social phenomena. In this regard the relatively recent incorporation of uncertainty and imperfect information into economic models has led to some rather exciting developments. We now have for example coherent economic models consistent with observations on investment activity, hiring-layoff decisions, job tenure, the rate of firm growth, and fluctuations in aggregate economic activity. The purpose of this essay is to provide a framework under which these and other models can be better understood and compared. We have attempted to minimize formalism and mathematical abstraction and to illustrate and generalize the constructs that have proven most useful in developing theories to explain economic phenomena for which uncertainty is the essential element.

We emphasize that the successful models to which we refer have in common a generalization of single agent statistical decision theory to market contexts and precisely defined concepts of equilibrium. These models also have in common a dynamic framework; stationary Markov processes, decision rules, and equilibrium concepts; and low dimensional state variables. We now elaborate on these themes.

In statistical decision theory a single agent chooses a decision from some specified set of possible decisions and receives a reward which is a function of that decision and some random outcomes (see Section 2 for a simple exposition). The decision problem of the agent is well defined in that the return or objective function and

the distribution of the random variables are known. Yet the distinguishing characteristic of economic environments is that there is more than one decision maker. In particular in market environments it is assumed, at least implicitly, that there are many agents. In such contexts then the reward to a single agent depends not only on his own decision but also on the decisions of the other agents. Thus to predict the decisions which agents will make in a multi-agent environment there is needed some notion of consistency. We emphasize here as did Rothschild [1973] that <u>the things which each agent takes as given in making his own decision must be consistent with maximizing behavior on the part of the other agents</u>. In short what is needed is a precisely stated definition of equilibrium with this property. We focus in this essay on the competitive rational expectations equilibrium and the Nash equilibrium.

The relationship between the rational expectations equilibrium and the Nash equilibrium is illustrated first in a simple (essentially static) framework of Muth [1961] (see Section 2). In a rational expectations equilibrium the distribution of prices which agents take as given, prior to making production decisions, is the distribution which is generated by these decisions. In contrast, in a Nash equilibrium, each agent takes as given the production decisions of the others. In the limit, however, as the output of each agent becomes negligible relative to the aggregate, the two equilibrium concepts are equivalent.

Frequently in statistical decision problems, an agent makes a decision after receiving some information, that is, after observing the realization of a random variable related to the underlying random outcome. One may then refer to the agent's decision rule or policy which specifies his decision as a function of the observation. In a multi-agent context then one may define a Nash equilibrium in the space of decision rules, a concept which allows for asymmetric information (see Section 3). This equilibrium concept is related to the notion of a Bayesian equilibrium as introduced by Harsanyi [1968].

Yet in markets with asymmetric information one frequently utilizes the notion of a rational expectations equilibrium with prices conveying information. In Section 3 of this essay we also analyze the relationship between the Nash equilibrium in the space of decision rules and such rational expectations equilibria in a simple version of an (essentially static) asset holding model of Grossman and Stiglitz [1976]. Again, an equivalence is established. Some have argued that the notion of a rational expectations competitive equilibrium with prices conveying information is conceptually flawed: is it not inconsistent that each agent obtains information from the prices which he takes as given prior to making his decision even though such prices are determined by the decisions of all agents? By viewing the rational expectations equilibrium as a Nash equilibrium in the space of decision rules, this apparent paradox is resolved.

One way of modeling dynamic economic phenomena, as would be suggested by the general equilibrium models of Arrow and Debreu, is to search for optimal actions conditional on the sequence of realizations of all past and present random variables or shocks. An alternative approach which has proven more useful in developing testable theories is to replace the attempt to locate equilibrium sequences with the search for equilibrium decision rules, where, in dynamic contexts, these decision rules are functions of a limited number of "state variables" (see Section 4). Such state variables include those elements which determine the technology available to the agents, such as the stocks of capital goods, and those elements which specify the effect of past decisions on contemporary preferences. State variables also include those elements which specify the relevant aspects of the agents' information sets. We emphasize that state variables should be of minimal dimension, indexing only those factors which can (potentially) change over time. A related point is that information sets (conditional distributions) are characterized by only a few parameters.

Consistent with the attempt to provide econometrically testable hypothesis, the models are time invariant or stationary. It is postulated that the stochastic process governing the distribution of next period's state variables, as a function of the current state variables and decisions, is stationary. Similarly the equilibrium concepts are stationary, either the recursive Nash or the recursive rational expectations equilibrium. Thus it turns out that the stochastic processes governing the relevant time series of the model are stationary and consequently subject to time series (econometric) analysis.

To reiterate, we hope to lay out in this essay the common elements of models which have proven successful in explaining economic phenomena. We do so in part by examining their role in a few selected structures. Section 5 presents a generalized version of the Lucas-Prescott (1971) model of investment under uncertainty. We note that the exposition here draws a distinction between individual and aggregate state variables. Section 5 also outlines the essential elements of Crawford's (1976) model of industry employment and hours decisions. Here we emphasize the important econometric feature that the state variables include a low dimensional statistic summarizing information on permanent and transient shocks, sufficient for forecasting future shocks. Section 6 presents a simple, discrete-time version of Jovanovic's (1978) model of job match and labor turnover. Section 7 presents Lucas's (1972) model of the Phillips curve, the essence of which is information asymmetries. Finally Section 8 presents a model of asymmetric information in which a few state variables characterize the distribution of beliefs of agents on an initial demand shock. The model is a modification of Townsend (1978), motivated by the search for stationary schemes and by the notion that agent's prior beliefs differ only if their information sets differ. The econometric implications of this last approach is the subject of ongoing research.

We would like to emphasize here that this essay certainly is not intended as a survey of existing literature on equilibrium under uncertainty. Rather we have chosen to concentrate on a few of the recent exciting contributions, a choice which we think best illustrates the points we have emphasized above.

2. Elements of Multi-Agent Decision Theory

In this section we offer a brief description of decision theory as it pertains to single agent problems and discuss extensions necessary for an application to multi-agent environments. We emphasize the need for a concept of equilibrium, and, in an essentially static context, define formally the Nash equilibrium and the rational expectations equilibrium. These equilibrium concepts are illustrated in the simple partial equilibrium framework adopted from Muth [1961].

The single agent decision problem is now described formally.[1/] An agent is to select a decision d from the space of all possible decisions φ. The reward to the decision maker for any such decision is assumed to depend on random variable ω with distribution $F(\omega)$.[2/] That is, the relationship among decisions, outcomes, and rewards can be expressed by a function r where $r(d,\omega)$ is the reward to the agent for decision d given outcome ω. The agent acts to maximize his expected reward $\int r(d,\omega)dF(\omega)$, by choice of his decision $d \in \varphi$.[3/]

The distinguishing characteristic of multi-agent environments is that the decision problem of any one agent cannot be considered independently of the decisions of others. To formalize this suppose there are n agents or decision makers. Each agent i can make a decision d_i in some set of possible decisions φ_i. Then the reward to agent i for decision d_i is assumed to depend on a random outcome ω as well as on the decisions of others, $D_i \equiv (d_1, d_2, \ldots d_{i-1}, d_{i+1}, \ldots d_n)$.

These relationships may be described by a function r_i where $r_i(d_i, D_i, \omega)$ is the reward to agent i.

If the decision problem of any agent i is to be well defined, then agent i must be able to view his reward as a function of the random outcome ω and his own decision d_i only. The natural way to do this in the space of decisions is for agent i to take the decisions D_i of others as fixed, independent of his own decision. Of course one would like to have the self-fulfilling property that the decisions which agent i takes as given <u>are</u> in fact the decisions employed by the others. This leads to the definition of a Nash equilibrium, which has this self-fulfilling property.

<u>Definition</u>. A <u>Nash equilibrium</u> is a specification of a decision $d_i^* \in \varphi_i$ for each agent $i \in \{1,2,\ldots,n\}$ such that

$$\int r_i(d_i^*, D_i^*, \omega)\,dF(\omega) \geq \int r_i(d_i, D_i^*, \omega)\,dF(\omega)$$

for all $d_i \in \varphi_i$.

Hence in a Nash equilibrium the decision d_i^* is maximal for each agent i given the decisions D_i^* of others.

An alternative way of suppressing the dependence of agent i's reward on the decisions of others is to focus attention on some random variable $\tilde{P}(\omega, D)$ which agents take as given even though it depends on the aggregate decisions of all agents $D = \Sigma_{j=1}^n d_j$. That is, it may be that there exists a random variable $P(\omega)$ and a reward function $\hat{r}_i$ such that the decision problem of each agent i is of the form

$$(2.1) \qquad \max_{d_i \in \varphi_i} \int \hat{r}_i[P(\omega), d_i] dF(\omega) \ .$$

In a rational expectations equilibrium $\tilde{P}(\omega, D)$ will have the self-fulfilling property.

