On the Shared Preferences of Two Bayesian Decision Makers

Teddy Seidenfeld; Joseph B. Kadane; Mark J. Schervish

The Journal of Philosophy, Vol. 86, No. 5. (May, 1989), pp. 225-244.

Stable URL:
http://links.jstor.org/sici?sici=0022-362X%28198905%2986%3A5%3C225%3AOTSPOT%3E2.0.CO%3B2-Y

The Journal of Philosophy is currently published by Journal of Philosophy, Inc..

THE JOURNAL OF PHILOSOPHY
VOLUME LXXXVI, NO. 5, MAY 1989

ON THE SHARED PREFERENCES OF TWO BAYESIAN DECISION MAKERS*

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? If rational parties to a negotiation can agree on collective actions merely by considering mutual gains, is it not possible to find a consensus Bayes model for their choices? 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?

Whatever else is to be required of a compromise, we suppose that a consensus Bayes model for the group preserves those strict preferences which the individuals already share. That is, we impose a *weak Pareto condition* on compromises. Whenever all parties to a decision have a common strict preference for one option over another, then any proposed Bayesian group model for their choice—any "neutral" position—must reflect this preference and assign higher expected utility to the Pareto dominating option.

Of course, the probabilities and utilities of any one of the agents satisfies this weak Pareto condition. That is, each agent on her own meets this condition—whatever strict preferences they all have, each has. But, it is hardly a compromise always to make the group decide all questions based on the preferences of a single member. We call

* We would like to thank John Broome, Robyn Dawes, Jay Goodman, Mark Kamlet, and Isaac Levi. This work was supported by the Buhl Foundation, ONR Contracts N00014-85-K-0539 and N00014-88-K-0013, and NSF Grant DMS-8805676.

0022-362X/89/8605/225–244

such Pareto solutions *autocratic*. What (nonautocratic) Bayesian compromises are there for the group decisions?

Imagine the dilemma that arises when two Bayesian agents agree on what to do by appeal to their unanimous preferences, but they find no neutral Bayes position (other than their own separate views) to endorse the rationality of their choices. It is our primary purpose in this essay to underscore the ubiquitous nature of this dilemma. Our central result, then, is this. When two Bayesian decision makers differ both in their beliefs (probabilities) and values (utilities), and subject to an assumption that they agree on the preference ranking of some two "constant" acts, the only candidates for a Bayes compromise are the two autocratic solutions. That is, there is no room for a Bayesian compromise. We develop this argument in section I.

In section II we contrast our results with several related theorems in social welfare theory. The question of finding a neutral Bayes model for the group's choices is seen at once to be a variety of Kenneth Arrow's[1] problem for social welfare rules. The Bayesian agents have individual preferences over social acts which are to be amalgamated into a single, Bayesian group preference over social acts. In order to dodge that "impossibility" theorem, some concession to Arrow's result is required.

The alternatives that have been investigated in connection with Arrow's problem follow several directions. One approach is to restrict the domain of applicability of a welfare rule to communities where the individual preferences are not too discrepant, e.g., where the individual preferences conform to Duncan Black's[2] "single-peakedness" condition. (This violates the "unrestricted domain" condition of Arrow's argument.) Our negative finding is unaffected by this consideration, since we show that the dilemma arises for all pairs of Bayesian agents who have even the slightest differences in

[1] *Social Choice and Individual Values* (New York: Wiley, 1951). Stated briefly, Arrow's *impossibility theorem* shows that it is inconsistent to posit a social choice rule for amalgamation of individual preferences into a group preference, subject to four conditions: (1) the rule has unrestricted domain—it applies with all varieties of individual preferences and all sets of feasible acts; (2) the rule satisfies the weak Pareto condition—if every one (strictly) prefers option A to option B, then the social ordering makes A strictly better than B; (3) the rule precludes *dictator* solutions—it cannot be that the social ordering is determined by one individual's preferences regardless the preferences of all others; (4) the rule conforms to an *independence of irrelevant alternatives* condition—the social ordering of a set of feasible options depends solely on the individuals' orderings for these options. (The fourth condition combines a prohibition on interpersonal utility comparisons with an assumption that the social choice rule induces an ordering of acts by revealed preferences.)

[2] "The decision of a committee using a special majority," *Econometrica*, XVI (1948): 245–261.

both their beliefs and their values. There is no gain made here by trying to bracket cases where the agents' probabilities or utilities reflect large, rather than small, discrepancies.

A second approach to avoiding Arrow's impossibility theorem is to add structure to the representation of individual preferences, to add interpersonal utility comparisons not allowed by Arrow's multipartite "independence of irrelevant alternatives" condition. The excellent paper of Kevin Roberts[3] summarizes how different classes of social welfare rules can be achieved by introducing alternative versions of positive interpersonal utility comparisons. Our negative finding about Bayesian social welfare rules applies, however, whether or not interpersonal utility comparisons are made. Thus, another familiar way around Arrow's result does not work when the problem is finding group compromises that are Bayes.

A third approach is to relax the "ordering" requirement, to liberalize the condition that the social choice rule induces a complete ordering of group options where any two social acts may be compared by the compromise social preference relation. In his seminal book on Bayesian decision theory and statistics, Leonard Savage[4] defends the minimax-regret rule for group deliberation by suggesting that the "ordering" postulate (*P*1) does not apply to group preference, though it does apply to individual preferences.

In this connection, Isaac Levi[5] offers an intriguing account, we think, why the norms on rational choice should be uniform between groups and individuals. The key assumption for a group decision problem is that there is only one agent, the (cooperative) group. Otherwise, there is not one decision problem but, instead, there are the several (noncooperative) deliberations of the separate individuals who, for prudential reasons, take an interest in each other's actions.

In section II.2 we explain how our result impacts on Levi's proposal for achieving a unified decision theory, unified across individual and group decisions. Levi's theory relaxes the "ordering" postulate of Bayesian expected utility. A rational agent (either an individual or a group) need not have a complete preference relation for comparing every two options. The explanation for this departure from Bayesian theory is his view that a rational agent (either an individual or a group) may experience *unresolved* conflicts.

For an individual, the conflicted preferences arise from uncertain beliefs (sets of probabilities) or from indeterminate values (sets of

[3] "Possibility Theorems with Interpersonally Comparable Welfare Levels," *Review of Economic Studies,* XLVII (1980): 409–420.

[4] *The Foundations of Statistics* (New York: Wiley, 1954), §13.5.

[5] "Conflict and Social Agency," this JOURNAL, LXXIX, 5 (May 1982): 231–247.

utilities). For a group, the conflicts among the individuals' beliefs and values generate the same uncertainties and indeterminacies when the agent is the cooperative group.

In Levi's theory, however, the compromise position on rational choice is achieved (at the expense of "ordering" for the preference relation) by constructing two *independent* "neutral" positions, one for conflicted beliefs (a set of probabilities) and one for conflicted values (a set of utilities). An incomplete preference relation is formed using inequalities in expected utilities with all pairs from these two sets. Based on our negative result, we show this rule leads to violations of the weak Pareto condition. That is, Levi's theory allows a group to choose a weak Pareto dominated option.

Nonetheless, we are in agreement with Levi on the desirability of relaxing the ordering postulate while preserving a respect for expected utility. Hence, we explore a theory in which preference need not be a complete relation, but also a theory in which preference does not compromise conflicts in beliefs and in values independently.

