

A Theorem for Bayesian Group Decisions

by

Ralph L. Keeney and Robert Nau

Fuqua School of Business

Duke University

Durham, NC 27708-0120

Published in *Journal of Risk and Uncertainty* (2011) 43:1-17

© Springer-Verlag 2011

Abstract: This paper presents a natural extension of Bayesian decision theory from the domain of individual decisions to the domain of group decisions. We assume that each group member accepts the assumptions of subjective expected utility theory with respect to the alternatives from which they must choose, but we do not assume, a priori, that the group as a whole accepts those assumptions. Instead, we impose a multiattribute utility independence condition on the preferences of the group with respect to the expected utilities of its actions as appraised by its members. The result is that the expected utility of an alternative for the group is a weighted average of the expected utilities of that alternative for its members. The weights must be determined collectively by the group. Pareto optimality is not assumed, though the result is consistent with Pareto optimality.

Keywords: group decisions, utility theory, probability aggregation, preference aggregation, impossibility theorems, Pareto assumption, constructing values, interpersonal comparison of preferences.

JEL Classifications: D70, D81

There is no unique definition of a group decision problem. In this paper, the group problem is defined as a decision in which two or more individuals must collectively choose from among a set of alternatives under conditions of uncertainty. Prototypical group decisions of this nature are made by boards of directors, partners in an organization, or teams. This definition excludes negotiations, where individual negotiators choose among different alternatives, and problems such as risk sharing, where consequences are split among the individual participants. It also excludes social planning decisions where the planner faces a decision and wishes to account for the preferences of many individuals affected by the decision.

Group decisions routinely occur; many are important and many are complex. If they are important, meaning that the differences in the possible consequences of the alternatives are significant, then it is important to make these decisions thoughtfully. If they are complex, then a structured approach that analyzes these complexities may provide insights well worth the effort. This paper develops a result for prescriptive analysis of group decisions.

The complexities of many group decisions that should be addressed by an informative analysis include the following:

1. At the time a decision must be made, there may be significant uncertainties concerning the events that will determine the consequences of each alternative.
2. Individuals may have uniquely personal beliefs about the likelihoods of the events and uniquely personal evaluations of the consequences.
3. The relative desirability (i.e., strength of preference) of the alternatives to each individual in the group is relevant to the decision, not only each individual's order of preferences for the alternatives.
4. Interpersonal comparison of utilities, meaning the differences in the relative desirabilities of alternatives to different members of the group, is relevant to the decision.
5. The relative importance or power of each individual in the group matters.

Section 1 of the paper provides a review of the literature and a discussion of the problem to be addressed, namely that the subjective expected utility (Bayesian) framework does not directly extend to group decisions without modification. Section 2 outlines a general group decision process that provides the context for our analysis. Section 3 reviews the fundamental assumptions of subjective expected utility preferences that are applicable in whole or in part to individuals and groups. Section 4 presents the conventional method for linking a state-independent subjective expected utility model of individual preferences to a state-dependent expected utility model of group preferences via a Pareto optimality condition. Section 5 presents our alternative derivation of a state-dependent subjective expected utility model of group preferences, which is based on explicit modeling of group preferences for assignments of expected utilities to its members, subject to a utility-independence condition. Section 6 discusses how a group might obtain the information needed to implement the group decision model for a specific decision, and Section 8 is a concluding summary.

1. Extending Bayesian Decision Making from Individuals to Groups

The axiomatic theory of subjective expected utility introduced by Savage (1954), which integrated and extended earlier work by Ramsey (1926), de Finetti (1937) and von Neumann and Morgenstern (1947), provides a foundation for the modeling of rational decisions that has been widely and successfully applied in theory and practice. The subjective expected utility model is called the “Bayesian” model insofar as it implies that posterior probabilities obtained by conditioning one’s beliefs on the prospective occurrence of an event should be obtained from prior probabilities by an application of Bayes’ rule.¹ Anscombe and Aumann (1963) provided an alternative axiom system for the subjective expected utility model in which events with objective probabilities are assumed to be available for purposes of constructing randomized acts and calibrating the measurements of subjective probability, which simplifies the application of the model.

¹ This does not mean that Bayes’ rule should be applied retrospectively in an automatic fashion after the event is observed, i.e., it is not a model of actual learning over time. Rather, it is a model of expected learning that can be used for developing and comparing hypothetical contingency plans and for evaluating sources of information.

The problem addressed by the Bayesian model is the following:

Individual Decision Problem. An individual decision maker must choose among a set of alternatives that may lead to different consequences depending on the outcomes of events. Specifically, there is a finite² set S of states of the world, that may occur and a finite set C of consequences that may be experienced. An alternative (“act”) is a function that maps events to consequences: if act f is chosen and state $s \in S$ then occurs, the consequence will be $f(s) \in C$.

The description of a consequence includes all aspects of the experience that might follow from the choice of an alternative that are valued by the decision maker.

Under the assumptions of the Bayesian model, an individual decision maker should evaluate an alternative f on the basis of its expected utility,

$$U(f) = \sum_s p(s) u(f(s)), \quad (1)$$

and should select the act with the highest expected utility, where p is a subjective probability distribution that represents her judgments about the likelihoods of events and u is a utility function representing her attitude toward risk and valuation of consequences. The latter function can be scaled without loss of generality so that

$$u(c^\circ) = 0 \text{ and } u(c^*) = 1, \quad (2)$$

where c° and c^* are her least-preferred and most-preferred consequences, respectively.

The problem of extending the Bayesian model for individual decisions to a model of group decisions has been studied by many authors. Some have searched for a group utility function and a group probability distribution that can be used analogously to (1) to calculate a group expected utility for each of the alternatives (Raiffa,1968; Hylland and Zeckhauser,1979; Seidenfeld et al.,1989; Mongin, 1995; Gilboa et al., 2004). In addition,

² Finiteness of the sets of states and consequences will be assumed here partly to simplify the exposition and partly because the problem is usually framed this way in applied decision analysis.