<u>Definition</u>. <u>A rational expectations equilibrium</u> is a random variable $P^*(\omega)$ and a specification of a decision d_i^* for each agent $i \in \{1,2,\ldots,n\}$ such that

(i) For every i, d_i^* solves problem (2.1) under $P^*(\omega)$.

(ii) $P^*(\omega) = \tilde{P}(\omega, D^*)$.

Hence in a rational expectations equilibrium the decision of each agent is maximal given $P^*(\omega)$, and $P^*(\omega)$ is generated by the decisions.

These two equilibrium concepts are now illustrated in a structure adopted from Muth [1961]. The objective function of each firm $i \in \{1,2,\ldots,n\}$ is expected profits, $E\{Pq_i - (1/2a)q_i^2\}$ a function which is quadratic in output $q_i \in \mathbb{R}$ and linear in the price P. This produced commodity is sold in a competitive market with price determined in accordance with stochastic inverse demand $\tilde{P}(\epsilon, Q) = \theta + \epsilon - (Q/n)$ where $Q = \Sigma_{j=1}^n q_j$ denotes the aggregate (industry) output, ϵ is a normal random variable with mean 0 and variance var(ϵ), and θ is a known parameter. Production decisions must be made prior to the realization of ϵ.

A rational expectations equilibrium in the Muth structure is a specification of an output choice q_i^* for each firm $i \in \{1,2,\ldots,n\}$ and a pricing function $P^*(\epsilon)$ such that for each i, q_i^* maximizes $E[P^*(\epsilon)q_i - (1/2a)q_i^2]$ and $P^*(\epsilon) = \tilde{P}(\epsilon, Q^*)$. It may be verified that $q_i^* = a\theta/(1+a)$, $i \in \{1,2,\ldots,n\}$, and $P^*(\epsilon) = [\theta/(1+a)] + \epsilon$ have the desired properties. A Nash equilibrium in the Muth structure is a specification of an output choice q_i^* for each firm $i \in \{1,2,\ldots,n\}$ such that q_i^* maximizes $E[\tilde{P}(\epsilon, \sum_{j \neq i} q_j^* + q_i)q_i - (1/2a)q_i^2]$. Here it may be verified that the $q_i^* = a\theta/[1+a+(a/n)]$, $i \in \{1,2,\ldots,n\}$, have the desired properties. Now it may be noted in the Nash equilibrium that firms take account of the influence of aggregate output $Q = \Sigma_{j=1}^n q_j$ on price. If the output of firm i were negligible relative to the aggregate Q, then the price P would be viewed by firm i as being independent of his individual action q_i, just as for the rational expectations equilibrium. In fact, as the number of firms gets large, i.e., as $n \to \infty$, it is clear that the Nash equilibrium converges to the rational expectations equilibrium.

An alternative way to establish this equivalence is to analyze directly the limit economy. That is, let the set of firms be the unit interval [0,1] and let $\tilde{P}(\epsilon, Q) = \theta + \epsilon - Q$. It can be verified that the appropriate elements of $q_i^* = a\theta/(1+a)$, $i \in [0,1]$, $P^*(\epsilon) = [\theta/(1+a)] + \epsilon$ constitute both a Nash and a rational expectations equilibrium.[4/]

3. Statistical Decision Theory with Asymmetric Information

In this section we consider statistical decision problems in which agents make decisions after observing the realization of a random variable related to the underlying random outcome. The concepts of the state variable, decision rule, and Nash equilibrium in the space of decision rules are introduced in an essentially static context. The relationship between the Nash equilibrium in the space of decision rules and the rational expectations equilibrium with prices conveying information is illustrated in a model of financial assets of Grossman and Stiglitz [1976] and certain simultaneity issues are resolved.

First the single agent statistical decision problem is modified. Prior to making a decision d, the agent observes a realization of a random variable x. The notion that x conveys information of the forthcoming outcome ω is formalized by postulating the existence of a joint distribution $F(\omega,x)$ of the random variables ω and x. The marginal distribution of x is denoted by $F(x)$, and the family of conditional (posterior) distributions of ω given x is denoted by $F(\omega|x)$. In this context a decision rule for an agent is a function δ where $\delta(x)$ is a decision in φ given an observation x.[5/] The observation x may also be referred to as the state; it should be noted that decisions depend entirely on the state. The agent seeks a best decision rule in the class of all feasible rules. That is, his problem is to choose some δ, with $\delta(x)\epsilon\varphi$ for every x, to maximize expected utility $\iint r[\delta(x),\omega]\,dF(\omega|x)\,dF(x)$.

Now returning to the multi-agent environment, recall that the reward to each agent $i\in\{1,2,\ldots,n\}$ depends on the random ω, on his own decision $d_i \in\varphi_i$ and on the decisions of others $D_i = (d_1,d_2,\ldots d_{i-1},d_{i+1},\ldots d_n)$, i.e., we have $r_i(d_i,D_i,\omega)$. It is now assumed that prior to making a decision, each agent i observes a realization of a random variable x_i. Let $F(\omega,x_1,x_2,\ldots x_n)$ denote the joint distribution of ω and the observations $(x_1,x_2,\ldots x_n)$. In this context a feasible decision rule for each agent i is a function δ_i where $\delta_i(x_i)$ is a decision in φ_i given the observation x_i. Given the decision rules of others, $\Delta_i = (\delta_1,\delta_2,\ldots\delta_{i-1},\delta_{i+1},\ldots\delta_n)$, the expected reward for agent i from decision rule δ_i is then

$$\iint r_i[\delta_i(x_i),\ \Delta_i(X_i),\omega]\,dF(\omega,X_i|x_i)\,dF(x_i) \equiv v_i(\delta_i,\Delta_i)$$

where $X_i = (x_1,x_2,\ldots x_{i-1},x_{i+1},\ldots x_n)$ and

$$\Delta_i(X_i) = [\delta_1(x_1),\delta_2(x_2),\ldots\delta_{i-1}(x_{i-1}),\delta_{i+1}(x_{i+1}),\ldots\delta_n(x_n)].$$

The natural candidate for an equilibrium in this context is the Nash equilibrium in the space of decision rules. A Nash equilibrium is a specification of feasible decision rules such that the decision rule of each agent is maximal for him given the decision rules of others. More formally we have

Definition. A Nash equilibrium is a specification of a feasible decision rule δ_i^* for each agent $i \in \{1,2,\ldots,n\}$ such that $v_i(\delta_i^*,\Delta_i^*) \geq v_i(\delta_i,\Delta_i^*)$ for all feasible rules δ_i.

In market contexts an alternative equilibrium concept is the rational expectations equilibrium with prices conveying information. This concept is first illustrated in a simple version of a model of financial assets by Grossman and Stiglitz [1976]. Its relationships to the Nash equilibrium is then established.

Consider the following two-period model. In the first period each agent i of a set of n agents is endowed with $\bar{M}_i$ units of a riskless asset and with $\bar{X}_i$ units of a risky asset. In the second period each unit of the riskless asset yields R units of the single consumption good of the model. Each unit of the risky asset yields u units of the consumption good where $u \sim N[E(u), \text{var}(u)]$. In the first period there exist competitive markets for the two assets. Let P denote the price of a unit of the risky asset with the price of the riskless asset at unity. Each trader i has a common utility function v for second-period consumption c_i of the form $v(c_i) = -\exp(-ac_i)$ where $a > 0$. The problem of each agent i is to choose asset holdings X_i and M_i to maximize the expected utility of consumption $c_i = RM_i + uX_i$ subject to the budget constraint $(M_i - \bar{M}_i) + P(X_i - \bar{X}_i) = 0$.

Expectations are taken with respect to all available information. The information structure is described as follows. A set I of informed traders, fraction λ of the population, receives a signal $\theta = u + \epsilon$, where $\epsilon \sim N[0, var(\epsilon)]$, prior to making a portfolio decision. Clearly θ conveys information on the forthcoming realization of u. In fact the posterior distribution of u conditional on the observation θ (cf. Zellner [1971] for such updating formulas) is normal with mean

$$E(u|\theta) = E(u) + \frac{var(u)}{var(u) + var(\epsilon)} [\theta - E(u)]$$

and with variance

$$var(u|\theta) = var(u) \left[1 - \frac{var(u)}{var(u) + var(\epsilon)}\right] .$$

Note that only the mean depends on θ. The set U of uninformed traders do not observe θ directly but do know the price P of the risky asset. Let us suppose that uninformed traders believe that P depends on θ, that is, that $P = \tilde{P}(\theta)$ is a realization of a random variable $\tilde{P}(\cdot)$. Then we may refer to the mean $E[u|P;\tilde{P}(\cdot)]$ and variance $var[u|P;\tilde{P}(\cdot)]$ of the posterior distribution of u of uninformed traders conditional on the observation P.