I. BAYESIAN COMPROMISES BETWEEN TWO BAYESIANS

Here we discuss in detail the case of a group composed of just two Bayesian decision makers. That is, we respond to the question posed in the opening sentences of this essay when only two Bayesian "experts" are involved and, for each, preference is a complete relation. We assume that each agent can compare every two acts, either by a (transitive, antisymmetric) relation of strict preference, or else by an equivalence relation of indifference.

The group decision involves amalgamating the preferences of two Bayesians, whom we shall call Dick and Jane. According to Bayesian decision theory, for simple problems, an agent chooses from among a set of feasible (state independent[6]) acts according to the principle of maximizing subjective expected utility.

[6] We assume acts and states are probabilistically independent. Later on where we use horse lotteries for acts, as explained in fn. 9, we assume also that states are value-neutral with respect to the lottery prizes.

Interesting discussions of the measurement problem for expected utility theory without either of these assumptions is found in J. Drèze, "Decision Theory with Moral Hazard and State–Dependent Preference," paper #8545, *C.O.R.E.* (1985), Universite Catholique de Louvain, Voie du Roman Pays, 34, B-1345, Louvain–la-Neuve, Belgium.

For additional, important commentary about the effects of state-dependent utilities on the measurement of probabilities, see E. Karni, D. Schmeidler, and K. Vind, "On State Dependent Preferences and Subjective Probabilities," *Econometrica,* LI, 4 (July 1983): 1021–1031; H. Rubin, "A Weak System of Axioms for 'Rational' Behavior and the Non-Separability of Utility from Prior," *Statistics and Decisions,* V (1987): 47–58; and J. B. Kadane and R. L. Winkler, "Separating Probability Elicitation from Utilities," *Journal American Statistical Association,* LXXXIII, 402 (1988): 357–363.

To be more precise, consider the familiar decision matrix:

Decision Matrix: Acts **X** States

	$S1$	$S2$			Sj			Sn
A_1	O_{11}	O_{12}			O_{1j}			O_{1n}
A_2	O_{21}	O_{22}			O_{2j}			O_{2n}
A_m	O_{m1}	O_{m2}			O_{mj}			O_{mn}

The columns denote a partition into (n) states of nature, $S_1, \ldots, S_n$, about which the agent is uncertain. The rows designate feasible acts, $A_1, \ldots, A_m$, whose outcomes in each state are denoted by the O_{ij}: the outcome of act A_i under state S_j. The agent's uncertainties about the states are given by a probability distribution, $P(S_j)$. The agent's values for outcomes are given by a (von Neumann-Morgenstern) utility function, $U(O_{ij})$. Then, according to the principle of maximizing expected utility, act A_1 is (strictly) preferred to act A_2 whenever

$$\sum_j P(S_j) \cdot U(O_{1j}) > \sum_j P(S_j) \cdot U(O_{2j}).$$

Suppose Dick's preferences over such acts are summarized by the pair (P_1, U_1) of his (personal) probability and utility, while Jane's preferences are depicted by the pair (P_2, U_2). What are the alternative Bayesian preference schemes that agree with the (strict) preferences shared by Dick and Jane? That is, for which pairs (P, U) is it the case that:

$$\sum_j P(S_j) \cdot U(O_{1j}) > \sum_j P(S_j) \cdot U(O_{2j})$$

whenever

$$\sum_j P_k(S_j) \cdot U_k(O_{1j}) > \sum_j P_k(S_j) \cdot U_k(O_{2j}) \qquad (k = 1, 2)?$$

In terms of social welfare rules, our question is this. In group choices where all agents receive the same outcomes (the social acts offer the same prospects to each agent), subject to a weak Pareto condition, what are the Bayesian social welfare rules? (Recall, a social choice rule satisfies the weak Pareto condition provided that an option is inadmissible whenever there is another feasible alternative that everyone strictly prefers. If everyone strictly prefers option B to option A, A is inadmissible whenever B is available.)

The answer is surprising. Stated informally, our result in the case of two agents is that no attractive compromises exist. Only autocratic solutions conform to the weak Pareto condition. That is, a Bayesian model for the group's collective decisions must use the beliefs and values of a single agent, thereby ignoring the preferences of everyone else whenever the weak Pareto condition does not apply.[7]

It is interesting to contrast the difference between the weak and strong Pareto conditions for group compromises. The strong Pareto condition for social choice rules requires that option A is socially inadmissible whenever there is a feasible option B, which everyone finds either strictly preferable or indifferent to A, and which someone strictly prefers to A. We find that with the imposition of a strong Pareto condition, there are *no* Bayes social welfare rules.

The proofs of these claims are contained in the following, simple example. The constructions in the example generalize to every pair, like Dick and Jane, that differs in both beliefs and values.

Example 1: Suppose, as before, Dick and Jane are Bayesian agents with preferences summarized by the probability and utility pairs (P_k, U_k), $k = 1, 2$. Assume that they have different beliefs, $P_1 \neq P_2$. That is, there is some event E with $P_1(E) \neq P_2(E)$. For instance, let Dick assign E a (personal) probability .1 while Jane assigns E a (personal) probability .3, i.e., $P_1(E) = .1$ and $P_2(E) = .3$. Also, suppose they have different values, $U_1 \neq U_2$.[8] For simplicity, suppose Dick and Jane

[7] This is not the same as Arrow's concept of a "dictator," since the determination of who is the autocrat may be a function of everyone's preferences, contrary to the requirements for an Arrovian dictator.

[8] This assumption is complicated by controversies over interpersonal utility comparisons. With or without concern for another's preferences, an individual's utility function is defined up to an affine transformation. Regarding this, we suppose that for any two agents there is some pair of v.N–M lotteries, L_* and L^*, which are ranked the same by both—each (strictly) prefers L^* to L_*. (This is mild, as we may introduce two new rewards with this property. We discuss the case of completely opposed preferences in fn. 11.) In case interpersonal utilities are recognized, consider the (perhaps, two) representation(s) of the agents' joint utilities:

$$U_a = \{U_{1a}, U_{2a}\} \text{ and } U_b = \{U_{1b}, U_{2b}\},$$

where

$$U_{1a}(L_*) = U_{2b}(L_*) = 0 \text{ and } U_{1a}(L^*) = U_{2b}(L^*) = 1.$$

Then, the assumption that Dick and Jane have different values is expressed by saying that there is a lottery, L, with $U_{1a}(L) \neq U_{2b}(L)$. That is, their (separate) utility functions do not coincide under affine transformations. Since we are concerned

Table 1: "Horse lotteries" used in fixing the upper and lower probabilities and utilities.

	E	$\neg E$
A_1	r^*	r_*
$A_{2\epsilon}$	$(.9 + \epsilon)(r_*) + (.1 - \epsilon)(r^*)$	$(.9 + \epsilon)(r_*) + (.1 - \epsilon)(r^*)$
$A_{3\epsilon}$	$(.7 - \epsilon)(r_*) + (.3 + \epsilon)(r^*)$	$(.7 - \epsilon)(r_*) + (.3 + \epsilon)(r^*)$
A_4	r	r
$A_{5\epsilon}$	$(.6 - \epsilon)(r_*) + (.4 + \epsilon)(r^*)$	$(.6 - \epsilon)(r_*) + (.4 + \epsilon)(r^*)$

have different (cardinal) utilities, in the following sense. They agree on the ranking of two particular rewards: each prefers r^* to r_*, though they differ in their valuation of a third reward r. Let r be a reward with $U_1(r) = .1$ and $U_2(r) = .4$, while there exist rewards r_* and r^* with $U_1(r_*) = U_2(r_*) = 0$ and $U_1(r^*) = U_2(r^*) = 1$.

