The Annals of Mathematical Statistics
1969, Vol. 40, No. 4, 1177–1186

BAYES' METHOD FOR BOOKIES

By David A. Freedman[1] and Roger A. Purves[2]

University of California, Berkeley

1. Introduction. Let Θ be a finite set, and for each $\theta \varepsilon \Theta$, let p_θ be a probability distribution on a finite set X. Consider three players: a master of ceremonies, a bookie, and a bettor. The master of ceremonies selects, at his pleasure, a θ belonging to Θ, and then an observation $x \varepsilon X$ at random, according to p_θ. He announces x to the bookie and the bettor. The bookie then posts odds on subsets of Θ, with the understanding that he must accept any combination of stakes the bettor might care to make. The bettor places his stakes. Finally, θ is revealed by the master of ceremonies and bookie and bettor settle up.

Before the game begins, how should the bookie plan to set the odds? One possibility is to choose a distribution on Θ, and when x is revealed, to calculate posterior odds by Bayes' rule. There is good reason for adopting this method. For any other procedure, there exists a system of bets with the following property: a bettor who places his stakes according to the system can expect to win money from the bookie, regardless of the θ chosen by the master of ceremonies. On the other hand, if the odds are calculated by Bayes' method, no such system exists. This is part of the content of Theorems 1 and 2 below. The two theorems extend a result of Bruno de Finetti (de Finetti, 1937, especially pages 6–8) which says (roughly) that someone who posts odds must do so on the basis of a finitely additive probability or else be certain to lose money to a clever bettor.

Section 2 of the paper treats the easier case where the odds are all finite and positive. The general case is developed in Section 3. Section 4, the final section, contains a theorem similar to the theorem of Section 2, but appropriate to situations involving prediction.

2. Finite, positive odds. Throughout this and the next section, the following assumptions will be in force. The sets Θ, X are finite and not empty. For each $\theta \varepsilon \Theta$, p_θ is a probability distribution on X. For each $x \varepsilon X$, the function $p.(x)$ is defined by the rule

$$p.(x): \theta \to p_\theta(x), \qquad \theta \varepsilon \Theta.$$

The complement of a subset A of Θ is written A^c, and A is said to be *proper* if neither A, A^c is empty. Everywhere A is a subset of Θ and x is a member of X.

Received 8 August 1967.

[1] Research sponsored by the Air Force Office of Scientific Research, Office of Aerospace Research, United States Air Force, under Grant AF-AFOSR-1312-67. This manuscript is submitted for publication with the understanding that the United States Government is authorized to reproduce and distribute reprints for governmental purposes.

[2] This investigation was prepared with the partial support of the National Science Foundation Grant GP-5059.

For each x and A, the bookie must be prepared to post odds $\lambda(x, A)$, where $\lambda(x, A)$ is either a nonnegative real number, or $+\infty$. The statement "$\lambda(x, A) = 7$" means that when x occurs the bookie will post odds of seven to one (7:1) against the occurrence of A. If the bookie happened to be setting the odds on the basis of a probability, the corresponding probability of A would be one-eighth. Odds determine the exchange of payments between the bettor and the bookmaker. To describe certain conventions to be adopted here, consider an event A with odds 7:1 posted. The bookie is assumed to be willing to permit the bettor to stake both on and against A. (Ordinary bookmakers are not so permissive.) When the bookie "accepts" stakes $s \geqq 0$ on A and $t \geqq 0$ against A, no money changes hands. Rather, bookie and bettor enter an agreement that the bookie will pay the bettor $7s - t$ if A occurs, and $(t/7) - s$ if A does not occur. Corresponding payments take place for any other positive odds. The bettor may, if he wishes, simultaneously place stakes on and against a number of events. In that case, the payment from the bookie to the bettor is simply the sum of the payments for the individual events.