It can be established that the demand $X_I(P|\theta)$ of any informed trader for the risky asset as a function of P and θ is of the form

$$X_I(P|\theta) = \frac{E(u|\theta) - R\,P}{a\ \mathrm{var}(u|\theta)}$$

and the demand $X_U[P|P;\tilde{P}(\cdot)]$ of any uninformed trader for the risky asset as a function of P is of the form

$$X_U[P|P;\tilde{P}(\cdot)] = \frac{E[u|P;\tilde{P}(\cdot)] - R\,P}{a\ \mathrm{var}[u|P;\tilde{P}(\cdot)]} .$$

Market clearing then requires that demand for the risky asset equal the total endowment. In addition it is assumed that the distribution of price $\tilde{P}(\cdot)$ which uninformed agents take as given is in fact the distribution generated in equilibrium. That is, a <u>rational expectations equilibrium</u> is a price function $P^*(\cdot)$ of the random variable θ such that, for every θ,

(3.1) $$\lambda\ nX_I[P^*(\theta)|\theta] + (1 - \lambda)nX_U[P^*(\theta)|P^*(\theta);P^*(\cdot)] = \sum_{j=1}^{n} \bar{X}_j .$$

It can be established that the equilibrium price function $P^*(\cdot)$ is of the form $P^*(\theta) = \alpha_1 + \alpha_2\ E(u|\theta)$ where $\alpha_1 > 0$ and $\alpha_2 > 0$ are real numbers. Clearly then $P^*(\theta)$ fully reveals to uninformed traders $E(u|\theta)$ which itself is a linear function of θ. Hence in equilibrium uninformed traders are as well informed as informed traders.

It remains to establish that this rational expectations equilibrium is also a Nash equilibrium in the space of decision rules if there is a continuum of traders. Let the set of agents be the unit interval.[6]

Let the decision element for each trader be a demand function for the risky asset, and let the final allocation be determined by these demand functions in accordance with market clearing. Now let the decision rule for an informed trader i be a mapping from the observed θ to a choice of a demand function $X_i(\cdot,\theta)$. Uninformed traders i choose a demand function $X_i(\cdot)$. Equilibrium (*) choices are specified as follows. For every informed agent i let $X_i^*(\cdot,\theta)$ be defined by

$$X_i^*(P,\theta) = \frac{E(u|\theta) - R\,P}{a\,\mathrm{var}(u|\theta)} \quad .$$

For every uninformed agent i let $X_i^*(\cdot)$ be defined by

$$X_i^*(P) = \frac{E[u|P;P^*(\cdot)] - R\,P}{a\,\mathrm{var}[u|P;P^*(\cdot)]}$$

where $P^*(\theta) = \alpha_1 + \alpha_2\, E(u|\theta)$ as before. Now consider some uninformed trader i who takes as given the demand functions $X_j^*(\cdot,\theta)$, $j\in I$ and $X_j^*(\cdot)$, $j\in U - \{i\}$. Given these functions agent i knows that the equilibrium price will be determined in accordance with (3.1), so in fact $P^*(\cdot)$ is the equilibrium price function.7/ But agent i can condition his demand on the price which will prevail, and he knows the relationship between prices and θ, namely $P^*(\cdot)$. This yields $X_i^*(\cdot)$ as his optimal decision. Clearly also $X_i^*(\cdot,\theta)$ is an optimal decision for any informed agent i by construction. This establishes the desired result. Note that what is crucial in the argument is that any one agent i has negligible weight, i.e., prices are determined entirely by the demand functions of the other agents. (The property that prices reveal completely all information is inessential to the argument.)

As was noted in the introduction, some have argued that the notion of a rational expectations competitive equilibrium with prices conveying information is conceptually flawed. It is inconsistent, they argue, for each agent to obtain information from the prices which he takes as given, prior to making his decision, even though such prices are determined by the decisions of all agents. By viewing the rational expectations equilibrium as a Nash equilibrium in the space of decision rules, this apparent paradox has been resolved.

4. Sequential Statistical Decision Theory and Economic Dynamics

There are important economic phenomena that are inherently dynamic and cannot be viewed as a sequence of static equilibria differing only in the value of certain exogenous shocks. Optimal current actions of the economic agents depend not only upon the current situation but upon their expectations of future events. This occurs for example whenever either preferences or technology fail to be time separable. One way of modeling dynamic economic phenomena is to search for optimal actions conditional upon all exogenous shocks observable at the time of the decision. This approach, which is how uncertainty is incorporated in Arrow-Debreu economies, has not proven very useful in characterizing equilibria and explaining the phemonena of interest. An alternative approach which has proven more useful is to replace the attempt to locate equilibrium sequences of decisions, given exogenous sequences of shocks, with the search for equilibrium decision rules, functions of a limited list of "state variables," which summarize both past decisions and current shocks.

A particularly useful class of structures are those of the recursive or time invariant variety, for the resulting equilibrium is a system of time invariant stochastic difference equations as assumed in most econometric analyses. There are three basic structures for preferences with this property. The first construct assumes an infinitely-lived family or firm with discounting. The investment under uncertainty example presented in Section 5 uses this construct. The second is the assumption of exponential life of agents as in the Javonovic job match equilibrium analysis in Section 6. The third is Samuelson's overlapping generations abstract which is used in the Lucas business cycle model summarized in Section 7.

The stationary statistical decision problem, which was rigorously analyzed by Blackwell [1965], is one for which[8/]

(i) The return function is time separable jointly in in the period t decision variable d_t and an appropriately defined state variable s_t, and returns are discounted by factor $\beta \in (0,1)$; that is, the return function has the form

$$\Sigma_{t=0}^{\infty} \beta^t \, r(d_t, s_t) \ .$$

(ii) The state variable is observed or is a function of observables at time t. The conditional distribution of s_{t+1} given current and past decisions depends only upon (d_t, s_t) and the distribution function $F(s_{t+1} | d_t, s_t)$ is time invariant or stationary. Frequently, s_{t+1} is a function of (d_t, s_t), and an identically and independently distributed random variable.

(iii) The decision d_t is constrained to belong to a correspondence $\varphi(s_t)$.

A <u>stationary decision policy</u> δ is a decision rule used in every period t specifying the current decision as a function of the current state:

$$d_t = \delta(s_t).$$

The stationary decision policy is feasible if $\delta(s_t) \in \varphi(s_t)$ for all possible s_t. Blackwell has shown that if an optimal policy exists in the more general class of not necessarily stationary decision policies, $\{d_t = \delta_t(s_1,\ldots,s_t)$ with $d_t \in \varphi(s_t)$ almost surely$\}$, then an optimal stationary policy exists.

Under some fairly weak continuity conditions that are typically satisfied in economic applications[9/], there is a unique bounded value function v satisfying Bellman's optimality equation:

$$v(s) = \sup_{d \in \varphi(s)} \{r(d,s) + \beta \int v(y) dF(y|d,s)\} .$$

Here $v(s)$ is the supremum of the (expected discounted) return over all feasible policies given state s. A feasible stationary policy $\delta(\cdot)$ is optimal if and only if for all s the above supremum is obtained for $d = \delta(s)$. Compactness of sets $\varphi(s)$ along with the continuity assumptions are sufficient to ensure the existence of such a policy.

The state variable s_t typically specifies stocks of assets or capital goods. Sometimes, however, s_t reflects the effect of past decisions upon current return. For example, if last period's labor supply decision affects the disutility of current labor supply, then last period's labor supply must be a component of s_t. Thus, a recursive structure does not require time separable utility functions any more than it requires time separable technologies.

The state variable must also include relevant aspects of the agents' information set. If, for example, earnings are subject to a second order autoregressive process, earnings in both the current and previous period are elements of the state variable, for they are needed to forecast future incomes of the utility-maximizing individual. If earnings are subject to unobserved permanent and transitory shocks as in Friedman's permanent income theory of consumption, an exponentially weighted sum of past incomes is sufficient for forecasting future incomes and is, therefore, a suitable information state variable.[10/]

The <u>multi-agent stationary decision problem with common information</u> is defined as follows:

(i) The return function of agent $i \in \{1,2,\dots,n\}$ has the form

$$E[\Sigma_{t=0}^{\infty} \beta^t r_i(d_{it}, s_{it}, D_{it}, S_{it})]$$

where (d_{it}, s_{it}) is the decision-state pair of agent i and (D_{it}, S_{it}) are the pairs for agents other than i.

(ii) The decision constraint correspondences are

$$d_{it} \in \varphi_i(s_{it}, S_{it}) \text{ for } i \in \{1,2,\ldots n\} .$$

(iii) The distribution of $(s_{1,t+1},\ldots,s_{n,t+1})$ depends only upon $(s_{1t},\ldots,s_{nt})$ and $(d_{1t},\ldots,d_{nt})$ and not upon time.