Next, we consider pairs of options where Dick and Jane hold common preferences. Examine the acts defined in table 1, above, on the binary partition formed by the event E.

These acts are "horse lotteries," in the language of F. J. Anscombe and R. J. Aumann.[9] We distinguish a von Neumann–Morgenstern (v.N–M) lottery from a horse lottery. A v.N–M lottery, L, is a specified probability distribution over a set of rewards. For example, a v.N–M lottery that yields the reward r_* with probability .6 and the reward r^* with probability .4 is denoted by: $L = .4r_* + .6r^*$. With horse lotteries as acts, the outcome O_{ij} of act A_1 in state s_j (see the decision matrix on page 229) is a v.N–M lottery L_{ij}.

---

with the set of Bayes compromises that preserve unanimous (strict) preferences, we may use the pair (U_{1a}, U_{2b}) to represent the agents' individual utilities while also respecting interpersonal utility comparisons. For convenience, abbreviate U_{1a} by U_1 and U_{2b} by U_2.

[9] "A definition of subjective probability," *Annals of Mathematical Statistics*, XXXIV (1963): 199–205. Their theory of "horse lotteries" is summarized by four axioms on preference. Informally, these four axioms require that: (A1) preference is a weak order; (A2) preference satisfies the "independence" condition (related to Savage's "sure-thing" postulate *P*2); (A3) preference obeys an Archimedean condition—to insure utilities are real-valued; and (A4) preference for v.N–M lotteries is state-independent.

Specifically, the final axiom requires this. Let H_1 be the "constant" horse lottery that awards the same v.N–M lottery L_1 in each state and let H_2 be the "constant" horse lottery that awards L_2 in each state. Let horse lotteries $H_{1'}$ and $H_{2'}$ differ solely in that, for some state s, $H_{1'}(s) = L_1$ and $H_{2'}(s) = L_2$. $H_{1'}$ and $H_{2'}$ have the same outcomes in every other state. Then axiom A4 says: H_1 is preferred to H_2 if and only if $H_{1'}$ is preferred to $H_{2'}$.

For example, the v.N–M lottery $(.9 + \epsilon)(r_*) + (.1 - \epsilon)(r^*)$ under event E for act $A_{2\epsilon}$ corresponds to the outcome: receive reward r_* with probability $(.9 + \epsilon)$ and reward r^* with probability $(.1 - \epsilon)$. When "horse$_j$" wins the hypothetical race, i.e., when state s_j obtains, the act A_i pays out the v.N–M lottery L_{ij}. Thus, horse lotteries are functions from states to v.N–M lotteries.

Since which state obtains (which "horse" wins) is uncertain, and how rewards are to be valued is not stipulated in advance, "horse lotteries" accommodate both uncertainty over states and (cardinal) utility over rewards.

First, as Dick and Jane agree that

$$.1 \leq P_k(E) \leq .3 \qquad (k = 1, 2)$$

they agree that act A_1 is preferred to each act $A_{2\epsilon}$, and they agree that each act $A_{3\epsilon}$ is preferred to A_1 $(.1 \geq \epsilon > 0)$. See table 1.

Likewise, they agree that

$$.1 \leq U_k(r) \leq .4 \qquad (k = 1, 2).$$

Then, also they are unanimous in their preference for A_4 over each act $A_{2\epsilon}$, while each act $A_{5\epsilon}$ is preferred to A_4. (All these preferences are "strict.")

In figure 1 we see the set of pairs of probabilities and utilities agreeing with these shared preferences. That is, figure 1 is based on the (strict) preferences involving the upper and lower probabilities of event E and the upper and lower utilities of reward r. These bounds for $P(E)$ and $U(r)$ box the family of Bayesian compromises between Dick and Jane with respect to their shared agreements for these (strict) preferences. Subject to the weak Pareto condition, it shows that the set of "neutral" Bayes' models (with respect to *all* their choices—not just for these few comparisons) is some subset of the cross product of weighted averages of their probabilities and weighted averages of their utilities.

Next, consider the set of pairs of acts defined in table 2, over which Dick and Jane also hold common preferences.[10]

For each value, $0 < \epsilon \leq .2$, Dick and Jane agree in their (strict) preference for $A_{7\epsilon}$ over $A_{6\epsilon}$. The set of Bayesian compromises for these preferences is *not* connected, however. See figure 2.

For $0 < \epsilon < .015$ the set of probability/utility pairs $(P(E), U(r))$ for which $A_{7\epsilon}$ has greater expected utility than $A_{6\epsilon}$, is bounded by a hyperbola centered at (.2, .25), which satisfies:

$$[P(E) - .2][U(r) - .25] = .015 - \epsilon.$$

[10] The particularly convenient form of these horse lotteries is due to Jay Goodman, formerly of the Statistics Dept. at Carnegie Mellon University.

Figure 1

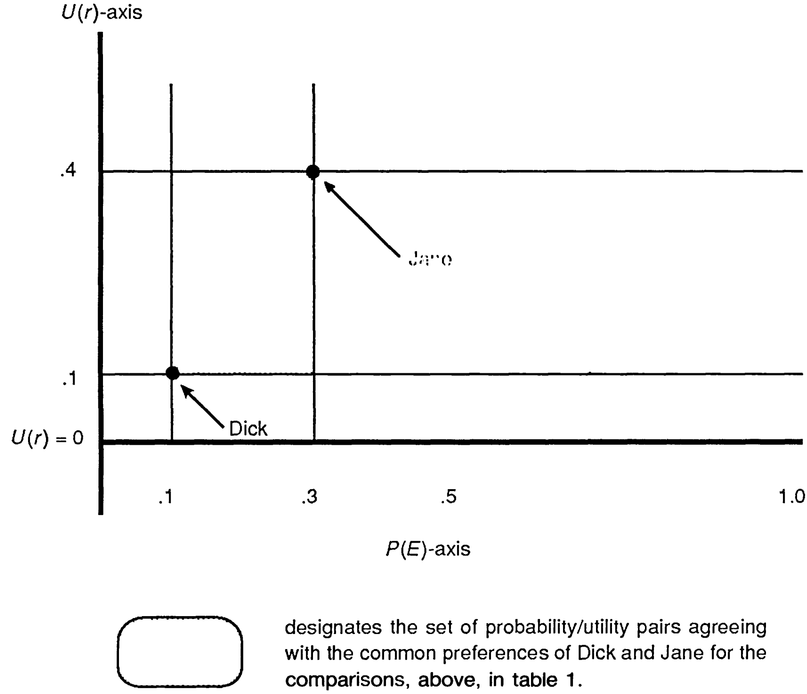

As $\epsilon \Rightarrow 0$, the hyperbola approaches the pair of points corresponding to Dick and Jane's preferences. When $\epsilon = 0$, the hyperbola intersects these two points.

If we superimpose the two figures, we obtain figure 3.

Figure 3 shows that the family of Bayesian agents who agree with these two, even for the few preferences already considered, consists exactly of the two themselves: (P_1, U_1) and (P_2, U_2). That is, since the hyperbolas all have negative slopes at the points interior to the box, as $\epsilon \Rightarrow 0$ the regions of overlap between the two figures collapse onto

Table 2: "Horse lotteries" used in separating the set of compromises between Dick and Jane.

