

BAYESIAN LEARNING IN DYNAMIC ECONOMIC MODELS

By

Ronald Britto

Discussion Paper Number 43
December 1973

Preliminary Report on Research in Progress
Not to be quoted without permission of the author.

BAYESIAN LEARNING IN DYNAMIC ECONOMIC MODELS*

I. Introduction

In recent years there has been a revival of interest in stochastic economic models where the basic variables are unobservable and must consequently be estimated by the economic decision-maker as part of the optimizing process.¹ Friedman's well-known "permanent income" hypothesis of consumption is a notable example. Nerlove (1967) and Taylor (1970) have successfully used such models to obtain conditions yielding estimators for the unobserved variables in the form of a distributed-lag function of past observations. Estimators with this property have been widely used in econometric work, making it a matter of some importance that a solid theoretical underpinning should be established.

My goals here are less specific than the derivation of the distributed lag function as an optimal estimator. While I shall use an economic model much like that of Nerlove and Taylor in that it is characterized by unobserved variables - states of the system - I intend to study its properties with the aid of the apparatus of decision theory and information theory. Rather than focusing on the conditions yielding distributed-lag estimators, I shall study how the probabilities the decision-maker attaches to the states of the system in each period are revised in the light of new information and what such revision implies for his behavior. Thanks to the work of Blackwell and Girshick (1954), Marschak (1971), and Marschak and Miyasawa (1968) there exist powerful theorems that can be applied for this purpose. It is worth mentioning

* I am deeply indebted to my colleague Joseph Ostroy, with whom I have had many helpful discussions, and whose perceptive comments on previous versions of the paper I have found always useful (and sometimes disturbing!) Mr. Rakesh Sarin has provided me with valuable research assistance in connection with the simulation work reported on in sections V and VI. I alone remain responsible for the views expressed and all errors.

¹Marschak (1963) has emphasized that the 2-step procedure of first estimating the parameter of interest and then taking an action based on this estimate is generally incorrect, because the optimal estimate is (in general) not independent of the payoff function.

here that the decision-theoretic approach casts the question of distributed-lag estimators in a different light, for one can now approach the problem from the point of view of the obsolescence of information, using the concept of the value of information developed by the above-mentioned writers.

The plan of the paper is as follows. Section II contains an example of the kind of economic situation we are modelling and lists some of the reasons why the model deserves serious consideration. In the next section we present the model formally and explain how the probability vector is revised sequentially. Section IV deals with the obsolescence of information. In section V we present a simulation of the working of the model for a simple case with only two states of the system and two "observations" (also known as "messages" or "signals"). The question investigated by means of this example is that of the number of periods it takes a decision-maker on the average to realize that the structure of the system has changed. In section VI the long-run properties of the model are discussed - more specifically, the long-run behavior of the subjective probability vector. We look at the formation of expectations in section VII. Our findings are summarized in section VIII. Finally, the Appendix contains mathematical proofs of some of the statements made in sections III and VII.

II. The Model and Its Scope

Consider the example of a seller of some good who must determine the demand curve facing him in the current period. He has a series of past observations of prices and quantities sold, which he can - and will - use in some way to obtain an estimate of the amounts that he will be able to sell at different prices today. However, there are two reasons why he does not know whether in fact the past observations have all come from the same demand curve. Firstly, the demand curve in any period may be one of several, depending upon what actions competitors have taken, the incomes of buyers, changes in tastes, etc. Although he can guess at the degree to which these forces have been operating, he does not know their effects with full certainty. Secondly, observations will not lie on the demand curve that generates them because of transient factors. Calling the various possible demand curves states of the system, we assume that the process generating them is a first-order Markov-process, known to the decision-maker. This means that the conditional probability of a particular state occurring in period t , given that some state occurred in period $t-1$, is independent of states occurring before $t-1$. The decision-maker also knows, we assume, the conditional probability distribution of observations, given the state of the system. The knowledge of these probability distributions, together with the history of past observations, is all he has to go by.

We hasten to remark that this model is very broad in scope and can be applied to a variety of problem. For example, the states of the system can be interpreted as equilibrium exchange rates (or prices of any sort), the observations being the current surpluses in the balance of payments of a country (or number of unfilled orders). In Taylor's paper, for example, the "permanent" demands for the goods produced by the firm are unobserved because of random disturbances and follow a first-order Markov process. The firm observes the

actual demands and uses these to estimate "permanent" demands with a view to maintaining stocks of the various goods close to demand levels. All the random variables involved are continuous.

The model is applicable in all situations in which it is desirable to distinguish the effects on observations of changes in the underlying economic structure from the effects due to transient factors. This distinguishes it from standard models used to analyze inventory problems, problems of determining optimal international reserves, etc., where the underlying economic structure is assumed to remain unchanged. Thus in all the literature on optimal international reserves the assumption is made that the equilibrium exchange rates do not change over the entire time-interval being considered, with all changes in the balance-of-payments surplus ascribed to random factors. In inventory models, the standard assumption is that the probability distribution of demands remains the same in each period.

Attractive features of the model outlined above are:

(i) It presents a picture of the world that is more realistic than the alternatives used hitherto for two reasons: a) The world as an economic decision-maker sees it is in a constant state of flux, where the past is an unreliable guide to the present, but is used because it is the only guide available. On the other hand, the world is not pure anarchy. Our assumption that the stochastic process is a Markov process is intended to capture some elements of stability of the underlying structure. b) Because one can allow the observation received by one individual in any period to be different from that received by another in the same period, even when the underlying structure is the same, individuals will in general have different subjective probability vectors for the states in any period.

(ii) It allows the decision-maker to form expectations about the future that

are much less rigid than a straightforward extrapolation of past observations would be. These expectations will be affected by past observations in a manner that depends on the transition probabilities.

(iii) Related to (i) above, the model suggests that the standard practice in economic theory of separating the analysis of equilibrium from that of the stability of this equilibrium is quite misleading. There is no such thing as static, deterministic equilibrium; instead we have a series of processes that never settle down.