For simplicity of exposition, it is assumed in this section that the bookie only accepts combinations of stakes on or against proper A, and Θ contains at least two members. Also, the odds $\lambda(x, A)$ are assumed to be positive and finite for all x and all proper A. In this situation, a *strategy* σ (for a bettor) is a function which assigns to each (x, A), with A proper, a pair of nonnegative real numbers,

$$s(x, A), \qquad t(x, A).$$

These are to be interpreted as stakes to be placed on and against A respectively, in the event that x occurs.

To express the amount received by a bettor using a particular strategy when (θ, x) occurs, let $w_A(\theta, x)$ be the amount he receives as a result of stakes placed on and against A only. Then

$$(1) \quad w_A(\theta, x) = [\lambda(x, A)s(x, A) - t(x, A)]1_A(\theta) + [t(x, A)/\lambda(x, A) - s(x, A)](1 - 1_A(\theta))$$

where $1_A(\theta) = 1$ for $\theta \varepsilon A$, $1_A(\theta) = 0$ for $\theta \not\varepsilon A$. The total amount received by the bettor when (θ, x) occurs is the sum of the terms (1) taken over all proper A. Given the occurrence of θ, the expectation of the total is

$$(2) \qquad \sum_x \sum_{A \text{ proper}} w_A(\theta, x)p_\theta(x).$$

Considered as a function of θ, this quantity will be referred to as a *payoff function*. The preceding calculation shows how each strategy σ leads to a real-valued function on Θ, the payoff function associated with σ. The set of all such functions, obtained by letting σ range over all strategies against a fixed λ, will be denoted $\mathcal{A}_\lambda$. Since Θ contains at least two points, $\mathcal{A}_\lambda$ contains a function which is not identically zero.

THEOREM 1. *Suppose Θ contains at least two points and $\lambda(x, A)$ is finite and positive for all proper A.*

If $\mathcal{A}_\lambda$ is not the set of all real-valued functions defined on Θ, there is a unique probability π on Θ such that

$$(1 + \lambda(x, A))^{-1} \textstyle\sum_\theta p_\theta(x)\pi(\theta) = \sum_{\theta \varepsilon A} p_\theta(x)\pi(\theta) \tag{3}$$

for all (x, A) with A proper. The probability π assigns positive mass to each member of Θ. Further, $\mathcal{A}_\lambda$ is the set of all functions f satisfying

$$\textstyle\sum_\theta f(\theta)\pi(\theta) = 0.$$

In other words, if the odds λ are not posterior odds (in the sense that no probability π on Θ satisfies (3) for all (x, A) with A proper), any real-valued function on Θ can be obtained as a payoff function by a suitable combination of stakes, including one that is positive for every θ. But if λ is the posterior odds determined by the prior π, every payoff function has π-expectation zero, and so at least one nonpositive value.

PROOF. The first step consists of establishing a more convenient expression for $w_A(\cdot, x)$. Omitting for the moment the dependence of λ, s, t on (x, A),

$$\begin{aligned} w_A(\cdot, x) &= (\lambda s - t)1_A + ((t/\lambda) - s)(1 - 1_A(\theta)) \\ &= (\lambda s - t)1_A + \lambda^{-1}(\lambda s - t)(1 - 1_A(\theta)) \\ &= (1 + \lambda^{-1})(\lambda s - t)[1_A - (1 + \lambda)^{-1}]. \end{aligned}$$

The final expression is of the form $\alpha(x, A)f_{x,A}$, where

$$f_{x,A} = 1_A - [1 + \lambda(x, A)]^{-1}$$

is a function of θ which does not depend on the strategy σ adopted by a bettor, and $\alpha(x, A)$ is a scalar which ranges over the real line as σ ranges over the set of all strategies. Incidentally, when $f_{x,A}$ is multiplied by $p.(x)$, the result is the payoff function associated with the strategy using only one positive stake, $[1 + \lambda(x, A)]^{-1}$, to be placed on A when x occurs. These simple payoff functions occur often in what follows.

