Bayesian Analysis (2014) **9**, Number 3, pp. 685–698

Toward Rational Social Decisions: A Review and Some Results

Joseph B. Kadane * and Steven N. MacEachern †

Abstract. Bayesian decision theory is profoundly personalistic. It prescribes the decision d that minimizes the expectation of the decision-maker's loss function $L(d, \theta)$ with respect to that person's opinion $\pi(\theta)$. Attempts to extend this paradigm to more than one decision-maker have generally been unsuccessful, as shown in Part A of this paper. Part B of this paper explores a different decision set-up, in which Bayesians make choices knowing that later Bayesians will make decisions that matter to the earlier Bayesians. We explore conditions under which they together can be modeled as a single Bayesian. There are three reasons for doing so:

1. To understand the common structure of various examples, in some of which the reduction to a single Bayesian is possible, and in some of which it is not. In particular, it helps to deepen our understanding of the desirability of randomization to Bayesians.
2. As a possible computational simplification. When such reduction is possible, standard expected loss minimization software can be used to find optimal actions.
3. As a start toward a better understanding of social decision-making.

Keywords: social decisions, compromise, randomization

A. Counterexamples in social decisions

1 Rational individual decisions

Savage (1954), preceded in part by Ramsey (1931) and deFinetti (1937), introduced axioms leading to the maximization of expected utility:

$$\underset{d}{\operatorname{argmax}} \int U(\theta, d)p(\theta)d\theta = \underset{d}{\operatorname{argmin}} \int L(\theta, d)p(\theta)d\theta \tag{1}$$

where d is a decision variable, θ is uncertain, and where $U(\theta, d) = -L(\theta, d)$. Savage's interpretation of U and p are personalistic (or subjective), meaning that they reflect, respectively, the values and opinions of the decision-maker. This perspective is supported by the Dutch-book argument that failing to make decisions according to (1) can

*Department of Statistics, Carnegie Mellon University kadane@stat.cmu.edu
†Department of Statistics, The Ohio State University snm@stat.osu.edu

 DOI:10.1214/14-BA876

$$\begin{array}{c} \Psi_1(A) > \Psi_1(B) > \Psi_1(C) \\ \Psi_2(B) > \psi_2(C) > \Psi_2(A) \\ \Psi_3(C) > \Psi_3(A) > \Psi_3(B) \end{array}$$

Table 1: Preferences of three voters among three alternatives

lead to sure loss. The ensuing development of this notion of decision-making has led it to be the dominant view among philosophers and economists and increasingly among statisticians. For the purposes of this paper, (1) is regarded as the rational way for an individual to make decisions. The issue addressed here is the extent to which it can be extended to decisions made by groups.

That social decisions are more complicated cannot come as a surprise. Many disciplines study aspects of it, including political science, economics, political economy, law, and public administration. None of these has a totally satisfying way of explaining how groups of people should make decisions that reflect the wishes and beliefs of the members of the group. The remainder of this section rehearses some of the well-known reasons why this is the case.

2 Groups and voting

If each member of the group has the same utility $U(d, \theta)$ and probability $p(\theta)$, then any member of the group can act for the group, and would achieve the same expected utility as would any other. Under this condition, the group problem essentially devolves into implementing (1).

When the group members do not necessarily have the same values and probabilities, things get more complicated. One frequently used method to decide matters is to vote. Suppose there are three voters, 1, 2, and 3 and three options, A, B and C. The preference function, $\Psi_i(\cdot)$, of each voter is given in Table 1, and there is no uncertainty at issue.

Suppose that each voter votes according to these preferences (*i.e.*, there is no strategic voting). Then, pairwise, A would defeat B (voters 1 and 3), B would defeat C (voters 1 and 2), and C would defeat A (voters 2 and 3). Thus the pairwise outcomes are intransitive. Suppose now that voter 1 gets to set the agenda, and knows the preferences of the other voters. Wanting A to be the outcome, and knowing C to be the threat, voter 1 proposes a vote between B and C, the winner to be paired against A. Then B defeats C, A defeats B, and A wins. Similarly, each other voter, given the agenda-setting power, could set an agenda so that his or her favored outcome wins.

