

Multi-Bayesian Statistical Decision Theory

By S. WEERAHANDI and J. V. ZIDEK

University of Sri Jayawardenapura

University of British Columbia

SUMMARY

A solution is proposed to the problem of finding a joint decision procedure for a group of n persons. It is any Pareto optimal solution which maximizes the generalized Nash product over the set of jointly achievable utility n -vectors. This result was originally proposed in the theory of bargaining but is readily adapted to the statistical context. The individuals involved need not have identical utility functions or identical prior (posterior) distributions. The solution may be a non-randomized rule but is randomized when the individual opinions or preferences are sufficiently diverse. Applications to hypothesis testing and estimation are included.

Keywords: BARGAINING; BAYES INFERENCE; GROUP DECISION ANALYSIS; HYPOTHESIS TESTING; LINEAR OPINION POOL; NASH PRODUCT

1. INTRODUCTION

THIS paper is concerned with the situation, described by Savage (1954), which confronts n Bayesians who are required to choose a mutually acceptable solution to a statistical decision problem. In Section 2, we survey some of the methods now available to these Bayesians. In particular, we include a description of very relevant results from the normative theory of n -person bargaining games which do not seem to have been considered previously in this context.

All of the methodology which has been developed in response to the fundamental issue raised here requires of these Bayesians that they be able and willing to state their subjective probability distribution on the appropriate event-set. One approach pursues a solution through *aggregation*; these n assessments are combined in some way into one which may then be used in a conventional uni-Bayesian analysis.

Difficulties will arise if, however, the decision-makers have widely divergent preferences, i.e. utility functions; if the use of threats is disallowed, it seems likely that no amount of discussion among these decision-makers will produce a unanimous decision. Thus a solution through *compromise* is called for.

In the next section various approaches involving “aggregation” and “compromise” are presented. Emphasis will be placed on the latter since the former has already been extensively treated in the statistical literature. In particular our presentation will focus on the Nash solution (see Nash, 1950) to the n -person bargaining game. It is the oldest such solution since it is embodied in Zeuthin’s original work (Zeuthin, 1930) and is easily the most celebrated. To reveal some of its shortcomings, several competitors will be briefly discussed.

In Section 4 we present two statistical examples in which the Nash solution is computed. The first is a bivariate normal estimation problem and the second, a hypothesis testing problem.

Nash’s theory, like its competitors, is normative, not descriptive. It proceeds logically from certain very weak assumptions to obtain a surprisingly strong conclusion. This conclusion ought to be acceptable to the decision-makers, if they agree that the assumptions are reasonable. This does not mean that all groups of decision-makers will find these assumptions reasonable and we are not saying that all such groups ought to find them reasonable. However, judging from the considerable attention the Nash theory has received, it would seem that they must generally be regarded as credible.

The n -Bayesian decision theory which derives from bargaining theory has two important features which usually are not mentioned in presentations of the corresponding uni-Bayesian theory. Firstly, account is taken of each Bayesian’s present (worth) utility. No decision-maker

would find acceptable a joint decision which on expectation would reduce his worth, and the theory provides the decision-maker with a veto against such decisions. Thus the group may well fail to agree on a decision. Secondly, randomized decision rules need to be reintroduced, as Savage (1954) points out. Of course, even in the uni-Bayesian theory such rules are present and some of them may well be optimal. However, they may be and are ignored for familiar reasons. In the n -Bayesian case they cannot be ignored; in many cases the solution is a randomized rule.

It should be noted that in assessing the worth or utility of a potential decision a decision-maker need not be entirely selfish. He may even be altruistic and assign relatively large utilities to decisions which would call on him to sacrifice some of his own physical well-being to the benefit of other individuals or society as a whole. Harsanyi (1977, Chapter 4) discusses this issue at some length.

2. MULTIPERSON DECISION PROBLEMS

There is a considerable literature dealing with group decision analysis and some notable contributions have not yet been published. Aspects of the theory are surveyed by Winkler (1968), Bacharach (1973), Hogarth (1975) and Madansky (1978).

As is clear from the work of Bacharach (1973, 1975) there cannot be a single solution to the multiperson decision problem. Consequently a variety of approaches are found in the contributions to this field.

In one such approach, it is hypothesized that the group reports to a single Bayesian who is to make the decision after incorporating the views of the group. The result is a conventional albeit complex, uni-Bayesian analysis (see Morris, 1977; Lindley, Tversky and Brown, 1979).

Alternatively, the group may enter into a “dialogue” in search of a consensus. Bacharach (1973) introduces the terminology “Bayesian dialogue” to distinguish the process in question from “bargaining”. In a dialogue preferences and opinions are openly expressed while in bargaining these things are deliberately misrepresented to gain strategic advantages.

A dialogue such as that embodied in the Delphi method (see Dalkey, 1967) may well lead to a consensus of opinion. Bacharach (1973) and De Groot (1974) propose models which entail, under certain circumstances, convergence of the group’s opinions given sufficient time. A variant of this method is presented by Press (1978) who also gives a comprehensive bibliography.