Keeney and Raiffa (1976) have worked on various approaches to develop a group utility function. Raiffa (1968), Aumann (1976), Clemen and Winkler (1999), and O'Hagan et al. (2006) have worked on various approaches to obtain group probabilities for the events that influence the consequences of the decision.

A variety of results have been proved regarding the “impossibility” of extending the Bayesian model of individual decision making to a corresponding Bayesian model of group decision making. In the special case of a common utility function, the group decision making problem reduces to one of pooling probability judgments, and there is no pooling formula that yields group probabilities that simultaneously satisfy the “marginalization property” (namely that marginalization of pooled probabilities should yield the same result as pooling of marginal probabilities) and the “external Bayesianity” property (namely that pooled probabilities should be updated according to Bayes’ rule when individuals agree on the likelihood function). These issues are discussed by Genest and Zidek (1986). Hylland and Zeckhauser investigate the problem of separately aggregating the individuals’ probability assessments into a group probability assessment and aggregating their state-independent utility functions into a state-independent group utility function, then calculating the group expected utility for alternatives as in (1). In addition, they require that the procedure should be weakly Pareto optimal, meaning that when all members of a group have a common preference for one alternative over another, then any proposed Bayesian group model for their choice must reflect this preference and assign a higher expected group utility to the commonly preferred alternative. They also exclude aggregations that are dictatorial, meaning identical to the set of probabilities and the utility function of only a single individual in the group. They then prove that there is no group decision procedure consistent with these assumptions.

Seidenfeld et al. (1989) state: "An outstanding challenge for ‘Bayesian’ decision theory is to extend its norms of rationality from individuals to groups. Specifically, can the beliefs and values of several Bayesian decision makers be amalgamated into a single Bayesian profile that respects their common preferences over options? ... In other words, can their shared strict preferences over acts be reproduced with a Bayesian rationale (maximizing

expected utility) from beliefs (probabilities) and desires (utilities) that signify a rational compromise between their rival positions?" They consider the implications of both weak and strong Pareto conditions, showing that the former leads to dictatorial solutions and the latter to non-existent solutions when both individual and group preferences satisfy all of the usual axioms. They go on to consider relaxing the completeness axiom for the preferences of the group, and they point out that incomplete group preferences *can* be determined from individual preferences via a Pareto condition, but not by considering probabilities and utilities separately. Rather, the group's preferences are determined by the set of "probability-utility" pairs of the members. The latter approach does not attempt to perform any weighting of individual preferences in pursuit of a group decision: a lack of consensus among the members merely leads to indecisiveness on the part of the group.

Mongin (1995) states that "The issue of Bayesian aggregation, as we shall refer to it, is obviously an important one in collective decision-making". His concern is with a group of individuals who accept Bayesian principles and also wish that the group will act consistently with these principles. Specifically, he investigates aggregation of the individual group member's probabilities over events, utilities over consequences, and expected utilities over alternatives to obtain representations of a group's probabilities, utilities, and expected utilities that are consistent with a Pareto principle for the individual's expected utilities. With Pareto indifference, Mongin concludes that the only consistent aggregation is that of a dictator (or inverse dictator) where the group aggregations are identical to (or the inverse of) those of one of the individuals in the group. With a strong Pareto condition, he proves that there are no consistent aggregations.

Gilboa et al. (2004), when considering social decisions, argue that a Pareto assumption is not appropriate to justify choices when individuals in the group (society) have different sets of beliefs (i.e. probabilities) to describe the possible consequences. They invoke a Pareto assumption only when all of the group members have common probabilities, which was the situation that Harsanyi (1955) analyzed using the Pareto assumption. With the Pareto assumption limited to comparisons with common probabilities, Gilboa et al.

prove that the group's utility function should be a linear combination of the individual's utility functions and that the group's probabilities should be a linear combination of the individual's probabilities. However, when individuals do not have identical probabilities, this result sometimes leads to evaluations of alternatives in which each individual prefers alternative A over alternative B and yet the group is advised to choose alternative B, in violation of Pareto optimality.

Many of the difficulties mentioned above—dictatorships, non-existence, indecisiveness, lack of a basis for interpersonal tradeoffs—can be circumvented if the preferences of the group are not required to satisfy one of the other conditions imposed by the Bayesian model, namely that utilities for consequences should be state-independent. Harsanyi (1955) showed that Arrow's famous impossibility theorem (Arrow, 1951) for group decisions does not apply when a cardinal expected utility model rather than an ordinal utility model is used to represent individual preferences among a set of consequences. With cardinal utility functions, it is meaningful to compute sums or averages of the utilities of different individuals. Harsanyi showed that if the preferences of both the individuals and the group have a cardinal expected-utility representation, and if group members share an objective probability measure, then the group preferences satisfy a Pareto condition with respect to the preferences of the individuals if and only if the group utility function is a weighted sum of the individual utility functions. This result can be proved using a separating hyperplane argument (Border, 1985).

An analogous result holds for subjective expected utility preferences, as shown by Mongin (1998) and Chambers and Hayashi (2006). If both the individuals and the group have subjective expected utility preferences, then the preferences of the group satisfy a Pareto condition with respect to the preferences of the individuals if and only if the expected utility of the group is a weighted sum of the expected utilities of the individuals. This fact also follows from a separating hyperplane argument, as will be shown in section 4. However, if the individuals are heterogeneous in both their probabilities for events and utilities for consequences, then a group utility function determined in this manner is not state-independent except in the trivial case of a dictatorship. So, the Pareto principle can

be respected, but only by giving up the state-independent-utility assumption that is usually applied to individual preferences. This weakens the meaning of “Bayesian” to the extent that the term is equated with the representation of beliefs by subjective probabilities, because subjective probabilities are not uniquely determined by choices when utilities are state-dependent.³

From a decision-analysis viewpoint, the problem with the standard derivation of the additive utility model is not so much the lack of determination of group probabilities. (As noted earlier, the full range of properties of subjective probabilities cannot be preserved by any preference-aggregation method even in the special case where only probabilities are heterogeneous.) Rather, the problem is that the Pareto argument assumes the conclusion, namely that the group has preferences among the same set of alternatives as the individuals and the group’s preferences satisfy the same axioms as those of the individuals, apart from state-independence. It does not explicitly address the fact that the substance of the group’s problem is to make tradeoffs among the preferences of its members. The weights that it applies to the members’ utility functions are determined implicitly from the relation between a group utility function and individual utility functions that have been assumed into existence separately.