(iv) Work done by Cyert and deGroot (1970) on the use of the Bayesian approach to duopoly theory is very similar in spirit to this model and suggests, as we have mentioned, that the model is applicable to a wide variety of problems in economic theory.

To summarize, the model is designed to study the adaptive behavior of economic units as they respond to changes in the stochastic environment.¹

1 In a recent paper, Cross (1973) has argued that valuable results can be obtained by using the psychologist's learning approach in economic theory. Among his results are the following two: a) the firm does not respond immediately to changes in the economic situation, its "speed of adjustment ... depending on... the potential gains from altered behavior and...the extent to which the individual is wedded to traditional forms of behavior" (p. 258); b) "a firm can be expected to spend much of its time out of 'equilibrium'" (p. 261). The latter result is essentially what we have argued in (iii) above. We shall see an example of the former in our paper in section V. Therefore, I would take the position that there is not much to be gained by adopting the learning approach of the psychologist in economics. Of course, Bayesian models are learning models - the point is that they are perfectly compatible with our basic postulate of the maximization of expected utility.

III. Formal Outline of the Model

Notation

$z_i(t)$ ($i = 1, \dots, m$) indicates the i 'th state in period t .

$x_j(t)$ ($j = 1, \dots, n$) indicates the j 'th observation in period t .

$y_k(t, s)$ ($k = 1, \dots, n^{t-s}$) is the k 'th member of the set of sequences of observations from the periods $s, \dots, t-1$, where s is a positive integer $< t$.

$P(t)$ is the $m \times m$ Markov transition matrix with elements $p_{ik}(t) = p(z_k(t+1) | z_i(t))$ ($i, k = 1, \dots, m$), where $p(\cdot)$ refers to the probability of the event in parenthesis.

$Q(t)$ is the $m \times n$ Markov matrix with elements $q_{ij}(t) = p(x_j(t) | z_i(t))$.

$a_i(t)$ is the i 'th action available in period t ($i = 1, \dots, \ell$).

p_0 is an m -component vector giving the a priori probability for the various states in the initial period.

$J(t, t_1)$ is an $n \times n$ Markov matrix whose elements $j_{ij}(t, t_1) = p(x_j(t_1) | p(x_i(t)))$, where t_1 is any positive integer $\neq t$.

Assumptions

- A1. The action in any period t must be taken before the observation $x_j(t)$ is available, although all past observations are known. The states of the system are never known.
- A2. The matrix $P(t)$ is constant over time and written as P . P is non-singular.
- A3. The matrix $Q(t)$ is constant over time and indicated by Q and T respectively.

Consider the situation facing our decision-maker in period t . He has a series of past observations $y_k(t, 0)$ and in addition knows the a priori probabilities of observations in the current and future periods. He will use the past

observations to revise these prior probabilities in Bayesian fashion, and make his decision with the aid of the revised probability vector. It is known that this revision can be made sequentially, i.e., before the observation of period t is known, all previous observations are already incorporated in the revised probability vector, which is then revised again according to Bayes' formula after observation t is obtained.¹ Let $p(z_i(t)|y_k(t-1,0))$ be the revised probability of state i in period t , given a particular sequence k of observations from the preceding periods. Then if observation x_j is made in period t , the new revised probability of state i is

$$\frac{\{q_{ij}p(z_i(t)|y_k(t-1,0))\}}{\{\sum_i q_{ij}p(z_i(t)|y_k(t-1,0))\}}$$

where $q_{ij} = p(x_j(t)|z_i(t))$ to recall. The revised probability vector for the different states in periods $t+1, t+2, \dots$, given the sequence of observations from periods $0, \dots, t$, is then obtained in straightforward fashion by applying the transition matrix P and its higher powers P^2, P^3 , etc. to the vector whose i 'th component is that given above.

The fact that the revised probability vectors depend on all previous observations implies of course that the action taken by the decision-maker in period t will depend, inter alia, on the entire set of past observations. In the simplest case, where the results of all past actions are embodied in the current value of some particular variable (like current stocks in inventory models), the optimal action in period t will then depend on both the value of this variable in period t as well as the set of past observations. Although this is also true in models where a particular parameter, assumed to be constant but unknown, has to be estimated from past observations, the difference in our model

¹See for example Bellman and Dreyfus (1962), Chapter 8.

is that these past observations are not necessarily given the same weights when used in the estimation procedure. In the limiting case, for example, where the observations are linked in a one-to-one fashion to states of the world, only one observation, that from the immediately preceding period, is required, for it identifies exactly the state of the system prevailing in that period: any information about states in previous periods is useless because of the assumption that they obey a first-order Markov process. A question of great interest is under what conditions the weights attached to past observations decline, the farther back in the past they come from; this is precisely the question investigated both by Nerlove and by Taylor. We turn to it now.

IV. The Obsolescence of Information

For the purpose of illustrating the concept of the value of information, we take the simplest case, where the decision-maker must act on the basis of one observation alone. Let us assume that the payoff depends on both the action taken and the state of the system. There is a payoff matrix T whose i, j 'th element, τ_{ij} , shows the payoff when action a_i is taken and the state of the system is z_j ($i = 1, \dots; j = 1, \dots, m$). The a priori vector of probabilities of states is p_0 , which is revised after the observation is made with the use of the matrix Q . The expected value to the decision-maker of the observation is then

$$\begin{aligned} & \sum_k [\max_i \sum_j \tau_{ij} p(x_k | z_j) p(z_j)] \\ &= \sum_k [\max_i \sum_j \tau_{ij} q_{jk} p(z_j)], \end{aligned}$$

which clearly depends on the elements of the matrix Q . For instance, should $Q = I$, the identity matrix, or a permutation thereof, there is perfect information, for each observation serves to identify a particular state with probability 1. At the other extreme is the case where all rows of Q are identical, i.e., the observations are statistically independent of the states of the system.¹

If there are two such information matrices Q_1 and Q_2 , then what Marschak has called Blackwell's theorem states that the value of information from the information structure represented by Q_1 will be not less than that from the information structure represented by Q_2 for all payoff matrices and all a priori probability vectors if and only if $Q_2 = Q_1 M$, where M is some Markov matrix.²

We return to the model outlined in the preceding section, assuming that

¹See Marschak (1971), pp. 199, 200.