	E	$\neg E$
$A_{6\epsilon}$	$.785(r_*) + .215(r^*)$	$\epsilon(r_*) + .2(r) + (.8 - \epsilon)(r^*)$
$A_{7\epsilon}$	$(.2 - \epsilon)(r_*) + .8(r) + \epsilon(r^*)$	$.165(r_*) + .835(r^*)$

Figure 2: Preferences which separate the family of agreeing probability/utility pairs

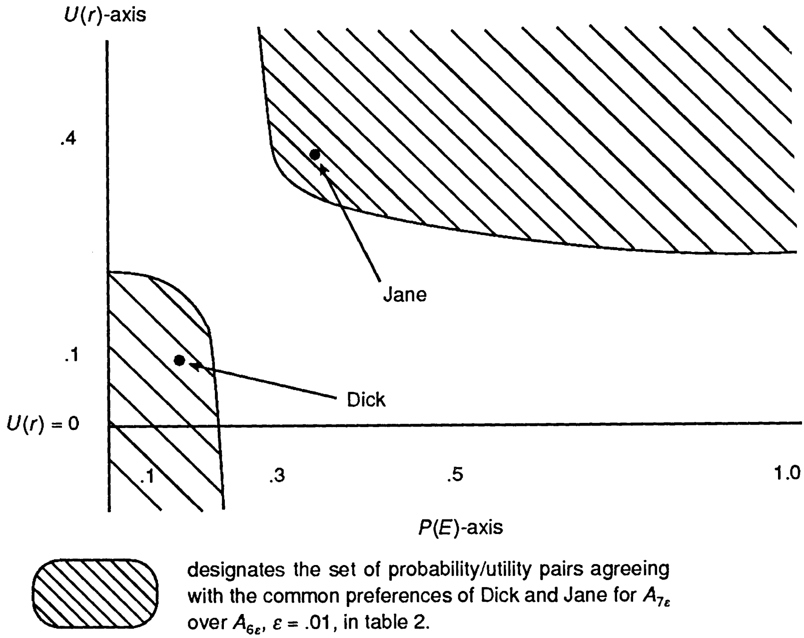

Figure 3: Preferences which separate the family of agreeing probability/utility pairs.

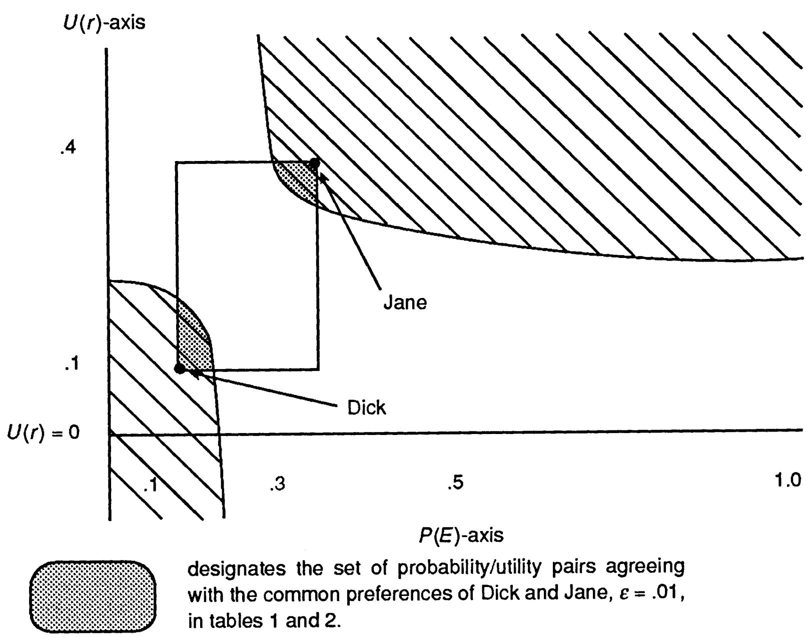

the two corner points: the points corresponding to Dick and Jane. We may express the lesson of the example as follows. The only Bayes models that preserve Dick and Jane's common strict preferences (a weak Pareto condition) are the two autocratic solutions: choose one of Dick and Jane.

Moreover, if we inquire about Bayes models that satisfy a strong Pareto condition ("When everyone is either indifferent or prefers act A_1 to A_2 and someone prefers A_1, the compromise is for the strict preference."), then there are *no* acceptable solutions. This follows since, e.g., Dick is indifferent between the bets B_1 and B_2 though Jane prefers B_1, and Jane is indifferent between the bets B_1 and B_3 though Dick prefers B_3 (table 3). Neither autocrat respects both of these strict preferences.

Table 3: Bets used with the strong Pareto condition.

	E	$\neg E$
B_1	r^*	r_*
B_2	$.9(r_*) + .1(r^*)$	$.9(r_*) + .1(r^*)$
B_3	$.7(r_*) + .3(r^*)$	$.7(r_*) + .3(r^*)$.

The various constructions in this example generalize easily to any pair of agents with different beliefs and values, subject to the common ranking for some pair of v.N–M lotteries, L^* and L_*.[11] (We choose $U_k(L_*) = 0$ and $U_k(L^*) = 1$, $k = 1, 2$.)

[11] This assumption cannot be removed. A result due to J. Kadane ["Opposition of Interest in Subjective Bayesian Theory," *Management Science*, XXXI (1985): 1586–1588, Theorem 1] says that two agents hold opposing strict preferences over all pairs of acts if, and only if, they share a common personal probability for the states and have opposite (cardinal) utilities for the rewards. Then, as a simple corollary, if two agents hold different probabilities and diametrically opposed (cardinal) utilities, there will be some pairs of horse lotteries which, by strict preference, they rank in common.

That is, suppose Susan and John have different personal probabilities (denoted P_S and P_J, with $P_S \neq P_J$) and strictly opposed preferences for all pairs of v.N–M lotteries, the "constant" horse-lottery acts. Thus, whenever John strictly prefers lottery L_1 to lottery L_2 Susan's preference goes the other way, and vice versa. There is no pair of lotteries which they rank order the same. Then, after standardizing their separate utility units with a single pair of v.N–M lotteries (as in fn. 8), where Susan prefers the first and, hence, where John prefers the second of the two lotteries, we have: $U_S = 1 - U_J$.

Nonetheless, as $P_S \neq P_J$, by Kadane's result, there are pairs of horse lotteries which they rank order in the same way. Let E be an event with

$$P_S(E) = p_S < p_J = P_J(E).$$

Theorem 1: Let (P_1, U_1) and (P_2, U_2) be the probability/utility pairs representing two decision makers with different beliefs and preferences over horse lotteries, then:

(i) The set of probability/utility pairs which agree with the strict preferences shared by these two decision makers (weak Pareto) consists exactly of the two pairs themselves—there are autocratic solutions only—no other Bayesian compromises exist; and

(ii) Subject to the strong Pareto condition, the set of probability/utility pairs that agree with these common preferences is empty—there are no Bayes models at all.