However, a consensus cannot always be achieved. Moreover, even in circumstances where the group is supervised by a Bayesian decision-maker, the task of evaluating the Bayes action may be exceedingly complex. For these reasons simple, direct ways of pooling opinions have been sought. Stone’s “linear opinion pool” is one such example (see Stone, 1961). Others are described by Madansky (1978).

When π_i denotes Π_i ’s density with respect to a fixed under-lying measure, the linear opinion pool consists of the prescription,

$$\pi_{AM} = \alpha_1 \pi_1 + \dots + \alpha_n \pi_n, \quad \alpha_i \geq 0, \quad \alpha_1 + \dots + \alpha_n = 1, \quad (2.1)$$

where we use the subscript “AM” to denote “arithmetic mean” and Π_i is Bayesian i ’s prior (or posterior) distribution.

When the decision-makers’ utility functions, say u_i , $i = 1, \dots, n$, are identical, the group’s optimal action may be chosen by entering π_{AM} into a conventional uni-Bayesian analysis (see Bacharach, 1975). Bacharach (1975) and Madansky (1978) present axioms which imply that the linear opinion pool is the one and only prescription for combining the π_i and Raiffa (1968) discusses these results.

Madansky (1978) points out that the linear opinion pool is not “externally Bayesian” if the α_i , $i = 1, \dots, n$, are fixed constants. This means that even if the group’s members share a common likelihood function, the joint posterior density functions will not be the result obtained by first adopting π_{AM} as the joint prior density function and then applying Bayes’ rule. Madansky (1978)

determines the unique manner in which the α_i 's must be updated using the data in order to achieve external Bayesianity.

The prescription for pooling opinions embodied in the linear opinion pool has several features which may in certain circumstances be viewed as drawbacks. Firstly, it does not lead to a decision theory when the u_i are not equal. In fact, Bacharach (1975) proves that no such theory exists if the group is required to be sufficiently coherent in its demands. Secondly, π_{AM} is typically, multi-modal on its domain and so may fail to identify a parameter which typifies its modes, i.e. the individual choices. Thirdly, π_{AM} is not scale invariant. Thus if any π_i is rescaled so that, for example, $\sup_{\theta} \pi_i(\theta) = 1$ instead of $\int \pi_i = 1$, π_{AM} will present a different expression of the group's relative preferences among the parameters, θ , no matter how it is normalized except when $n = 1$ even though individual opinions remain unaltered. This criticism would not be a source of concern perhaps to any Bayesian who demanded that the π_i 's be probabilities in order that each Bayesian's expression of his degree of belief be coherent (cf. di Finetti, 1974). It might be of concern, however, in a situation where the decision-makers shared a common utility function and one of them insisted on adopting a diffuse, i.e. improper prior measure. In the uni-Bayesian theory of course the optimal decision cannot depend on the normalization chosen for the prior measure because the utility functions themselves are determined only to positive affine transformations. This contrasts with multi-Bayesian theory derived in this case from the linear opinion pool where the optimal decision rule would depend on these normalizations.

This third objection given above provides the starting point for the analysis of Weerahandi and Zidek (1978). By requiring, axiomatically, that a pooling prescription be scale invariant, they derive an alternative to the linear opinion pool, namely,

$$\pi_{GM} = \pi_1^{\alpha_1} \dots \pi_n^{\alpha_n}, \quad \alpha_i \geq 0, \quad \alpha_1 + \dots + \alpha_n = 1, \quad (2.2)$$

where "GM" stands for "geometric mean".

This last prescription is not new. Bacharach (1973) calls this the "logarithmic opinion pool" and attributes it to Hammond. Dalkey (1975) proposes the same rule on an *ad hoc* basis as a descriptive summary of the diverse opinions of the group. McConway (personal communication) axiomatically derives a more general formula than (2.2) which includes the latter as a special case.

Madansky (1978) points-out that the prescription embodied in (2.2) is externally Bayesian. As well, this prescription may be employed, more generally, even when the π_i 's are not prior probability densities. For example, as noted above, the π_i 's may be non-integrable, i.e. vague priors, or they might be fiducial density functions which need not integrate to 1 (Wilkinson, 1977), or belief functions (Dempster, 1968; Schafer, 1976).

In the methods described above, the n opinions are aggregated into a single expression of the group's joint opinion and this may then be used in a uni-Bayesian analysis. As pointed out in the Introduction such methods will prove unsatisfactory when there is a marked divergence in the decision-makers' preferences. In this case some compromise solution must be sought. One method of determining such a compromise is provided by Savage (1954). In the remainder of this section alternatives to Savage's method, i.e. the Nash solution and certain of its variants, will be presented. The Nash solution in particular provides a multi-Bayesian decision theory which, in form, more strongly resembles the univariate theory than does Savage's solution.

