Economic Theory

A 20981

Volume 9 Number 2 1997

ECTHEA 9 (2) 197–384 February 1997
Printed on acid-free paper

Springer

0938-2259(199702)9:2;1-1

Economic Theory 9, 309–324 (1997)

Economic Theory

The rationality and efficacy of decisions under uncertainty and the value of an experiment*

Stephen Morris[1] **and Hyun Song Shin**[2]

[1] Department of Economics, University of Pennsylvania, Philadelphia, PA 19104-6297, USA

[2] Nuffield College, Oxford OX1 1NF, ENGLAND

Received: March 17, 1995; revised version September 11, 1995

Summary. A decision maker faces a known prior distribution over payoff relevant states. We compare the expected utility of this individual under two scenarios. In the first, the decision maker makes a choice without further information. In the second, the decision maker has access to an experiment before choosing an action. However, the decision maker does not know the true joint distribution over states and messages. The value of the experiment as measured by the difference in the two utility levels can be negative as well as positive. We give a condition which is necessary and sufficient for the experiment to be valuable in our sense, for any decision problem.

JEL Classification Number: D81.

1 Introduction

Actions are motivated by beliefs and preferences. The judgement as to whether an action is rational is one based on the consistency of ends and means given beliefs. However, if the decision maker does not know all relevant facts about the world, it is possible to draw a distinction between the *rationality* of an action and its *efficacy* in furthering the goals of the decision maker. The judgement as to the efficacy of an action is based on whether the action *in fact* furthers the goals of the decision maker. It is concerned with how the world is, rather than with what the decision maker believes.

* An earlier version of this paper was circulated under the title "Noisy Bayes Updating and the Value of Information." We have gained from the comments of Stephen Coate, John Geanakoplos, Larry Samuelson, Timothy Van Zandt and seminar participants at Harvard Business School, Princeton, Boston University, the international conference on game theory at Stony Brook 1992 and the Winter meeting of the Econometric Society at Anaheim 1993. The first author received support for this project from NSF grant #SES-9308515 and a University of Pennsylvania Research Foundation Grant.

Correspondence to: S. Morris

Although the distinction between rationality and efficacy is of little comfort to the decision maker, it is of considerable interest to an outside observer who is concerned with the value to the decision maker of receiving additional information. To explain this point, consider the plight of two decision makers, A and B. Both decision makers know the true prior distribution over a set of payoff relevant states, have identical preferences, and are facing identical decision problems. Moreover, they have identical beliefs concerning the joint distribution over the payoff relevant states and a set of messages yielded by an experiment. However, although they share the same beliefs concerning the experiment, these beliefs may not coincide with the true probabilities over state-message pairs. Crucially, while A observes the outcome of the experiment, B does not. Hence, A will condition his action on the message observed, based on his subjective beliefs, while B will choose an action which is optimal given the known prior distribution over the payoff relevant states alone. Who is better off?

This question is posed in terms of the ex ante expected utilities of the two decision makers. Given that A will follow an optimal decision rule given beliefs, an outside observer who knows the true joint distribution over states and messages will be able to calculate the ex ante expected utility of A. The ex ante expected utility of B is simply the expected utility from choosing an optimal action in the absence of any experiment. We may then regard the *value* of the experiment as the difference between the ex ante expected utility of A over the expected utility of B.

The value of an experiment in our sense can be positive or negative. Its sign depends on the form which the decision problem takes, such as the set of available actions and the corresponding payoffs. It also depends on the relationship between the true distribution over state-message pairs governing the experiment and the *subjective* beliefs over state-message pairs. If beliefs wander 'too far' from the truth, B may be better off, so that the experiment has negative value. Then the question becomes how far is 'too far'?

This is the issue addressed in this paper. For the record, we should emphasize that the question addressed in this paper is different from that addressed by Blackwell [1]. Blackwell envisaged a decision maker who knows his environment, including the nature of the noise affecting his experiment. In such a setting, with expected utility maximization, an experiment can never have a negative value. The decision maker in our paper is someone who does not know every detail of his environment, and the value of information can be negative as well as positive. (Our approach appears to be closer to that of Lehrer [4] who examines the value of accepting advice; however, our results do not appear to be directly comparable).

Our question will be posed in the following framework. A decision maker faces a finite set of payoff relevant states with an objective, and known prior distribution over these states. We now consider the expected utility of this individual under two scenarios. In the first, the decision maker makes a choice without further information. The action chosen will be optimal given the prior distribution over the payoff relevant states. In the second scenario, the decision maker has access to an experiment, and observes the outcome of this experiment before choosing an action. However, the decision maker does not know the true joint distribution over states and messages and has to rely on his beliefs. We do not model the process through which these beliefs are formed, but it is the crucial assumption of this paper that these beliefs may diverge from the true distribution over state-message pairs. In each of the two scenarios, we can calculate the (true) expected utility of the decision maker.

Within this framework, we characterize the relationship between the prior distribution over the payoff relevant states and the subjective beliefs over states and messages that is necessary and sufficient for the decision maker never to be made worse off under the second scenario as compared to the first, in *any* decision problem. In other words, we characterize the relationship between two probability distributions which ensures that the experiment has positive value for any decision problem. We emphasize that the question is not whether the decision maker would choose to observe the signal, but rather whether he will be made better off ex ante, from the viewpoint of the analyst who knows the true distribution.

The applicability of the ideas in this paper is at its greatest when we have the conjunction of two features: an innovative (and perhaps controversial) experiment which is put forward as an aid to decision in a common and well-understood problem. Such a combination would come closest to the assumption in this paper that, while the prior distribution over the payoff relevant states is known, the joint distribution governing the experiment is unknown. New diagnostic tests for common, well-understood illnesses would be one prime example from the field of medicine. Indeed, we shall pursue this medical theme in our worked example, below. The example will help in posing our question in a sharper way, and in relating our question to other standard questions in decision theory.