²See Blackwell and Girshick (1954), p. 328 (Theorem 12.2.2), or Marschak and Miyasawa (1968), p. 152 (Theorem 8.1).

the payoff in each period depends on the action taken and on the observation made in that period. (Of course the action must be taken before the observation is known.) It is shown in the Appendix that, with mild restrictions on the matrices Q and P, the value of an observation from any past period t_1 will be not less than that from an earlier period t_2 ($t_2 < t_1$). For example, one restriction that will suffice is that the stochastic process governing the movement from one state to another over time is stable, the probability being high that the system will stay in the same state from one period to the next. The same restrictions serve to ensure that a batch of observations from the immediate past will be at least as valuable in making the current decision as another batch from the more distant past.

What is the significance of this obsolescence of information? As we mentioned at the outset, it suggests that the weights attached to past observations for the purpose of taking today's action will decline, the farther back from the past they come. Care must be taken in interpreting this statement, for it is not true for any individual decision-maker. The concept of the value of information applies to an entire information structure, hence, we are asserting that the phenomenon of declining weights will occur for the average individual. This is to be contrasted with the Nerlove-Taylor approach, where the goal is to obtain estimates in the form of distributed lags for each individual.¹

¹Because of their more stringent requirements, both Nerlove and Taylor placed restrictions on their payoff functions. In Nerlove's case this function took the form of minimizing the mean-square-error. Taylor used a payoff function quadratic in the action and state variables, amounts produced and demanded respectively. On the other hand, Taylor's objective function was a sum of inventory costs and production costs, and the former clearly depends on past actions (amounts produced) as well as the current action.

V. Lags in Adjustment to Changes in the Environment

We study in the section the case where $n=m=2$ and the matrix

$$P = \begin{bmatrix} p_1 & 1-p_1 \\ 0 & 1 \end{bmatrix}$$

State 2 is thus an absorbing state. We shall assume that the system has moved to state 2 in period 1 and that it was known by every individual that the system was in state 1 prior to that. Thus the prior subjective probability for state 1 in period 1 is p_1 for each individual. The question of interest is how long it takes the average individual to learn that the change has occurred, and how this depends on the matrix Q . It is to be expected that it will normally take a few periods before the revised probability for state 1 will have dropped to a reasonably low level (say .10) and that the greater the value of the information structure Q the shorter will be this lag.

For our example we put

$$p_1 = .9, \quad Q_1 = \begin{bmatrix} .8 & .2 \\ .1 & .9 \end{bmatrix} \quad \text{and} \quad Q_2 = \begin{bmatrix} .6 & .4 \\ .3 & .7 \end{bmatrix}$$

It is easily verified that Q_2 is less informative than Q_1 , i.e., that $Q_2 = Q_1 M$ for some Markov matrix M . With each of these information matrices the model was simulated for 20 runs, the length of a run being 20 periods. The results were as follows.

For each run a sequence of revised probabilities was obtained. The average of these revised probabilities in each period was then calculated. With the information structure of Q_1 the average revised probabilities for state 1 in periods 2, 3, 4 and 5, given the observations from the preceding periods, were .63, .27, .14 and .065 respectively. The average individual had revised his probability to below .10 after just three observations. The corresponding probabilities with Q_2 were .78, .63, .58 and .44. (For period 11 it was still as high as .19). When no observations are used at all the sequence of probabilities is of course $(.9)^2 = .81$, $(.9)^3 = .73$, $(.9)^4 = .66$ and $(.9)^5 = .59$.

The chief point we are trying to make by means of this simple example is that the introduction of uncertainty into models of economic behavior is sufficient to justify the use of lags even when a) large numbers of participants are involved and b) there are no adjustment costs. When there has been a change in the state of the system, it takes a few periods before this change is discerned by the average individual. The length of this lapse will be determined by the information pattern available to individuals.

VI. Some Long-run Features of the Model

We look in this section at the long-run behavior of the revised probability vector for states in period t when t is large. The i 'th component of this vector (when the j 'th sequence of observations from the preceding $t-1$ periods has occurred) is $p(z_i(t)|y_j(t,1))$, $i = 1, \dots, m$; $j = 1, \dots, n^{t-1}$. We assume that the process described by the transition matrix P is acyclic and irreducible, so that $\lim_{t \rightarrow \infty} p_0 P^t = p^*$ independently of p_0 . Because of this, the dependence on p_0 of the revised probability vector, which shall be indicated by $p_r(z(t))$, will clearly diminish as t increases. The question of interest is whether $\lim_{t \rightarrow \infty} p_r(z(t))$ exists under these circumstances.

It is easy to see that the answer is generally in the negative. Assume that $n=m$ and put $Q = I$, the identity matrix. Then for any t the state in period $t-1$ is known exactly by every individual. It follows that $p_r(z(t)) = e_i P$, if the system was in state i in period $t-1$, where e_i is the i 'th unit vector. If p^* contains more than one positive component, the system will switch from one state to another and $p_r(z(t))$ will fluctuate accordingly. More specifically, for large t $p_r(z(t))$ will equal $e_i P$ with probability p_i . This is true for each individual and hence also holds for the average individual.