Consider a statistical decision problem with parameter space Θ and action space \mathcal{A} . For the i th member of a group of n individuals, Π_i will denote the posterior or prior distribution according as data are or are not available. The corresponding utility function will be denoted by $u_i(a, \theta)$, $a \in \mathcal{A}$, $\theta \in \Theta$. The domain of u_i is extended in the usual way to include randomized rules, δ , which may depend on the data when the latter is available. Thus

$$u_i(\delta, \theta) = \int u_i(a, \theta) \delta(da).$$

We regard the multi-Bayesian decision problem as a co-operative n -person, non zero-sum game. This implies that the $\{u_i\}$ and the $\{\Pi_i\}$ are known to each member of the group. Furthermore, joint randomization is permissible, so, for example, the group may agree to choose between alternative rules, δ_1 and δ_2 say by means of a coin-toss.

It will be assumed that the $\{u_i\}$ are bounded, that the i th Bayesian's current worth (utility) is c_i and that \mathcal{A} includes "no action" as a possible action. The use of bluffs and threats is not permitted, and, in fact, we intend that \mathcal{A} should consist of conventional statistical decisions except for the "no action" decision which has been appended.

Nash's solution is implied by certain weak assumptions which will now be stated. Some of these assumptions are non-controversial within the ambit of the normative theory of bargaining and these we will label as *axioms* 1 and 2. The remainder are open to criticism and will be designated as *postulates* 1, 2 and 3; criticism leads to modifications of this list of postulates and variants of Nash's solution. While Nash's solution cannot, therefore, be declared, inarguably, the winner, none of the variants seems as yet to have brought it down from its distinguished position. Moreover, Harsanyi's recent monograph (Harsanyi, 1977) gives Nash's solution additional plausibility by developing it from even weaker and more qualitative assumptions than those of Nash (1950). In presenting below criticisms and variants our aim is to throw more light on the nature of the Nash solution itself rather than to provide a comprehensive account (which it is not) of the theory of bargaining games (cf. Luce and Raiffa, 1957; Cross, 1969; Young, 1975).

Let $\mathbf{u}(\delta) = (u_1(\delta), \dots, u_n(\delta))$ where $u_i(\delta) = \int u_i(\delta, \theta) \Pi_i(d\theta)$; our assumptions imply that $S_n = \{[\mathbf{u}(\delta) - \mathbf{c}]_+ : \delta \text{ randomized}\}$ is a compact, convex subset of $[0, \infty)^n$, where if $\mathbf{x} = (x_1, \dots, x_n)$, $[\mathbf{x}]_+ = (\max\{x_1, 0\}, \dots, \max\{x_n, 0\})$.

Let us now recentre for convenience all of the u_i by the translation, $u_i \rightarrow u_i - c_i$. Then $S_n = \{[\mathbf{u}(\delta)]_+ : \delta \text{ randomized}\}$ represents the set of utility vectors which are attainable by adopting a (jointly) randomized decision rule. Observe that $\mathbf{0} = (0, \dots, 0) \in S_n$.

A *multi-Bayesian decision rule*, δ_n , relative to S_n may now be defined. To this end, let \mathcal{B}_n designate the class of all compact convex S_n 's which contain $\mathbf{0}$ and are subsets of $[0, \infty)^n$. Then δ_n is any rule satisfying $\mathbf{u}(\delta_n) = \boldsymbol{\mu}(S_n)$ where $\boldsymbol{\mu} : \mathcal{B}_n \rightarrow [0, \infty)^n$ is a mapping which satisfies the following conditions:

Axiom 1. (Feasibility). $\boldsymbol{\mu}(S) \in S$ for all $S \in \mathcal{B}_n$.

2. (*Pareto optimality*). There is no $\mathbf{u} = (u_1, \dots, u_n) \in S$, $\mathbf{u} \neq \boldsymbol{\mu}(S)$ such that $\mu_i(S) \leq u_i$ for all i .

Postulate 1. (Total invariance). If $T(\mathbf{u}) = (a_1 u_1, \dots, a_n u_n)$ for all $\mathbf{u} \in [0, \infty)^n$, $a_i > 0$, $i = 1, \dots, n$, then $(T \circ \boldsymbol{\mu})(S) = \boldsymbol{\mu}(T(S))$.

2. (*Symmetry*). If S is symmetric, i.e. closed under permutations of the co-ordinates of its elements, then $\mu_i(S) = \mu_j(S)$ for all i and j .

3. (*Independence of irrelevant alternatives*). $\boldsymbol{\mu}(S) = \boldsymbol{\mu}(U)$ whenever $S \subset U$ and $\boldsymbol{\mu}(U) \in S$.