2 An example

Our decision maker is the unfortunate victim of a dog bite, and faces the real prospect of developing rabies. It is commonly acknowledged that our decision maker will develop rabies with probability $\pi > 0$ if untreated. A course of treatment is available if the decision maker acts immediately, but the treatment is long and painful. However, this inconvenience is small compared to the cost inflicted by untreated rabies. The treatment incurs a cost of c in utility terms, while the disease itself entails a cost of k in utility terms. The utility of not developing rabies and not undergoing the treatment is set at zero.

In the absence of any further information, our decision maker will seek treatment if $\pi > \frac{c}{k}$. However, a new test for rabies has been developed, and our decision maker has access to the test results before making a decision. Because the test is new, the accuracy of the test is uncertain. The clinic estimates that the test has a false "positive" rate of ε_p and a false "negative" rate of ε_n. In other words, the clinic estimates that proportion ε_p of patients who do not have rabies test "positive", while proportion ε_n of patients who do have rabies test "negative". However, these numbers are merely estimates, and may not correspond to the actual false positive and false negative rates.

Based on the clinic's estimates, the probability that our decision maker has rabies conditional on a positive result is:

$$\gamma_p = \frac{\pi(1-\varepsilon_n)}{\pi(1-\varepsilon_n)+(1-\pi)\varepsilon_p}. \tag{2.1}$$

The probability of rabies conditional on a negative result is:

$$\gamma_n = \frac{\pi\varepsilon_n}{\pi\varepsilon_n+(1-\pi)(1-\varepsilon_p)}. \tag{2.2}$$

We will assume that $\varepsilon_p + \varepsilon_n < 1$. This ensures that $\gamma_p > \pi > \gamma_n$, since $\gamma_p/\gamma_n > 1$ if and only if $(1-\varepsilon_n)(1-\varepsilon_p) > \varepsilon_n\varepsilon_p$ and π is a convex combination of γ_p and γ_n. This assumption is essentially without loss of generality: we can always label "p" the result which is believed to be correlated with rabies.

To the extent that the clinic's estimate of the false positive and false negative rates may diverge from the actual probabilities, it is an open question as to whether the decision maker will be made better off as a result of heeding the test results. Say that the *actual* false positive and false negative rates are given by ε_p^* and ε_n^* respectively. Then, the true probability of rabies conditional on a positive result is $\gamma_p^* = \pi(1-\varepsilon_n^*)/(\pi(1-\varepsilon_n^*)+(1-\pi)\varepsilon_p^*)$ while the true probability of rabies conditional on a negative result is $\gamma_n^* = \pi\varepsilon_n^*/(\pi\varepsilon_n^* + (1-\pi)(1-\varepsilon_p^*))$. We allow for the possibility that $\gamma_p^* < \gamma_n^*$, so that the result identified as "positive" actually makes rabies less likely.

We shall now argue that our decision maker is better off in ex ante terms for all values of c and k by having access to the test (as compared to acting on the prior π alone) if and only if

$$\gamma_n^* \leq \gamma_n \leq \pi \leq \gamma_p \leq \gamma_p^*. \tag{2.3}$$

Thus it is necessary that the decision maker is qualitatively correct in his interpretation of information, but he puts less weight on the information than he should according to Bayes rule. We shall first present an argument which is *ad hoc* to this example. It is designed to draw out some suggestive features of the problem. However, we will have a chance to return to this example in section 4 after we have presented the general result, where we can be much more explicit about how this argument fits into the general framework.

Let us first consider the 'if' part of the claim. Suppose that $\gamma_n^* \leq \gamma_n < \pi < \gamma_p \leq \gamma_p^*$. If a positive result is observed, our decision maker will make the same choice as he would have made according to the actual probabilities *unless* $\gamma_p \leq \frac{c}{k} < \gamma_p^*$. However, in this case he would not have been better off without the test, since the prior probability π would have recommended the same course of action as the posterior belief γ_p. Similarly, if the test result is negative, our decision maker will make the correct decision unless $\gamma_n^* < \frac{c}{k} \leq \gamma_n$.

But in this case, he would not have been better off without the test, since the prior probability π would have recommended the same course of action as the belief γ_n. Thus, in ex ante terms, the expected utility of heeding the test result is at least as high as that from acting in accordance with π alone.

Conversely, suppose that (2.3) fails to hold. We will show that for some values of $\frac{c}{k}$, the decision maker is worse off in ex ante terms *with* the test than without. Since π is a convex combination of γ_n^* and γ_p^*, there are four cases to consider.

If $\gamma_n^* \leq \gamma_n \leq \pi \leq \gamma_p^* < \gamma_p$, then for those values of $\frac{c}{k}$ such that $\gamma_p^* < \frac{c}{k} < \gamma_p$, our decision maker will accept the treatment following a positive test result $\left(\frac{c}{k} < \gamma_p\right)$ even though he should have rejected it $\left(\gamma_p^* < \frac{c}{k}\right)$. However, the prior π gives the *correct* recommendation whereas the posterior belief γ_p gives the wrong recommendation. Hence, conditional on observing a positive result, the decision maker would have been better off without the test. If the test is negative, the decision maker rejects the treatment. Moreover, he would have rejected the treatment anyway if he were acting in accordance with π alone. Hence, in ex ante terms, the decision maker is better off *without* the test.

Similarly, if $\gamma_n < \gamma_n^* \leq \pi \leq \gamma_p \leq \gamma_p^*$, then for those values of $\frac{c}{k}$ such that $\gamma_n < \frac{c}{k} < \gamma_n^*$, the decision maker will reject the treatment following a negative test result $\left(\frac{c}{k} > \gamma_n\right)$ even though he should have accepted it $\left(\gamma_n^* > \frac{c}{k}\right)$. However, the prior π gives the *correct* recommendation whereas the posterior belief γ_n gives the wrong recommendation. Hence, conditional on observing a negative result, the decision maker would have been better off without the test. If the test is positive, the decision maker accepts the treatment. Moreover, he would have accepted the treatment anyway if he were acting in accordance with π alone. In ex ante terms, therefore, he is better off without the test.