A stationary policy for agent i is a decision rule,

$$d_{it} = \delta_i(s_{it}, S_{it}) ,$$

specifying the current decision of agent i as a function of his and other agents' state variables. If agents other than i are using stationary decision rules $\Delta_i = \{\delta_1,\ldots,\delta_{i-1},\delta_{i+1},\ldots,\delta_n\}$, the decision problem facing agent i is a stationary statistical decision problem. For a feasible decision policy δ_i there will be a return,[11/]

$$v_i(\delta_i,\Delta_i) = E\{\Sigma_{t=0}^{\infty}\beta^t r_i(d_{it},s_{it},D_{it},S_{it}) | \delta_i,\Delta_i\} .$$

A <u>stationary sequential Nash equilibrium</u> is an n-tuplet of mutually feasible stationary decision policies $(\delta_1^*, \delta_2^*,\ldots,\delta_n^*)$ such that

$$v_i(\delta_i^*,\Delta_i^*) \geq v_i(\delta_i,\Delta_i^*)$$

for all feasible policies δ_i for each $i\in\{1,2,\ldots n\}$.[12/]

At this point we do not have a general definition of a stationary competitive equilibrium that is suitable for all the existing examples of stationary competitive equilibrium analyses. Instead we shall focus on three structures and be precise in the definition of equilibrium for each. For the first two, the investment under uncertainty and job-match market analyses, the agents are faced with a discounted stationary statistical decision problem. The business cycle example is stationary in the equilibrium decision rules of each generation.

5. Investment Under Uncertainty Example

A competitive industry with constant returns to scale, increasing costs of rapid adjustment, and demand uncertainty will be used to illustrate the equilibrium constructs developed here. The constant returns to scale assumption greatly simplifies the analysis as well as being consistent with percentage growth rates being uncorrelated with firm size (Gibrat's Law) and the observed great variability of firm sizes. The simplification occurs because the distribution of capital or capacity over price-taking firms is irrelevant; only total industry capital or capacity matters.

A representative firm's production possibility set is a closed convex cone constraining current period capital $k_t \in \mathbb{R}$, next period capital $k_{t+1} \in \mathbb{R}$, and the current period commodity vector $x_t \in \mathbb{R}^n$. With the constant returns to scale assumption, (x_t, k_{t+1}) is feasible given k_t if

$$(5.1) \qquad (x_t/k_t,\ k_{t+1}/k_t) \in T \subset \mathbb{R}^{n+1}$$

for some closed convex technology set T. A component of x_t is nonnegative if an output and nonpositive if an input. With this convention, often used in general equilibrium theory, there is no need to distinguish between inputs and outputs and all prices are nonnegative. The set T is assumed to have an interior.

The industry is subject to demand shocks z_t. Letting X_t be industry "output," the downward sloping inverse "demand" function is

(5.2) $$p_t = D(X_t, z_t) .$$

The demand shocks are subject to a stationary Markov process and initially we assume shocks observed by all firms. The probability distribution function for z_{t+1} conditional on z_t is $F(z_{t+1}|z_t)$. The objective of the firm is to maximize its expected present value,

(5.3) $$E\{\Sigma_{t=0}^{\infty} \beta^t p_t \cdot x_t\}$$

given its initial capital, k_0. The discount factor β equals $1/(1+r)$ where r is the time invariant interest rate.

The state of the industry is the beginning-of-period distribution of capital among firms and the current demand shock. The former completely specifies current and future production possibilities of all firms, while the latter is sufficient relative to the entire history for predicting subsequent shocks. With the constant returns to scale and price-taking assumptions, however, only total industry capital,

(5.4) $$K_t = \sum_{\text{all firms}} k_t ,$$

matters for market variables and not how it is distributed over firms. Therefore the pair (K_t, z_t) is an appropriate set of state variables for this <u>industry</u> equilibrium analysis.

<u>Definition</u>. A recursive competitive equilibrium is a pricing function $p_t = p(K_t, z_t)$, a law of motion for industry capacity $K_{t+1} = f(K_t, z_t)$, and a feasible, stationary firm investment-output policy

$$x_t = k_t\ \delta_1(K_t,z_t) \text{ and } k_{t+1} = k_t\ \delta_2(K_t,z_t) \tag{5.5}$$

such that

(i) Decision policy $\delta = (\delta_1,\delta_2)$ maximizes value, that is, maximizes (5.3), given functions p and f.

(ii) For decision policy δ_2, the law of motion for industry capacity will be f; that is for all (K_t,z_t)

$$f(K_t,z_t) = \sum_{\text{all firms}} k_t\delta_2(K_t,z_t) = K_t\delta_2(K_t,z_t)\ . \tag{5.6}$$

(iii) For decision policy δ_1

$D(X_t,z_t) = p(K_t,z_t)$ for all (K_t,z_t) where

$$X_t \equiv \sum_{\text{all firms}} x_t = K_t\ \delta_1(K_t,z_t)\ .$$

These equilibrium conditions warrant further discussion. The firm faces the following stationary statistical decision problem in state variable (k_t,K_t,z_t) and decision (k_{t+1},x_t). The objective function whose expected value is being maximized is

$$\Sigma_{t=0}^{\infty}\ \beta^t p_t\cdot x_t = \Sigma_{t=0}^{\infty}\ \beta^t\ p(K_t,z_t)\cdot x_t\ . \tag{5.7}$$

Period t return depends only upon the state and decision vector as required. The industry state (K_t,z_t) is subject to a first order time invariant Markov process defined by the conditional distribution $F(z_{t+1}|z_t)$ and the law of motion for industry capital $K_{t+1} = f(K_t,z_t)$. The element k_{t+1} is as selected in period t so is trivially determined by decision in period t. The constraint on period t decision (x_t,k_{t+1}) depends only upon state variable k_t and not on time. This constraint is $(x_t/k_t,k_{t+1}/k_t)\in T$. Because of the special form of the constraint set and the objective function, if an optimal policy exists it is of the form.

$$x_t = k_t\ \delta_1(K_t,z_t) \text{ and } k_{t+1} = k_t\ \delta_2(K_t,z_t) \tag{5.8}$$

as required in our definition of industry equilibrium.

Conditions (ii) and (iii) are that these optimal decision rules are consistent with the equilibrium pricing function $p_t = p(K_t,z_t)$ and equilibrium motion of industry capital $K_{t+1} = f(K_t,z_t)$. This is the condition that predicted and actual distributions of future prices be equal or that expectations are rational.

In Lucas and Prescott (1971) it was shown for a special case of the above model that a unique recursive competitive equilibrium exists. The equilibrium is such that the expected discounted consumer surplus function,

$$E\{\Sigma_{t=0}^{\infty}\ \beta^t s(X_t,z_t)\},\text{ where } s(X_t,z_t) \equiv \Sigma_{j=1}^{n} \int_0^{X_{tj}} D_j(y,z_t)\ dy, \tag{5.9a}$$

is maximal over all feasible programs. Analytically one considers the functional equation

$$(5.9b) \qquad v(K,z) = \max_{(X/K,K'/K)\in T} \{s(X,z) + \beta\int v(K',z')\, dF(z'|z)\} ,$$

where the prime indicates the next period's variable. The stochastic process governing the industry in equilibrium is Markovian and has stationary transition probabilities. In Lucas and Prescott, additional conditions are imposed which ensure the stationarity of this process. The stationarity of the equilibrium process governing the industry behavior is, we think, important for it is necessary for testing hypotheses using standard time series methods.

Crawford (1976) uses recursive equilibrium analysis to model employment, output and hours decisions for an industry. He assumes costly labor force adjustment so the workforce, as well as beginning of period inventories, are components of the "capital" vector k. He estimates output, hours, and workforce change equations as a function of industry capital K and the information set relevant for forecasting future demand shocks.

For the Crawford analysis the process governing the industry demand shock is not first order Markovian and z_t is not a sufficient statistic for forecasting future shocks. The same problem arises in Friedman's permanent income consumption theory as observed income is the sum of unobserved permanent and transitory components. The econometric resolution of this problem is to develop a low dimensional statistic m_t which is sufficient relative to the entire history of shocks for forecasting future demand shocks. It, along with z_t and the capital stocks, are the elements of the state vector. To obtain the Markov process governing (m_t, z_t) one develops the conditional distribution of z_{t+1} given m_t and the functional dependency $m_{t+1} = m(m_t, z_{t+1})$. These can be used to determine

$$(5.10) \qquad F(m_{t+1}, z_{t+1} | m_t) = F(m_{t+1}, z_{t+1} | z_t, z_{t-1}, \ldots)$$

The pair (m_t, z_t) are thus subject to a time invariant Markov process and the problem has the stationary structure required.