It appears that the behavior of $p_r(z(t))$ for large t will depend on the information matrix Q and one is tempted to conjecture that the fluctuations will be smaller in amplitude as the value of Q decreases, since it is clear that $\lim_{t \rightarrow \infty} p_r(z(t)) = p^*$ if Q has zero information value. For the case where only one observation is used the above conjecture is easily validated. It is not difficult to show that the matrix $K(t, t_1)$, whose i, j 'th element is $p(x_j(t_1)|z_i(t))$, $= (P_0(t))^*{}^{-1} (P^{t-t_1})^* Q$. Let us take two information matrices Q_1 and Q_2 such that $Q_2 = Q_1 M$ for some Markov matrix M . It follows easily that $K(t, t_1)_{Q_2} = K(t, t_1)_{Q_1} M$, where the subscript denotes the specific information matrix used to obtain $K(t, t_1)$.

Let $N(t, t_1)_Q$ be the $n \times m$ matrix whose i, j 'th element is

$$p(z_i(t) | x_j(t_1)), \quad i = 1, \dots, m; \quad j = 1, \dots, n,$$

where Q represents the information matrix. Then it follows from a theorem proved by Marschak and Miyasawa (1968 Theorem 8.2, p. 154) that $N(t, t_1)_{Q_2} = M^* N(t, t_1)_{Q_1}$, where M^* is some $n \times n$ Markov matrix. Since the rows of $N(t, t_1)_Q$ give the revised probabilities, this implies that the revised probabilities obtained with information structure Q_2 are convex linear combinations of those obtained with Q_1 . Thus, the set of revised probability vectors obtained with Q_2 in any period is a subset of that obtained with Q_1 for the same period. This means that the amplitude of the fluctuations of the revised probability vector will be greater as the value of the information matrix increases.

We conjecture that this result is also valid in the general case, where all past observations are used. To gain an idea of the extent of the fluctuations the model was simulated for 100 periods with the following matrices:

$$P = \begin{bmatrix} .8 & .2 \\ .3 & .7 \end{bmatrix} \quad Q_1 = \begin{bmatrix} .8 & .2 \\ .1 & .9 \end{bmatrix} \quad Q_2 = \begin{bmatrix} .6 & .4 \\ .3 & .7 \end{bmatrix}$$

In each case, the system started in state 2. When Q_1 was used, the revised probability of state 1 fluctuated between .783 and .354. On the other hand, the fluctuations ranged between .716 and .471 when Q_2 was used. Note that $p^* = (.6, .4)$ and that therefore, were Q equal to I , the revised probability would be .8 roughly 60% of the time and .3 the remaining 40% of the time.

What are the implications of this result? The most important one appears to be the following. If the action of a decision-maker in period t depends on $p_T(z(t))$ alone, one would expect more frequent changes in the optimal action as the matrix Q becomes more informative, not merely for the individual decision-maker, but also for the group as a whole. To put the matter in a "common-sense" form, economic responses will tend to follow fluctuations in the underlying

ing structure more closely if the pattern of information available to decision-makers is good and will remain more or less constant in the opposite case. An interesting possibility emerges as a corollary: if a public agency attempts to improve the flow of information to decision-makers, this might have the unexpected effect of increasing the fluctuations of the economic variables that the latter control!¹

¹Of course our model presupposes that the environment itself is subject to shifts. Also to be stressed is the hidden assumption that the Markov process described by the matrix P is unaffected by the actions of the individuals involved. In a full-fledged general equilibrium model P itself would be endogenous.

VII. The Formation of Expectations

We turn now to the revision of the probability vector of observations in future periods with the use of past and current observations. Let $F(t_2, t_1)$ be the $n^{t-t_2} \times n$ Markov matrix whose i, j 'th element is $p(x_j(t_1) | y_i(t, t_2))$, i.e., the conditional probability of observation x_j in period t_1 (where $t_1 > t$), given the particular sequence i of observations from the preceding $t-t_2$ periods ($t_2 < t$). Then it is shown in the Appendix that conditions similar to the conditions ensuring the obsolescence of information mentioned in section IV yield the result that $F(t_2, t_1) = F(t_2, t_1')M$, where M is some Markov matrix and $t_1 > t_1' > t$. What are the implications of this result? Let $f_i(t_2, t_1)$ be the i 'th row of $F(t_2, t_1)$ and x be a random variable. Then $f_i(t_2, t_1)x$, where x is the vector (x_1, \dots, x_n) of observations, is the expected value of the random variable x in period t_1 , given the sequence of observations $y_i(t, t_2)$. Let $\bar{x} = \max(x_1, \dots, x_n)$ and $\underline{x} = \min(x_1, \dots, x_n)$. If $F(t_2, t_1) = F(t_2, t_1')M$, then $F(t_2, t_1)x = F(t_2, t_1')Mx$. Since M is a Markov matrix any component of the vector Mx is a weighted average of the components of x and lies therefore between \underline{x} and \bar{x} . Thus, the range of the conditional expected value of x in period t_1 is not greater than its range in period t_1' . This suggests the inelasticity of expectations concerning future values of the mean of the random variable x . More precisely, two persons sharing the same assumptions about the transition matrix P and the information matrix Q but receiving different observations will tend to be in closer agreement about the expected mean of the random variable x for periods farther out in the future than for periods closer to the present. If these expectations affect their current decisions, we have a factor inducing individuals to take similar actions in the present, despite different subjective probability vectors about current states.

VIII. Summary

The model studied in this paper is essentially the same as those used in earlier studies by Nerlove (1967) and Taylor (1970) to justify the use of distributed lags in econometrics. The unknown states of the system are assumed to be governed by a known Markov process; information about these states is provided by a set of observations more or less closely related to the states. The matrix whose typical element is the conditional probability of a particular observation, given a particular state, is called the information matrix.

The result discussed in section IV. bears on the obsolescence of information, meaning that knowledge of an observation is less valuable on the average for current decision-making, the farther back in the past it lies. Conditions are given in terms of the matrix of transition probabilities and the information matrix that yields this result for the case where the number of observations is less than or equal to that of states. The value of the result lies in the suggestion that the model can justify the use of distributed lags with declining weights for any payoff functions, provided one is looking at the behavior of the typical individual.