The preceding calculation, in conjunction with (2), implies that $\mathcal{A}_\lambda$ is the linear subspace (of the space V of all real-valued functions on Θ) spanned by the set of functions

$$[1_A - (1 + \lambda(x, A))^{-1}]p.(x) \tag{4}$$

where $x \varepsilon X$ and A is proper. If $\mathcal{A}_\lambda \neq V$, there is a $\pi \neq 0$ in V such that π is orthogonal to every member of $\mathcal{A}_\lambda$. In particular, π, and any scalar multiple of π, satisfy the equations

$$\textstyle\sum_\theta \pi(\theta)[1_A(\theta) - (1 + \lambda(x, A))^{-1}]p_\theta(x) = 0 \tag{5}$$

for all (x, A) with A proper. The remainder of the proof consists in showing that, up to a scalar multiple, π is a probability distribution. Once this has been established, the equation (5), which is equivalent to (3), shows λ to be the posterior odds for the prior obtained by multiplying π by a suitable scalar.

There is a $\tau \varepsilon \Theta$ such that $\pi(\tau) \neq 0$. Setting $A = \{\tau\}$ in (5) and rearranging slightly,

$$(6) \qquad p_\tau(x)\pi(\tau) = (1 + \lambda(x, \tau))^{-1} \sum_\theta p_\theta(x)\pi(\theta)$$

for all x. If an x is chosen for which $p_\tau(x)$ is positive, the right factor on the right side of (6) is nonzero. For any other τ', the equation (6), with τ' replacing τ, remains true, and since $\lambda(x, \tau')$ is positive and finite, both $\pi(\tau')$ and $p_{\tau'}(x)$ are nonzero. Dividing one equation by another shows $\pi(\tau')/\pi(\tau)$ to be positive. It follows that by scalar multiplication, π can be taken to be a probability distribution on Θ with $\pi(\theta)$ positive for all $\theta \varepsilon \Theta$. A similar argument shows that probabilities π_1, π_2 satisfying (5) are proportional to one another and hence equal.

Since every function (4) satisfies

$$\sum_\theta f(\theta)\pi(\theta) = 0,$$

$\mathfrak{A}_\lambda$ is a subspace of the orthogonal complement of π. But the uniqueness of π implies that the orthogonal complement of $\mathfrak{A}_\lambda$ has dimension 1. This completes the proof.

3. Zero and infinite odds. To establish an analogue of Theorem 1 when zero and infinite odds are posted would seem to require certain conventions. Rather than work with a payoff of $+\infty$, we have found it convenient to interpret infinite odds on an event A in this way: if a bettor stakes a positive amount s on A, and A occurs, he may ask the bookie for any nonnegative real amount u. If A does not occur, the bettor loses s. For a stake $t \geqq 0$ against A, the bettor receives zero if A does not occur. In total, the bettor receives $u - t$ when A occurs, and $-s$ otherwise, with the proviso that $u = 0$ when $s = 0$. In case that the odds are 0:1 against A, analogous conventions will be observed.

With these stipulations, the strategy of a bettor includes not only the choice of stakes for events with finite odds, but for each (x, A) for which $\lambda(x, A) = 0$ or ∞, the choice of a triple of nonnegative numbers

$$s(x, A), \qquad t(x, A), \qquad u(x, A)$$

satisfying the following conditions: if $\lambda(x, A) = \infty$, $s(x, A) = 0$ implies $u(x, A) = 0$; if $\lambda(x, A) = 0$, $t(x, A) = 0$ implies $u(x, A) = 0$.