Strategic voting would apply to voter 2 in the first vote. Foreseeing that a vote for B against C is defacto a vote for A (which is voter 2's least preferred alternative), voter 2 could vote for C instead of B. Then C would be the outcome in the second vote. Now voter 1, foreseeing this possibility sets A vs. C in the first vote with the winner to be paired against B. Voter 3, seeing the eventual results of a B vs. C contest, votes for A.

A wins the first vote and also wins against B.

This example illustrates in general why the exact manner in which decisions are taken can matter a lot, and must be specified in order to determine the outcome. (It also shows the importance of the agenda-setting power of the Speaker of the US House of Representatives and the Majority Leader of the US Senate.)

3 Compromise

If voting has such paradoxes, perhaps Bayesians can come to a Bayesian compromise. To explain what is meant by a Bayesian group decision, we mean only that it is characterized by a utility function $U(d, \theta)$ and a probability distribution $p(\theta)$. It is a compromise if it satisfies a weak Pareto condition: if each Bayesian strictly prefers decision d_1 to decision d_2, then so will their compromise. A compromise is autocratic if the "compromise" utility and probability are those of one of the parties. The parties can be offered choices that depend on both the θ that obtains and the utility consequence of making a particular decision. Then the following theorem shows that in general such compromises do not exist:

Theorem 4. *There exist non-autocratic, weak Pareto compromises for two Bayesians if and only if either*

a) they agree in probability, but not in utility or

b) they agree in utility, but not in probability.

(For proof, see Kadane (2011, Theorem 11.8.1) *and Seidenfeld et al. (1989)).*

The implication of this theorem is that when two Bayesians disagree in both probability and utility, there is no compromise Bayesian position for them.

4 Simultaneous moves

This is the decision-structure of traditional game theory. Two or more decision-makers each make decisions without knowing what decisions the others have made. It is in this context, and in particular for a two-person game in which the parties' interests are diametrically opposed (*i.e.*, zero, or constant sum), that Von Neumann and Morgenstern (1944) advocate the minimax strategy. It is supported by the minimax theorem (due to Von Neumann), that says that in a two-person zero sum game in which mixed strategies are permitted, if a player is sure that the opponent will play his minimax strategy, the best the opponent can do is also to play a minimax strategy.

But why should a player be sure that his opponent will in fact play a minimax strategy? A Bayesian with a probability distribution other than minimax will be able to maximize utility and do better. See Kadane and Larkey (1982) and Harsanyi (1982) for debate on this point.

An especially interesting game is the prisoner's dilemma. The story is that two criminal suspects are separated, and given the choice of whether to confess. Each wishes to minimize the length of prison sentence they may get. If one confesses and the other does not, the confessor gets only one year in jail (say on a weapons charge), while the other prisoner gets 10 years. If neither confesses, they each get 2 years in prison. If both confess, they each get 5 years, What are the rational decisions for them?

If prisoner A knows that prisoner B is going to confess, prisoner A faces a choice between 10 years for failing to confess, and 5 years for confessing. If prisoner A knows that prisoner B is not going to confess, prisoner A faces a choice between 2 years for failing to confess, and 1 year for confessing. Hence, whatever prisoner B is going to do, prisoner A does better confessing. The same is true for prisoner B, so they both confess, and get 5 years each.

There are several things to notice about this game. First, the decisions that are optimal do not depend on the probabilities that each player assigns to the move of the other player. In this case, the decision to confess is said to dominate the decision to refuse to confess. Second, the individually rational decisions, to confess, with the result of 5 years each, is worse for both of them than the jointly rational decision not to confess, with the result of 2 years each. If they could make an enforceable contract not to confess, they would both be better off. This illustrates the social importance of the ability to make such contracts, and of the court's role in enforcing them. It also explains why contracts contrary to social policy, such as those in restraint of trade, are not enforceable.