In section V we argue that the uncertainty about states of the world is responsible for lags in adjustment to a shift to a new state, the length of the lag depending on the information structure, and demonstrate this by some simple examples. The long-run revised probability vectors for states of the world are shown to fluctuate in section VI, even when the prior probabilities converge to some limiting values, the extent of these fluctuations again depending on the information matrix. Finally, in section VII the same considerations giving rise to the obsolescence of information are shown to imply that the elasticity of expectations declines, the farther out in the future they lie.

REFERENCES

- Bellman, R. and S. Dreyfus, Applied Dynamic Programming. Princeton: Princeton University Press, 1962.
- Blackwell, D. and M.S. Girshick, Theory of Games and Statistical Decisions. New York: John Wiley & Sons, 1954.
- Cross, J.G., "A Stochastic Learning Model of Economic Behavior", Quarterly Journal of Economics, May 1973, Vol. 87, No. 2; pp. 239-66.
- Cyert, R.M. and M.H. deGroot, "Bayesian Analysis and Duopoly Theory", Journal of Political Economy, September-October 1970, Vol. 78(5), pp. 1168-84.
- Marschak, J., "On Adaptive Programming", Management Science, July 1963, Vol. 9(4), pp. 517-26.
- _____, "Economics of Information Systems", Journal of the American Statistical Association, March 1971, Vol. 66 (333), pp. 192-219.
- _____, and K. Miyasawa, "Economic Comparability of Information Systems", International Economic Review, June 1968, Vol. 9 (2) pp. 137-74.
- Nerlove, M., "Distributed Lags and Unobserved Components in Economic Time Series", in Ten Economic Studies in the Tradition of Irving Fisher, New York: John Wiley & Sons, 1967.
- Taylor, L., "The Existence of Optimal Distributed Lags", Review of Economic Studies, January 1970, Vol. 37(1), pp. 95-106.

In addition to assumptions A1 - A4 in section III we need:

- A5. The rank of Q is $\min(m,n)$;
- A6. The decision-maker maximizes the discounted sum of the expected value of the payoffs over all periods.
- A7. The vector $p_0 P^t$, which gives the a priori probabilities of the states in period t , is strictly positive for all t .

Because of A5 and A7, the vector $p_0 P^t Q$, which gives the a priori probabilities of observations in period t , is strictly positive for all t .

Our assumptions allow us to break down the maximization of the objective function into a series of independent maximizations, one for each period. In what follows we shall accordingly focus on the decision to be made in some specific period t alone. Let us commence with the simplest case, where there is only one past observation $x_j(t_1)$ ($t_1 < t$). Since the decision-maker's current payoff depends only on $a_i(t)$ and $x_j(t)$, the matrix $J(t, t_1)$ will determine the value of an observation from the period t_1 , as we saw in section III of the text.

$$\begin{aligned} \text{Now } p(x_j(t_1) | x_i(t)) &= p(x_j(t_1), x_i(t)) / p(x_i(t)) \\ &= \sum_{rs} p(z_r(t_1), z_s(t)) p(x_j(t_1) | z_r(t_1)) p(x_i(t) | z_s(t)) / p(x_i(t)) \\ &= \sum_{rs} p(z_r(t_1)) p(z_s(t) | z_r(t_1)) p(x_j(t_1) | z_r(t_1)) p(x_i(t) | z_s(t)) / p(x_i(t)). \end{aligned}$$

The numerator of this last expression is, for fixed value of i and j , a quadratic form in the variables $p(z_r(t_1)) p(x_j(t_1) | z_r(t_1))$ and $p(x_i(t) | z_s(t))$, where r and $s = 1, \dots, m$. It can consequently be expressed as the i, j 'th element of a certain matrix, which is the product of three matrices. Without presenting the derivation we proceed immediately to write the product matrix, which is

$$Q' (P^{t-t_1})' P_0(t_1) * Q,$$

where $P_0(t_1) *$ is the $m \times m$ diagonal matrix whose main diagonal is the vector $p_0 P^{t_1}$, and the primes indicate the transposes of the corresponding matrices.

$p(x_i(t))$ is of course the i 'th component of the vector $p_0 P^t Q$. Let $Q_0(t)^*$ denote the $n \times n$ diagonal matrix whose main diagonal is this vector. $(Q_0(t)^*)^{-1}$ exists by A7. Then it can be verified that the matrix whose i, j 'th element is $p(x_j(t_1) | x_i(t))$,

$$J(t, t_1) = (Q_0(t)^*)^{-1} Q' (P^{t-t_1})' P_0(t_1)^* Q, t_1 = 1, \dots, t-1 \quad (1)$$

The process of deriving the elements of the matrix $J(t, t_1)$, where t_1 is any positive integer $> t$, is similar and yields

$$J(t, t_1) = (Q_0(t)^*)^{-1} Q' P_0(t)^* P^{t_1-t} Q, t_1 = t+1, \dots \quad (2)$$

Suppose we wish to compare the value of an observation from period $t_1 < t$ with one from some preceding period t_2 , where $t_2 < t_1$. The two matrices of conditional probabilities,

$$J(t, t_1) = A(t) P_0(t_1)^* Q, \quad (3)$$

$$\text{and } J(t, t_2) = A(t) (P^{t_1-t_2})' P_0(t_2)^* Q, \quad (4)$$

where $A(t) = (Q_0(t)^*)^{-1} Q' (P^{t-t_1})'$.

The expected value to the decision-maker of an observation from any past period t_1

$$v(t, t_1) = \sum_k [\max_j \sum_j \tau_{1j} p(x_k(t_1) | p(x_j)) p(x_j(t))]. \quad (5)$$

How does $v(t, t_1)$ compare with $v(t, t_2)$? Blackwell's theorem states that

$$v(t, t_1) \geq v(t, t_2) \quad (6)$$

if and only if

$$J(t, t_2) = J(t, t_1) M_1, \quad (7)$$

where M_1 is some $n \times n$ Markov matrix.