In this paper we present an alternative and more constructive derivation of the additive model of group utility, in which the preferences of the group are not assumed a priori to satisfy the axioms of subjective expected utility with respect to the alternatives in the original problem, with or without state-independence. Instead, the group is assumed to have expected-utility preferences among gambles whose outcomes are assignments of expected utility to its individual members. Thus, interpersonal comparisons enter into the model at the level of fundamental measurements by the group. A utility independence

³ When utilities are state-dependent, the unique determination of subjective probabilities requires either that they should be measured verbally rather than inferred from preferences (DeGroot, 1970) or that they should be measured by assessing preferences among larger sets of acts involving higher degrees of counterfactualism, for example, acts in which states of nature are imagined to have objective probabilities or lotteries in which each outcome is a combination of a prize and a state (Karni et al., 1983; Karni, 1985; Schervish et al., 1990; Karni 2007). However, models of the latter kind, which invoke preferences that do not correspond to the sort of choices that are really available, kind give rise to their own impossibility theorems.

condition, similar to one introduced by Fishburn (1965) for multiple objectives, is then used to link the group's preferences for the risks taken on behalf of its members with the members' own preferences for the risky alternatives, which turns out to yield a state-dependent expected utility representation of group preferences among the same alternatives.

2. The Group Decision Process

Within the general framework of Bayesian decision analysis, it would be natural to use the individuals' models of their own preferences on behalf of the group in constructing a model for the preferences of the group itself. To understand why this is a natural approach for group decisions of the nature defined in this paper, consider how a group might make such a decision. Imagine that it is an important decision that will require a number of meetings over time in order to develop a decision frame, gather and exchange information, create and discuss alternatives, and compare probability judgments, attitudes toward risk, and valuations of outcomes. The greater the differences among the members' prior knowledge and expertise and points of view, the greater their potential for learning from each other and updating their own judgments and preferences. This process of communication, together with the gathering of new information from mutually respected sources, may lead to some convergence. That being said, the individual's updated judgments and preferences are still individual judgments and preferences. Using these, each of the group members should evaluate the alternatives and decide what he or she thinks is the best alternative for the group to select. Situations in which the group members completely agree on all judgments or on preferences or on both are special and fortunate cases. Even when the members do not end up agreeing on judgments and preferences, the group still must collectively make a decision on everyone's behalf.

Hence, the decision process can be viewed in two stages, and our formulation focuses on them separately. In the first stage, the group members develop and refine a common understanding of the decision problem and then each of them individually evaluates the alternatives from the perspective of what he or she thinks is best for the group. In the

second stage, the group collectively evaluates the alternatives, using their individual evaluations of those alternatives as inputs. The second stage is the real group decision, in which the group must use some formal or informal procedure to choose a single alternative on behalf of its members in light of everything that is known. The procedure might involve continuing the discussion in hope of reaching a satisfactory level of consensus, or it might terminate in one of many types of voting procedures. Or, it might involve the application of an analytical procedure in order to construct the preferences of the group from comparisons of the preferences of its members, and such is the approach that will be presented here.

3. Foundation for the Group Decision Theorem

The decision faced by a group in this paper is the following:

Group Decision Problem. A group of M members must choose among a set of alternatives that may lead to different consequences depending on the outcomes of events. Specifically, there is a finite set S of states of the world, that may occur and a finite set C of consequences that may be experienced. Each combination of a chosen alternative and a realized state leads to an element from the set of consequences. An alternative (“act”) is a function that maps events to consequences: if act f is chosen and state $s \in S$ then occurs, the consequence will be $f(s) \in C$.

The first stage of the group decision problem is the evaluation of alternatives by the individual members in order to construct their own preferences. Under the assumptions of the Bayesian model for individual decision making under uncertainty, member m should evaluate alternative f using the expected utility index

$$U_m(f) = \sum_s p_m(s) u_m(f(s)), \quad (3)$$

where p_m is her own subjective probability distribution that represents judgments about the likelihoods of events and u_m is her own utility function that represents risk preferences and values for consequences. These may be scaled in the same manner as (2) by

$$u_m(c_m^\circ) = 0 \text{ and } u_m(c_m^*) = 1, \quad (4)$$

where c_m° and c_m^* are the consequences that member m judges to be worst and best, but they not be the same for everyone, i.e., the members may disagree on what is best and worst as well as on their relative strengths of preference for what lies in between.

In order to calibrate the preferences of an individual in terms of real-valued probabilities and utilities, it is necessary to embed the concrete decision problem in a larger problem in which the objects of choice are acts that include hypothetical and even counterfactual mappings from states to consequences, as first proposed by Savage (1954). The approach that will be used here follows the later version of Anscombe and Aumann (1963) in which the set of acts includes not only hypothetical mappings from states to consequences, but also hypothetical mappings from states to objective lotteries over consequences. The incorporation of objective probabilities into the very definition of the acts makes it possible to form probabilistic mixtures of them (“gambles”) in a natural way, which simplifies the formulation of axioms that rational preferences should satisfy. In particular, it allows von Neumann and Morgenstern’s version of the independence axiom to be used, and it also fits nicely with the use of objective randomization devices (e.g., “probability wheels”) for calibrating subjective probability judgments in decision analysis applications.