To calculate the amount received by a bettor when (θ, x) occurs, recall that $w_A(\theta, x)$ is given by (1) if $\lambda(x, A)$ is positive and finite. If $\lambda(x, A) = +\infty$,

$$(7) \qquad w_A(\theta, x) = [u(x, A) - t(x, A)]1_A(\theta) - s(x, A)(1 - 1_A(\theta)).$$

If $\lambda(x, A) = 0$,

$$(8) \qquad w_A(\theta, x) = -t(x, A)1_A(\theta) + [u(x, A) - s(x, A)](1 - 1_A(\theta)).$$

The *payoff function* arising from a strategy is the function which assigns to each $\theta \varepsilon \Theta$ the number

$$(9) \qquad \sum_x \sum_A w_A(\theta, x)p_\theta(x)$$

where the w_A's are determined by the strategy through equations (1), (7), (8) and the inner sum is taken over all subsets, proper or not, of Θ. The set of all such payoff functions will be denoted $\mathcal{B}_\lambda$. The following theorem extends only the first part of Theorem 1.

THEOREM 2. *If $\mathcal{B}_\lambda$ is a proper subset of the set of all real-valued functions defined on Θ, there is a probability π on Θ such that* (3) *holds for every* (x, A).

In (3) (and elsewhere) the quantity $(1 + \lambda(x, A))^{-1}$ is taken to be zero when $\lambda(x, A)$ is infinite, and conversely. Unlike Theorem 1, the π of Theorem 2 need not be unique. An example will be given in the final remark following the proof of Theorem 3.

LEMMA 1. *$\mathcal{B}_\lambda$ is closed under pointwise addition and multiplication by non-negative scalars.*

In the next lemma, a linear subspace L of functions is introduced. In the proof of Theorem 2, L will play the role played by $\mathcal{A}_\lambda$ in the proof of Theorem 1.

LEMMA 2. *The closure of $\mathcal{B}_\lambda$ includes the linear space L spanned by the set of functions*

$$[1_A - (1 + \lambda(x, A)^{-1}]p.(x) \tag{10}$$

where $x \varepsilon X$, A is a subset of Θ.

PROOF. It is sufficient to show, that for each (x, A), the result $r_{x,A}$ of multiplying (10) by a real scalar c, can be realized as a limit of functions in $\mathcal{B}_\lambda$. If $0 < \lambda(x, A) < \infty$, this is obvious, because $r_{x,A}$ is a member of $\mathcal{B}_\lambda$. If $\lambda(x, A) = +\infty$, the function

$$(u(x, A) - t(x, A))1_A - s(x, A)(1-1_A)$$

when multiplied by $p.(x)$, becomes a payoff function. If c is not positive, set $s(x, A) = 0$ and $t(x, A) = -c$ to obtain $r_{x,A}$ as a member of $\mathcal{B}_\lambda$. If c is positive, set $t(x, A) = 0$, $u(x, A) = c$. Then as $s(x, A)$ approaches zero, the payoff functions so obtained approach $r_{x,A}$. The case $\lambda(x, A) = 0$ can be handled in a similar way.

PROOF OF THEOREM 2. If $\mathcal{B}_\lambda$ is a proper subset of V, so is L. This follows from the two lemmas and the fact that a convex dense subset of V coincides with V. Following the proof of Theorem 1, there must be a nontrivial π satisfying (5) for all (x, A), $A = \phi$ or Θ included. The main obstacle to continuing with the proof of Theorem 1 is that $p_\theta(x)$ may be zero for some x and θ. To overcome this, certain equivalence relations will be defined and the sign of π changed on appropriate equivalence classes.

If Γ is a nonempty subset of Θ, let $\tilde{\Gamma}$ be the set of pairs (θ, τ) satisfying: for some $n \geqq 0$, there are $\theta_0, \theta_1, \cdots, \theta_{n+1}$, in Γ such that $\theta_0 = \theta$, $\theta_{n+1} = \tau$, and the supports of p_{θ_i}, $p_{\theta_{i+1}}$ are not disjoint for $i = 0, 1, \cdots, n$. The relation $\tilde{\Gamma}$ is an equivalence relation on Γ. The argument given in the proof of Theorem 1 to show the $\pi(\theta)$'s are of the same sign is easily adapted to establish the following fact. If Γ_1 is an equivalence class of $\tilde{\Gamma}$ and $\pi(\theta) \neq 0$ for all $\theta \varepsilon \Gamma_1$, the numbers $\pi(\theta)$ have the same sign for all θ in Γ_1.