From (3) and (4), it is sufficient for (7) that for some Markov matrix M_1 ,

$$(P^{t_1-t_2})' P_0(t_2)^* Q = P_0(t_1)^* Q M_1,$$

or $BQ = QM_1$, where $B = P_0(t_1)^*^{-1} (P^{t_1-t_2})' P_0(t_2)^*$.

Note that $(P_0(t_1)^*)^{-1}$ exists because of A7.

Lemma 1

B is a Markov matrix, i.e., its elements $b_{ij} \geq 0$ ($i, j = 1, \dots, m$), and $Bu = u$, where u is the m -component vector $(1, \dots, 1)$.

Proof: It is clear that B , the product of three non-negative matrices, is non-negative.

$$\text{Now } uP_o(t_2)^* = p_o P^{t_2}.$$

$$\text{Therefore } uP_o(t_2)^* P^{t_1-t_2} = p_o P^t P^{t_1-t_2} = p_o P^{t_1} = uP_o(t_1)^*.$$

$$\text{Therefore } uP_o(t_2)^* P^{t_1-t_2} (P_o(t_1)^*)^{-1} = u.$$

Taking transposes, we obtain immediately

$$(P_o(t_1)^*)^{-1} (P^{t_1-t_2})' P_o(t_2)^* u = u$$

Let $\hat{Q} = \{q \mid \sum_1 w_i q^i = q\}$, where $w_i \geq 0$ and q^i is the i 'th column of Q .

We are now ready to state

Theorem 1

Let $n \geq m$. A necessary condition that the value of information from period t_1 should be not less than the value of information from an earlier period t_2 , where $t_2 < t_1$, is: $Bq^i \in \hat{Q}$ for all i ($i=1, \dots, n$).

Proof: Since $n \geq m$ and the rank of P is m , it follows from A5 that the rank of the matrix $A(t)$ in (4) is m . $J(t, t_2) = J(t, t_1)M_1$ therefore implies that $QM_1 = BQ$, where M_1 is an $n \times n$ Markov matrix. Take the i 'th column of BQ , Bq^i . This equals Qm_1^i , where m_1^i is the i 'th column of M_1 . Since $m_1^i \geq 0$, $Bq^i \in \hat{Q}$ ($i=1, \dots, n$).

Theorem 2

Let $n \leq m$. Then a sufficient condition that the value of information from period t_1 should be not less than the value of information from an earlier period t_2 , where $t_2 < t_1$, is: $Bq^i \in \hat{Q}$ for all i ($i=1, \dots, n$).

Proof: Since the rank of Q is n by A5, there exists an $n \times m$ matrix Q^+ such that $Q^+Q = I$. $Q^+ = (Q'Q)^{-1}Q'$. Also $Q^+u_m = Q^+Qu_n = Iu_n = u_n$, where u_n (u_n) is

the n -component (m -component) vector $(1, \dots, 1)$ and we have used the fact that $Qu_n = u_m$, since Q is a Markov matrix. Now there exists by assumption a non-negative $n \times n$ matrix M_1 such that $QM_1 = BQ$. Also $BQu_n = Bu_m = u_m$. Therefore $Q^+BQu_n = Q^+u_m = u_n$. But $Q^+BQu_n = Q^+QM_1u_n = IM_1u_n$. Therefore $M_1u_n = u_n$ and M_1 is the required Markov matrix that satisfies $QM_1 = BQ$ and it follows easily that $J(t, t_2) = J(t, t_1)M_1$.

From Theorems 1 and 2 we obtain immediately

Theorem 3

Let $m = n$. Then a necessary and sufficient condition that the value of information from period t_1 should be not less than that from an earlier period t_2 is: $Bq^i \in \hat{Q}$ for all i ($i=1, \dots, m$).

Suppose we wished to compare the Markov matrices whose i, j 'th elements are $p(x_j(t_1) | z_i(t))$ and $p(x_j(t_2) | z_i(t))$ for the value of the information they provide. The two matrices, it can be easily shown, are $(P_o(t)^*)^{-1}(P^{t-t_1})'P_o(t_1)^*Q$ and $(P_o(t)^*)^{-1}(P^{t-t_2})'P_o(t_2)^*Q$ respectively, and we shall indicate them by $K(t, t_1)$ and $K(t, t_2)$. The proof of the following theorem is omitted because it parallels that of Theorems 1 and 2.

Theorem 4

(i) A necessary condition that $K(t, t_2) = K(t, t_1)M_2$ for some $n \times n$ Markov matrix M_2 is: $Bq^i \in \hat{Q}$ for all i ($i=1, \dots, n$).

(ii) This condition is sufficient if $n \leq m$.

We proceed now to the general case. Let $L(t, t_1)$ be the $n \times n^{t-t_1}$ Markov matrix whose i, j 'th element is $p(y_j(t, t_1) | x_i(t))$, where $y_j(t, t_1)$, to repeat, is a particular member of the set consisting of all possible sequences of observations from the preceding $t-t_1$ periods. Let $\hat{y}_j(t_1, t_2)$ represent the j 'th member from the set consisting of past observations from the preceding $t-t_1-1$ periods plus the observation from the t_2 'th period, where $t_2 \leq t_1$. $\hat{y}_j(t_1, t_2)$ is thus

obtained from $y_j(t, t_1)$ by replacing the t_1 'th observation by the t_2 'th one.

Let $\hat{L}(t, t_1, t_2)$ be the $n \times n^{t-t_1}$ Markov matrix whose i, j 'th element is $p(\hat{y}_j(t_1, t_2) | x_i(t))$.

Lemma 2

If $K(t, t_2) = K(t, t_1)M_2$ for some Markov matrix M_2 , then $K(t_0, t_2) = K(t_0, t_1)M_2$, where $t_1 < t_0 < t$.

Proof: By assumption $(P_0(t))^*{}^{-1}(P^{t-t_2})'P_0(t_2)^*Q = (P_0(t))^*{}^{-1}(P^{t-t_1})'P_0(t_1)^*QM_2$.