If $\gamma_n < \gamma_n^* \leq \pi \leq \gamma_p^* < \gamma_p$, then there are *two* intervals of value of $\frac{c}{k}$ for which the above argument holds.

Since we know that $\gamma_n < \pi < \gamma_p$ and π is a convex combination of γ_n^* and γ_p^*, it remains only to check what happens if $\gamma_p^* < \pi < \gamma_n^*$. But then we must have $\gamma_p > \pi > \gamma_p^*$ and $\gamma_n^* > \pi > \gamma_n$. In this case, the decision maker's informed choices are never better than his uninformed ones; and, for $\gamma_p^* < \frac{c}{k} < \pi$, the decision maker is made strictly worse off observing a positive result; and, for $\gamma_n^* > \frac{c}{k} > \pi$, the decision maker is made strictly worse off observing a negative result.

Thus we have exhausted all relevant violations of (2.3). In each case, there is a value of $\frac{c}{k}$ for which the decision maker is better off in ex ante terms without the test than with it, thereby demonstrating that (2.3) is necessary and sufficient for the test to be of value for every value of $\frac{c}{k}$. Our objective in this paper is to see how far this type of reasoning may be generalized.

3 Value of information

Let us now pose the problem in a more general framework in a way which parallels the example we have examined above. Let Ω be a finite set of payoff relevant states. The true probability distribution over Ω is denoted by π, and it is assumed that the decision maker knows π. There is a signal which takes values in a finite set S. We call S the set of *messages*. The true probability of the message $s \in S$ conditional on the state being $\omega \in \Omega$ is denoted by

$$\beta^*(s|\omega).$$

However, the decision maker does not know the joint distribution over states and messages. Instead, he believes that the probability of the message s conditional on state ω is governed by a distribution β, say, which may diverge from β^*. We do not model the process by which the beliefs β are formed, but it is crucial for our analysis that β may diverge from β^*. The decision makers' posterior probability of ω given the message s is then given by:

$$\gamma(\omega|s) = \frac{\pi(\omega)\beta(s|\omega)}{\sum_{\omega'} \pi(\omega')\beta(s|\omega')}. \tag{3.1}$$

A *decision problem* is a pair (A, u) consisting of a finite set of actions A, and a utility function $u: A \times \Omega \to \Re$, where $u(a, \omega)$ is the utility of action a at state ω. A *decision rule* is a mapping $f: S \to A$ which associates an action for each message.

We say that a decision rule f is *optimal given beliefs* γ if:

$$\sum_{\omega} \gamma(\omega|s)u(f(s), \omega) \geq \sum_{\omega} \gamma(\omega|s)u(a, \omega), \quad \forall a \in A, \forall s \in S. \tag{3.2}$$

In other words, the decision rule maximizes subjective expected utility for every message. We contrast this with our final definition. We say that the beliefs γ are *valuable given* π if, for every decision problem (A, u), every decision rule f which is optimal given beliefs γ, and every action $a \in A$, we have:

$$\sum_{\omega} \pi(\omega) \sum_{\omega} \beta^*(s|\omega)u(f(s), \omega) \geq \sum_{\omega} \pi(\omega)u(a, \omega). \tag{3.3}$$

In other words, the beliefs γ are valuable given π if, for every decision problem, the (true) expected utility from following the optimal decision rule given beliefs γ is no lower than choosing an action based solely on π. This condition is a formalization of the idea that the decision maker cannot be made worse off in ex ante terms by acting in accordance with the posterior beliefs γ generated by the experiment. Our result characterizes the relationship between π and γ which is necessary and sufficient for γ to be valuable given π.

Theorem 3.1. *The beliefs γ are valuable given π if and only if there exist $\lambda: S \to \Re_+$ and $\mu: S^2 \to \Re_+$ such that, for all $\omega \in \Omega$ and $s \in S$,*

$$\pi(\omega)\beta^*(s|\omega) = \lambda(s)\gamma(\omega|s) + \sum_{s' \neq s} (\mu(s, s')\gamma(\omega|s) - \mu(s', s)\gamma(\omega|s')). \tag{3.4}$$

Postponing the interpretation of this condition until section 5, we first present a proof. Let us take the proof in stages. The basic strategy behind the proof is to formulate the problem in such a way that checking for the valuable information condition can be reduced to checking whether a non-negative vector lies in a closed convex cone. The proof is then completed by appeal to Farkas' lemma. We begin with the following preliminary result.

Lemma 3.2. *There exists a decision problem (A, u), a decision rule f, and an action $a^* \in A$ such that*

$$\sum_{\omega} \gamma(\omega|s)u(f(s), \omega) \geq \sum_{\omega} \gamma(\omega|s)u(a, \omega), \quad \forall s \in S, \forall a \in A \tag{3.5}$$

$$\sum_{\omega} \pi(\omega) \sum_{s} \beta^*(s|\omega)u(f(s), \omega) < \sum_{\omega} \pi(\omega)u(a^*, \omega) \tag{3.6}$$

if and only if there exists a function $x: \Omega \times S \to \Re$ such that

$$\sum_{\omega} \gamma(\omega|s)x(\omega, s) \geq \sum_{\omega} \gamma(\omega|s)x(\omega, s'), \quad \forall s, s' \in S \tag{3.7}$$

$$\sum_{\omega} \gamma(\omega|s)x(\omega, s) \geq 0, \quad \forall s \in S \tag{3.8}$$

$$\sum_{\omega} \pi(\omega) \sum_{s} \beta^*(s|\omega)x(\omega, s) < 0. \tag{3.9}$$