Proof: Since $P_1 \neq P_2$ and $U_1 \neq U_2$, there is an event E and a v.N–M lottery L such that $P_1(E) = p_1 < P_2(E) = p_2$, and $U_1(L) = u_1 < U_2(L) = u_2$. Then the agents agree that:

$$p_1 \leq P_k(E) \leq p_2$$

$$u_1 \leq U_k(L) \leq u_2 \qquad (k = 1, 2).$$

(*i*) Table 1 is modified by these bounds, creating a "box of compromises," as in figure 1. The pairs of acts used to "separate" the set of compromises, analogous to those of table 2, are defined by the hyperbolic equations (provided by Jay Goodman):

$$(P(E) - [p_1 + p_2]/2) \cdot (U(E) - [u_1 + u_2]/2) = (p_2 - p_1)(u_2 - u_1)/4 - \epsilon$$

for $0 < \epsilon < (p_2 - p_1)(u_2 - u_1)/4$. Combining these sets of preferences, letting $\epsilon \Rightarrow 0$, we discover that only the original pairs (P_1, U_1) and (P_2, U_2) remain. By holding fixed the binary partition $\{E, \neg E\}$ and varying the lottery L, we see that the group utility function must be the utility for one of the two agents. Then, by varying the partition, since utility is state independent, we see that the group probability also must be taken from the same agent.[12] Hence, only autocratic solutions agree with the weak Pareto condition on strict preferences.

Consider two horse lotteries $H_1 = \{L_{11}, L_{12}\}$ and $H_2 = \{L_{21}, L_{22}\}$, defined on the binary partition $\{E, \neg E\}$, where the L_{ij} $(i, j = 1, 2)$ are v.N–M lotteries. John prefers H_1 to H_2 if, and only if, $cp_J + d > 0$, where
$c = [U_J(L_{11}) + U_J(L_{22}) - U_J(L_{12}) - U_J(L_{21})]$
and $d = [U_J(L_{12}) - U_J(L_{22})]$. Because $U_S = 1 - U_J$, Susan also prefers H_1 to H_2 if, and only if, $cp_S + d < 0$. Hence, with $d < 0 < c$ and $-p_Jc < d < -p_Sc$, John and Susan have common strict preferences for H_1 over H_2.

Other Bayesians will agree with John and Susan in preferring H_1 to H_2 provided they have preferences of the form: (P^*, U_J) or (P_*, U_S) where $P^*(E) > P_J(E)$ and $P_S(E) > P_*(E)$ whenever $P_J(E) > P_S(E)$. That is, there exist nonautocratic, Bayes models for the set of weak Pareto agreements between Susan and John. These models correspond to Bayesian agents with *more extreme* degrees of belief (combined with the respective utility)—models that exaggerate the expected utility differences between the two of them. None of these more extreme Bayes models strikes us as a serious compromise between Susan and John. Thus, even in this exceptional case of diametrically opposed utilities, we cannot locate a viable Bayes compromise for their unanimous preferences.

[12] The proof depends upon the "state independence" of utilities. Without it, i.e., if we require only axioms A1–A3 (see fn. 9) there are a continuum of (weak) Pareto,

(*ii*) Upon adding a strong Pareto requirement for compromises, we may modify the bets from table 3, to read as follows:

Modified Table 3: Bets used with the strong Pareto condition.

	E	$\neg E$
B_1	L^*	L_*
B_2	$(1 - p_1)(L_*) + p_1(L^*)$	$(1 - p_1)(L_*) + p_1(L^*)$
B_3	$(1 - p_2)(L_*) + p_2(L^*)$	$(1 - p_2)(L_*) + p_2(L^*)$.

Then, as before, the first agent is indifferent between B_1 and B_2 though the second prefers B_1, and the second agent is indifferent between B_1 and B_3 though the first prefers B_3. Neither "autocrat" respects both these strict preferences. Hence, under a strong Pareto condition on compromises, there are none. □

II. GROUP DECISION MAKING AND A SEPARATION OF BELIEF AND UTILITY

II.1. Social Welfare Theory: An immediate consequence of theorem 1 is the impossibility of a general, nonautocratic Bayesian social welfare function (subject to the weak Pareto condition). That is, even when interpersonal utility comparisons are admitted, and provided the domain of social acts include the simple varieties discussed in the previous section, there is no interesting Bayesian solution to the social welfare problem.

That corollary to theorem 1 is obvious when the population consists in two agents. For larger communities, of size n, consider populations where the beliefs and preferences of different agents fall into one of two camps, i.e., where the n-agents are clones of two agents.[13] Thus, the requirement that group welfare satisfies the strong Pareto condition (or else a conjunction of "weak Pareto" and "no autocrats") is inconsistent with the rationality postulates of subjective

Bayesian compromises using state-dependent utilities. This follows from two other results: (1) Theorem 13.1, p. 176, of P. Fishburn's *Utility Theory for Decision Making* (New York: Krieger, 1979); and (2) The existence of a (convex) set of preferences, satisfying axioms A1, A2, and A3, each of which extends the strict partial order formed by the (weak) Pareto condition. Result (2) is proven in our unpublished technical report, "A Representation for Preference as a Strict Partial Order in Terms of Sets of Pairs of Probabilities and Utilities," in preparation, Dept. of Statistics, Carnegie Mellon Univ.

[13] This mimics the technique used in the two papers we discuss, by A. Hylland and R. Zeckhauser, "The Impossibility of Bayesian Group Decisions with Separate Aggregation of Beliefs and Values," *Econometrica*, XLVII (1979): 1321–1336, p. 1330, equation 7; and by P. Hammond, "Ex-*ante* and Ex-*post* Welfare Optimality under Uncertainty," *Economica*, XLVIII (1981): 235–250, p. 241, for larger communities of n-many decision makers.

expected utility theory. Moreover, this result is resistant to the standard cures for Arrow's impossibility theorem. Neither considerations of interpersonal utility nor restrictions on the extent of the discrepancies between the beliefs and values of two Bayesians avoids the dilemma.

The impossibility of Bayes solutions to group decisions involving uncertainty is discussed also in two important papers, one by A. Hylland and R. Zeckhauser, and one by P. Hammond.[14] We detail a contrast between theorem 1 and these earlier results in an appendix. Briefly stated, those accounts are limited in two ways that leave theorem 1 unaffected.

(1) The earlier findings establish the impossibility of a general, Bayesian welfare rule subject to the condition that group preference amalgamates probability and utility independently. In other words, they require that the Bayes model for the group has the group probability defined solely in terms of the individuals' probabilities, and has the group utility defined solely in terms of the individuals' utilities. Theorem 1 applies without this limitation.

(2) The earlier findings show that with each (Bayesian) social welfare rule for amalgamation of individual preferences into a group preference (and subject to the independent amalgamation of belief and desire discussed in the previous paragraph), *there is some* profile of individual preferences leading to a failure of the (strong) Pareto condition. Theorem 1, however, shows a failure of the (strong) Pareto condition with each rule and for *every* pair of agents (who differ in both beliefs and desires). Thus, theorem 1 derives a universal statement where the other derives an existential.

II.2. A comparison with Levi's quasi-Bayesian decision theory and a problem of independent compromises of beliefs and desires: In important papers and books, Levi advocates a unified theory of rational decision making under unresolved conflict, unified between individual and group decisions.[15] Both for individuals and groups,

[14] Cf. fn. 13. and see John Broome, "Utilitarianism and Expected Utility," this JOURNAL, LXXXIV, 8 (August 1987): 405–422, who comments on the relationship of Hammond's theorem with utilitartianism. Also, he discusses the problem in two unpublished essays, "Bolker-Jeffrey Decision Theory and Axiomatic Utilitarianism," and "Should Social Preferences Be Consistent?" For a survey of some related issues, see C. Genest and J. Zidek, "Combining Probability Distributions: A Critique and an Annotated Bibliography," with discussion, *Statistical Science,* I (1986): 114–148.