The probability structure of the industry demand shock in Crawford's analysis is as follows: The observed or deducible shock z_t is the sum of unobserved transient shock y_{1t} and unobserved permanent shock y_{2t}:

$$(5.11) \qquad z_t = y_{1t} + y_{2t} \; .$$

The transient and permanent shocks satisfy

$$(5.12) \qquad y_{1,t+1} = \rho y_{1t} + \epsilon_{1t}, \; 0 < \rho < 1 \; ,$$

$$y_{2,t+1} = y_{2t} + \epsilon_{2t} \; .$$

The ϵ_{1t} and ϵ_{2t} are identically and independently distributed normal zero mean variates over time and are independent of each other.

This probability structure is an example of the Kalman filter model which is proving to be as useful in economic studies as it has in engineering analyses. In general the observed z_t, which can be a vector, is a known linear combination of the unobserved components

$$z_t = B\, y_t \ . \tag{5.13}$$

It is required that z_t have lower dimensionality than unobserved y_t and that B be of full row rank. For the Crawford analysis z_t has dimensionality one and $B = (1 \;\; 1)$.

The vector y_t is subject to a first order vector autoregressive process with identically and independently distributed normal zero mean errors ϵ_t:

$$y_{t+1} = A\, y_t + \epsilon_t \ . \tag{5.14}$$

For the Crawford problem

$$A = \begin{bmatrix} \rho & 0 \\ 0 & 1 \end{bmatrix} \tag{5.15}$$

and Σ_ϵ, the covariance matrix of ϵ_t, is diagonal. Crawford finds that this probability structure fits the data well and estimates parameters ρ and Σ_ϵ.

If y_t were known, it would be sufficient relative to the entire history of observables for forecasting future demand shocks and would be a valid information state variable. But y_t is not observed, and it is the prior distribution for y_t at time t which is sufficient relative to the history for forecasting future shocks. We will see for this structure that the prior distribution for y_t is completely specified by m_t, the conditional expectation of y_t given the history, and that m_{t+1} is a linear function of m_t and z_{t+1}. The pair (m_t, z_t) are subject to a Markov process with time invariant transition probabilities as required by the theory We now establish these results.

If at the time of the period t production decisions the prior distribution for y_t (i.e., the conditional distribution of y_t given $z_t, z_{t-1}, \ldots$) is normal with mean m_t and covariance Σ_t, the predictive probability distribution function of the pair (z_{t+1}, y_{t+1}) is from (5.13) and (5.14) (singular) normal with

(5.16)
$$E\{z_{t+1} | m_t, \Sigma_t\} = B\ A\ m_t$$

$$E\{y_{t+1} | m_t, \Sigma_t\} = A\ m_t$$

$$Var(z_{t+1} | m_t, \Sigma_t\} = B\ (A\ \Sigma_t A + \Sigma_\varepsilon)\ B' = Var(z_{t+1} | \Sigma_t)$$

$$Var(y_{t+1} | m_t, \Sigma_t) = A\ \Sigma_t\ A' + \Sigma_\varepsilon = Var(y_{t+1} | \Sigma_t)$$

$$Cov(y_{t+1}, z_{t+1} | m_t, \Sigma_t) = (A\ \Sigma_t\ A + \Sigma_\varepsilon)\ B' = Cov(y_{t+1}, z_{t+1} | \Sigma_t).$$

Letting m_{t+1} be the mean of the prior distribution for y_{t+1} at time $t+1$, then

$$(5.17) \qquad m_{t+1} = A\, m_t + \mathrm{Cov}(y_{t+1}, z_{t+1}|\Sigma_t)[\mathrm{Var}(z_{t+1}|\Sigma_t)]^{-1}\,(z_{t+1} - B\,A\,m_t)\ .$$

In addition

$$(5.18) \qquad \Sigma_{t+1} = \mathrm{Var}(y_{t+1}|\Sigma_t) - \mathrm{Cov}(y_{t+1}, z_{t+1}|\Sigma_t)[\mathrm{Var}(z_{t+1}|\Sigma_t)]^{-1}$$

$$\mathrm{Cov}(z_{t+1}, y_{t+1}|\Sigma_t)\ .$$

Assuming the modulus of the largest eigenvalue of A is in or on the unit circle, Σ_t converges exponentially fast to the unique positive definite matrix Σ satisfying (5.18).[13/]

Crawford sets Σ_0 equal to Σ so all Σ_t are equal to Σ. As Σ_t is time invariant it need not and is not included in the state vector. Further with $\Sigma_t=\Sigma$, the difference equation governing m_t, (5.17), is stable and the effect of m_0 upon m_t becomes negligible after a few periods.

Remaining to be determined is $F(m_{t+1}, z_{t+1}|m_t)$. The distribution of z_{t+1} given m_t is normal with mean and variance given in (5.16). With $\Sigma_t = \Sigma$, from (5.17) m_{t+1} is a linear function of m_t and z_{t+1}; that is

$$(5.19) \qquad m_{t+1} = L_1\, m_t + L_2\, z_{t+1}$$

where matrices L_1 and L_2 are functions of Σ, Σ_ε, A and B. Therefore it can be established that the joint distribution of (m_{t+1}, z_{t+1}) conditional on m_t is singular normal with mean $[A\ m_t,\ B\ A\ m_t)$ and covariance matrix

$$\begin{bmatrix} L_2 Var(z_{t+1}|\Sigma)L_2' & L_2 Var(z_{t+1}|\Sigma) \\ Var(z_{t+1}|\Sigma)L_2' & Var(z_{t+1}|\Sigma) \end{bmatrix} \tag{5.20}$$

This establishes that a Markov process with stationary transition probabilities governs the evolution of (m_t, z_t) as required for the analysis.

With this approach, the demand shock is decomposed into an expected permanent component m_{2t} and an expected transient component m_{1t} the effect of which declines at rate ρ. Crawford explains industry output, hiring-layoff, and hour decisions as a function of capital stock state variables, industry inventory and workforce, and of the information state variables, m_{1t} and m_{2t}. The decomposition of z_t into expected effects facilitates testing restrictions imposed by the theory. As predicted by the theory, Crawford finds that permanent demand shocks have a greater impact than transient shocks on the hire-layoff decision, and that transient shocks have the greater impact on hours of employment.

6. A Market Analysis of Job Matching

The study of Jovanovic [1978] is another illustration of the usefulness of stationary equilibrium constructs. His analysis is in continuous time but, in order to match better with the notation previously developed and to avoid some difficult mathematical issues, a discrete time version of his model is considered here.

There are many firms and many workers. Workers and firms are matched pairwise, with the unknown job characteristic θ of a worker under a specified match determined by a draw from a known distribution. That is, the output x_t of an employed worker in a time period t is a normal random variate with unobserved worker-job match mean θ and known variance 1. At the beginning of the first period of a worker-job match, the prior distribution of θ is normal with mean μ and precision π. Thus the information about θ, the parameter of the match, changes subsequent to each observed product of the match x. For this probability structure the distribution on θ at time t is normal with mean m_t and precision h_t. Next period job-match state variables, assuming the match continues, satisfy

$$m_{t+1} = (h_t + 1)^{-1}(m_t h_t + x_t) \tag{6.1}$$

and

$$h_{t+1} = h_t + 1 \, . \tag{6.2}$$

The distribution of x_t conditional on m_t and h_t is normal with mean m_t and precision $h_t/(h_t + 1)$. Thus, the distribution of next period job-match state (m_{t+1}, h_{t+1}) given the current state (m_t, h_t) is well defined.

If the match is terminated in period t, the worker is unemployed for the period and his net output is a possibly negative constant reflecting value of leisure and moving costs. The mean and precision of such a worker's next period match are $m_{t+1} = \mu$ and $h_{t+1} = \pi$.

Worker's Problem

The objective of the worker is to maximize the present value of earnings. Let Q be the maximal obtainable present value of earnings for a new worker or one who has terminated a job-match and was unemployed in the previous period. If the stationary wage policy of the current employer is to pay wage

$$(6.3) \qquad w_t = w(m_t, h_t) ,$$

the problem facing the worker is a stationary statistical decision problem. A stationary policy or stopping rule is defined by a set S; the worker terminates the match at the first t for which

$$(6.4) \qquad (m_t, h_t) \in S.$$

Assuming the worker behaves optimally subsequent to the termination, his present value of earnings $v(S,w,Q)$ will depend upon the separation set S, the wage policy for the current match, and parameter Q. Results cited in Section 4 justify restricting the search for optimal separation policies to the stationary set. The worker's problem is to select S given the function w and the parameter Q so as to maximize $v(S,w,Q)$. The correspondence of optimal separation sets (there may be more than one) is denoted by $\sigma(w,Q)$.