Nash (1950) shows that these assumptions imply that $\boldsymbol{\mu}(S)$ is the unique point in S which maximizes over S the so-called Nash product, say $P(\mathbf{u})$, given by

$$P(\mathbf{u}) = \prod_{i=1}^n [u_i]^{1/n}, \quad \mathbf{u} \in S. \tag{2.3}$$

The exponent, $1/n$, could have been omitted in this expression since the resulting product would yield the same solution. However, as we point out below, dropping Postulate 2 leads to an alternative product with $1/n$ replaced by $\alpha_i \geq 0$, $i = 1, \dots, n$ with $\sum \alpha_i = 1$. Nash's product is stated in the above form so that it will then be a special case of the more general result. It then follows that a multi-Bayesian decision rule, δ_n , may be found by maximizing $P([\mathbf{u}(\delta)]_+)$ with respect to δ . If the n decision-makers jointly agree that the basic assumptions are reasonable, then they ought to agree to accept the logical implication, the Nash solution to their problem.

Various authors discuss, criticize or offer inconsistent alternatives to Postulates 1, 2 or 3 (cf. Luce and Raiffa, 1957; Bishop, 1963; Owen, 1968; Kalai and Smorodinsky, 1975; Nydegger and Owen, 1975; Kalai, 1977a, b; Roth, 1977).

Kalai (1977a) drops Postulate 2 and shows that the remaining assumptions imply that $\mu(S)$ is the unique point in S which maximizes the Non-symmetric Nash Product,

$$\prod_{i=1}^n [u_i]^{\alpha_i}, \quad \mathbf{u} \in S.$$

Here $\alpha_i \geq 0$ for all i and $\alpha_1 + \dots + \alpha_n = 1$; these constants are not arbitrary and, in fact, $\alpha = (\alpha_1, \dots, \alpha_n)$ is the solution these decision-makers would choose if S were the set, $\{(s_1, \dots, s_n) : s_i \geq 0, s_1 + \dots + s_n \leq 1\}$. Thus the decision-makers would be called upon to agree in preliminary discussions on a choice of α in a hypothetical but intuitively simple situation. This would permit an individual i who lacked confidence in his own judgements to defer to the group by agreeing to accept a small value for $\alpha_i < 1/n$. This same individual might well insist on a large value for $\alpha_k > 1/n$ if he regarded individual k as an expert. Kalai (1977a) does not, however, discuss the choice of α and hence the exact method for its determination is unclear.

Owen (1968) argues that if $S \subset T$ and both share the same disagreement point, \mathbf{c} , then $\mu(S) \leq \mu(T)$ should be a reasonable requirement. While reasonable, this condition is incompatible with Postulates 1–3 (if the Axioms are maintained). To accommodate this condition two alternatives have been pursued.

Kalai and Smorodinsky (1975) replace Postulate 3 by an assumption which Riddell (1978) calls the axiom of individual monotonicity:

Postulate 4. (Individual monotonicity). If $S \subset T$ share the same disagreement point \mathbf{c} , and $b_i(S) = b_i(T)$, $i \leq n-1$, where $b_i(A) = \sup \{s_i : \mathbf{s} \in A\}$ for any $A \subset [c_1, \infty) \times \dots \times [c_n, \infty)$, then $\mu(S) \leq \mu(T)$.

They then show that Axioms 1 and 2 and Postulates 1, 2 and 4 imply that $\mu(S)$ is the unique point on S 's Northeastern frontier which is on the line joining \mathbf{c} (i.e. $\mathbf{0}$ in our presentation) to $(b_1(S), \dots, b_n(S))$. This solution was first proposed by Raiffa (1953).

An alternative approach is adopted by Kalai (1977b) who replaces Postulate 1 by an assumption of homogeneity (in Riddell's 1978 terminology):

Postulate 5. (Homogeneity). $\mu(kS) = k\mu(S)$ for all $k > 0$.

He then shows that $\mu(S)$ must provide a "proportional solution", i.e.

$$\mu(S) = \max \{t : t\mathbf{p} \in S\} \cdot \mathbf{p},$$

where $\mathbf{p} = (p_1, \dots, p_n)$, $p_i > 0$, for all i . His conclusion follows from Axioms 1 and 2 and Postulates 2–5 and one significant additional assumption: S is "comprehensive", i.e. if $\mathbf{s}_1 \in S$ and $\mathbf{0} < \mathbf{s}_2 \leq \mathbf{s}_1$ then $\mathbf{s}_2 \in S$. The constants, p_i , are not determined by Kalai's assumptions. In practice it may well be reasonable to choose the p_i 's so that in every bargaining situation the decision-makers would obtain equal gains in utility.