[15] A selection of these works are his "On Indeterminate Probabilities," this JOURNAL, LXXI, 13 (July 18, 1974): 391–418; *The Enterprise of Knowledge* (Cambridge MIT, 1980); "Conflict and Social Agency," *op. cit.;* several entries in *Decisions and Revisions* (New York: Cambridge, 1984), and in *Hard Choices* (New York: Cambridge, 1986).

there is no requirement that an agent's preferences induce an ordering of options—not all acts need be compared by preference.

Levi's decision theory is a liberalization of (Bayesian) expected utility theory. In his theory, an agent's beliefs are represented by a convex set of (personal) probabilities, $\mathcal{P}$, and preferences for outcomes are represented by a convex set of (cardinal) utilities, $\mathcal{U}$. In brief, an option o is an admissible choice from a set of feasible options $\mathcal{O}$, provided that it satisfies a (lexicographically ordered) sequence of maximizations. The first of these is E-admissibility, which requires that an option maximize expected utility for some probability/utility pair (P, U), where $P \in \mathcal{P}$ and $U \in \mathcal{U}$.

Definition: o is *E-admissible* if and only if

$$\exists(P \in \mathcal{P}, U \in \mathcal{U})\forall(o' \in \mathcal{O})\ E_{P,U}(o) \geq E_{P,U}(o').$$

A secondary decision test for narrowing the set of E-admissible options is maximizing a "security" index among those options which are E-admissible. Illustrations of "security" include (a) a vacuous standard—all options have equal security; (b) security indexed by worst ($\inf[\mathcal{U}]$) payoff; and (c) security indexed by least ($\inf[\mathcal{P} \times \mathcal{U}]$) expected utility. In addition to these two, Levi[16] entertains ternary, etc., maximization requirements, reflecting added structure in the agent's system of values.

Because of the first condition, E-admissibility, an admissible option is "Bayes," i.e., maximizes expected utility for some probability/utility pair. Hence, admissibility takes expected utility theory as a special case, when both $\mathcal{P}$ and $\mathcal{U}$ are unit sets and, e.g., all other value considerations are vacuous. Even when other value considerations are vacuous, however, if either $\mathcal{P}$ or $\mathcal{U}$ is not a singleton, admissibility fails to induce an ordering.[17]

A central theme in Levi's account of choice under unresolved conflict is that a "neutral" position among conflicting beliefs and desires is a position that preserves the shared agreements between the rivals, yet introduces no judgments over which there is disagreement. If an agent experiences a value conflict between, e.g., two (cardinal) utilities U_1 and U_2 then, according to Levi's theory, the convex combinations of these two,

$$\mathcal{U} = \{\alpha U_1 + (1 - \alpha)U_2 \colon 0 \leq \alpha \leq 1\},$$

[16] *Hard Choices,* §5.7.

[17] Then admissible choices fail to satisfy A. K. Sen's property γ for choice rules, "Social Choice Theory: A Re-examination," *Econometrica,* XLV (1977): 53–89. Hence, it is not even a *normal* choice rule in Sen's terminology. For details, see Seidenfeld's Discussion of A. P. Dempster "Probability, Evidence, and Judgment," in *Bayesian Statistics 2,* Bernardo, DeGroot, Lindely, and Smith, eds. (Amsterdam: North-Holland, 1985), pp. 127–129.

represents the consensus position of unresolved value conflict from which he makes rational decisions. For example, if John is conflicted between his value assessments of options to acquire art objects—the options rank differently under comparisons of their economic outcomes (U_1) and their aesthetic outcomes (U_2)—nonetheless, by representing his conflicted values with the set $\mathcal{U}$, he may proceed to make rational decisions without first having to resolve the value conflict. The analysis for a consensus position with conflicted beliefs is similar. The "neutral" position for belief is the (convex) set $\mathcal{P}$ of conflicted probabilities.

As we remarked in our introduction, Levi[18] offers arguments that rational deliberation should follow the same standards regardless whether the decision is by an individual or by a cooperative group. The intrapersonal conflicts of values and beliefs for individual decision making should be treated in a like fashion with the parallel interpersonal conflicts for group decisions. Our understanding of his theory is that the "neutral" position for decision under unresolved conflict is found by analyzing desires and beliefs *independently*. That is, the consensus for conflicted preferences over outcomes is a (convex) set of utilities, the consensus for conflicted beliefs is a (convex) set of probabilities, and the independence between them is built into E-admissibility with the appeal to the Cartesian product of these two sets.

If we apply his method to the group decisions faced by Dick and Jane, our example 1, we see that the E-admissible social acts are those which maximize expected utility for some probability/utility pair in the rectangle pictured in figure 1. Then, both acts $A_{6\epsilon}$ and $A_{7\epsilon}$ are E-admissible in a pairwise choice between them. (The expected utility of act $A_{6\epsilon}$ is greater for each probability/utility pair within the unshaded region of the rectangle, depicted in figure 3.) This challenges the claim to the "neutrality" of the E-admissible options, we believe, since E-admissibility fails to respect the shared strict preference for $A_{7\epsilon}$ over $A_{6\epsilon}$ E-admissibility conflicts with the weak Pareto condition.

To emphasize this point, Levi's method of group decision making makes identical the E-admissible options in the feasible sets $\mathcal{O}_\epsilon = \{A_{6\epsilon}, A_{7\epsilon}\}$ for the following three groups of Bayesians. Group$_1$ consists of Dick and Jane. Group$_2$ consists of Tom and Mary; where Tom's preferences are summarized by the probability/utility pair (P_1, U_2) and Mary's by (P_2, U_1). Group$_3$ is composed of all four agents: Tom, Dick, Jane, and Mary. For each of the three groups, both options in the sets $\mathcal{O}_\epsilon$ are E-admissible. Group$_1$ declares a unani-

[18] "Conflict and Social Agency," *op. cit.*

mous strict preference for $A_{7\epsilon}$ over $A_{6\epsilon}$, however. Group_2 declares the opposite. And the members of Group_3, of course, find no common preferences over these choices.

Levi's E-admissibility is a normative theory of consensus in *reasons* for preference, the rational causes of decisions. It treats beliefs and values for outcomes as the independent springs for our rational actions. But we discover that this account of consensus in reasons is at odds with the conservation of shared (strict) preferences among options. Then Pareto agreements are taken to be superficial unless they are supported from below by consensus in reasons. Levi's theory makes it a serious question between two exclusive strategies facing group_1: either (merely) appeal to the existing agreements on what to do—choose $A_{7\epsilon}$ without giving reasons or, instead, agree first on the "neutral" reasons for the group's choices and let that consensus dictate admissibility—in which case both options are admissible. How is this higher-order decision problem resolved?

At the expense of denying an independent consensus for conflicts of beliefs and values, we avoid this dilemma. In "Decisions without Ordering,"[19] we propose a theory of choice in which preference over pairs of horse lotteries is a strict partial order. In our theory, preference is represented by a set $\mathcal{S}$ of pairs of probabilities and utilities, as follows. Option o_1 is preferred to option o_2 according to the partial order if and only if the expected utility of o_1 is greater than that of o_2 for each probability/utility pair in $\mathcal{S}$.[20] When the strict partial order characterizes a consensus among different (Bayesian) agents, we take the elements of $\mathcal{S}$ to be all the probability/utility pairs that agree with the shared (weak Pareto) preferences of those agents. Trivially, the weak Pareto condition is respected with this consensus set. Each element of $\mathcal{S}$, each potential compromise among the agents, preserves their shared agreements. Levi's theory does not satisfy this condition for potential compromises.