The prisoner's dilemma gets more interesting if one regards it sequentially, so each party faces many such decisions, one at a time, against the same opponent/partner. A standard argument runs by backward induction, as follows. On the last play, each knows that the other will confess. So on the next-to-last play, each should confess, etc. Thus, this argument goes, the only rational strategy is to confess at each play. Suppose, however, that a Bayesian believes that if he confesses on any given play (except the last), his counterpart will confess forever afterward, but if he refuses to confess, his counterpart will do so on the next game. Then it would be expected-utility maximizing to refuse to confess.

5 Sequential decision making

In sequential decision making, each player makes decisions knowing the decision made before. Chess and checkers are familiar examples of this kind of structure.

The example given here comes from a real situation in which attorneys choosing a jury can exercise peremptory challenges to exclude certain potential jurors. The example here is simplified in many ways to reveal a surprising aspect of this decision problem. The rules for the simplified game are:

1. There are three kinds of jurors, H, M, and L.

2. Both sides know that they will be considered in the order H, M, L, L.

3. A jury with two members is to be chosen.

4. The defense will choose whether to exercise a challenge first, then the prosecution decides.

5. Both sides know the preferences of both parties, which are, respectively:

$$\Psi_D(HM) > \Psi_D(LL) > \Psi_D(ML) > \Psi_D(HL)$$
and
$$\Psi_P(HL) > \Psi_P(HM) > \Psi_P(ML) > \Psi_P(LL)$$

6. The defense gets d challenges and the prosecution p, where $p + d = 2$.

We consider two cases. First, suppose that $p = d = 1$. The defense goes first. If the defense chooses not to challenge H, the prosecution will accept and challenge M, leading to an HL jury, which is the worst outcome for the defense. Therefore, the defense challenges H, and the result is an ML jury.

Now suppose $d = 2$ and $p = 0$. Then the defense will not exercise any challenges, and the result is an HM jury.

So by giving its challenge to the defense, the prosecution achieves a jury more to its liking than it gets otherwise. This example is somewhat similar to the prisoner's dilemma, in that if the two sides could make an enforceable agreement, they would both do better. This example is also notable in that uncertainty plays no part here. It is the sequential structure of the decision problem alone that leads to this anomaly. For more on the problem, see Roth et al. (1977), DeGroot and Kadane (1980), and Kadane et al. (1999).

B. A partial positive result

Given the many ways in which Bayesian ideas cannot be extended to social decisions, a conclusion might be drawn that the problem is totally hopeless. However, there is one additional possibility to explore. Is it possible, in the setting of sequential decisions, for the (we'll say two, at first) Bayesians to make decisions that are consistent with some probability and utility function, even if these may not be those of either player? This section explores this possibility. In contrast to Theorem 4, the choices offered depend only on the decision d, and not on θ.

6 Decision-equivalent Bayesian formulations

A *formulation* of a decision problem specifies a weak ordering of the decisions $\mathcal{D}$ (i.e., if d_1 and d_2 are elements of $\mathcal{D}$, then one of the following holds: d_1 is strictly preferred to d_2, d_2 is strictly preferred to d_1, or d_1 and d_2 are indifferent). *Decision-equivalent formulations* operate on the same space of decisions $\mathcal{D}$, and specify the same weak order. For example, a Bayesian with utility function $U(d, \theta)$ and prior $\pi(\theta)$, and who maximizes expected utility orders decisions in the same way as a Bayesian who minimizes expected loss, where $L(d; \theta) = -U(d, \theta)$, with the same prior $\pi(\theta)$. Similarly, an expected-utility maximizing Bayesian with utility $aU(d, \theta) + b$ $(a > 0)$ and prior $\pi(\theta)$ orders decisions in the same way as an expected utility maximizing Bayesian with utility $U(d, \theta)$ and prior $\pi(\theta)$.

If two formulations are decision-equivalent, the set of optimal decisions (whether empty, a single decision, or a set of more than one decision) is the same. Similarly, provided that the weak ordering comes from a real-valued function (as is true of Bayesian decision theory), the set of ϵ-optimal decisions (for $\epsilon > 0$) for the first formulation is the same as the set of ϵ'-optimal decisions in the second formulation, for some $\epsilon' > 0$.