Let ∂ be an equivalence class of $\tilde{\Theta}$ and Δ be the set of $\theta \varepsilon \partial$ such that $\pi(\theta) \neq 0$. Assuming that Δ is nonempty, let C be an equivalence class of $\tilde{\Delta}$. Then $\pi(\theta)$ has the same sign for all $\theta \varepsilon C$. If this were negative, $\hat{\pi}$, defined by

$$\begin{aligned} \hat{\pi}(\theta) &= -\pi(\theta), \qquad \theta \varepsilon C, \\ &= \pi(\theta), \qquad \theta \varepsilon \Theta - C, \end{aligned}$$

also satisfies (5) for all (x, A). There are two cases to check.

Fix an x in X and let us suppose that for some $\tau \varepsilon C$, $p_\tau(x)$ is positive. Then

$$p_\theta(x)\pi(\theta) = p_\theta(x)\hat{\pi}(\theta) = 0$$

for all $\theta \not\varepsilon C$. Indeed, if $\theta \varepsilon (\partial - \Delta)$, $\pi(\theta) = \hat{\pi}(\theta) = 0$; and if $\theta \not\varepsilon C \cup (\partial - \Delta)$, $p_\theta(x) = 0$ by the definition of C and ∂. Thus, for any A,

$$\begin{aligned} (11) \quad \textstyle\sum_A p_\theta(x)\hat{\pi}(\theta) &= \textstyle\sum_{A \cap C} p_\theta(x)\hat{\pi}(\theta) = -\textstyle\sum_{A \cap C} p_\theta(x)\pi(\theta) \\ &= -\textstyle\sum_A p_\theta(x)\pi(\theta). \end{aligned}$$

Using (5), this equality may be continued.

$$\begin{aligned} &= (1 + \lambda(x, A))^{-1}(-\textstyle\sum_\theta p_\theta(x)\pi(\theta)) \\ &= (1 + \lambda(x, A))^{-1}(\textstyle\sum_\theta p_\theta(x)\hat{\pi}(\theta)). \end{aligned}$$

The last step is (11) with $A = \Theta$.

The other case, when $p_\tau(x)$ is zero for all $\tau \varepsilon C$, is clear since $\hat{\pi} = \pi$ off C.

It follows immediately that π may be modified to give a solution of (5) which is also a probability distribution.

The next theorem, which corresponds to the second part of Theorem 1, characterizes the set of payoff functions available to the bettor when the odds offered to him are posterior odds.

THEOREM 3. *Suppose that $p_\theta(x)$ is positive for all θ, x. If π is a probability distribution on Θ which is not a point mass, and λ satisfies* (3) *for all* (x, A), *then $\mathcal{B}_\lambda$ is the half-space of functions f satisfying*

$$\textstyle\sum_\theta f(\theta)\pi(\theta) \leqq 0.$$

PROOF. (i) *If $f \varepsilon \mathcal{B}_\lambda$, the inequality holds.* For

$$\textstyle\sum_\theta w_A(\theta, x)p_\theta(x)\pi(\theta)$$

is either

$$0, \qquad -s(x, A)(\textstyle\sum_\theta p_\theta(x)\pi(\theta)), \qquad -t(x, A)(\textstyle\sum_\theta(p_\theta(x)\pi(\theta)),$$

depending on whether $\lambda(x, A)$ is nonzero and finite, infinite, or zero. This follows upon replacing $w_A(\theta, x)$ by its definition and applying (3). In particular, in the last two cases, the calculation depends on two consequences of (3): if $\lambda(x, A)$ is infinite,

$$\textstyle\sum_{\theta \varepsilon A} p_\theta(x)\pi(\theta) = 0;$$

and if $\lambda(x, A)$ is zero,

$$\sum_{\theta \varepsilon A} p_\theta(x)\pi(\theta) = \sum_\theta p_\theta(x)\pi(\theta).$$

(ii) *The subspace L is a subset of $\mathcal{B}_\lambda$.* It is clear that the subspace K spanned by the set of functions

$$[1_A - (1 + \lambda(x, A))^{-1}]p.(x),$$

where (x, A) satisfies $0 < \lambda(x, A) < \infty$, is a subset of $\mathcal{B}_\lambda$. Thus, it suffices to show $K = L$. If $\lambda(x, B) = \infty$, it will be argued that the function $1_B p.(x)$ is in K; the rest of the proof of (ii) is then routine.