The sets $\mathcal{S}$ are "conditionally convex": If pairs (P_1, U_1) and (P_1, U_2) belong to $\mathcal{S}$, so too do all pairs of the form (P_1, U_3), where $U_3 = \beta U_1 + (1 - \beta)U_2$ $(0 \leq \beta \leq 1)$. Likewise $\mathcal{S}$ is closed under mixtures of pairs that share a common utility. $\mathcal{S}$ need not be a (convex) cross product of a set of probabilities and a set of utilities, however, as required in Levi's theory. Nor need $\mathcal{S}$ be convex, or even

[19] In W. Sieg, ed., *Acting and Reflecting* (Dordrecht: Reidel, in press). Also available as Tech. Report #391 (1987), Dept. of Statistics, Carnegie Mellon University.

[20] If the weak Pareto condition fails to specify action, i.e., where there are conflicted recommendations based on the expectation inequalities taken from $\mathcal{S}$, then one possibility is Levi's proposal to deploy second-tier considerations in order to choose among those options admissible at the first tier. For example, one might use "security" to compromise choice among the "Pareto" admissible alternatives.

connected—as illustrated in figure 3. According to this account, the shared preferences for the three groups (above) are represented by three different consensus sets:

$$\mathcal{S}_1 = \{(P_1, U_1), (P_2, U_2)\}$$

$$\mathcal{S}_2 = \{(P_1, U_2), (P_2, U_1)\}$$

$$\mathcal{S}_3 = \{(P, U): P = \alpha P_1 + (1 - \alpha)P_2,\ U = \beta U_1 + (1 - \beta)U_2,\ 0 \leq \alpha, \beta \leq 1\}.$$

Only the third group has (partially ordered) strict preferences agreeing with Levi's E-admissible choices, since $\mathcal{S}_3$ is the cross product of two (convex) sets of probabilities and utilities. In general, the consensus set $\mathcal{S}$ is *not* arrived at by analyzing individual beliefs and values independently.

In another study,[21] we explore a representation for strict partial orders over pairs of horse lotteries in terms of such sets $\mathcal{S}$. It remains for us an open question, in general, what set $\mathcal{S}$ is generated this way by the preferences of several (quasi-) Bayesian agents—agents each of whose preferences is given by a strict partial order of this very sort. With two Bayesian agents, the quasi-Bayesian group preference generated by the weak Pareto condition is reported in theorem 1. Then, $\mathcal{S}$ is the set consisting of the two points. What strict partial order corresponds in this way to the shared preferences of n-Bayesians? To know the answer is to know the Pareto decisions of a panel of Bayesian experts.[22]

TEDDY SEIDENFELD, JOSEPH B. KADANE, and MARK J. SCHERVISH
Carnegie-Mellon University

APPENDIX

In their 1979 article, Hylland and Zeckhauser pursue a heuristic argument [due to Zeckhauser, "Group Decision and Allocation," Discussion Paper #51 (1968), Harvard Institute of Economic Research] to establish the impossibility of a rule for amalgamating a set of Bayesian preferences into a single Bayesian preference, provided that the rule: (1) applies to all potential sets of Bayesian agents; (2) respects the weak Pareto condition; (3) avoids "Dictators" for the group-probability; (4) combines probabilities and utilities separately; and (5) preserves unanimity of the agents' probabilities in case they all have the same degrees of belief. Also, they show that (2) and (3) may be replaced by the single condition (6): the strong Pareto requirement.

We note, first, that regarding "impossibility," the two results share an

[21] "A Representation for Preference as a Strict Partial Order in terms of Sets of Pairs of Probabilities and Utilities," *op. cit.*

[22] In his doctoral thesis, "Existence of Compromises in Simple Group Decisions" (Dept. of Statistics, Carnegie Mellon University), Goodman extends the negative finding of theorem 1 to groups of three agents (n = 3). Also, he gives sufficient conditions that only autocratic "compromises" exist with groups of n-many Bayesian agents.

equivalence between the conjunction of conditions (2) weak Pareto and (3) no (probability) dictators/no autocrats, and condition (6) the strong Pareto requirement. (The weaker condition, "no (probability) dictators," suffices in their argument because they have condition (4). Instead, we use "no autocrats.") Unlike their result, however, theorem 1 does not place restrictions on the form of the amalgamation, i.e., neither condition (4) nor (5) is required for our analysis. Thus, it is insufficient for finding a Bayes compromise merely to abandon independent amalgamations of probability and utility.

Also, theorem 1 shows that there are none but autocratic, weak Pareto "compromises" for *every* pair of agents with differing beliefs and values. This is in marked contrast with their result which states that, for each Bayesian social welfare rule satisfying 3–5, there is some pair of Bayesian agents whose preferences are amalgamated in violation of (2): the weak Pareto condition.

A somewhat different treatment of the Bayesian group decision problem is found in Hammond's 1981 paper. Hammond's interesting work is concerned with issues of welfare optimality in dynamic (intertemporal) social decisions. Abstracting away from the dynamic features of his analysis, we find the following result about Bayesian amalgamations:

"*Theorem 2* " (Hammond, *op. cit.*, p. 241): Suppose there are n Bayesian agents, whose preferences over their own gains from social acts are represented by the probability/utility pairs (P_k, U_k), $k = 1, \ldots, n$. Then, a Bayesian (Bergson) social welfare function, W, satisfying the strong Pareto condition exists, provided it is of the form (P_W, U_W), where: (i) $P_W = P_k$ $(k = 1, \ldots, n)$, i.e., the n agents all have the same personal probabilities, and (ii) $U_W = \sum_k \gamma_k U_k$ $(k = 1, \ldots, n$ and $0 < \gamma_k)$, i.e., social utility is a convex combination of the n individual utilities. Thus, according to this result, there can be no Bayesian amalgamation (subject to the strong Pareto condition) whenever different agents hold different personal probabilities, regardless of the nature of their personal values for outcomes.

Theorem 1 is stronger than this result because, as in the previous comparison, Hammond's conclusion depends upon a restricted form of group amalgamation—a Bergson social welfare rule makes the group utility a function of the individual (interpersonal) utilities, independent of their personal probabilities. A Bergson amalgamation treats beliefs and values separately. (There is, in addition, an assumption that the amalgamation of individual utilities is differentiable, i.e., Hammond's argument requires a smoothly changing Bergson social welfare.) Also paralleling the comparison with Hylland and Zeckhauser's result, Hammond's argument shows that, for every Bergson social welfare rule (unless the agents share a common personal probability), there is some configuration of personal utilities which leads to a violation of the strong Pareto condition. Theorem 1, by contrast, establishes the violation of the strong Pareto condition for each Bayes model and for every pair of agents with different beliefs and values.

Finally, it is worth explaining away an apparent conflict between Hammond's theorem 2 and other findings we have made, concerning Bayesian compromises when values are shared (there is a common utility) and beliefs differ (cf. "Decisions without Ordering," §4). "Theorem 2" asserts that no Bergson social welfare rule is possible in this case. Our discussion indicates that there exist weak Pareto compromises created by the following pairs: Let the Bayes model use the (assumed) common utility and any convex combination of the personal probability distributions. In the case of two agents, each nonextreme convex combination of the individual probabilities (no one individual is autocrat) creates a compromise that satisfies the strong Pareto condition. [This generalizes to n agents. For conditions under which the strong Pareto compromises require all positive weights, see, e.g., P. C. Fishburn, "On Harsanyi's Utilitarian Cardinal Welfare Theorem," *Theory and Decision,* XVII (1984): 21–28.]