Proof. ('if') Suppose there is a function x satisfying (3.7) to (3.9). Consider the decision problem (A, u), where $A \equiv S \cup \{a^*\}$, $u(s, \omega) \equiv x(\omega, s)$, and $u(a^*, \omega) \equiv 0$ for any ω. Consider the decision rule $f: S \to A$ which is the identity map. That is, $f(s) = s$ for any s. We claim that (A, u), f and a^* satisfy (3.5) and (3.6). To verify the former, note that

$$\begin{aligned} \sum_{\omega} \gamma(\omega|s)u(f(s), \omega) &= \sum_{\omega} \gamma(\omega|s)x(\omega, s), && \text{(by construction)} \\ &\geq \sum_{\omega} \gamma(\omega|s)x(\omega, s'), && \text{(by (3.7))} \\ &= \sum_{\omega} \gamma(\omega|s)u(f(s'), \omega), && \text{(by construction)} \end{aligned}$$

so that (3.5) holds for actions in $A \backslash \{a^*\}$. To see that (3.5) holds for a^* also, note

$$\begin{aligned} \sum_{\omega} \gamma(\omega|s)u(f(s), \omega) &= \sum_{\omega} \gamma(\omega|s)x(\omega, s), && \text{(by construction)} \\ &\geq 0 && \text{(by (3.8))} \\ &= \sum_{\omega} \gamma(\omega|s)u(a^*, \omega), && \text{(by construction)} \end{aligned}$$

Thus, (3.5) holds generally. We can verify that (3.6) holds from:

$$\begin{aligned}\sum_{\omega}\pi(\omega)\sum_{s}\beta^*(s|\omega)u(f(s),\omega) &= \sum_{\omega}\pi(\omega)\sum_{s}\beta^*(s|\omega)x(\omega,s) && \text{(by construction)}\\ &< 0 && \text{(by (3.9))}\\ &= \sum_{\omega}\pi(\omega)u(a^*,\omega). && \text{(by construction)}\end{aligned}$$

('only if') Suppose there is a decision problem (A, u), a decision rule f, and an action $a^*\in A$ satisfying (3.5) and (3.6). We define the function $x: \Omega\times S\to\Re$ as $x(\omega, s)\equiv u(f(s),\omega)-u(a^*,\omega)$, so that

$$u(f(s),\omega)=x(\omega,s)+u(a^*,\omega). \tag{3.10}$$

We claim that x satisfies (3.7) to (3.9). First, since the function f takes its value in the set A, (3.5) implies:

$$\sum_{\omega}\gamma(\omega|s)u(f(s),\omega)\geq\sum_{\omega}\gamma(\omega|s)u(f(s'),\omega). \quad \forall s,s'\in S \tag{3.11}$$

Substituting (3.10) yields

$$\sum_{\omega}\gamma(\omega|s)x(\omega,s)\geq\sum_{\omega}\gamma(\omega|s)x(\omega,s'). \quad \forall s,s'\in S,$$

which is (3.7). To verify (3.8), note that since $a^*\in A$, (3.5) implies:

$$\sum_{\omega}\gamma(\omega|s)(u(f(s),\omega)-u(a^*,\omega))\geq 0, \quad \forall s\in S \tag{3.12}$$

and substituting (3.10) gives $\sum_{\omega}\gamma(\omega|s)x(\omega,s)\geq 0$ for all s, which is (3.8). Finally, to verify (3.9), note that (3.6) implies

$$\sum_{\omega}\pi(\omega)\sum_{s}\beta^*(\omega|s)(u(f(s),\omega)-u(a^*,\omega))<0,$$

so that substituting (3.10) yields $\sum_{\omega}\pi(\omega)\sum_{s}\beta^*(s|\omega)x(\omega,s)<0$, which is (3.9). This completes the proof of lemma 3.2.

The second plank in our proof of theorem 3.1 is Farkas' lemma, a well-known result from the theory of convex sets. See, for example, Gale [2, page 44].

Lemma 3.3. *Let Q be an $a\times b$ matrix and y be a row vector with b components. Then, there exists either a non-negative a-vector α such that $y=\alpha Q$, or a b-vector x such that $yx<0$ and $Qx\geq 0$, but not both.*

Farkas' lemma has the following geometric interpretation. Consider the convex cone in $\Re^b$ generated by the rows of the matrix Q, and interpret y as a point in $\Re^b$. Then, either y is an element of the convex cone (in which case $y=\alpha Q$ for some $\alpha\geq 0$), or y is *not* an element of the convex cone, in which case there is a separating hyperplane through the origin with normal x which separates y from the convex cone (which is the "or" clause). Clearly, these possibilities are mutually exclusive.

Armed with these two preliminary results, we can now complete the proof of theorem 3.1. Suppose Ω has n elements, and denote by ω_i the ith state. Suppose that S has m elements, and denote by s_j is the jth message. We denote by $\gamma(s_i)$ the row vector:

$$\gamma(s_i)\equiv(\gamma(\omega_1|s_i),\gamma(\omega_2|s_i),\ldots,\gamma(\omega_n|s_i)) \tag{3.13}$$

and denote by G the block diagonal matrix:

$$G=\begin{bmatrix}\gamma(s_1) & & & \\ & \gamma(s_2) & & \\ & & \ddots & \\ & & & \gamma(s_m)\end{bmatrix} \tag{3.14}$$

G has m rows and mn columns. Next, we denote by $x(s_i)$ the column vector:

$$x(s_i)=\begin{bmatrix}x(\omega_1,s_i)\\ x(\omega_2,s_i)\\ \vdots\\ x(\omega_n,s_i)\end{bmatrix} \tag{3.15}$$

and denote by x the stacked vector:

$$x=\begin{bmatrix}x(s_1)\\ x(s_2)\\ \vdots\\ x(s_m)\end{bmatrix} \tag{3.16}$$

So, x is a column vector with mn components. Using this notation, we can express the set of inequalities given by (3.8) as:

$$Gx\geq 0. \tag{3.17}$$

The set of inequalities given by (3.7) can be expressed in a similar way. We shall denote by B_i the $(m-1)\times mn$ matrix which satisfies the following four conditions.