As this last comment suggests, Postulate 5 unlike Postulate 1 admits the possibility of interpersonal comparisons of utility and so yields a qualitatively different solution than does the latter postulate. Whether or not such comparisons are feasible is controversial (cf. Luce and Raiffa, 1957; Owen, 1968). However, Nash's solution may seem objectionable because such comparisons are inadmissible. As Riddell (1978) points out, it would compel subjects to split 100 chips equally regardless of the respective monetary values of the chips to the subject. For example, suppose the first subject can cash-in any chips he receives as a result of any potential agreement, for 1 dollar each while the second subject is able to obtain 99 dollars per chip. Then an intuitively natural division of the chips would give the first subject 99 chips and the second, just 1 chip. However, the Nash solution splits the chips evenly between the two subjects. This is because of Postulate 1 and the resulting form of the Nash product (see equation 2.3). The

factors, 1 dollar and 99 dollars, respectively, which would be inserted in an effort to “adjust” the two utilities to take account of the differing monetary values of the chips to the subjects would simply “factor out” of the product and in no way influence the computation of the Nash solution.

This solution of Kalai is of very limited appeal in statistical applications because the set of feasible utility vectors, S , may well fail to be comprehensive. For example, in the degenerate case of identical preferences and opinions, S will be a radial line segment emanating from the disagreement point along the 45° line.

3. EXAMPLES

3.2. *Normal Theory Example*

Here the problem is that of estimating the mean, θ , of a normally distributed random variable. The prior (or posterior if data are available) distributions are normal, i.e. $\pi_i(\theta) = \exp[-\frac{1}{2}(\theta - \theta_i)^2]$, $-\infty < \theta < \infty$, $i = 1, \dots, n$. A joint action is to be chosen from $\mathcal{A} = (-\infty, \infty) \cup \{a_0\}$ where a_0 represents the “no action decision”.

For simplicity in this illustrative example, assume that the gain-in-utility is (approximately) a 0–1 function, i.e.

$$u(a, \theta) = \begin{cases} 1, & a = \theta, \\ 0, & a \neq \theta, \end{cases}$$

$\infty < \theta < \infty$. This is formally equivalent to assuming (cf. Weerahandi and Zidek, 1980a, a gain-in-utility function of conjugate form—see Lindley, 1976), i.e. $u(a, \theta) = \exp[-\frac{1}{2}(a - \theta)^2]$ or 0 according to whether $a \neq a_0$ or $a = a_0$.

Since the n utility functions are equal it may be reasonable to aggregate the n -opinions, for example by pooling them in π_{GM} or π_{AM} as discussed in Section 2 and then to perform a conventional Bayesian analysis. Since the gain-in-utility function is the 0–1 function the best supported group decision would be the action which maximizes π_{GM} or π_{AM} if either of these is adopted. It is easily shown that $\pi_{GM}(\theta) \propto \exp[-\frac{1}{2}(\theta - \bar{\theta})^2]$, $-\infty < \theta < \infty$, where $\bar{\theta} = \sum \alpha_i \theta_i$. Thus, the point estimate of θ which is best supported by $\pi_{GM}(\theta)$ is $\theta = \bar{\theta}$.

The point estimate, $\bar{\theta}$, is unsatisfactory in cases where the θ_i are widely separated. If, for example, $n = 2, \alpha_1 = \alpha_2 = \frac{1}{2}, \theta_1 = 0, \theta_2 = 100$, then $\bar{\theta} = 50$. However, $\bar{\theta}$ is not well supported by either of the two Bayesians and so choosing a point near θ_1 or θ_2 might well seem preferable, in this case, to choosing $\bar{\theta}$ itself.

The linear pool, π_{AM} (see (2.1)), has no simple expression and the point estimates best supported by it are difficult to evaluate. If the θ_i are widely separated they will be modes of π_{AM} and hence equally good choices.

Let us now consider the Nash solution. To keep the analysis simple assume $n = 2$. Our goal then reduces approximately to the maximization over δ of

$$\prod_{i=1}^2 \left[\int \pi_i(\theta) \delta(d\theta) \right]^{\frac{1}{2}} \tag{3.1}$$

For convenience let $\pi_i(\delta) = \int \pi_i(\theta) \delta(d\theta)$. Consider

$$S = \{(x, y): x = \pi_1(\theta), y = \pi_2(\theta), \text{ all } \theta\}.$$

Its convex hull is

$$\bar{S} = \{(x, y): x = \pi_1(\delta), y = \pi_2(\delta), \text{ all } \delta\}.$$

It is clear that $[\pi_1(\delta)\pi_2(\delta)]^{\frac{1}{2}}$ is maximized on the Northeast boundary of \bar{S} at say δ^* . By symmetry, δ^* is either degenerate at $\bar{\theta} = \frac{1}{2}\theta_1 + \frac{1}{2}\theta_2$, i.e. $\delta^* = [\bar{\theta}]$, if $(\pi_1(\bar{\theta}), \pi_2(\bar{\theta}))$ is on the boundary of \bar{S} or else δ^* is the two point distribution, $\frac{1}{2}[(1 - \beta^*)\theta_1 + \beta^*\theta_2] + \frac{1}{2}[\beta^*\theta_1 + (1 - \beta^*)\theta_2]$, which assigns half of its mass to each of the points