To see that $1_B p.(x) \varepsilon K$, it is enough to verify that $\delta_\tau p.(x) \varepsilon K$ for each $\tau \varepsilon B$, where $\delta_\tau = 1_{\{\tau\}}$. By assumption, there is an $\omega \varepsilon \Theta$ such that $0 < \pi(\omega) < 1$. Let $F = \{\omega, \tau\}$ and $E = \{\omega\}$.

The next step is to see that

$$(12) \qquad [1_F - (1 + \lambda(x, F))^{-1}]p.(x), \qquad [1_E - (1 + \lambda(x, E))^{-1}]p.(x)$$

are both in K. Indeed, on setting $A = E$ in (3), and using $p_\theta(x) > 0$ and $0 < \pi(\omega) < 1$, it follows that $0 < \lambda(x, E) < \infty$. But

$$(13) \qquad \lambda(x, F) = \lambda(x, E),$$

because $\lambda(x, B) = \infty$ means $\sum_{\theta \varepsilon B} p_\theta(x)\pi(\theta) = 0$ by (3), so $p_\tau(x)\pi(\tau) = 0$; using (3) again,

$$p_\tau(x)\pi(\tau) + p_\omega(x)\pi(\omega) = (1 + \lambda(x, F))^{-1} \sum_\theta p_\theta(x)\pi(\theta),$$

which proves (13). In particular, both functions in (12) are in K. In view of (13), their difference is $\delta_\tau p.(x)$, which is therefore in K. This completes the proof of (ii).

(iii) *If f satisfies $\sum_\theta f(\theta)\pi(\theta) = 0$, f is a member of L and conversely.* The converse is established by a calculation similar to the one described in (i). The assumption that $p_\theta(x)$ is everywhere positive implies that the system of equations (5) (where A is also permitted to be ϕ or Θ) has, up to a scalar multiple, at most one solution. Thus the orthogonal complement of L coincides with the set of scalar multiples of π.

(iv) *The set U of all functions f satisfying $\sum_\theta f(\theta)\pi(\theta) < 0$ is a subset of $\mathcal{B}_\lambda$.* If h is a fixed member of U it is easy to verify that the set of functions of the form $g + \alpha h$ where g ranges over L and α over positive scalars, coincides with U. To exploit this, let $h = -p.(x)$, where x is some member of X. To show h is in $\mathcal{B}_\lambda$, let $s(x, \phi) = u(x, \phi) = 1$ and all other stakes be zero. The payoff function resulting from this strategy is h. Since $\mathcal{B}_\lambda$ is a convex cone, and every $g \varepsilon L$ is in $\mathcal{B}_\lambda$ by (ii), U is a subset of $\mathcal{B}_\lambda$.

REMARKS. 1. If the assumption of Theorem 3 that π is not a point mass is replaced by the assumption that π is a point mass giving probability one to

$\omega \varepsilon \Theta$, say, $\mathcal{B}_\lambda$ becomes the set of functions f satisfying either

$$f(\omega) < 0, \qquad \text{or}$$
$$f(\omega) = 0 \qquad \text{and} \qquad f(\theta) \leqq 0 \qquad \text{for all} \quad \theta \neq \omega.$$

The proof is similar to the proof of Theorem 3. In this case, then, $\mathcal{B}_\lambda$ is not a half space.