(i) For each row of B_i, the entries from $(i-1)n+1$ to in are given by $\gamma(s_i)$.

(ii) For the jth row of B_i, where $j<i$, the entries from $(j-1)n+1$ to jn are given by $-\gamma(s_j)$.

(iii) For the jth row of B_i, where $j\geq i$, the entries from $jn+1$ to $(j+1)n$ are given by $-\gamma(s_j)$.

(iv) All other entries of B_i are zero.

The matrix B_i can be represented as follows.

$$B_i = \begin{bmatrix} -\gamma(s_i) & & \gamma(s_i) & & & \\ & \ddots & \vdots & & & \\ & & -\gamma(s_i) \quad \gamma(s_i) & & & \\ & & \gamma(s_i) & -\gamma(s_i) & & \leftarrow i\text{th row} \\ & & \vdots & & \ddots & \\ & & \gamma(s_i) & & & -\gamma(s_i) \end{bmatrix} \tag{3.18}$$

Then the set of inequalities in (3.7) can be expressed as:

$$B_i x \geq 0, \quad \text{for all } i \in \{1, 2, \ldots, m\}. \tag{3.19}$$

Finally, denote by $\pi(s_i)$ the row vector:

$$\pi(s_i) \equiv (\pi(\omega_1)\beta^*(s_i|\omega_1), \pi(\omega_2)\beta^*(s_i|\omega_2), \ldots, \pi(\omega_n)\beta^*(s_i|\omega_n)) \tag{3.20}$$

and denote by π the row vector:

$$\pi \equiv (\pi(s_1), \pi(s_2), \ldots, \pi(s_m)). \tag{3.21}$$

Then, inequality (3.9) can be expressed as: $\pi x < 0$. Taken together, there is no function $x: \Omega \times S \to \Re$ satisfying (3.7), (3.8) and (3.9) if and only if there is no vector x such that:

$$B_i x \geq 0 \quad \text{for all } i, \tag{3.22}$$

$$Gx \geq 0, \tag{3.23}$$

$$\pi x < 0. \tag{3.24}$$

We can now appeal to Farkas' lemma. There is no such x if and only if π is a non-negative combination of the rows of the matrices $\{G, B_1, B_2, \ldots, B_m\}$. By lemma 3.2, the non-existence of such an x is equivalent to the statement that γ is valuable given π. Hence, γ is valuable given π if and only if π is a non-negative combination of the rows of the matrices $\{G, B_1, B_2, \ldots, B_m\}$.

Now, visualize a non-negative combination of the rows of $\{G, B_1, B_2, \ldots, B_m\}$, where $\lambda(s_i)$ is the weight on the ith row of G, and for $i > j$, $\mu(s_i, s_j)$ is the weight on the jth row of B_i, while for $i < j$, $\mu(s_i, s_j)$ is the weight on the $(j-1)$th row of B_i. This non-negative combination results in a row vector with mn entries:

$$y \equiv (\underbrace{y_{11}, y_{12}, \ldots, y_{1n}}_{n} | \underbrace{y_{21}, y_{22}, \ldots, y_{2n}}_{n} | \cdots | \underbrace{y_{m1}, y_{m2}, \ldots, y_{mn}}_{n}),$$

where

$$y_{ij} = \lambda(s_i)\gamma(\omega_j|s_i) + \sum_{k \neq i} (\mu(s_i, s_k)\gamma(\omega_j|s_i) - \mu(s_k, s_i)\gamma(\omega_j|s_k)). \tag{3.25}$$

Thus, γ is valuable given π if and only if $\pi = y$ for some non-negative $\lambda(\cdot)$ and $\mu(\cdot, \cdot)$. That is, $y_{ij} = \pi(\omega_j)\beta^*(s_i|\omega_j)$. However, this is precisely the statement (3.4). In conclusion then, γ is valuable given π if and only if there are non-negative λ and μ such that (3.4) holds. This proves theorem 3.1.

4 The rabies example revisited

Now that we have identified the general condition for γ to be valuable given π, let us return to the example with which we started, and see how the *ad hoc* argument given earlier matches the condition identified in theorem 3.1. First, note that in the example we considered the restricted class of decision problems parameterized by the cost c of the test and the cost k of having rabies. In fact, this class of decision problems is without loss of generality given the two state payoff relevant states. Suppose that more generally the payoff is A if sick and treated, B if sick and untreated, C if well and treated and D if well and untreated. Consider generic problems with no ties. In the non-trivial case, we have $A > B$ and $D > C$ i.e. one wants treatment if and only if sick. Then, the whole problem boils down to what the critical probability α is when the decision maker is indifferent, where α solves $\alpha A + (1-\alpha)B = \alpha C + (1-\alpha)D$. This α plays exactly the same role as $\frac{c}{k}$.

Now if the *ad hoc* argument is correct, we ought to have an equivalence between the requirement that $\gamma_n^* \leq \gamma_n \leq \pi \leq \gamma_p \leq \gamma_p^*$ and the condition identified by theorem 3.1. Let us demonstrate that this is indeed the case.