$(1 - \beta^*)\theta_1 + \beta^*\theta_2$ and $\beta^*\theta_1 + (1 - \beta^*)\theta_2$, where $0 < \beta^* < 1$ is the point at which

$$g(\alpha) \triangleq \exp\left[\frac{1}{2}(\theta_1 - \theta_2)^2 \alpha^2\right] + \exp\left[-\frac{1}{2}(\theta_1 - \theta_2)^2 (1 - \alpha)^2\right]$$

is maximized or equivalently is the point corresponding to the maximum of $f(u) \triangleq \exp[-\gamma u^2] \cosh u$, where $u = \frac{1}{2}(\theta_1 - \theta_2)^2 (\alpha - \frac{1}{2})$ and $\gamma = 2/(\theta_1 - \theta_2)^2$. Thus δ^* is the non-randomized point estimate $\bar{\theta}$ or else $(1 - \beta^*)\theta_1 + \beta^*\theta_2$ or $\beta^*\theta_1 + (1 - \beta^*)\theta_2$ according to whether a fair coin-toss yields heads or tails. To determine which of these alternatives is correct we need only notice that $f(u)$ has a single maximum at 0 or has two global maximums at u^* and $-u^*$ (and one minimum at 0) according as $2\gamma \geq 1$ or $0 \leq 2\gamma < 1$; i.e. according to whether $|\theta_1 - \theta_2| \leq 2$ or $|\theta_1 - \theta_2| > 2$, where u^* is the solution of $\tan h(u^*)/u^* = 2\gamma$ and $u^* = 2(\theta_1 - \theta_2)^2 (\beta^* - \frac{1}{2})$, $\gamma = 2/(\theta_1 - \theta_2)^2$. Hence δ^* is non-randomized or randomized according to whether $|\theta_1 - \theta_2| \leq 2$ or $|\theta_1 - \theta_2| > 2$.

The last result may be described as saying that when given two normal prior distributions with modes θ_1 and θ_2 , a point estimate for θ may be formed by performing a preliminary test of significance. If $|\theta_1 - \theta_2| > 2$, reject the null hypothesis of similar knowledge and toss a fair coin to choose between $\bar{\theta} + u^*/(\theta_2 - \theta_1)$ and $\bar{\theta} + u^*/(\theta_1 - \theta_2)$. If, on the other hand, $|\theta_1 - \theta_2| \leq 2$ then $\bar{\theta} = \frac{1}{2}(\theta_1 + \theta_2)$ is sufficiently well supported by both sets of prior knowledge as to render it an acceptable compromise between θ_1 and θ_2 .

The multi-Bayesian normal-mean estimation theory, which derives from the adoption of the Nash solution involves randomized rules in the manner of the special case considered here. The multi-Bayesian procedure is randomized unless there is fairly strong agreement in beliefs and preferences among the decision makers. An analytical theory is possible in the case of just $n = 2$ decision-makers and some aspects of this theory are presented by Weerahandi and Zidek (1980a). The general case is largely computational. Results relevant to these computations are given in Weerahandi and Zidek (1980b).

3.2. Hypothesis testing

Suppose $\Theta = \Theta_0 \cup \Theta_1$, u_i is the 0–1 utility function, and \mathcal{A} is the two action space, with the “no action” decision adjoined.

To obtain the multi-Bayesian testing procedure for Θ_0 against Θ_1 consider first the non-randomized rule “select Θ_0 ”. The non-symmetrical Nash product of Kalai (1977a) for this rule is $\prod p_i^{\alpha_i}$ where $p_i = \Pi_i(\Theta_0)$, $i = 1, \dots, n$. The corresponding quantity for Θ_1 is $\prod q_i^{\alpha_i}$ where $q_i = 1 - p_i$. If only non-randomized rules were admissible then the best jointly supported action would be “select Θ_0 ” if and only if $\prod (p_i/q_i)^{\alpha_i} \geq 1$, i.e. in terms of log-odds ratios if and only if $\sum \alpha_i \ln (p_i/q_i)^{\alpha_i} \geq 0$. If $n = 1$, we obtain the usual Bayes testing procedure.

To find the best jointly supported randomized rule, say δ^* , note that S_n (see Section 2) consists of convex hull of $\mathbf{0}$ and the line segment joining (p_1, \dots, p_n) and (q_1, \dots, q_n) . The optimal rule is found by maximizing the Nash product over this line segment, i.e. by maximizing $g(\delta) \triangleq \prod [\delta p_i + (1 - \delta) q_i]^{\alpha_i}$, $0 \leq \delta \leq 1$. It is straightforward to show that g is strictly concave. Thus $\delta^* = 0$ if $g'(0) \leq 0$, 1 if $g'(1) \geq 0$ and the unique solution of $g'(\delta) = 0$ if $g'(1) < 0 < g'(0)$. Note that, disregarding an always positive factor,

$$g'(\delta) \propto \sum_{i=1}^n \alpha_i [\delta + q_i/(p_i - q_i)]^{-1}.$$

Thus $g'(1) = 1 - \sum \alpha_i (q_i/p_i)$ so $g'(1) \geq 0$ becomes $\sum \alpha_i (q_i/p_i) \leq 1$. The best supported rule then chooses the action, “select Θ_0 ”. This may be interpreted to mean that Θ_0 is chosen because the average of the odds in favour of Θ_1 is no larger than one. The decision, “select Θ_1 ” has a similar interpretation in terms of the average of odds in favour of Θ_0 .