In this setting, an act can be denoted by a vector f whose elements are doubly indexed by states and consequences, with $f(s,c)$ denoting the objective probability it assigns to consequence c in state s . An act f is a “constant act” if it yields the same objective lottery in every state, i.e., if $f(s,c) = f(s',c)$ for any two states s and s' and any consequence c . If f and h are acts and λ is a number between 0 and 1, then $\lambda f + (1-\lambda)h$ denotes the “objectively mixed” act that yields consequence c with probability $\lambda f(s,c) + (1-\lambda)h(s,c)$ in state s . If E is any event (a state or set of states), then $Ef + (1-E)h$ denotes the “subjectively mixed” act that yields the same lotteries as f in states where E is true and yields the same lotteries as h in states where E is false. Let \succsim and \succ denote weak and strict preference, respectively, between pairs of acts. In these terms, the axioms of subjective-expected-utility preferences can be given as follows:

- A1: (weak order) \succsim is complete, transitive, and reflexive,
A2: (non-triviality) there exist f and g such that $f \succ g$ (i.e., NOT $g \succsim f$),
A3: (continuity) for fixed f , the set of g such that $f \succsim g$ is closed, and vice versa,
A4: (independence) $f \succsim g \Rightarrow \lambda f + (1-\lambda)h \succsim \lambda g + (1-\lambda)h$ where $0 < \lambda < 1$,
A5: (state-independence): $f \succsim g \Rightarrow Ef + (1-E)h \succsim Eg + (1-E)h$ for all *constant* acts f , g , h , and every event E .

Axioms A1 through A4 are the axioms of expected utility, merely applied to acts whose outcomes depend on events as well as objective randomization, and A5 is the additional assumption needed to separate subjective probabilities for events from utilities for consequences. From these assumptions it follows that there exist a unique subjective probability distribution p and a state-independent von Neumann-Morgenstern utility function u , unique up to positive affine transformations, with respect to which $f \succsim g$ if and only if the expected utility of f is greater than or equal to the expected utility of g .⁴

The function of states and consequences that is defined by the product of the probability distribution p and utility function u , will be called the member's *expected utility function*. It can be represented by a vector v of length $|S| \times |C|$ in which $v(s,c) = p(s)u(c)$. Then the subjective expected utility of f can be conveniently expressed as the inner product $f \square v = \sum_s \sum_c f(s,c) v(s,c)$, in terms of which $f \succsim g$ if and only if $f \square v \geq g \square v$.

Preferences which satisfy all of the same axioms *except* for the state-independence axiom (A5) are also representable by an expected utility function v , unique up to positive affine transformations, such that $f \succsim g$ if and only if $f \square v \geq g \square v$, but in this case v need not have the decomposition $v(s,c) = p(s)u(c)$ for some probability distribution p and some state-independent utility function u . Rather, $v(s,c)$ must simply be interpreted as the

⁴ From the first four axioms, it follows by a separating hyperplane argument that preferences have (at least) a state-dependent expected utility representation, and the fifth axiom allows it to be decomposed into a product of a unique probability distribution and a state-independent utility function, although the uniqueness of the probabilities depends on a conventional implicit assumption that utilities which are state-independent in relative terms are also state-independent in absolute terms..

contribution to expected utility that is associated with the receipt of consequence c in state s . It can be considered as the product of a probability $p(s)$ and a state-*dependent* utility $u(s,c)$, but neither the probability nor the state-dependent utility is uniquely determined, absent additional measurements that are verbal rather than behavioral and/or which involve even-more-counterfactual alternatives.

There is a subtlety concerning the interpretation of the consequences that are used in the definitions of acts for the individuals and for the group. Strictly speaking, a consequence is a personal experience of an individual decision maker at a requisite level of description, so different individuals do not necessarily experience the “same” consequences of events when actions are taken on their behalf. A different interpretation is that a consequence is “whatever happens” when a given real alternative is chosen and a given real event occurs, regardless of how it may be perceived and valued by different individuals. Under the latter interpretation, which is adopted here, the set of consequences is the same for all individuals and for the group, and the set of acts is therefore the same for the individuals and for the group.

4. The Original Group Decision Theorem

The individual members of the decision-making group are henceforth assumed to have preferences consistent with the Bayesian model:

Assumption 1: Each of the group members has state-independent subjective-expected-utility preferences among acts, according to axioms A1–A5.

It follows from Assumption 1 that the preferences of member m are represented by a state-independent expected utility function v_m that is unique up to positive affine transformations, for $m = 1, \dots, M$.

An additive representation of the group’s preferences can be derived in one of two ways. The approach used by Mongin (1998) and Chambers and Hayashi (2006) follows that of

Harsanyi, generalized to a setting of subjective expected utility rather than expected utility. A version of this derivation is presented below to provide a basis for comparing its modeling assumptions and method of proof with those of our theorem in Section 5.

The conditions that are assumed to be satisfied by the group preferences are the following:

Assumption 2: The group has state-dependent subjective-expected-utility preferences among the same acts, according to axioms A1-A4.

It follows that the preferences of the group are represented by a not-necessarily-state-independent expected utility function v_G , also unique up to a positive affine transformation. Finally, the individual and group preferences are connected by:

Assumption 3: The preferences of the group satisfy a strict Pareto condition with respect to the preferences of its members: if all the members weakly prefer f to g , then so does the group, and if, in addition, any one member strictly prefers f to g , then so does the group.

In these terms, the original version of the group decision theorem is as follows:

Theorem 1: Given Assumptions 1 and 2, Assumption 3 is satisfied if and only if there exist positive weights $\{\alpha_1, \dots, \alpha_M\}$, unique up to a common scale factor, with respect to which the expected utility function of the group can be expressed as a weighted sum of the expected utility functions of the individual members, i.e., $v_G = \sum_m \alpha_m v_m$.

Proof: Let p_m and u_m denote the subjective probability distribution and state-independent utility function that represent the preferences of member m according to Assumption 1; let v_m denote the corresponding expected-utility function that is determined by $v_m(s,c) = p_m(s)u_m(c)$; and let v_G denote a state-dependent expected-utility function that represents the preferences of the group, according to Assumption 2.