In the terminology of the general framework, the set Ω has two elements, corresponding to whether the decision maker does or does not develop rabies. Let $\Omega \equiv \{Y, N\}$, where Y is the state in which he does, and N is the state in which he does not. The set of message S also has two elements, corresponding to a positive or negative test result. Let $S \equiv \{p, n\}$, with the obvious interpretation. Then, the valuable information condition is a statement concerning the existence of non-negative weights $\lambda(\cdot)$ and $\mu(\cdot, \cdot)$ which satisfy four equations, corresponding to each element of $\Omega \times S$. These equations are:

(Y, p) $$\pi(1 - \varepsilon_n^*) = \lambda(p)\gamma_p + \mu(p, n)\gamma_p - \mu(n, p)\gamma_n \tag{i}$$

(Y, n) $$\pi\varepsilon_n^* = \lambda(n)\gamma_n + \mu(n, p)\gamma_n - \mu(p, n)\gamma_p \tag{ii}$$

(N, p) $$(1 - \pi)\varepsilon_p^* = \lambda(p)(1 - \gamma_p) + \mu(p, n)(1 - \gamma_p) - \mu(n, p)(1 - \gamma_n) \tag{iii}$$

(N, n) $$(1 - \pi)(1 - \varepsilon_p^*) = \lambda(n)(1 - \gamma_n) + \mu(n, p)(1 - \gamma_n) - \mu(p, n)(1 - \gamma_p) \tag{iv}$$

We show that, under the assumption $\gamma_p > \gamma_n$, we have $\gamma_n^* \leq \gamma_n \leq \pi \leq \gamma_p \leq \gamma_p^*$ if and only if there exist non-negative λ and μ for which conditions (i) to (iv) hold.

Proof. ('if') Suppose there are λ and μ for which (i) to (iv) hold. By summing these conditions, we have:

$$(i) + (ii) \quad \pi = \lambda(p)\gamma_p + \lambda(n)\gamma_n \tag{4.1}$$

$$(i) + (iii) \quad \pi(1 - \varepsilon_n^*) + (1 - \pi)\varepsilon_p^* = \lambda(p) + \mu(p, n) - \mu(n, p) \tag{4.2}$$

$$(ii) + (iv) \quad \pi\varepsilon_n^* + (1 - \pi)(1 - \varepsilon_p^*) = \lambda(n) + \mu(n, p) - \mu(p, n) \tag{4.3}$$

$$(i) + (ii) + (iii) + (iv) \quad 1 = \lambda(p) + \lambda(n) \tag{4.4}$$

From (4.1) and (4.4), π is a convex combination of γ_n and γ_p, so that:

$$\gamma_n \leq \pi \leq \gamma_p. \tag{4.5}$$

Rearranging (i) gives:

$$\gamma_n = \frac{\pi(1-\varepsilon_n^*)}{\lambda(p)+\mu(p,n)} + \frac{\mu(n,p)}{\lambda(p)+\mu(p,n)}\gamma_n. \tag{4.6}$$

From (i), (iii) and the fact that $\gamma_p^* = \pi(1-\varepsilon_n^*)/(\pi(1-\varepsilon_n^*)+(1-\pi)\varepsilon_p^*)$, we obtain:

$$\gamma_p = \frac{\lambda(p)+\mu(p,n)-\mu(n,p)}{\lambda(p)+\mu(p,n)}\gamma_p^* + \frac{\mu(n,p)}{\lambda(p)+\mu(p,n)}\gamma_n, \tag{4.7}$$

so that γ_p is the convex combination of γ_p^* and γ_n. Since we have assumed that $\gamma_p > \gamma_n$, we have $\gamma_p \le \gamma_p^*$. From an analogous argument, we can show that $\gamma_n^* \le \gamma_n$. Thus, together with (4.5), we have $\gamma_n^* \le \gamma_n \le \pi \le \gamma_p \le \gamma_p^*$.

('only if') Suppose that $\gamma_n^* \le \gamma_n \le \pi \le \gamma_p \le \gamma_p^*$, and $\gamma_n < \gamma_p$. Define $\lambda(p)$ and $\lambda(n)$ as:

$$\lambda(p) \equiv \frac{\pi - \gamma_n}{\gamma_p - \gamma_n} \quad \text{and} \quad \lambda(n) \equiv \frac{\gamma_p - \pi}{\gamma_p - \gamma_n}. \tag{4.8}$$

and define $z \equiv \pi(1-\varepsilon_n^*) + (1-\pi)\varepsilon_p^*$. Note that $0 \le z \le 1$, since $\pi = z\gamma_p^* + (1-z)\gamma_n^*$. The function $\mu(\cdot,\cdot)$ is defined as:

$$\mu(p,n) \equiv \frac{1-z}{\gamma_p - \gamma_n}(\gamma_n - \gamma_n^*) \quad \text{and} \quad \mu(n,p) \equiv \frac{z}{\gamma_p - \gamma_n}(\gamma_p^* - \gamma_n). \tag{4.9}$$

By construction, both λ and μ have non-negative values. It can now be verified that when λ and μ are defined thus, equations (i) to (iv) hold. For example, to see that (i) holds,

$$\begin{aligned}
\lambda(p)\gamma_p + \mu(p,n)\gamma_p - \mu(n,p)\gamma_n &= \frac{1}{\gamma_p - \gamma_n}\{(\pi-\gamma_n)\gamma_p + (1-z)(\gamma_n - \gamma_n^*)\gamma_p - z(\gamma_p^* - \gamma_p)\gamma_n\} \\
&= \frac{1}{\gamma_p - \gamma_p}\{\gamma(z\gamma_p^* + (1-z)\gamma_n^*) - (1-z)\gamma_n^*\gamma_p - z\gamma_p^*\gamma_n\} \\
&= z\gamma_p^* \\
&= \pi(1-\varepsilon_n^*),
\end{aligned}$$

which is (i). (In the second equality, we have used the fact that $\pi = z\gamma_p^* + (1-z)\gamma_n^*$). Equations (ii), (iii) and (iv) can be shown in a similar way by straightforward substitution.