7 Single stage decisions

Suppose two Bayesians face a decision problem with the same space of decisions $\mathcal{D}$ and the same space indexing their uncertainties Ω. Suppose their loss functions are $L_1(d, \theta)$ and $L_2(d, \theta)$, and their opinions over Ω are respectively $\pi_1(\theta)$ and $\pi_2(\theta)$. Consider the following representation of the first Bayesian's expected loss for decision $d \epsilon \mathcal{D}$:

$$\begin{aligned} \int_\Omega L_1(d, \theta)\pi_1(\theta)d\theta &= \int_\Omega L_1(d, \theta)\frac{\pi_1(\theta)}{\pi_2(\theta)}\pi_2(\theta)d\theta \\ &= \int_\Omega L_2(d, \theta)\pi_2(\theta)d\theta \end{aligned} \tag{2}$$

where

$$L_2(d, \theta) = L_1(d, \theta)\pi_1(\theta)/\pi_2(\theta). \tag{3}$$

This formal representation is valid provided $\pi_2(\theta) = 0 \Rightarrow \pi_1(\theta) = 0$, *i.e.*, provided that θ's impossible to Bayesian 2 are impossible to Bayesian 1, or that Bayesian 1's opinion about θ is absolutely continuous with respect to Bayesian 2's opinion. Under this condition, L_2 in (3) is well-defined. These results are summarized below:

Theorem 5. *Bayesian 1 with loss $L_1(d, \theta)$ and opinion $\pi_{1(\theta)}$ suffers the same expected loss for each decision $d \epsilon \mathcal{D}$ as does Bayesian 2 with loss function $L_2(d, \theta)$ defined in (3) provided π_1 is absolutely continuous with respect to π_2.*

Note that if π_2 is absolutely continuous with respect to π_1, the roles of Bayesian 1 and 2 can be interchanged.

Ratios such as (3) are familiar objects in Bayesian statistics, particularly in checking detailed balance for Markov Chain Monte Carlo methods.

8 Single party two-stage decisions

This section establishes notation, and rehearses the Bayesian approach to experimental design (Lindley 1972), which is a paradigmatic special case of a general two-stage decision problem. The decision-maker chooses a design $\delta \epsilon \Delta$. This leads to the observation of data x, which in turn leads to some action $a \epsilon \mathcal{A}$ by the decision-maker. Of interest to but uncertain to the decision maker is a parameter $\theta \epsilon \Omega$. How should a Bayesian order choices of δ and a?

Suppose that the Bayesian has a loss function $L(a, \theta, \delta, x)$ which determines the losses suffered if design δ and action a are taken in the situation in which data x are observed and the parameter takes the value θ. The Bayesian wishes to minimize an appropriate version of expected loss.

After choosing design δ and observing data x, the Bayesian's information about θ is summarized by the posterior distribution of θ given x and δ. Hence for choice of a, a Bayesian would order elements of $\mathcal{A}$ to minimize

$$\int_\Omega L(a, \theta, \delta, x) p(\theta \,|\, x, \delta) d\theta. \tag{4}$$

Note that in this ordering a is a function of δ and x.

After choice of δ, but before observing x, the Bayesian has the probability distribution on x of

$$p(x \,|\, \delta) = \int_\Omega p(x, \theta \,|\, \delta) d\theta. \tag{5}$$

Hence a Bayesian would order Δ, preferring small values of

$$\int_X \min_a \int_\Omega d\theta dx \; L(a, \theta, \delta, x) p(\theta \,|\, x, \delta) p(x \,|\, \delta) \tag{6}$$

and the resulting expected loss is

$$\min_\delta \int_X \min_a \int_\Omega d\theta dx \; L(a, \theta, \delta, x) p(\theta \,|\, x, \delta) p(x \,|\, \delta). \tag{7}$$