Firm's Problem

The contribution of a match to a firm's valuation (the expected discounted difference between the product of the match and the wages paid the worker) depends upon its wage policy $w(m,h)$ and the separation set S chosen by the worker. It is denoted by

$$(6.5) \qquad \pi(w,S) \ .$$

The firm maximizes this quantity subject to the constraints that the worker chooses S optimally and that the expected present value of the match for the worker is at least Q; that is subject to

$$(6.6) \qquad S \in \sigma(w,Q) \quad \text{and} \quad v(S,w,Q) \geq Q \ .$$

If there is no wage contract with non-negative profits, the demand for workers will be zero. If a wage contract exists which yields strictly positive profits, the demand will be infinite. As for all constant returns to scale technologies, profits must be zero in equilibrium.

Definition of Equilibrium

An equilibrium is a triplet (Q^*, S^*, w^*) such that

(i) $v(S,w^*,Q^*) \leq v(S^*,w^*,Q^*)$ or equivalently $S^* \in \sigma(w^*,Q^*)$.

[utility maximization by workers]

(ii) $\pi(w,S) \leq \pi(w^*,S^*) = 0$ for all (w,S) such that $S \in \sigma(w,Q^*)$ and $v(S,w,Q^*) \geq Q^*$. [value maximization by firms]

(iii) $Q^* = v(S^*,w^*,Q^*)$ [The maximal obtainable utility is indeed Q^*.]

Jovanovic establishes that paying workers their expected product, $w^*(m,h) = m$, is an equilibrium wage policy. The equilibrium set S^* is the unique one which results in society's product being maximized. The set S^* has the property that the more precisely a match's θ is known, that is the larger precision h, the greater must mean m be for the match to continue. As h approaches infinity, the m required for continuation approaches an asymptote.

This model has the implication, consistent with the data, that there is a negative correlation between labor turnover and job tenure with most job separation occurring in the first few periods of employment. Other important phenomena explained are the positive correlation between wages and job tenure, and, when tenure is held constant, the negative relationship between the wage and the probability of subsequent separation.

7. Business Cycle Example

Lucas considers a variant of Samuelson's model of pure consumption loans. In each of an infinite number of time periods N people are born. Each person lives two periods. Hence in each period there are N young people and N old. Each young person has a linear technology for transforming labor effort n into output y of the single nonstorable consumption good in accordance with $y = n$. Each young person values current consumption c and leisure in accordance with a common utility function $u(c,n), u_c > 0$, $u_n < 0$. The old have no labor and value consumption c' in accordance with a common utility function $v(c')$.

There is one other good in this economy--fiat money, issued by the government. At the beginning of the second period of his life each old person receives a transfer of such money which is proportional to his pretransfer holdings λ. Let the proportionality factor x be random with known density f on $(0,\infty)$. There are also two competitive markets in which the consumption good can be exchanged for money. Fraction $\theta/2$ of the young are distributed to the first market and $1 - (\theta/2)$ to the second. Here θ is random with known density g on $(0,2)$. (Both x and θ are independent of each other and independent over time.) In each of these markets each old person will supply his money balances, λx, inelastically, and each young person will exchange for money some of the consumption good which he produces, the amount depending on the price p of the consumption good. The old are distributed to markets in such a way as to equate the supply of money balances in each market.

The decision problems of the agents are now stated formally. It should be noted first that the state variables of this economy are m, the per capita pre-transfer holdings of the old, and the realizations x and θ. The old know m and x, and the problem of the old is trivial, namely maximize $v(c')$ by choice of consumption c' and money supply m^s subject to the budget constraint $p\,c' \leq m^s$ where $m^s \leq \lambda x$. Clearly $m^s = \lambda x$ and $c' = \lambda x/p$ are maximizing choices. Each young person is assumed to know m, but observes neither x nor θ. Yet each believes p depends on m, x and θ, that is, $p = \tilde{p}(m,x,\theta)$ is a realization of the random variable $\tilde{p}(\cdot)$. Hence the price p conveys information on x and θ. Similarly the young believe the future price $p' = \tilde{p}(m',x',\theta')$ is a realization of $\tilde{p}(\cdot)$. Finally each young person takes into account his maximizing behavior when old, i.e., $c' = (\lambda x')/p'$ where again λ denotes money balances acquired when young. This yields the following problem:

$$\max_{\lambda,c,n \geq 0} \left\{u(c,n) + \int v\left[\frac{\lambda\, x'}{\tilde{p}(m\,x,\, x',\theta')}\right] dG[x,x',\theta' | m,p;\tilde{p}(\cdot)]\right\}$$

subject to the budget constraint $p\,c + \lambda \leq p\,n$. Again note that since p' is a realization of $\tilde{p}(\cdot)$, expectations can be taken with respect to the known distributions of x, x', and θ'. Note also that p' depends on x through $m' = m\,x$ and under $\tilde{p}(\cdot)$, p conveys information on x. Hence the distribution G is conditional on m and p and depends on $\tilde{p}(\cdot)$. Let $\lambda[p|m,p;\tilde{p}(\cdot)]$ denote the maximizing choice of λ.

In a rational expectations equilibrium of the first market the demand for money of the $N(\theta/2)$ young must equal the supply of money of the old. Moreover it is postulated that the pricing function $\tilde{p}(\cdot)$ which agents take as given is also realized. That is, a rational expectations equilibrium of the first market is a pricing function $p^*(\cdot)$ of m, x and θ such that, for all m, x and θ,

$$N(\theta/2)\lambda[p^*(m,x,\theta)|m,p^*(m,x,\theta);p^*(\cdot)] = (N/2)mx .$$

Lucas establishes under specified assumptions on the utility functions $u(\cdot)$ and $v(\cdot)$ and on the densities f and g that the equilibrium pricing function $p^*(\cdot)$ of the first market is of the form $p^*(m,x,\theta) = m\,\varphi(x/\theta)$ where $\varphi(z)$ is monotone increasing in z. (Similarly in the second market $p^*(m,x,2-\theta) = m\,\varphi[x/(2-\theta)]$.) Thus in equilibrium the young discern the ratio z of x to θ. It is also shown that the maximizing labor supply decision is an increasing function of real money balances acquired, λ/p. Also equilibrium real balances $mx/\theta p$ are shown to be an increasing function of z. Hence a relatively large z will be associated with relatively large output in the first market as well as high prices. Now, roughly speaking, if z is high because θ is low, then there is an effect of opposite directions in the second market. On the other hand if

z is large because x is large, then the effect in the second market is in the same direction. This is shown by Lucas to yield a Phillips curve in which aggregate real output is positively associated with the rate of inflation--an empirical observation not explained by an equilibrium theory prior to Lucas' work. It should be emphasized that this result turns on information asymmetries. Because only x/θ is known, monetary disturbances are interpreted in part as real disturbances and have real consequences. If in contrast there are no real disturbances, i.e., $\theta=1$ with probability one, then under $p^*(\cdot)$, x would be made known to the young, and labor supply decision would be some constant independent of x.

8. The Distribution of Beliefs as State Variable

The models considered thus far essentially have the property that informed agents trade with uninformed agents only once. The state variable then is naturally the vector of observables as well as those variables which determine the exchange opportunities available to the agents. In contrast agents may get repeated observations on variables which are determined in part by the actions of those with additional information. This is true for example of Lucas' [1975] model of the business cycle. It is established now by the way of an example that the natural candidate for the state variable in such environments is the distribution of beliefs of the agents. The exposition here is a modified version of Townsend [1978].

Returning to the partial equilibrium structure of Section 2, it will be convenient to take the set of firms as the unit interval. Each firm $i \in [0,1]$ maximizes expected profits at time t by choice of output q_{it}, yielding best decision $q_{it} = a\, E_{it}(P_t)$, a linear function of the price expected by firm i. These decisions in conjunction with inverse market demand at time t, $\tilde{P}_t(\epsilon_t, Q_t) = \theta + \epsilon_t - Q_t$ yield the price at time t. Here industry output is $Q_t = \int_0^1 q_t(i)\,d\mu(i)$. (See footnote 4.) Also, the $\{\epsilon_t\}_{t=1}^{\infty}$ are independently and identically distributed with $\epsilon_t \sim N[0, \mathrm{var}(\epsilon)]$. It is assumed that parameter θ is unknown initially and viewed by all firms as drawn from a normal distribution with mean $E(\theta)$ and variance $\mathrm{var}(\theta)$. The price P_t is observed at time t. Industry output Q_t and the realization of ϵ_t are unobserved.