If in the above analysis the p_i 's represent posterior probabilities then we readily perceive an optional stopping-at-consensus rule which may be used if randomized rules are deemed to be objectionable: continue sampling until either $g'(1) \geq 0$ or $g'(0) \leq 0$ and then take the appropriate terminal action. We have not investigated the properties of this rule.

4. CONCLUSION

The Nash (or possibly non-symmetrical Nash) solution has been proposed in Section 2 for the problem facing n Bayesian decision-makers who are required to make a joint statistical decision even though their beliefs and preference may be quite diverse. Whether or not to adopt the Nash solution is itself a choice confronting the group. This decision is facilitated by the fact that it depends on the acceptability to the group of certain weak and intuitively natural assumptions of which the Nash solution is a logical consequence. Criticisms of these assumptions have led to certain alternatives to the Nash solutions but none of these seems to have been so widely accepted. Therefore, as a multi-person extension of uni-Bayesian statistical decision theory, that provided by the Nash solution would seem to be the most promising candidate.

When the decision-makers' preferences are very similar, a broad alternate approach is to aggregate in some way the beliefs expressed by the decision-makers and thereafter to rely on a conventional Bayesian analysis. Methods for aggregating opinions were summarized in Section 2.

The multi-Bayesian decision rule we propose in Section 2 is that, possibly randomized, $\delta = \delta_n$ which maximizes as a function of δ

$$P([\mathbf{u}(\delta)]_+) = \prod_{i=1}^n \left[\int \int u_i(a, \Theta) \delta(da) \Pi_i(d\Theta) - c_i \right]_+^{\alpha_i};$$

here $\alpha_i = 1/n$ for all i if Nash's symmetry postulate holds and is otherwise determined as a preliminary step in the analysis, subject to the conditions $\alpha_i \geq 0, \alpha_1 + \dots + \alpha_n = 1$. The constant, c_i , represents, essentially, the decision-maker's current worth in utility and is the amount he shall receive if the group fails to agree on a solution.

We conclude with a series of observations.

It is of interest to consider equation (2.2) when $\alpha_i = n^{-1}$. A Bayesian federal statistics bureau in designing a survey might wish to use as a prior distribution, the "democratic" combination of its n citizens' prior densities which is proportional to $[\pi_1 \dots \pi_n]^{1/n}$. Or, more realistically, the Bureau might use a random sample of prior densities, $\tilde{\pi}_1, \dots, \tilde{\pi}_n$. Assuming these are positive, independent and identically distributed, and that the law of large numbers applies, it would follow that

$$\tilde{\pi} \propto [\tilde{\pi}_1 \dots \tilde{\pi}_n]^{1/n} \simeq \exp [E(\ln \tilde{\pi}^*)], \tag{4.1}$$

where $\tilde{\pi}^*$ has the common marginal distribution of the $\{\tilde{\pi}_i\}$. If, for example, $\pi_i(\theta) \propto \exp(-\frac{1}{2}(\theta - Z_i)^2)$, where $Z_i \sim N(\theta_0, 1)$ and θ_0 denotes the actual (but unknown) mean, equation (4.1) would then imply that, approximately, for large n ,

$$\tilde{\pi}(\theta) \propto \exp[-\frac{1}{2}(\theta - \theta_0)^2]. \tag{4.2}$$

Thus in this case the mode of the combined prior would be situated approximately at the correct value. Of course, the problem of designing the survey to collect the sample of priors remains.

Formula (2.2) may be used in an entirely different spirit than that which has motivated its derivation. Suppose there is just one individual whose prior, π_1 , is obtained. The decision-maker may well elect to throw into the analysis a second, very diffuse prior, π_2 , and then combine them according to (2.2). In this way he can respond to one of the frequent criticisms levelled at the Bayesian methodology, namely, that it offers no means by which the surmised quality of a prior can be reflected in the analysis. He may choose $\alpha_2 \simeq 1$ if he deems π_1 to be of low quality and $\alpha_2 \simeq 0$ if he believes it to be of high quality.

ACKNOWLEDGEMENTS

We are indebted to Professors A. P. Dawid, D. V. Lindley, A. Madansky and M. Stone and to Dr K. McConway for their comments and for supplying references. Our thanks as well to Professor A. Madansky and to Dr K. McConway for providing accounts of their own recent

contributions to this area. The suggestions of a referee led to substantial improvements in our presentation and for these suggestions we are very grateful. This work was supported by the Natural Sciences and Engineering Research Council of Canada.