Now suppose that Bayesian 2 is to make these decisions instead of Bayesian 1. After choosing design δ and observing data x, Bayesian 2 would order $\mathcal{A}$, preferring small values of

$$\int_\Omega L_2(a, \theta, \delta, x) p_2(\theta \,|\, x, \delta) d\theta = \int L_1(a, \theta, \delta, x) p_1(\theta \,|\, x, \delta) d\theta \tag{8}$$

provided

$$L_2(a, \theta, \delta, x) = L_1(a, \theta, \delta, x) \frac{p_1(\theta \,|\, x, \delta)}{p_2(\theta \,|\, x, \delta)} \tag{9}$$

and $p_1(\theta \,|\, x, \delta)$ is absolutely continuous with respect to $p_2(\theta \,|\, x, \delta)$. Furthermore, such L_2 could be multiplied by an arbitrary positive function of x and δ without changing the ordering by expected loss of actions $a \epsilon \mathcal{A}$.

Next, suppose that Bayesian 2 is to choose the design as well. Then Bayesian 2 will order Δ, preferring small values of

$$\int_X \min_a \int_\Omega d\theta dx \; L_2(a,\theta,\delta,x) p_2(\theta \mid x,\delta) p_2(x \mid \delta) \tag{10}$$

which will have the same expected loss to Bayesian 2 as they would to Bayesian 1 provided

$$L_2(a,\theta,\delta,x) = L_1(a,\theta,\delta,x) \frac{p_1(\theta \mid x,\delta)}{p_2(\theta \mid x,\delta)} \; \frac{p_1(x \mid \delta)}{p_2(x \mid \delta)}. \tag{11}$$

Can (9) and (11) be reconciled? Note that $p_1(x \mid \delta)/p_2(x \mid \delta)$ depends only on x and δ. Hence, with L_2 satisfying (11), and provided $p_1(\theta, x \mid \delta)$ is absolutely continuous with respect to $p_2(\theta, x \mid \delta)$ for all δ, Bayesians 1 and 2 will order both Δ and $\mathcal{A}$ in the same way.

Theorem 6. *Bayesian 1 with loss $L_1(a,\theta,\delta,x)$ and opinion $p_1(\theta,x \mid \delta)$ orders $\mathcal{A}$ and Δ in the same way as Bayesian 2 with loss $L_2(a,\theta,\delta,x)$ and opinion $p_2(\theta,x \mid \delta)$ provided that p_1 is absolutely continuous with respect to p_2 for all δ, and provided L_2 is defined as in (11). Bayesian 1 has the same expected loss for each δ as does Bayesian 2; however at the estimation stage of ordering $\mathcal{A}$ given δ and x, Bayesian 2's expected loss is a factor of $p_1(x \mid \delta)/p_2(x \mid \delta)$ times Bayesian 1's expected loss for each $a \epsilon \mathcal{A}$.*

9 Two-party two-stage decisions

There is no particular reason why the two decisions, a and δ, have to be taken by the same decision maker, and there are many applications in which it is convenient to suppose they are not. In this case, it is necessary to specify whose probabilities and losses are under consideration. Since the Bayesians are no longer in symmetric positions, we distinguish them with letters and names instead of numbers. Suppose Dan, the designer, chooses δ, while Edward, the estimator, chooses a. Probabilities and losses reflecting Dan's beliefs are subscripted "d", while those reflecting Edward's are subscripted e.

After Dan chooses design δ, and data x are observed, Edward's information about θ is summarized by his posterior distribution of θ given x and δ. Hence, Edward orders $\mathcal{A}$, preferring small values of

$$\int_\Omega L_e(a,\theta,\delta,x) p_e(\theta \mid x,\delta) d\theta. \tag{12}$$

Note again that the ordering of $\mathcal{A}$ is a function of δ and x. Suppose the optimal choice for Edward is $a^*(\delta,x)$.

For choice of δ, and knowing Edward's ordering of $\mathcal{A}$, Dan is uncertain about both x and θ, and has probability distribution over them of $p_d(x,\theta \mid \delta) = p_d(x \mid \delta) p_d(\theta \mid x,\delta)$. Hence, Dan suffers expected loss

$$\int_x \int_\Omega L_d(a^*(\delta,x),\theta,\delta,x) p_d(\theta \mid x,\delta) d\theta \; p_d(x \mid \delta) dx \tag{13}$$

and orders Δ to minimize (13).