2. If X has exactly one member, the bookie has to specify only one set of odds λ. Essentially, this is the problem of posting odds for an experiment about to be performed, no other information given. Bruno de Finetti has considered this situation and our three theorems are very close to his discussion of the theorem of total probabilities (given on pages 6–8 of de Finetti, 1937). For example, Theorem 1 says that if λ (positive and finite) does not satisfy

$$(1 + \lambda(A))^{-1} = \textstyle\sum_{\theta \varepsilon A} \pi(\theta),$$

for all proper A, where π is some probability on Θ, then every real-valued function is available as a payoff function (no expectations here) to a bettor.

3. Under the assumptions of Theorem 3, if a bettor decides to bet only when a particular x occurs, that is to adopt strategies with

$$s(x', A) = t(x', A) = 0$$

for all $x' \neq x$ and all A, the set of all payoff functions

$$\textstyle\sum_A w_A(\cdot, x)p.(x)$$

available to him remains the entire half-space of functions f with nonpositive π-expectation. This follows easily from a special case of Theorem 3, already noted in the preceding remark. Letting $\hat{X}$ be the set consisting of x alone, $\hat{p}_\theta(x) = 1$ for all θ, and $\hat{\pi}$ be the posterior distribution given x, Theorem 3 says that the set $\mathcal{B}_\lambda$ of all functions of the form

$$\textstyle\sum_A w_A(\theta, x)$$

is the half-space of functions with nonpositive $\hat{\pi}$-mean. Now if f has nonpositive π-expectation, the quotient $f/p.(x)$ has nonpositive $\hat{\pi}$-expectation and so $f/p.(x)$ is in $\mathcal{B}_{\hat{\lambda}}$. Since members of $\mathcal{B}_{\hat{\lambda}}$, when multiplied by $p.(x)$, become payoff functions available to the restricted bettor, f is available to the restricted bettor.

4. If the assumption that $p_\theta(x)$ is everywhere positive is dropped, every member of $\mathcal{B}_\lambda$ will still have nonpositive, π-expectation, but the converse may not hold. For example, suppose $\Theta = X = \{1, 2\}$; $p_1(1) = p_2(2) = 1$, $p_1(2) = p_2(1) = 0$; and $\pi(1) = \pi(2) = \frac{1}{2}$. Then the set $\mathcal{B}_\lambda$ is the set of functions on Θ which are everywhere nonpositive (the lower left quadrant of the plane). In general, $\mathcal{B}_\lambda$ did not seem to us to admit of a simple description.

5. In the example described in the preceding remark, where λ is the posterior odds for the uniform prior, any probability on Θ will be a solution of (3). Thus the probability of Theorem 2 need not be unique. When $p_\theta(x)$ is positive for all

θ, x, however, it is unique, as already mentioned in part (iii) of the proof of Theorem 3.

4. Prediction. Let Θ, X, Y be finite, nonempty sets and for each $\theta \varepsilon \Theta$, p_θ a probability distribution on $X \times Y$. In this situation, the master of ceremonies selects a θ belonging to Θ, and then, according to the distribution p_θ, a point of $X \times Y$. He reveals the first coordinate of this point to the bookie and the bettor. The bookie proceeds to post odds on events depending only on the second coordinate (the future) and the bettor places bets on those events. The game ends with the master of ceremonies revealing the second coordinate of the point and bookie and bettor settling up.

The notation of Section 3 can be carried over to this section, with the exception that A here will be a subset of Y rather than Θ. The amount received by a bettor as the result of stakes placed on and against A only, if θ, x, y occurs, is no longer a function of θ. If $\lambda(x, A)$ is positive and finite, this quantity is given by

$$w_A(x, y) = (\lambda(x, A)s(x, A) - t(x, A))1_A(y) \\ + (t(x, A)/\lambda(x, A) - s(x, A))(1 - 1_A(y)).$$

If $\lambda(x, A)$ is either infinite or zero, the substitution of 'y' for 'θ' in the right hand side of (7) and (8) will give the correct expression for $w_A(x, y)$.