REFERENCES

- BACHARACH, M. (1973). Bayesian dialogues. Unpublished manuscript. Christ Church, Oxford.
- (1975). Group decisions in the face of differences of opinion. *Manag. Sci.*, **22**, 182–191.
- BISHOP, R. L. (1963). Game theoretical analysis of bargaining. *Quart. J. Econ.*, **77**, 559–602.
- CROSS, J. G. (1969). *The Economics of Bargaining*. New York.
- DALKEY, N. C. (1969). The Delphi method: an experimental study of one group decision. RAND Memo RM–5888–PR.
- (1975). Toward a theory of group estimation. In *The Delphi Method: Techniques and Applications* (H. A. Linstone and M. Turoff, eds). Reading, Mass: Addison-Wesley.
- DEGROOT, M. (1974). Reaching a consensus. *J. Amer. Statist. Ass.*, **69**, 118–121.
- DI FINETTI, B. (1974). *Theory of Probability*. New York: Wiley.
- DEMPSTER, A. P. (1968). A generalization of Bayesian inference. *J. R. Statist. Soc. B*, **30**, 205–247.
- HARSANYI, J. (1955). Cardinal welfare, individualistic ethics and interpersonal comparisons of utility. *J. Polit. Econ.*, **63**, 309–321.
- (1977). *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*. New York: Cambridge University Press.
- HOGARTH, R. M. (1975). Methods for aggregating opinions. In *Decision Making in Human Affairs* (H. Jungerman and G. de Zeeuw, eds), pp. 231–255. Dordrecht, Holland: D. Reidel Publishing Co.
- KALAI, E. (1977). Nonsymmetric Nash solutions and replications of 2-person bargaining. *Int. J. Game Theory*, **6**, 129–133.
- (1977). Proportional solutions to bargaining situations: interpersonal utility comparisons. *Econometrica*, **45**, 1623–1630.
- KALAI, E. and SMORODINSKY, M. (1975). Other solutions to Nash's bargaining problem. *Econometrica*, **43**, 513–518.
- LINDLEY, D. V. (1976). A class of utility functions. *Ann. Statist.*, **1**, 1–10.
- LINDLEY, D. V., TVERSKY, A. and BROWN, R. V. (1979). On the reconciliation of probability assessments (with Discussion). *J. R. Statist. Soc. A*, **142**, 146–180.
- LUCE, D. and RAIFFA, H. (1957). *Games and Decisions*. New York: Wiley.
- MADANSKY, A. (1978). Externally Bayesian groups. Unpublished manuscript. University of Chicago.
- MORRIS, P. A. (1977). Combining expert judgements: a Bayesian approach. *Manag. Sci.*, **23**, 679–693.
- NASH, J. F., Jr (1950). The bargaining problem. *Econometrica*, **18**, 155–162.
- NYDEGGER, R. and OWEN, G. (1975). Two person bargaining: an experimental test of the Nash axioms. *Int. J. Game Theory*, **3**, 239–249.
- OWEN, G. (1968). *Game Theory*. Philadelphia: W. B. Saunders Co.
- PRESS, S. J. (1978). Qualitative controlled feedback for forming group judgements and making decisions. *J. Amer. Statist. Ass.*, **73**, 526–535.
- RAIFFA, H. (1953). Arbitration schemes for generalized two person games. In *Annals of Mathematical Studies* (H. W. Kuhn and A. W. Tucker, eds), Vol. 28. Princeton: Princeton University Press.
- RIDDELL, W. C. (1978). Bargaining under uncertainty. Res. Paper No. 78–15, Department of Economics, University of Alberta.
- ROTH, A. F. (1977). Independence of irrelevant alternatives and solutions to Nash's bargaining problem. *J. Econ. Theory*, **16**, 247–251.
- SAVAGE, L. J. (1954). *The Foundations of Statistics*. New York: Wiley.
- SCHAFFER, G. (1976). *A Mathematical Theory of Evidence*. Princeton: Princeton University Press.
- STONE, M. (1961). The opinion pool. *Ann. Math. Statist.*, **32**, 1339–1342.
- WEERAHANDI, S. and ZIDEK, J. V. (1978). Pooling prior distributions, Technical Report 78–34. Institute of Applied Mathematics and Statistics, University of British Columbia.
- (1980a). Bi-Bayesian estimators of multinormal means. (Submitted.)
- (1980b). Multi-Bayesian estimators of the normal mean. (In preparation.)
- WILKINSON, G. N. (1977). On resolving the controversy in statistical inference. *J. R. Statist. Soc. B*, **39**, 119–142.
- WINKLER, R. L. (1968). The consensus of subjective probability distributions. *Manag. Sci.*, **2**, B65–B75.
- YOUNG, O. R. (1975). *Bargaining*. Chicago: University of Illinois Press, 1975.
- ZEUTHEN, F. (1930). Du monopole bilateral. *Révue d'Economie Politique*, **XLVII**, 1651–70.