Now suppose that

$$L_d(a, \theta, \delta, x) = L_e(a, \theta, \delta, x) \frac{p_e(\theta, x \mid \delta)}{p_d(\theta, x \mid \delta)}. \tag{14}$$

From the arguments given above, Edward's ordering of final actions $\mathcal{A}$ would be the same as Dan's, and Dan's choice of design would be the same as Edward's.

Theorem 7. *In the Two-Party Two-Stage case, if Dan's utility function is given by (14) and $p_e(\theta, x \mid \delta)$ is absolutely continuous with respect to $p_d(\theta, x \mid \delta)$, then Dan's ordering of Δ and $\mathcal{A}$ are the same as Edward's. Dan's expected loss is the same as Edward's at the design stage; however, at the stage of choice of final $a \epsilon \mathcal{A}$ action, given δ and x, Dan's expected loss is a factor of $p_e(x \mid \delta)/p_d(x \mid \delta)$ times Edward's expected loss for each action $a \epsilon \mathcal{A}$.*

10 Dan doesn't know Edward's loss function and prior

The analysis in Section 9 assumes that Dan knows both Edward's loss function and his prior distribution. What happens if this is not the case? Suppose that Dan's uncertainty about Edward's loss function and prior distribution are indexed by $\psi \epsilon \Psi$. The value of ψ is known to Edward but not to Dan. In this respect it is just like the data x, which, again, Dan does not know but Edward does know when he chooses his action a. Therefore the analysis of Section 9 can be reinterpreted, where what had been data x is now pairs (x, ψ). With this reinterpretation, Theorem 7 applies to this case. Such a reinterpretation is also used in game theory, where ψ is called Edward's "type" (Harsanyi (1967/1968)).

11 Several Bayesians

The argument so far has been limited to two Bayesians, Dan and Edward, respectively a designer and an estimator. Can it be extended to three or more Bayesians? Suppose there is a third Bayesian, Sabrina, a super designer, who decides what design problem to pose to Dan. As Theorem 7 shows, provided its conditions are satisfied, Dan can make the design-and-estimation decisions for both himself and Edward, with a modified loss function. Now to Sabrina, Dan looks like a new Edward, a Bayesian with a loss function, a likelihood, a prior, and some decisions to make. Once again, Theorem 7 applies provided its conditions are met. Hence, any number of Bayesians can be collapsed to a single Bayesian via repeated use of Theorem 7.

12 Examples

Theorems 5 through 7 provide a framework under which one Bayesian can emulate another. The key condition in the theorems is that the emulator must be as open-minded as the emulated – that a set of states of nature judged impossible by the emulator is also judged impossible by the emulated. The additional conditions describe knowledge of or beliefs about the emulated's loss function. Theorem 4 stands in sharp contrast to these results, as the following example shows.

Example 1:

All parties agree that there are two states of nature, θ_1 and θ_2. There are three actions, a_1, a_2, and a_3. There is no data. Three Bayesians have the following prior distributions and loss functions. The tables provide the values of $L(a, \theta)$.

	Alice			*Bob*			*Clarice*		
	a_1	a_2	a_3	a_1	a_2	a_3	a_1	a_2	a_3
θ_1	0	1	0	0	2	0	0	4/3	0
θ_2	2	0	4	1	0	2	4/3	0	8/3
$\pi(\theta_1)$		2/3			1/3			1/2	
$\pi(\theta_2)$		1/3			2/3			1/2	

These three prior-loss pairs satisfy the conditions of Theorem 5 for all assignments of the participants to the roles of Bayesians 1 and 2. Each of the three can emulate each of the others. Furthermore, emulation extends to a modified version of the problem where data x follows a universally agreed upon sampling distribution $p(x|\theta)$ is observed and decisions $d(x)$ replace the no-data actions.