At the beginning of the first period, each firm of a set of informed firms I, fraction ρ_I of the set of all firms, observe a random variable u where $u = \theta + v$, $v \sim N[0, var(v)]$. The distribution of v is known by everyone. The observation u conveys to informed firms information about θ. In particular after observing u informed firms view θ as distributed normally with mean $E_{It}(\theta)$ and variance $var_{It}(\theta)$, $t = 1$, where

$$E_{I1}(\theta) = E(\theta) + \frac{var(\theta)}{var(\theta)+var(v)} [u - E(\theta)]$$

$$var_{I1}(\theta) = var(\theta) - \frac{[var(\theta)]^2}{var(\theta)+var(v)} .$$

None of those in the set of uninformed firms U, fraction ρ_U of the set of all firms, observes u. Hence initially uninformed firms view θ as distributed normally with mean $E_{Ut}(\theta)$ and variance $var_{Ut}(\theta)$, $t = 1$, where

$$E_{U1}(\theta) = E(\theta) \text{ and } var_{U1}(\theta) = var(\theta) .$$

Now from the point of view of informed firms, $E_{I1}(\theta)$ is known. Of course θ is unknown. Uninformed firms, however, do not observe u and hence $E_{I1}(\theta)$ is not known. But from their point of view, given the above updating formulas, $E_{It}(\theta)$ can be viewed as an unknown parameter π_t, $t = 1$, which is jointly normally distributed with the unknown parameter θ, i.e., $[\theta,\pi_t]$ is distributed normally with mean $[E_{Ut}(\theta), E_{Ut}(\pi_t)]$ and variance-covariance matrix $\Sigma_{Ut}(\theta,\pi_t)$, $t = 1$, where $[E_{U1}(\theta), E_{U1}(\pi_1)] = [E(\theta), E(\theta)]$,

$$\Sigma_{U1}(\theta,\pi_1) = \begin{bmatrix} \mathrm{var}(\theta) & C\ \mathrm{var}(\theta) \\ C\ \mathrm{var}(\theta) & C^2[\mathrm{var}(\theta) + \mathrm{var}(v)] \end{bmatrix}$$

given $C \equiv \mathrm{var}(\theta)/[\mathrm{var}(\theta) + \mathrm{var}(v)]$.

Motivated by this discussion it may be guessed that the state variable of this model are the beliefs of informed and uninformed firms. In particular let the state variables at time t be specified by $s_t = [\pi_t, \mathrm{var}_{It}(\theta), E_{Ut}(\theta), E_{Ut}(\pi_t), \Sigma_{Ut}(\theta,\pi_t)]$ where $\pi_t = E_{It}(\theta)$. One may guess from the linearity that the Nash equilibrium decision rule of each firm j, specifying output as a function of the state variables <u>known</u> to firm j, is a stationary (linear) function of mean beliefs alone. That is $q_{jt} = \delta_U[E_{Ut}(\theta), E_{Ut}(\pi_t)]$, $j \in U$, and $q_{jt} = \delta_I[E_{It}(\theta), E_{Ut}(\theta), E_{Ut}(\pi_t)]$ $j \in I$. It is understood here that at any time t informed firms know all the state variables and uninformed firms do not know π_t. Moreover it may be guessed that, given the Nash equilibrium in decision rules each period, the state variables s_t evolve in accordance with a stationary Markov process. Finally, as information variables, the state variable s_{t+1} should be expressible, in equilibrium, as a deterministic, time invariant function of the state variables s_t and the observable P_t; i.e., $s_{t+1} = g(s_t,P_t)$. Of course all these conjectures turn out to be correct, as is now shown.

First postulate that Nash equilibrium decision rules at time t are of the form

(8.1) $$q^*_{jt} = \alpha_1 E_{Ut}(\theta) + \alpha_2 E_{Ut}(\pi_t) , \quad j \in U$$

(8.2) $$q^*_{jt} = \beta_0 E_{It}(\theta) + \beta_1 E_{Ut}(\theta) + \beta_2 E_{Ut}(\pi_t) , \quad j \in I .$$

If this is so, then

$$Q^*_t = \rho_U[\alpha_1 E_{Ut}(\theta) + \alpha_2 E_{Ut}(\pi_t)] + \rho_I [\beta_0 E_{It}(\theta) + \beta_1 E_{Ut}(\theta) + \beta_2 E_{Ut}(\pi_t)] .$$

Given that $\tilde{P}_t(\epsilon_t, Q^*_t) = \theta + \epsilon_t - Q^*_t$, it follows that each uninformed trader j has maximizing decision

(8.3) $$q_{jt} = a\{E_{Ut}(\theta) - \rho_U[\alpha_1 E_{Ut}(\theta) + \alpha_2 E_{Ut}(\pi_t)] - \rho_I[\beta_0 E_{Ut}(\pi_t) + \beta_1 E_{Ut}(\theta) + \beta_2 E_{Ut}(\pi_t)]\}$$

and each informed trader j has maximizing decision

(8.4) $$q_{jt} = a\{E_{It}(\theta) - \rho_U[\alpha_1 E_{Ut}(\theta) + \alpha_2 E_{Ut}(\pi_t)] - \rho_I[\beta_0 E_{It}(\theta) + \beta_1 E_{Ut}(\theta) + \beta_2 E_{Ut}(\pi_t)]\} .$$

(Note again that $\pi_t = E_{It}(\theta)$ is unknown by uninformed firms.) Equations (8.1)-(8.4) can be made consistent by appropriate choice of the parameters $\alpha_1, \alpha_2, \beta_0, \beta_1$, and β_2. Hence there exists a Nash equilibrium at time t with the desired properties.

It remains to consider the evolution of the state variables and verify the conjectured properties. Note first that upon observing P_t informed firms view

$$X_{It} \equiv P_t + \rho_U [\alpha_1 E_{Ut}(\theta) + \alpha_2 E_{Ut}(\pi_t)] + \rho_I [\beta_0 E_{It}(\theta)$$

$$+ \beta_1 E_{Ut}(\theta) + \beta_2 E_{Ut}(\pi_t)]$$

as a realization of $\theta + \varepsilon_t$. (It is assumed here that all state variables are known by informed firms at time t.) It follows that upon observing P_t,

(8.5a) $$E_{I,t+1}(\theta) = E_{It}(\theta) + C_0[X_{It} - E_{It}(\theta)]$$

where $$C_0 = \text{var}_{It}(\theta) / [\text{var}_{It}(\theta) + \text{var}(\varepsilon)]$$

and

(8.5b) $$\text{var}_{I,t+1}(\theta) = \text{var}_{It}(\theta) - \frac{[\text{var}_{It}(\theta)]^2}{\text{var}_{It}(\theta) + \text{var}(\varepsilon)} .$$

Thus in equilibrium $E_{I,t+1}(\theta)$ and $\mathrm{var}_{I,t+1}(\theta)$ are time-invariant functions of the price P_t and the state variables at time t. Moreover in equilibrium the price P_t is a time invariant function of the state variables at time t and the random variable ϵ_t, and hence $E_{I,t+1}(\theta)$ and $\mathrm{var}_{I,t+1}(\theta)$ evolve in a time invariant stochastic way. Uninformed firms view

$$X_{Ut} \equiv P_t + \rho_U[\alpha_1 E_{Ut}(\theta) + \alpha_2 E_{Ut}(\pi_t)] + \rho_I[\beta_1 E_{Ut}(\theta) + \beta_2 E_{Ut}(\pi_t)]$$

as a realization of $\theta + \epsilon_t - \rho_I\beta_0\pi_t$. It follows that for uninformed firms

(8.6a) $$E_{U,t+1}(\theta) = E_{Ut}(\theta) + C_1\,[X_{Ut} - E_{Ut}(\theta) + \rho_I\beta_0 E_{Ut}(\pi_t)]$$

(8.6b) $$-\rho_I\beta_0 E_{U,t+1}(\pi_t) = -\rho_I\beta_0 E_{Ut}(\pi_t) + C_2[X_{Ut} - E_{Ut}(\theta) + \rho_I\beta_0 E_{Ut}(\pi_t)]$$

(8.6c) $$\Sigma_{U,t+1}(\theta,\pi_t) = \psi[\Sigma_{Ut}(\theta,\pi_t)]$$

where C_1 and C_2 can be expressed in terms of the entries of $\Sigma_{Ut}(\theta,\pi_t)$ and ψ maps $\Sigma_{Ut}(\theta,\pi_t)$ into $\Sigma_{U,t+1}(\theta,\pi_t)$.

Now at the beginning of period $t+1$ uninformed firms care about $\pi_{t+1} = E_{I,t+1}(\theta)$, not $\pi_t = E_{It}(\theta)$. (See 8.1.) But from (8.5a)

$$\pi_{t+1} = \pi_t + C_0\{P_t + \rho_U[\alpha_1 E_{Ut}(\theta) + \alpha_2 E_{Ut}(\pi_t)] + \rho_I[\beta_0\pi_t + \beta_1 E_{Ut}(\theta) + \beta_2 E_{Ut}(\pi_t)] - \pi_t\} . \tag{8.7}$$

Therefore taking the expectation in (8.7) and substituting from (8.6b), $E_{U,t+1}(\pi_{t+1})$ can be expressed as a time-invariant function of the price and state variables of time t. It is also clear from (8.6a), (8.6c) and (8.7) that $E_{U,t+1}(\theta)$ and $\Sigma_{U,t+1}(\theta,\pi_{t+1})$ have the same property. Hence all the state variables will be known by informed firms at time t+1, and all the state variables evolve in a time invariant stochastic way.