Without loss of generality, assume that the individual state-independent utility functions $\{u_m\}$ are normalized by the addition of constants so that their values sum to zero across consequences, and similarly assume that the group's state-dependent expected utility function v_G is normalized by the addition of a constant within each state so that its values sum to zero across consequences within each state. (Normalization of a utility function in this fashion merely adds an identical constant to the expected utility of every act and therefore does not change the representation of preferences.) Under this normalization, $\{v_m\}$ and v_G are all unique up to positive scaling. Let V denote the open convex cone generated by the vectors $\{v_m\}$, i.e., the set of all positively weighted sums of the individual expected utility functions. By the finite dimensional separating hyperplane theorem, v_G and the relative interior of V are disjoint sets if and only if they can be properly separated by a hyperplane, which in this case may be taken to pass through the origin because V is a cone that recedes from the origin. Proper separation under the latter condition means that there exists a hyperplane with normal vector w such that $w \cdot v_m \geq 0$ for every m , and $w \cdot v_G \leq 0$, and either $w \cdot v_m > 0$ for some m or else $w \cdot v_G < 0$. Without loss of generality it may be assumed that the elements of w sum to zero within each state because, in view of the normalization of the utility functions, the addition of the same constant to the elements of w within any state does not change any of the inner products $w \cdot v_m$ or $w \cdot v_G$. Any such vector w is proportional to the difference between two acts, i.e., it can be written in the form $w = \alpha(f - g)$ for some acts f and g and $\alpha > 0$. The existence of two such acts is precisely a violation of the strict Pareto condition, because it means that either (a) $f \cdot v_m \geq g \cdot v_m$ for every m but meanwhile $f \cdot v_G < g \cdot v_G$, or else (b) $f \cdot v_m \geq g \cdot v_m$ for every m , with $f \cdot v_m > g \cdot v_m$ for some m , but meanwhile $f \cdot v_G \leq g \cdot v_G$. Hence the strict Pareto condition is satisfied if and only if v_G is *not* disjoint from the relative interior of V , in which case it is a positively weighted sum of $\{v_m\}$. ■

The fact that the utility weights of the individual members are uniquely determined without any explicit interpersonal comparisons seems rather remarkable on the surface, but it merely follows from the fact that the group expected utility function is assumed into existence at the beginning along with the individual expected utility functions. So, the question of how the weights “ought” to be determined is not addressed by the setup of

the theorem. It is still up to the group members to decide how to weight each others' utility functions if they want the preferences of the group to conform to the axioms of expected utility, minus the state-independence requirement.⁵

In the special cases in which the group members have identical probabilities for events or identical utilities for consequences, the group expected utility function reduces to a state-independent form in which the weighted average is performed with respect to the non-identical parameters by themselves. In particular, if the members have a common subjective probability distribution \mathbf{p} (or if the probabilities of the events are objective, as in Harsanyi's model), then the group expected utility function has the factorization $v_G(s,c) = p(s)u_G(c)$, where $u_G(c) = \sum_m \alpha_m u_m(c)$ is an aggregate utility function for consequences. If they have a common utility function \mathbf{u} for consequences (for example, if they are all risk neutral), then the group expected-utility function has the factorization $v_G(s,c) = p_G(s)u(c)$ where $p_G(s) = \sum_m \alpha_m p_m(s)$ is an aggregate subjective probability distribution.

⁵ Mongin (1998) observes that the additive model is a "mathematically trivial resolution" of the group decision problem insofar as the Pareto condition is obviously satisfied by any positive weighted sum of the individual utility functions, although the uniqueness of the weights and the "only if" part of the theorem are not quite so trivial, requiring an appeal to the separating hyperplane argument.

5. A New Group Decision Theorem

An alternative formulation of the additive model of group preferences will now be given. In this version, the group is not assumed a priori to have state-dependent subjective expected utility preferences over the common set of acts, nor is the usual Pareto condition invoked, i.e., Assumptions 2 and 3 are not used. In fact, the objects of the group's preferences are not the same acts as those experienced by its members. Rather, the group is assumed to have preferences among different possible assignments of expected utilities to its members. In particular, the group is assumed to have expected-utility preferences among gambles over such assignments, as if the group contemplates making randomized choices among alternatives that may yield different profiles of expected utilities for its members.

To provide a rationale for this approach to modeling the group decision problem, note that every act f (a mapping from states to objective lotteries over consequences) can be associated with the vector $(f \square v_1, \dots, f \square v_M)$ whose elements are the expected utilities on the basis of which it is evaluated by its members from their own individual perspectives, as developed in the first stage of the analysis. Every such vector is a point in the hypercube $[0,1]^M$, and the set of all such vectors generated by the stage-1 analysis is convex, with extreme points corresponding to non-randomized acts.

In modeling the decision problem from the perspective of the group as a whole, we imagine that the group explicitly considers risks and tradeoffs with respect to the expected utilities that its decisions yield for different members. In so doing, we enlarge the canvas in two ways. We first expand the stage-1 problem so that the set of possible expected-utility vectors consists of the entire hypercube, as if there are hypothetical alternatives available that yield all possible matchups of least-preferred and most-preferred outcomes for different members, i.e., all possible 0-1 profiles of expected utilities, as well as their convex combinations. The elements of the hypercube will henceforth be called "pure group acts." A pure group act can be denoted by $F = (F_1, \dots, F_M)$, where F_m is the expected utility it yields for member m , measured on a scale of 0 to

1. In the special case where F corresponds to the choice of an act f that is available in the real decision problem, we have $F_m = U_m(f)$ for every m .