Choosing any two of the three, say Alice and Bob, we ask whether there is a non-autocratic weak Pareto compromise. Clarice differs from both Alice and Bob in prior distribution (probability) and loss function (utility), and so is a natural candidate for a compromise. However, Theorem 4 tells us that there is no non-autocratic compromise, in spite of agreement on the part of Alice, Bob, and Clarice on the value of all integrals of the form (2).

The difference between the emulation of Theorems 5 through 7 and the compromise of Theorem 4 is in the set of decisions that are considered. Theorem 4 entertains a larger set of decisions, for example, allowing the Bayesian to choose between (θ_2, a_1) and (θ_1, a_2) as well as lotteries of such state-of-nature, action pairs. This style of choice sidesteps the integral over θ. In the larger set of choices we find disagreement between the three Bayesians, and it is these disagreements that prevent Clarice from representing a compromise.

This section discusses two additional examples. The first falls under the conditions of Theorem 7; the second does not.

Example 2: (Lindley and Singpurwalla 1991)

A manufacturer $\mathcal{M}$ produces items which the consumer $\mathcal{C}$ may accept or reject. The manufacturer pays for binomial sampling to convince the (more skeptical) consumer to accept the lot. Both have beta priors on θ, the proportion of items that are non-conforming. With specific specifications for the utility functions of both parties, optimal choices of sample size for the manufacturer are computed.

Because both parties have specified beta priors on the same parameter θ, the priors are absolutely continuous with respect to one-another. Hence, Theorem 7 applies, and (14) specifies a loss function with respect to which the manufacturer's optimal sample size can be computed.

Similarly, Etzioni and Kadane (1993) study a normal-normal model in which the designer, Dan, decides on a sample size knowing that the estimator, Edward, will use Edward's prior, together with the normal likelihood, to produce an estimate of the mean. Unlike the Lindley/Singpurwalla case, the optimal sample size can be found explicitly, as it depends on the largest real root of a certain cubic equation. This model falls into the orbit of Theorem 7, as both Dan and Edward have (possibly differing) normal distributions on the mean of the normal likelihood, so these distributions are mutually absolutely continuous.

Example 3:

Berry and Kadane (1997) study a Bayesian model in which it is optimal to randomize. To explain why this is a bit of a surprise, it is useful first to go back to the single Bayesian model of Section 8. In that model, each design that might be chosen has some expected loss to the Bayesian. Suppose designs d_1 and d_2 are such that one of them, say d_1, has smaller expected loss. Then it would be the better choice. If d_1 and d_2 have the same expected loss, then choosing either or a randomized choice between them are all equally good. But this is not a case for randomization, since choosing either deterministically is just as good.

The Berry/Kadane model extends the scenario of Etzioni and Kadane by inserting a new character, Phyllis (a physician) between Dan and Edward. In this extended scenario, Dan chooses the design, Phyllis implements it with patients, and Edward produces the final estimate. What is different is that Phyllis sees, or senses, a covariate, the health of the patient, which Dan suspects might be there but doesn't know, and Edward is blind to. (For simplicity, this model is discrete: the two treatments either succeed or fail, and "health" has only two values.) Let $p_{jh}(j = 0, 1, h = 0, 1)$ be the probability that a patient has a successful outcome if treatment j is assigned to a patient with health status h. If w is the proportion of patients in the population with health status 1, Dan wants Edward's estimates to converge (in the large-sample sense) to

$$p_j^* = wp_{j1} + (1 - w)p_{j0}, j = 1, 2. \tag{15}$$

Suppose that λ_1 is the probability that a healthy patient is assigned to treatment

1, and λ_0 is the probability that an unhealthy patient is assigned to treatment 1. The values of λ_1 and λ_0 could be chosen by Dan, or he could leave the decisions to Phyllis. In either case, Dan's view of the value of treatment j is $P\{\text{success}|\text{treatment}\}$. After some algebra, the following results:

$$p_j^* - P\{\text{success}|\text{treatment } j\} = \frac{w(1-w)(p_{j1}-p_{j0})(\lambda_0-\lambda_1)}{w\lambda_1+(1-w)\lambda_0}. \tag{16}$$

Equation (16) must equal zero for Dan's goal for Edward's estimates to be met. There are three cases in which this happens:

(i) $w(1-w)=0$. In this case there are no covariates.