Finally it may be noted that this recursive Nash equilibrium may also be viewed as a recursive competitive equilibrium in the state variables s_t. Here a recursive competitive equilibrium is a price function $P_t = p(s_t,\epsilon_t)$, a law of motion for the state variables, $s_{t+1} = g(P_t,s_t)$, and decision rules $q_{it} = \delta_I[E_{It}(\theta), E_{Ut}(\theta), E_{Ut}(\pi_t)]$ for $i\in I$, $q_{it} = \delta_U[E_{Ut}(\theta), E_{Ut}(\pi_t)]$ for $i\in U$ such that (i) decision rules are maximal given the price function p, and, (ii) given the price function p, the law of motion g, and updating in accordance with Bayes' rule, the state variables evolve by the law of motion g. For having established above the existence of Nash equilibrium decision rules at time t in $[\pi_t, E_{Ut}(\theta), E_{Ut}(\pi_t)]$, it is clear that the price P_t can be expressed as a linear function of these variables and ϵ_t in such a way as to have the self-fulfilling property. That is, given

$$P_t = \theta + \epsilon_t - \rho_U[\alpha_1 E_{Ut}(\theta) + \alpha_2 E_{Ut}(\pi_t)] - \rho_I[\beta_0 E_{It}(\theta)$$

$$+ \beta_1 E_{Ut}(\theta) + \beta_2 E_{Ut}(\pi_t)] ,$$

where the parameters $\alpha_1, \alpha_2, \beta_0, \beta_1$, and β_2 are determined in the Nash equilibrium, each firm will choose the previous Nash equilibrium decision rule and P_t will be realized. Now suppose that each firm believes that the state variables evolve in accordance with the inference process outlined above, i.e., $s_{t+1} = g(P_t, s_t)$. Then, by the arguments given above, the state variables will evolve in fact in this way. That is, given the equilibrium price function, the same as in the Nash equilibrium, each informed firm forms posterior beliefs on θ in accordance with (8.5a) and (8.5b). Taking this updating as given, each uninformed firm forms beliefs on θ and π_{t+1}, as before. As before, informed firms take this updating as given. Thus the state variables at time t+1 are determined as before. Again at time t+1 there exists a self-fulfilling price function of these variables, and so on.

9. Concluding Remarks

We would like to reiterate in these concluding remarks the success of a rigorous application of statistical decision theory to economic modeling of markets subject to uncertainty. Statistical decision theory suggests that the prior beliefs of agents concerning as yet unobserved random shocks or unknown parameters are not arbitrary. Rather agents are assumed to know the joint distribution of these random shocks and observed random variables. Then priors are defined by the conditional distributions in accordance with Bayes' rule. In short one must be precise about the information structure of the model -- who observes what and when -- and this in turn gives the model some content. The message we wish to convey in this essay is that the restrictions on observations implied by these models are consistent with certain stylized facts.

In closing we would like to broach some questions concerning the optimality of the equilibrium notions we have commended. Unfortunately little can be said in general, though there has been some work on particular structures. One can assert that with common information sets and under miminal assumptions as specified by Debreu, competitive equilibrium allocations are necessarily optimal. But even here one must assert that the set of markets is sufficiently rich, that is, that the modeler has not exogenously limited the set of markets in a binding way. Much less can be said of models with asymmetric information, though there have been some recent advances. Harris and Townsend [1978] have argued that with asymmetric information one can not define optimal allocations independent of a consideration of resource allocation mechanisms. Employing the notion that sequential Nash equilibrium in the space of decision rules, they establish that the equilibrium final allocations of well defined mechanisms must necessarily satisfy certain self-selection properties. They are then able to show, in the context of a pure exchange economy with one informed and one uninformed agent, that a particularly simple mechanism is optimal (in their sense) even though the equilibrium final allocations are not optimal relative to full information. Subsequent work generalizes this result. Yet there remain many open questions concerning the optimality market and non-market arrangements under uncertainty.

FOOTNOTES

1. The exposition here is motivated by DeGroot [1970].

2. Both here and below random variables may be viewed as having realizations in n dimensional Euclidean space. This space is not made explicit.

3. Alternatively one may view the reward to the agent as the expected utility of consumption, in which case r is a composite function.

4. With a continuum of firms the correct way to sum is by integration so that $Q = \int_0^1 q(j)d\mu(j)$ where $q:[0,1] \to \mathbb{R}$ is a Lebesgue (μ) integrable function and $q(j) = q_j$ is the output of firm j. Thus $Q = \int_0^1 q(j)\, d\mu(j)$ is independent of $q(i)$ as firm i is of Lebesgue measure zero.

5. Here the rule δ is assumed to be a measurable function of x.

6. The definition of a equilibrium must be modified of course to allow for the continuum.

7. With a continuum (3.1) must be modified. In particular set $n = 1$ and integrate on the right side.

8. Measurability of functions is assumed as required for expectations to be well defined.

9. These conditions are the continuity of the return function r, the correspondence $\varphi(s)$, and the probability s_{t+1} belongs to any measurable subset of the state space as a function of (d_t, s_t). If in addition for any s, $\varphi(s)$ is compact, then the supremum is obtained and the max can be substituted. If the function $r(d,s)$ is concave

and the set $\{(d,s)\colon\ d\in\varphi(s)\}$ is convex, then the value function v is concave. All these results are rigorously established in Lucas and Prescott [1971]. Without these continuity properties, as Blackwell [1965] has shown, optimal or even almost optimal (ϵ-optimal) policies may not exist.

10. See Lucas [1976] for an exposition of this theory in statistical decision theoretic terms.

11. This return functional depends upon the initial conditions $(s_{1,0},\ldots,s_{n,0})$ as well. Optimal δ_i given Δ_i does not, so this dependency can be suppressed.

12. For dynamic games there are many Nash equilibria depending upon how the decision space is defined (see Friedman [1977], Kydland [1975]); thus the predicate sequential is needed.

13. The norm used is the standard one for symmetric matrices, namely, the largest absolute value of the eigenvalues.

References

Blackwell, David [1965], "Discounted Dynamic Programming," Annals of Mathematical Statistics, v. 36:871-90.

Crawford, Robert G. [1976], An Empirical Investigation of a Dynamic Model of Labor Turnover in U.S. Manufacturing Industries, Carnegie-Mellon Dissertation.

DeGroot, Morris H. [1970], Optimal Statistical Decisions, McGraw-Hill, New York.

Friedman, James W. [1977], Oligopoly and the Theory of Games, North-Holland, Amsterdam.

Grossman, Sanford and Stiglitz, Joseph E. [1976], "Information and Competitive Price Systems," American Economic Review, v. 66, pp. 246-53.

Harris, Milton and Townsend, Robert M. [1978]. "Allocation Mechanisms for Asymmetrically Informed Agents," Carnegie-Mellon working paper #35-76-77, June 1978.

Harsanyi, John C. [1967-68], "Games with Incomplete Information Played by 'Bayesian' Players," Management Science, v. 14, pp. 159-82, 320-34 and 481-502.

Kydland, F. [1975]. "Noncooperative and Dominant Player Solutions in Discrete Dynamic Games," International Economic Review, 16, 321-335.

Jovanovic, Boyan [1978], Job Matching and the Theory of Turnover, University of Chicago Dissertation.

Lucas, Robert E., Jr. [1972], "Expectations and the Neutrality of Money," Journal of Economic Theory, v. 4, pp. 103-124. Reprinted in P. Diamond and M. Rothschild, eds., Uncertainty in Economics: A Book of Readings, [1974].

____________ [1975], "An Equilibrium Model of the Business Cycle," Journal of Political Economy, v. 83, pp. 1130-44.

____________ [1976], "Econometric Policy Evaluation: A Critique," in The Phillips Curve and Labor Markets, K. Brunner and A.H. Meltzer, eds., North-Holland, Amsterdam, pp. 19-46.

____________ and Prescott, Edward C. [1971], "Investment Under Uncertainty," Econometrica, v. 39, pp. 659-81.

Muth, John F. [1961], "Rational Expectations and the Theory of Price Movements," Econometrica, v. 29, pp. 315-35.

Rothschild, Michael [1973], "Models of Market Organization with Imperfect Information," Journal of Political Economy, v. 81, pp. 1283-1308.

Townsend, Robert M. [1978], "Market Anticipations, Rational Expectations and Bayesian Analysis," International Economic Review, v. 19, pp. 481-94.

Zellner, Arnold [1971], An Introduction to Bayesian Inference in Econometrics, John Wiley, New York.