In stage 2 of the analysis, where the preferences of the group as a whole are modeled, the canvas is further enlarged by assuming that the group contemplates not only pure group acts but also objectively randomized acts, i.e., gambles among vectors of expected utilities, and the group's preferences among such gambles are required to satisfy the axioms of expected utility. Under these assumptions, it is conceivable that the group's preferences among gambles could depend in complicated ways on the joint distributions that they yield for its members' expected utilities. In particular, it does not (yet) follow that the group must reduce compound lotteries on a member-by-member basis when evaluating gambles among expected-utility vectors, nor must its preferences align exactly with those of any one of its members, *ceteris paribus*. For example, in a 2-person scenario, it could be the case that a 50-50 gamble between the pure group acts $(0, 0)$ and $(1, 1)$ is strictly preferred by the group to a 50-50 gamble between $(1,0)$ and $(0,1)$, on grounds of *ex post* equity, despite the fact that from the perspective of either member as an individual, both gambles yield an expected utility of $\frac{1}{2}$ and are therefore equally preferred. As another example, a 2-person group could regard the pure group act $(\frac{1}{2}, x)$ to be strictly preferred to a 50-50 gamble between the pure group acts $(0, x)$ and $(1, x)$ on the grounds that the group is averse to risks that it takes with respect to the expected utility of member 1, even though, speaking only for herself, member 1 would be indifferent between these two choices. The key additional assumption that we impose on group preferences in order to rule out such examples is an independence condition that is an extension of Fishburn's (1965) concept of "mutual independence in the utility sense," which suffices to yield a group utility function that is a weighted sum of its members' utility functions.

Formally, we proceed as follows. Let $\mathcal{F} = \{F^1, \dots, F^N\}$ denote a finite set of pure group acts and let $(\mathcal{F}, q) := ((q^1, F^1), \dots, (q^N, F^N))$ denote a "finite gamble" in which the pure group act F^n is chosen with objective probability q^n , where $q^1 + \dots + q^N = 1$.

Assumption 4: The group has expected-utility preferences among finite gambles, i.e., there exists a von Neumann-Morgenstern utility function φ on the set of pure group acts with respect to which one finite gamble is weakly preferred to another if and only if it has greater or equal expected utility. In particular, the expected utility of $(\mathcal{F}, \mathbf{q})$ is $q^1\varphi(\mathbf{F}^1) + \dots + q^N\varphi(\mathbf{F}^N)$.

Note that each outcome of the finite gamble is a vector of expected utilities assigned to different members, i.e., \mathbf{F}^n stands for (F_1^n, \dots, F_M^n) where F_m^n is the expected utility assigned to member m . In the construction of such a gamble, it might be the case that member m receives the same expected utility in two or more outcomes. In such a case we can compute a marginal probability for the assignment of a given expected utility to member m . For example, let F_m^n and $F_m^{n'}$ denote the expected utilities assigned to member m by two pure group acts \mathbf{F}^n and $\mathbf{F}^{n'}$ that are among the outcomes of a finite gamble $(\mathcal{F}, \mathbf{q})$. If $F_m^n = F_m^{n'} = u$ and $F_m^{n''} \neq u$ for $n'' \neq n, n'$, then the marginal probability of expected utility u for member m is $q^n + q^{n'}$. In these terms, the key additional assumption is that only the marginal distributions of expected utility are relevant to the group's preferences among gambles on behalf of its members, and the group's preferences must be consistent with those of any one member in choices affecting only that member's marginal distribution of expected utility.

Assumption 5: If two finite gambles yield identical marginal probabilities for the expected utilities of all members other than m , then the group is indifferent between the gambles if and only if member m is indifferent between them.

This assumption has an individual-sovereignty aspect, namely that each member has sovereignty over group choices in which only her own marginal distribution of expected utility is at issue, and in such choices, because the individual reduces compound lotteries in her own personal evaluation of any objective gamble, the group is required to do likewise. This is weaker than the usual strict Pareto condition insofar as it applies only

when the expected utilities of the other $M-1$ group members have the same marginal distributions under the two gambles, not merely when the other members have the same directions of preference between the two gambles.

Theorem 2: Given Assumptions 1, 4, and 5, the group has state-dependent expected-utility preferences among acts, and there exist positive weights $\{\alpha_1, \dots, \alpha_M\}$ summing to 1 with respect to which the expected utility function of the group can be expressed as a weighted sum of the expected utility functions of the individual members, i.e., $\nu_G = \sum_m \alpha_m \nu_m$.

Proof: As a special case of Assumption 5, the group must be indifferent between any two finite gambles that yield, for each member, exactly the same marginal distribution of expected utilities. This means that the group's preferences among finite gambles satisfy the assumption of mutual independence in the utility sense introduced by Fishburn (1965). From Fishburn's theorem 4, it follows that the von Neumann-Morgenstern utility function φ that represents the group's preferences among finite gambles has the additive representation:

$$\varphi(\mathbf{F}) = \varphi_1(F_1) + \dots + \varphi_m(F_m) + \dots + \varphi_M(F_M) \quad (5)$$

for some set of functions $\{\varphi_m\}$. By applying Assumption 5 again to choices among finite gambles in which only member m 's marginal distribution of expected utility is varied, it follows that $\varphi_m(\mathbf{f} \square \nu_m)$ is a von Neumann-Morgenstern utility function that represents member m 's preferences among acts, which means that (up to the addition of an arbitrary constant) it is an increasing linear function of $\mathbf{f} \square \nu_m$. We can therefore write $\varphi_m(\mathbf{f} \square \nu_m) = \alpha_m (\mathbf{f} \square \nu_m)$ for some positive α_m , and without loss of generality we can scale the α 's so that $\sum_m \alpha_m = 1$. Hence, a pure group act \mathbf{F} that corresponds to an act \mathbf{f} in the original decision problem is evaluated by the group according to the utility function $\sum_m \alpha_m (\mathbf{f}_m \square \nu_m)$, which reduces to $\mathbf{f} \square \nu_G$ where $\nu_G = \sum_m \alpha_m \nu_m$. ■

Hence we arrive at the same representation of group preferences as in Theorem 1, which has the same special cases as those mentioned at the end of the preceding section, but we

arrive there via a different route. The key difference is that in Theorem 2 the group explicitly considers risks and tradeoffs among hypothetical experiences of its different members. In the latter setup, the group is assumed a priori to have expected-utility preferences among gambles whose outcomes are vectors of expected utilities assigned to its members, but the fact that the group effectively has state-dependent subjective-expected-utility preferences among the original acts emerges as a conclusion when the independence condition of Assumption 5 is also applied.