Premultiplying both sides by the matrix $(P_0(t_0))^*{}^{-1}((P^{-1})^{t-t_0})'P_0(t)^*$ we obtain $(P_0(t_0))^*{}^{-1}(P^{t_0-t_2})'P_0(t_2)^*Q = (P_0(t_0))^*{}^{-1}(P^{t_0-t_1})'P_0(t_1)^*QM_2$.

Theorem 5

If $K(t, t_2) = K(t, t_1)M_2$ for some Markov matrix M_2 , then

(i) $J(t, t_2) = J(t, t_1)M_2$, and

(ii) $\hat{L}(t, t_1, t_2) = L(t, t_1)M_3$ for some $n \times n^{t-t_1}$ Markov matrix M_3 .

Proof:

(i) By assumption

$$(P_0(t))^*{}^{-1}(P^{t-t_2})'P_0(t_2)^*Q = (P_0(t))^*{}^{-1}(P^{t-t_1})'P_0(t_1)^*QM_2.$$

Premultiplying both sides by the $n \times m$ matrix $(Q_0(t))^*{}^{-1}Q_0'P_0(t)^*$ we obtain the desired result.

(ii) The i, j 'th elements of $\hat{L}(t, t_1, t_2)$ and $L(t, t_1)$ are

$p(x_{j_{t_2}}(t_2), x_{j_{t_1+1}}(t_1+1), \dots, x_{j_{t-1}}(t-1) | x_i(t))$ and

$p(x_{j_{t_1}}(t_1), \dots, x_{j_{t-1}}(t-1) | x_i(t))$ respectively.

Now $p(x_{j_{t_2}}(t_2), x_{j_{t_1+1}}(t_1+1), \dots, x_{j_{t-1}}(t-1) | x_i(t))$

$$= p(x_{j_{t_2}}(t_2), x_{j_{t_1+1}}(t_1+1), \dots, x_{j_{t-1}}(t-1), x_i(t)) / p(x_i(t))$$

$$\begin{aligned}
&= \left[\sum_{r_{t_2}} \sum_{r_{t_1+1}} \dots \sum_{r_t} p\{(x_{j_{t_2}}(t_2), x_{j_{t_1+1}}(t_1+1), \dots, x_{j_{t-1}}(t-1), x_i(t) | \right. \\
&\quad \left. z_{r_{t_2}}(t_2), z_{r_{t_1+1}}(t_1+1), \dots, z_{r_t}(t)\} \right. \\
&\quad \left. \times p(z_{r_{t_2}}(t_2), z_{r_{t_1+1}}(t_1+1), \dots, z_{r_t}(t)) \right] / p(x_i(t)) \\
&= \left[\sum_{r_{t_2}} \sum_{r_{t_1+1}} \dots \sum_{r_t} p(x_{j_{t_2}}(t_2) | z_{r_{t_2}}(t_2)) p(x_{j_{t_1+1}}(t_1+1) | z_{r_{t_1+1}}(t_1+1)) \dots \right. \\
&\quad \left. \dots p(x_i(t) | z_{r_t}(t)) p(z_{r_{t_2}}(t_2), z_{r_{t_1+1}}(t_1+1), \dots, \right. \\
&\quad \left. z_{r_t}(t)) \right] / p(x_i(t)) \\
&= \left[\sum_{r_{t_2}} \sum_{r_{t_1+1}} \dots \sum_{r_t} p(x_{j_{t_2}}(t_2) | z_{r_{t_2}}(t_2)) p(x_{j_{t_1+1}}(t_1+1) | z_{r_{t_1+1}}(t_1+1)) \dots \right. \\
&\quad \left. \dots p(x_i(t) | z_{r_t}(t)) p(z_{r_{t_2}}(t_2) | z_{r_{t_1+1}}(t_1+1)) p(z_{r_{t_1+1}}(t_1+1) | \right. \\
&\quad \left. z_{r_{t_1+2}}(t_1+2)) \dots p(z_{r_{t-1}}(t-1) | z_{r_t}(t)) p(z_{r_t}(t)) \right] / p(x_i(t)) \\
&= \left[\sum_{r_{t_1+1}} \dots \sum_{r_t} p(x_{j_{t_2}}(t_2) | z_{r_{t_1+1}}(t_1+1)) \{p(x_{j_{t_1+1}}(t_1+1) | z_{r_{t_1+1}}(t_1+1)) \dots \right. \\
&\quad \left. \dots p(x_i(t) | z_{r_t}(t)) p(z_{r_{t_1+1}}(t_1+1) | z_{r_{t_1+2}}(t_1+2)) \dots p(z_{r_t}(t)) \} \right] \\
&\quad / p(x_i(t))
\end{aligned}$$

Now, from Lemma 2

$$p(x_{j_{t_2}}(t_2) | z_{r_{t_1+1}}(t_1+1)) = \sum_k p(x_k(t_1) | z_{r_{t_1+1}}(t_1+1)) m_{kj},$$

where m_{kj} is the k, j 'th element of the matrix M_2 . Therefore, substituting in the last but one expression for $p(x_{j_{t_2}}(t_2) | z_{r_{t_1+1}}(t_1+1))$, we obtain

$$p(x_{j_{t_2}}(t_2), x_{j_{t_1+1}}(t_1+1), \dots, x_{j_{t-1}}(t-1) | x_i(t))$$

$$\begin{aligned}
&= \left[\sum_{r_{t_1+1}} \dots \sum_{r_t} \sum_k p(x_k(t_1) | z_{r_{t_1+1}}(t_1+1))^{m_{kj}} \right] / p(x_1(t)) \\
&= \left[\sum_{r_{t_1}} \dots \sum_{r_t} \sum_k p(x_k(t_1) | z_{r_{t_1}}(t_1)) p(z_{r_{t_1}}(t_1) | z_{r_{t_1+1}}(t_1+1))^{m_{kj}} \right] \\
&\quad / p(x_1(t)) \\
&= \sum_k \{ p(x_k(t_1), x_{j_{t_1+1}}(t_1+1), \dots, x_{j_{t-1}}(t-1) | x_1(t)) \}^{m_{kj}}
\end{aligned}$$