(ii) $p_{j1}=p_{j0}(j=1,2)$. In this case, there are covariates, but they don't matter.

(iii) $\lambda_0=\lambda_1$. To achieve this, Dan must randomize the patients, and insist that Phyllis obey the randomization.

What relationship has this model to Theorem 7? Under cases (i) and (ii), there is a single relevant probability for success for each treatment. Theorem 7 applies and randomization is not strictly optimal. Under case (iii), because randomization is strictly optimal, the reduction to a single Bayesian model cannot occur. The reason it falls outside of Theorem 7 is that Phyllis knows the health of the patients, which neither Dan nor Edward knows, so absolute continuity fails.

13 Conclusion

The "possibility" theorems of Part B all have the same character, namely that the Bayesians can agree if the product of their relevant prior or posterior or joint probabilities and their loss functions are identical (up to a constant scale factor). This general result is not surprising: it merely says that differences in personal probabilities between agents can be compensated for by inversely proportional differences in personal utilities. The main reason we find these results enlightening is that they shed further light on how it is that although a Bayesian in a single-party problem has no incentive to randomize, this is not the case in the multi-party scenario of Example 3. Since sampling and randomization are such an important part of statistical practice, it behooves us to understand why.

References

Berry, S. and Kadane, J. (1997). "Optimal Bayesian randomization." *Journal of the Royal Statistical Society: Series B*, 59(4): 813–819. 695

deFinetti, B. (1937). *"La Prèvision: ses lois logique, se sources subjectives," Annales de L'Institut Henri Poincaṙe*. Translated in "Studies in Subjective Probability". John Wiley and Sons. 685

DeGroot, M. and Kadane, J. (1980). "Optimal challenges for selection." *Operations Research*, 28: 952–968. 689

Etzioni, R. and Kadane, J. (1993). "Optimal experimental design for another's analysis." *Journal of the American Statistical Association*, 88: 1404–1411. 695

Harsanyi, J. (1967/1968). "Games with incomplete information played by Bayesian players." *Management Science*, 14(3), 159-183 (Part I), 14(5), 320-344 (Part II), 14(7), 486-502 (Part III). 693

— (1982). "Subjective probability and the theory of games: Comments on Kadane and Larkey's paper." *Managment Science*, 28: 120–124. 687

Kadane, J. (2011). *Principles of Uncertainty*. Boca Raton: Chapman and Hall. 687

Kadane, J. and Larkey, P. (1982). "Subjective probability and the theory of games." *Management Science*, 28: 113–120. 687

Kadane, J., Stone, C., and Wallstrom, G. (1999). "The donation paradox for peremptory challenges." *Theory and Decision*, 47: 139–151. 689

Lindley, D. (1972). *Bayesian Statistics*. Society for Industrial and Applied Mathematics. 691

Lindley, D. and Singpurwalla, N. (1991). "On the evidence needed to reach action between adversaries, with application to acceptance sampling." *Journal of the American Statistical Association*, 86: 933–937. 695

Ramsey, F. (1931). *The Foundations of Mathematics and Other Logical Essays*, chapter Truth and Probability, Chapter VII. Routledge. 685

Roth, A., Kadane, J., and DeGroot, M. (1977). "Optimal peremptory challenges in trial by juries: A bilateral sequential process." *Operations Research*, 25: 901–919. 689

Savage, L. (1954). *Foundations of Statistics*. New York: John Wiley and Sons. 685

Seidenfeld, T., Kadane, J., and Schervish, M. (1989). "On the shared preferences of two Bayesian decision-makers (reprinted in **The Philosopher's Annual**, Vol. XII (1989)." *Journal of Philosophy*, 5: 225–244. 687

Von Neumann, J. and Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton: Princeton University Press. 687

Acknowledgments

Dr. MacEachern acknowledges grant support from NSF grants DMS-1007682 and DMS-1209194.