6. Information Needed for Implementation

Displaying the result of Theorem 2 with the members' utilities in (3) yields

$$U_G(f) = \sum_m \alpha_m U_m(f) = \sum_m \sum_s \alpha_m p_m(s) u_m(f(s)). \quad (6)$$

To determine the group utility function (6), it is necessary to assess the members' subjective probability distributions $\{p_m\}$, their utility functions $\{u_m\}$, and the weights $\{\alpha_m\}$ that their individual expected utilities should be given. To compute their personal probabilities and utilities, individual members can use the theory and procedures of Bayesian decision analysis (Savage, 1954; Pratt et al., 1964; Raiffa, 1968). The responsibility for these tasks rests entirely with the individual group members; it is not a group exercise. However, as mentioned earlier in this paper, individual members may wish to take into account the facts, information, reasoning, and judgments that were communicated by other members of the decision-making group.

The group as a whole must determine the weights $\{\alpha_m\}$, assuming that they accept the conditions of Assumptions 4 and 5 concerning the properties of the group's risk preferences. In principle, the weights are determined by two issues: the members' own relative strengths of preference between consequences and the relative importance (i.e. power) of the members within the group. In practice it might be more reasonable to consider these issues separately in constructing the preferences of the group, rather than asking the group to make holistic judgments. For example, suppose that both members of a two-member group feel that the range of utility for member 1 from $U_1 = 0$ to $U_1 = 1$ (i.e., the difference between worst-case and best-case outcomes) is twice as significant to him as the range of utility for member 2 from $U_2 = 0$ to $U_2 = 1$ is for her. This would

mean that the ratio of α_1 to α_2 should be 2-to-1 based only on an interpersonal comparison of relative strengths of preference. However, suppose the group also feels that in relative terms member 2's opinion is three times as important as the opinion of member 1 for the decision being faced. It could be, for example, that member 2 owns 75% of their joint company. In this situation, the ratio of α_1 to α_2 should be 1-to-3 based only on the relative importance of the individuals. In order to obtain the weights that should be used to construct a group utility function $U_G(\mathbf{f}) = \alpha_1 U_1(\mathbf{f}) + \alpha_2 U_2(\mathbf{f})$ that takes both issues into account, the relative widths of the personal utility scales and the relative importance of the members could be multiplied to get the overall ratio appropriate for the utility weights. In this case, the result would be a ratio of (2×1) to (1×3) , so, with the normalization $\alpha_1 + \alpha_2 = 1$, the appropriate weights would be $\alpha_1 = 0.4$ and $\alpha_2 = 0.6$.

Conceptually, it is easier to specify the relative importance of the group members than to make interpersonal comparisons of the group members' utilities. In many cases, the group may specify that all members' opinions are equally important, using the spirit of "one person, one vote" or that the group is a team and team members should be treated equally. In fact, without ever explicitly considering the relative importance issue, the group may de facto assume that the relative importance of all group members is the same. When it is explicitly stated or implicitly assumed that the opinions of all group members are equally important, the implication for the α -weights is that they should depend only on the interpersonal comparisons of preferences.

In some cases, such as the one in which the group members own different percentages of the organization making a decision, their relative importance may be specified to be proportional to their ownership percent. For other cases, the relative importance may be specified for all decisions made by a standing group. For example, a co-op might acknowledge loyalty or a work organization might respect experience, so the relative importance might be proportional to the years in the organization or to the years in the organization with a leveling off at (say) 10 years. In specifying the relative importance of group members, the group has a decision, but this decision is likely to be less complex than many subsequent decisions which the group must face about actions it should take.

Both the logical foundations and complexity of making interpersonal comparisons of utilities have a long and contentious history (Harsanyi, 1955; Sen, 1970). However the fact remains that the interpersonal comparison of utilities is essential to many group decisions, even those without the support of any analysis, and they are essential to the analytical results of this paper. Here we outline how one could prescriptively address this issue in the group expected utility framework.

When making interpersonal comparisons of utility, it is useful to first try to rank the relative significance of the ranges of utilities from 0 to 1 for different group members. The group must compare, for example, the significance of the change in utility from $U_1 = 0$ to $U_1 = 1$ for member 1 to the significance of the change in utility from $U_2 = 0$ to $U_2 = 1$ for member 2. The group may conclude that one is greater than the other or that they are both the same. If such a ranking can be agreed upon, the relative ratings of interpersonal comparison of utilities can often closely be specified for groups of more than five or six members. In groups with just a few members, it may be feasible to pursue the interpersonal comparison of utilities more directly. For example, if a three-member group feels that the change in utility from $U_1 = 0$ to $U_1 = 1$ for member 1 is more significant than the change in utility from $U_2 = 0$ to $U_2 = 1$ for member 2, then the group can try to determine the level of U_1 , call it U_1' , such that the change in utility from $U_1 = 0$ to $U_1 = U_1'$ for member 1 is equally significant to the change in utility from $U_2 = 0$ to $U_2 = 1$ for member 2.

In practice, the weights $\{\alpha_m\}$ may sometimes be chosen without the benefit of clearly considering, or perhaps even recognizing the relevance of, the members' interpersonal comparison of preferences and relative importance. Indeed, for many decisions such as those faced by a board of directors or university department faculty, the spirit may be that all group members are equal, so quite simply they would assume that $\alpha_m \equiv 1/M$. However, it is essential to recognize that the issues of interpersonal comparison of preferences and the relative importance (i.e. power) of members are an inherent part of the complexity of group decisions. The group, or individuals in the group, can avoid

explicitly thinking about these issues, but by doing so just renders this critical part of the group's decision to the subconscious.