The payoff function arising from a strategy is now defined by

$$\sum_{x \varepsilon X} \sum_{y \varepsilon Y} \sum_A w_A(x, y)p.(x, y)$$

where A ranges over all subsets of Y. Letting $\mathfrak{B}_\lambda$ be the set of all such payoff functions, we have

THEOREM 4. *Either $\mathfrak{B}_\lambda$ contains a function which is positive for every $\theta \varepsilon \Theta$, or there is a probability distribution π on Θ for which*

$$(14) \quad (1 + \lambda(x, A))^{-1}\left(\sum_y \sum_\theta p_\theta(x, y)\pi(\theta)\right) = \sum_{y \varepsilon A} \sum_\theta p_\theta(x, y)\pi(\theta)$$

for every $x \varepsilon X$ and A a subset of Y.

SKETCH OF PROOF. It can be verified that $\mathfrak{B}_\lambda$ is a convex cone of real-valued functions defined on Θ. If this cone contains no everywhere positive function, it is disjoint from the convex set of all such functions. By a well-known separation theorem for convex sets (Dunford, (1958), page 412) there is a nontrivial linear functional Π and a real number b such that if f is everywhere positive, $b \leqq \Pi(f)$; and if $f \varepsilon \mathfrak{B}_\lambda$, $\Pi(f) \leqq b$. It is easily argued that b is zero and there is a probability distribution π for which

$$\Pi(f) = \sum_\theta f(\theta)\pi(\theta)$$

for all real-valued f on Θ.

By choosing certain f in $\mathfrak{B}_\lambda$ and examining the inequality $\Pi(f) \leqq 0$, it becomes clear that λ is the odds computed on the basis of π, that is (14) holds. For example, let $\lambda(x, A)$ be positive and finite and consider the strategy with

$$s(x, A) = 1, \qquad t(x, A) = 0.$$

and

$$s(x', A') = t(x', A') = 0, \qquad \text{for all} \quad (x', A') \neq (x, A).$$

For this strategy, the payoff function is given by

$$f = \sum_y [\lambda(x, A)1_A(y) - (1 - 1_A(y))]p.(x, y).$$

The inequality $\Pi(f) \leqq 0$ becomes

$$\sum_\theta \sum_y [\lambda(x, A)1_A(y) - (1 - 1_A(y))]p_\theta(x, y)\pi(\theta) \leqq 0.$$

This implies the left side of (14) is no smaller than the right side. By considering a stake against A, the inequality can also be established in the other direction.

If $\lambda(x, A)$ is infinite, taking $s(x, A)$ to be one and every other stake zero, a similar argument leads to an inequality which implies the right hand side of (14) also is zero. This completes the sketch of the proof.

REMARKS. 1. Theorem 4 does not say that there is a unique π. This need not be the case. For example, consider the situation where X has one member, Y has two members and Θ has many members. Then a prior distribution has only to satisfy five linear equations.

2. To return to the setting of Section 3, if $Y = \Theta$ and

$$\begin{aligned} p_\theta(x, y) &= p_\theta(x) \qquad \text{for} \qquad y = \theta \\ &= 0 \qquad\qquad \text{for} \qquad y \neq \theta, \end{aligned}$$

Theorem 4 gives a weaker result than Theorem 2.

REFERENCES

DE FINETTI, BRUNO (1937). La Prevision: ses lois logigues, ses sources subjectives. *Annales de l'Institut Henri Poincaré* 7 1–68.

DUNFORD, NELSON and SCHWARZ, JACOB T. (1958). *Linear Operators, Part I: General Theory*. Interscience, New York.

KEMENY, JOHN G. (1955). Fair bets and inductive probabilities. *J. Symbolic Logic* **20** 263–273.

LEHMAN, R. SHERMAN, (1955). On confirmation and rational betting. *J. Symbolic Logic* **20** 251–262.

SHIMONY, ABNER (1955). Coherence and the axioms of confirmation. *J. Symbolic Logic* **20** 1–28.