Therefore, if the columns of the matrices $L(t, t_1)$ and $\hat{L}(t, t_1, t_2)$ are arranged so that each set of n columns contains in numerical order the set of n observations from periods t_1 and t_2 respectively, the matrix M_3 has the form

$$\begin{bmatrix}
M_2 & 0 & \dots & 0 \\
0 & M_2 & \dots & 0 \\
\dots & \dots & \dots & \dots \\
0 & 0 & & M_2
\end{bmatrix}$$

with $n^{t-t_1-1} M_2$ matrices along the principal diagonal and null matrices elsewhere. Since M_2 is a Markov matrix, M_3 will clearly be a Markov matrix as well.

From Theorems 4 and 5 we obtain immediately

Theorem 6

Let $n \leq m$ and $B_q^i \in \hat{Q}$, $i = 1, \dots, n$.

Then $\hat{L}(t, t_1, t_2) = L(t, t_1) M_3$ for some Markov matrix M_3 .

Theorem 6 provides a sufficient condition that the value of the observations from the preceding $t-t_1$ periods should be not less than that from the preceding $t-t_1-1$ periods plus the observation from period t_2 , where $t_2 < t_1 < t$.

This condition involves the matrices Q and B , the latter being $(P_0(t_1))^*^{-1}(P^{t_1-t_2})P_0(t_2)^*$. Since B (besides depending on P) depends not just on the difference t_1-t_2 , but on t_1 and t_2 , it is desirable to have a condition free of the latter dependence.

We first note that, if $J(t, t_1)M_{t_1}$ for every $t_1 < t$, then $J(t, t_2) = J(t, t_1)M$ for some Markov matrix M for any $t_2 < t_1$, a result implicit in the transitivity of the ordering by value of information. Thus it suffices to look at B for $t_2 = t_1 - 1$, i.e. $(P_0(t_1))^*^{-1}P_0(t_1-1)^*$. If $(P_0(t_1))^* \approx P_0(t_1-1)^*$, a condition that will obtain for Markov processes whose limiting probabilities exist when t is large, the elements on the principal diagonal of B will be approximately equal to those on the principal diagonal of P . Now let $B = \lambda I + (1-\lambda)C$, where $0 \leq \lambda \leq 1$ and C is any $m \times m$ Markov matrix all of whose rows are equal.¹ Then $BQ = \lambda IQ + (1-\lambda)BC = \lambda Q + (1-\lambda)D$, where D is a Markov matrix with the same property as C . Thus the columns of D are scalar multiples of the vector $u_m = (1, \dots, 1)$. But $u_m \in \hat{Q}$ because $Qu_m = u_m$; therefore the columns of D , and hence those of $(1-\lambda)D \in \hat{Q}$. Thus $B_q^i \in \hat{Q}$ for all i . Taking now λ close to 1, we find that B is a matrix with elements on the principal diagonal close to 1.

Since these elements are approximately equal to elements on the principal diagonal of P , we have a condition on the latter matrix, viz. $p_{ii} \approx 1 - \delta_i$ (δ_i small), $i = 1, \dots, m$, that should ensure that $B_q^i \in \hat{Q}$, $i = 1, \dots, n$, independently of Q . In terms of our model, this condition means that the stochastic process governing the movement from one state to another is stable, the probability being high that the environment will stay in the same state from one period to the next.

It is also clear that, for $n = m$, the condition $B_q^i \in \hat{Q}$ ($i = 1, \dots, n$) is likely to hold for $Q \approx I$ independently of B , since it is trivially satisfied in

¹It is easy to show that any 2×2 matrix $B = \begin{vmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{vmatrix}$ can be written in this form provided $b_{11} > b_{21}$ and $b_{22} > b_{12}$.

the limiting case where $Q = I$. Indeed, in this limiting case the only observation of value is that from the preceding period, for it serves to identify with probability 1 the state prevailing in that period.

We turn now to the revision of the probability distribution of observations in future periods as a result of past and current observations. The theorems here are essentially similar to the preceding ones, and we shall therefore dispense with proofs in general. If $K(t, t_1)$ is the $m \times n$ Markov matrix whose i, j 'th element is $p(x_j(t_1) | z_i(t))$, where $t_1 > t$, then

$$K(t, t_1) = P^{t_1 - t} Q, \quad t_1 = t+1, t+2, \dots$$

Similar to Theorem 4 is

Theorem 7

Let $t_1 > t_1' > t$.

(i) A necessary condition that $K(t, t_1) = K(t, t_1') M_4$ for some Markov matrix M_4 is:

$$P^{t_1 - t_1'} q^i \in \hat{Q} \text{ for all } i \text{ (} i=1, \dots, n \text{)}.$$

(ii) If $n \leq m$, a sufficient condition that $K(t, t_1) = K(t, t_1') M_4$ for some Markov matrix M_4 is:

$$P q^i \in \hat{Q} \text{ for all } i.$$

Let $t_2 < t < t_1$. In section VII the matrix $F(t_2, t_1)$ was defined as the $n^{t-t_2} \times n$ Markov matrix whose i, j 'th element is $p(x_j(t_2) | y_i(t, t_2))$. A theorem similar to Theorem 5 (ii) is

Theorem 8

If $K(t, t_1) = K(t, t_1') M_4$ for some Markov matrix M_4 , where $t_1 > t_1'$, then $F(t_2, t_1) = F(t_2, t_1') M_4$.

From Theorem 7(ii) and Theorem 8 we obtain

Theorem 9

If $n \leq m$, a sufficient condition that $F(t_2, t_1) = F(t_2, t_1') M_4$ for some Markov matrix is:

$$P q^i \in \hat{Q} \text{ for all } i.$$