The two claims are not contradictory: they deal with different domains of social acts. In Hammond's presentation, the class of social acts includes the (smaller) domain of our analysis. For Hammond's "Theorem 2," the domain of social acts arises by feasible market transactions. In the restricted domain of our analysis, each agent receives the same outcome as every other agent, in every state. Thus, when we suppose that agents share a common utility for individual rewards, perforce it is also a common utility over the "constant" social acts in this limited variety of social choices.

Instead, the social acts used in Hammond's analysis include the commonplace situation where different individuals receive different rewards. The personal utilities of Hammond's analysis reflect the agent's preferences for her individual rewards only. Thus, for Hammond's argument, two agents with the same individual utility will differ in their preferences for some "constant" social acts. If they receive different rewards, one from another, their preferences between two social acts may be in direct opposition despite the common utility for individual rewards. For example, they can have the same (cardinal) utility for money but differ in their ranking of two social options, depending upon which of them receives the greater monetary reward under which of the two acts.

We can extend theorem 1 to the larger domain of Hammond's social choices by recasting our notion of a social option. Let us mean, rather, that an act is an n-tuple of horse lotteries, one for each agent. This is inclusive of the restricted class of social acts in which each agent receives the same horse lottery. Then, as we suppose when arguing that a convex combination of personal probabilities creates a weak Pareto compromise, in order for two agents to share a common utility over outcomes—over the "constant" acts in this extended sense of "option"—they must agree on the ranking of all social acts which award distinct (von Neumann–Morgenstern) lotteries to different individuals. If they have the same utility for these "constant" social acts then [in J. C. Harsanyi's sense; cf. Rational Behavior and Bargaining Equilibrium in Games and Social Situations (New York: Cambridge, 1977), ch. 4] they share one "moral" utility over such options. In that case, an agent's preference over "constant" social acts does not depend upon his identity; it does not depend upon which (lottery) reward is his. But when agents have a common utility for personal rewards and evaluate social acts solely in terms of their own gains (as in Hammond's analysis), their preferences do *not* correspond to a "moral" utility. Thus, the appearance of a conflict between the two claims is illusory.

LINKED CITATIONS

On the Shared Preferences of Two Bayesian Decision Makers
Teddy Seidenfeld; Joseph B. Kadane; Mark J. Schervish
The Journal of Philosophy, Vol. 86, No. 5. (May, 1989), pp. 225-244.
Stable URL:
http://links.jstor.org/sici?sici=0022-362X%28198905%2986%3A5%3C225%3AOTSPOT%3E2.0.CO%3B2-Y

[Footnotes]

[2] **The Decisions of a Committee Using a Special Majority**
Duncan Black
Econometrica, Vol. 16, No. 3. (Jul., 1948), pp. 245-261.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28194807%2916%3A3%3C245%3ATDOACU%3E2.0.CO%3B2-1

[3] **Possibility Theorems with Interpersonally Comparable Welfare Levels**
Kevin W. S. Roberts
The Review of Economic Studies, Vol. 47, No. 2. (Jan., 1980), pp. 409-420.
Stable URL:
http://links.jstor.org/sici?sici=0034-6527%28198001%2947%3A2%3C409%3APTWICW%3E2.0.CO%3B2-6

[5] **Conflict and Social Agency**
Isaac Levi
The Journal of Philosophy, Vol. 79, No. 5. (May, 1982), pp. 231-247.
Stable URL:
http://links.jstor.org/sici?sici=0022-362X%28198205%2979%3A5%3C231%3ACASA%3E2.0.CO%3B2-M

[6] **On State Dependent Preferences and Subjective Probabilities**
Edi Karni; David Schmeidler; Karl Vind
Econometrica, Vol. 51, No. 4. (Jul., 1983), pp. 1021-1031.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28198307%2951%3A4%3C1021%3AOSDPAS%3E2.0.CO%3B2-M

NOTE: *The reference numbering from the original has been maintained in this citation list.*

LINKED CITATIONS

[9] **A Definition of Subjective Probability**
F. J. Anscombe; R. J. Aumann
The Annals of Mathematical Statistics, Vol. 34, No. 1. (Mar., 1963), pp. 199-205.
Stable URL:
http://links.jstor.org/sici?sici=0003-4851%28196303%2934%3A1%3C199%3AADOSP%3E2.0.CO%3B2-8

[11] **Opposition of Interest in Subjective Bayesian Theory**
Joseph B. Kadane
Management Science, Vol. 31, No. 12. (Dec., 1985), pp. 1586-1588.
Stable URL:
http://links.jstor.org/sici?sici=0025-1909%28198512%2931%3A12%3C1586%3AOOIISB%3E2.0.CO%3B2-H

[13] **The Impossibility of Bayesian Group Decision Making with Separate Aggregation of Beliefs and Values**
Aanund Hylland; Richard Zeckhauser
Econometrica, Vol. 47, No. 6. (Nov., 1979), pp. 1321-1336.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28197911%2947%3A6%3C1321%3ATIOBGD%3E2.0.CO%3B2-H

[13] **Ex-Ante and Ex-Post Welfare Optimality under Uncertainty**
Peter J. Hammond
Economica, New Series, Vol. 48, No. 191. (Aug., 1981), pp. 235-250.
Stable URL:
http://links.jstor.org/sici?sici=0013-0427%28198108%292%3A48%3A191%3C235%3AEAEWOU%3E2.0.CO%3B2-8

[14] **Utilitarianism and Expected Utility**
John Broome
The Journal of Philosophy, Vol. 84, No. 8. (Aug., 1987), pp. 405-422.
Stable URL:
http://links.jstor.org/sici?sici=0022-362X%28198708%2984%3A8%3C405%3AUAEU%3E2.0.CO%3B2-A

[14] **Combining Probability Distributions: A Critique and an Annotated Bibliography**
Christian Genest; James V. Zidek
Statistical Science, Vol. 1, No. 1. (Feb., 1986), pp. 114-135.
Stable URL:
http://links.jstor.org/sici?sici=0883-4237%28198602%291%3A1%3C114%3ACPDACA%3E2.0.CO%3B2-W

NOTE: *The reference numbering from the original has been maintained in this citation list.*

LINKED CITATIONS

[15] **Conflict and Social Agency**
Isaac Levi
The Journal of Philosophy, Vol. 79, No. 5. (May, 1982), pp. 231-247.
Stable URL:
http://links.jstor.org/sici?sici=0022-362X%28198205%2979%3A5%3C231%3ACASA%3E2.0.CO%3B2-M

[17] **Social Choice Theory: A Re-Examination**
Amartya Sen
Econometrica, Vol. 45, No. 1. (Jan., 1977), pp. 53-88.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28197701%2945%3A1%3C53%3ASCTAR%3E2.0.CO%3B2-3

[18] **Conflict and Social Agency**
Isaac Levi
The Journal of Philosophy, Vol. 79, No. 5. (May, 1982), pp. 231-247.
Stable URL:
http://links.jstor.org/sici?sici=0022-362X%28198205%2979%3A5%3C231%3ACASA%3E2.0.CO%3B2-M