7. Summary

Over the last 40 years, there have been many attempts to extend the Bayesian decision analysis framework to group decisions. Many of these efforts have searched for a logically sound method to combine the individual group members' probability functions into a compromise group probability function and to combine the individual group members' utility functions into a compromise group utility function. Raiffa (1968) illustrates specific cases for both of these combinations that do not lead to satisfactory results. In different circumstances, Hylland and Zeckhauser (1979), Seidenfeld al. (1989), and Mongin (1995) each prove that there is no way, consistent with Pareto optimality, to construct group probability functions or a group utility function from the individual members' corresponding functions that can be used in a group Bayesian decision model logically equivalent to the individual Bayesian decision model.

This paper takes a different approach to derive a group decision model within the general Bayesian framework. It assumes that each member of the decision-making group has state-independent subjective expected utility preferences over the original set of acts, but it does not assume a priori that the group also has subjective-expected-utility preferences, state-independent or otherwise, over the same set of acts. Rather, the group is assumed to have expected utility preferences with respect to gambles whose outcomes are specified in terms of the vectors of expected utilities that they induce for the group members. If the group preferences also satisfy an independence condition stating that only the marginal distributions of the members' utilities are relevant and that any one member has sovereignty over group choices that affect only herself, then it follows that the group effectively has state-dependent expected utility preferences over the original set of acts, which are represented by a weighted sum of the state-independent expected-utility functions of the individual members. The values of the weights depend on both the members' relative strengths of preference among the alternatives and the relative importance of the members in the group. If the group members are in agreement on these

parameters, then appropriate values for the weights can, in principle, be determined by separate consideration of these issues.

Acknowledgment. Helpful comments on various drafts were provided by David Bell, Jay Kadane, Howard Raiffa, and Richard Zeckhauser.

References

- Arrow, K. (1951) (2nd ed. 1963). *Social Choice and Individual Values*. Wiley, New York.
- Anscombe, F. and R. Aumann (1963). A definition of subjective probability. *Annals of Mathematical Statistics*, **34**, 199–205.
- Aumann, R. (1976). Agreeing to disagree. *Annals of Statistics*, **4**, 1236-1239.
- Border, K. (1985). More on Harsanyi's utilitarian cardinal welfare theorem. *Social Choice and Welfare*, **1**, 279-281.
- Chambers, C. and Hayashi, T. (2006). Preference aggregation under uncertainty: Savage vs. Pareto. *Games and Economic Behavior* **54**, 430-440.
- Clemen, R., and Winkler, R. (1999). Combining probability distributions from experts in risk analysis. *Risk Analysis*, **19**, 187-203.
- DeGroot, M. (1970). *Optimal Statistical Decisions*. Wiley, New York
- Fishburn, P. (1965). Independence in utility theory with whole product sets. *Operations Research*, **13**, 28-45.
- Genest, C., and Zidek, J. (1986) Combining probability distributions: a critique and annotated bibliography. *Statistical Science*, **1**, 114-148
- Gilboa, I., Samet, D., and Schmeidler, D. (2004). Utilitarian aggregation of beliefs and tastes. *Journal of Political Economy*, **112**, 932-938.
- Harsanyi, J. (1955). Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of Political Economy*, **63**, 309-321.
- Hylland, A., and Zeckhauser, R. (1979). The impossibility of Bayesian group decision making with separate aggregation of the beliefs and values. *Econometrica*, **47**, 1321-1336.

- Karni, E., D. Schmeidler, and K. Vind (1983). On state-dependent preferences and subjective probabilities. *Econometrica*, **51**, 1021-1032
- Karni, E. (1985). *Decision Making Under Uncertainty: The Case of State Dependent Preferences*. Harvard University Press, Cambridge
- Karni, E. (2007). Foundations of Bayesian Theory. *Journal of Economic Theory*, **132**, 167-188
- Keeney, R., and Raiffa, H. (1976). *Decisions with Multiple Objectives*. Wiley, New York. Reprinted in 1993, Cambridge University Press, New York.
- Mongin, P., (1995). Consistent Bayesian aggregation. *Journal of Economic Theory*, **66**, 313-351.
- Mongin, P. (1998). The paradox of the Bayesian experts and state-dependent utility theory. *Journal of Mathematical Economics*, **29**, 331-361.
- O'Hagan, A., Buck, C., Daneshkhan, A., Eiser, J., Garthwaite, P., Jenkinson, D., Oakley, J., and Rakow, T. (2006). *Uncertain Judgments: Eliciting Experts' Probabilities*, John Wiley and Sons, Chichester, England.
- Pratt, J., Raiffa, H., and Schlaifer, R. (1964). The foundations of decision under uncertainty: an elementary exposition. *Journal of the American Statistical Association*, **59**, 353-375.
- Raiffa, H. (1968). *Decision Analysis*. Addison-Wesley, Reading, Massachusetts.
- Ramsey, F. (1926). Truth and probability. Reprinted in 1964 in *Studies in Subjective Probability*, Kyberg, H.E. and Smokler H.E., editors, Wiley, New York.
- Savage, L. (1954). *The Foundations of Statistics*, Wiley, New York.
- Schervish, M., T. Seidenfeld, and J. Kadane (1990). State-dependent utilities. *Journal of the American Statistical Association*, **85**, 840-847
- Sen, A. (1970). *Collective Choice and Social Welfare*. Holden-Day, San Francisco.
- Seidenfeld, T., Kadane, J., and Schervish, M. (1989). On the shared preferences of two Bayesian decision makers. *Journal of Philosophy*, **86**, 225-244.
- Von Neumann, J. and Morgenstern, O. (1947). *Theory of Games and Economic Behavior*, 2nd edition, Princeton University Press, Princeton