Let us make one final observation about the example. It turns out to be the case that whenever the decision maker is made better off in ex ante terms by observing the experiment, he is also made better off conditional on each signal he might observe. This was a special property of the example.

5 An interpretation

The condition for an experiment to be valuable identified in theorem 3.1 has some unfamiliar features. Of particular interest are the expressions involving the weights $\mu(\cdot,\cdot)$ in (3.4). To our knowledge, such expressions have not been seen previously in the theory of choice under uncertainty. It will be useful, therefore, to trace our steps and see where these expressions come from, and what role they play in the analysis.

First of all, we should emphasize that (3.4) also gives us some fairly familiar and easily interpretable conditions, too. For instance, if we sum (3.4) over the messages s, the expressions involving the μ cancel out, leaving us with:

$$\pi(\omega) = \sum_s \lambda(s)\gamma(\omega|s). \tag{5.1}$$

Hence, the ex ante distribution over the states Ω is a convex combination of the posterior distributions conditional on the set of messages. Thus the ex ante distribution and the posterior distributions could have been derived from the same joint distribution. (5.1) also makes clear that the weights $\lambda(s)$ are not only non-negative, but that they add up to 1.

In understanding where the weights $\mu(\cdot,\cdot)$ come from, let us go back to lemma 3.2. Now, γ *fails* to be valuable given π if, there is a decision problem (A,u), decision rule f and action a^* such that (3.5) and (3.6) hold. It is instructive to restate these conditions, except that we split (3.5) into two separate conditions, as follows.

$$\sum_\omega \gamma(\omega|s)u(f(s),\omega) \ge \sum_\omega \gamma(\omega|s)u(a,\omega), \quad \forall s \in S, \forall a \ne a^* \tag{3.5(i)}$$

$$\sum_\omega \gamma(\omega|s)u(f(s),\omega) \ge \sum_\omega \gamma(\omega|s)u(a^*,\omega), \quad \forall s \in S \tag{3.5(ii)}$$

$$\sum_\omega \pi(\omega)\sum_s \beta^*(s|\omega)u(f(s),\omega) < \sum_\omega \pi(\omega)u(a^*,\omega) \tag{3.6}$$

If we were now to drop (3.5)(i) altogether, we would lose the weights $\mu(\cdot,\cdot)$ from the analysis. This can be seen from the proof of lemma 3.2. Equation (3.8) has its counterpart in (3.5)(ii), while (3.9) has its counterpart in (3.6). In effect, by dropping (3.5)(i), we lose (3.7).

Let us suppose that we were to define a new relation between γ and π called valuable* as follows. Say that γ is valuable* given π if, for any decision problem (A,u), decision rule f and action $a^* \in A$, (3.5)(ii) holds and (3.6) fails. Then, by selectively pruning the proof of lemma 3.2, we have the proposition that γ is valuable* given π if and only if there is no function $x: \Omega \times S \to \mathfrak{R}$ such that:

$$\sum_\omega \gamma(\omega|s)x(\omega,s) \ge 0, \quad \forall s \in S \tag{3.8}$$

$$\sum_\omega \pi(\omega)\sum_s \beta^*(s|\omega)x(\omega,s) < 0. \tag{3.9}$$

In terms of the matrices defined in section 3, this implies, by Farkas' lemma that the vector (3.21) is a non-negative combination of the rows of the matrix G. We have lost all the matrices $\{B_1, B_2, \ldots, B_m\}$ from the analysis, since they originate from the deleted condition (3.5)(i). Thus, γ is valuable* given π if and only if, there exists $\lambda: S \to \mathfrak{R}_+$ such that for all ω and s,

$$\pi(\omega)\beta^*(s|\omega) = \lambda(s)\gamma(\omega|s). \tag{5.2}$$

This is a modified version of the valuable information condition of theorem 3.1, in which all the expressions involving μ have disappeared.

It should be evident from (5.2) that the notion of being valuable* is a very stringent one. The left hand side of (5.2) is the (true) joint distribution over $\Omega \times S$, while the right hand side is the imputed subjective joint distribution over $\Omega \times S$. In effect, the notion of being valuable* reduces to the trivial requirement that the decision maker knows the true joint distribution. This is unsurprising, since by removing (3.5)(i), we are requiring the beliefs γ to yield an expected utility no lower than that given by the prior π for *any* decision rule, not just the optimal ones.

When stated in these terms, the role of the weights μ become clearer. These weights relax the requirement that the decision maker knows the joint distribution by allowing some room for manoeuvre in the beliefs of the decision maker. This room for manoeuvre comes from the fact that we restrict the decision rule f to be *optimal* given beliefs γ. Thus, the requirement which the beliefs γ must satisfy is less stringent. The relevant comparison applies only to the set of *rational* decision rules, rather than the set of *all* decision rules, rational and irrational.

The weights μ give the decision maker some freedom of movement in holding beliefs which diverge from the true distribution. However, theorem 3.1 also sets limits on how far these beliefs may diverge from the truth. This line of reasoning gives us another way to interpret the valuable information condition. If we regard the beliefs γ as having been perturbed from the true distribution by the addition of some 'noise' in the manner of Blackwell [1], we can set limits on how large, and of what sort, this noise can be. To pursue this line of reasoning, it is illuminating to recast the valuable information condition using matrix notation.

Thus, let us denote by Π the $n \times m$ matrix, whose (i, j)th entry is given by

$$\pi(\omega_i)\beta^*(s_j|\omega_i). \tag{5.3}$$

Π is the matrix of the true probabilities over the set $\Omega \times S$. In a similar way, let us denote by Γ then $n \times m$ matrix whose (i, j)th entry is:

$$\gamma(\omega_i|s_j). \tag{5.4}$$

We define three more matrices, Λ, Δ, and Θ. Λ is an $m \times m$ diagonal matrix, where

$$\Lambda = \begin{pmatrix} \lambda(s_1) & & & \\ & \lambda(s_2) & & \\ & & \ddots & \\ & & & \lambda(s_m) \end{pmatrix}. \tag{5.5}$$

Δ is also an $m \times m$ diagonal matrix. The ith entry along the diagonal is the number:

$$\sum_{k \neq i} \mu(s_i, s_k). \tag{5.6}$$

Finally, the matrix Θ is an $m \times m$ stochastic matrix, whose (i, j)th entry is given by:

$$\mu(s_i, s_j) \Big/ \sum_{k \neq i} \mu(s_i, s_k), \tag{5.7}$$

when $i \neq j$, and is zero along the diagonal. Then, the condition identified in theorem 3.1 in terms of (3.4) can be expressed as

$$\Pi = \Gamma(\Lambda + \Delta(I - \Theta)). \tag{5.8}$$

There are two key features of the matrix $\Lambda + \Delta(I - \Theta)$. The first is that all of its off-diagonal entries are zero or negative. This follows from the fact that Θ is a stochastic matrix, and the fact that Λ, Δ and I are all diagonal matrices with non-negative entries. The second feature is that the sum of any row of $\Lambda + \Delta(I - \Theta)$ is non-negative. This follows from the fact that the sum of any row of $I - \Theta$ is zero, so that the sum of any particular row of $\Lambda + \Delta(I - \Theta)$ in the entry $\lambda(s)$ in the matrix Λ.

Matrices which have the above pair of features are called *dominant diagonal* matrices, and are familiar to economists from the analysis of stability of linear system. (See, for example, McKenzie [3] or Nikaoido [5, chapter 2]). One of the features of such matrices is that its inverse is non-negative. From (5.8), then, we have:

$$\Gamma = \Pi M, \tag{5.9}$$

where M is a non-negative matrix. In other words, the beliefs are such that they are obtained from the true distribution by means of adding noise, in the shape of M. Thus, it is a necessary condition of the valuable information condition that the beliefs are obtained from the true distribution by adding noise. However, it is not sufficient. The reason is that the noise M must be the inverse of a dominant diagonal matrix. Not all non-negative matrices will have a dominant diagonal inverse.

In comparing Π and Γ, we have glossed over the fact that Π is a matrix representing the joint distribution over $\Omega \times S$, while Γ is a matrix representing the conditional beliefs. Thus, to be accurate in our talk of adding noise, we ought to normalize one of Π and Γ so that they are comparable objects. One way of doing this would be to post-multiply (5.9)) by a diagonal matrix whose jth entry is the marginal probability of the message s_j. If we denote this matrix by N, then (5.9) gives

$$\Gamma N = \Pi M N, \tag{5.10}$$

so that ΓN is now a matrix representing a joint distribution over $\Omega \times S$, so that it is comparable with Π. Then, the 'noise' is given by the matrix MN.

It is tantalizing that the apparently 'correct' sort of noise is that linked to dominant diagonal matrices. It is tantalizing for us, since we have not been able to figure out whether this is merely a superficial coincidence, or whether there is something deeper going on in our analysis which ties it in with the analysis of stability in linear systems. In an earlier version of this paper, we had used the dominant diagonal condition to give a characterization of the valuable information condition in terms of the long-run distribution associated with an absorbing Markov chain. We hope that future research will uncover further insights.

References

1. Blackwell, D.: The comparison of experiments. Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability, pp. 93–102. University of California Press 1951
2. Gale, D.: The theory of linear economic models. New York: McGraw-Hill 1960
3. McKenzie, L.: Matrices with dominant diagonals. In Mathematical Methods in the Social Sciences, Proceedings of the First Stanford Symposium. Arrow, K., Karlin, S., Suppes, P. (eds.) Standard: Stanford University Press 1960
4. Lehrer, E.: Comparison of experts. Unpublished working paper, Northwestern University (1994)
5. Nikaido, H.: Convex structures and economic theory. New York: Academic Press 1968

Economic Theory 9, 325–339 (1997)

Economic Theory

Reconstructing dynamics from intertemporal economic data

Venkatesh Bala★

Department of Economics, McGill University, 855 Sherbrooke Street W., Montreal, Quebec H3A 2T7, CANADA

Received: April 25, 1994; revised version December 29, 1995

Summary. This paper is concerned with the relationship between a continuous dynamical system and the trajectory generated by such a system. The main result provides necessary and sufficient conditions for an infinite data stream to be rationalized as the output of a continuous law of motion. The paper develops concepts of informativeness of a given set of intertemporal data and shows that informativeness is maximal when the data is chaotic. It also demonstrates that with probability one the sample paths from a non-trivial independent and identically distributed stochastic process cannot be rationalized as the output of a continuous deterministic system. Two impossibility results are discussed which show that even with an infinite amount of data the hypothesis that the data has been generated by a non-monotonic function cannot be ruled out. An application concerning the recovery of the excess demand function from a sequence of price observations from the tatonnement process is also given.

JEL Classification Number: D83.

1 Introduction

In many dynamic economic models, the evolution of the system can be succinctly characterized by a transition rule of the form $x' = f(x)$, where x represents a state variable such as prices, quantities or capital stocks and f is a mapping on the state space. As is well-known, research over the last fifteen years has demonstrated that in many simple economic models the evolution of the system is chaotic, i.e. resembles the output of a stochastic process (see, e.g. Benhabib and Day (1982), Day (1982), Grandmont (1985), Boldrin and Montrucchio (1986), Deneckere and Pelikan (1986), Day and Shafer (1987),

★ I would like to thank Professors Bent Christensen and Mukul Majumdar for valuable discussions and an anonymous referee for helpful comments. Financial support from SSHRC and Quebec's Fonds FCAR is gratefully acknowledged.