

argument), methinks so much protest—here and elsewhere—signifies the futility of the search for a weak link in the Bayesian argument. A theory which does not expect a coherent decision maker to stick to a strategy chosen in advance will certainly be unattractive for everyday normative use, if not chaotic. Defining consequences inadequately clouds the argument but does not refute it.

9. Savage (*Foundations*, Section 5.5) explicitly recognized that a small-world consequence depends on grand-world decisions, probabilities, and more fundamental consequences.

10. On one strictly peripheral point I disagree with both Savage and Shafer: people are regularly taken in by pseudomicrocosms that focus on one risk when others, even negatively correlated ones, are present but unmentioned. For example, to someone negotiating for the right to use a patented production process, a fixed payment may seem less risky than royalties, but the picture reverses when profits are looked at, because higher sales accompany higher royalties.

### CONCLUSIONS

Talking about the behavior of a mythical ideally consistent person may still be the best way to convince people—and many still need convincing—that subjected expected utility is uniquely normative. Resisting this idea plays only a regressive role, and obstructs a sound understanding and appraisal of alternative tools. The Bayesian view helps one to distinguish what's important, trivial, ad hoc, fundamental, non-

sensical, misleading, irrelevant, or misguided in areas of statistics from sequential stopping to ridge regression to hypothesis testing to unbiased or parameterization-invariant estimation. In problems of decision and inference under uncertainty, other arguments may sometimes be simpler and good enough, but they are never more cogent.

No new rationality has found widespread acceptance since Savage, nor should have. It is no revision of rationality to adopt short cuts, approximations, or even deliberate irrationality according to taste and circumstances, or to recognize that the main concerns often lie elsewhere. Other routes to Bayesian rationality may have advantages, but once it is accepted, even with amendments, the jig is up and the rest is tactics (or strictly for philosophical specialists).

Read literally, Shafer does not contradict most of my numbered remarks. But if he accepts them, and accepts that they are far from novel, what does all his sound and fury signify? If he does not, we live in different worlds.

I am sorry to sound so nasty. For some reason, statisticians who work in the foundations of the field often seem nicer in person than in writing. Shafer does, and I hope I do too.

### ADDITIONAL REFERENCES

- BELL, D. E. (1982). Regret in decision making under uncertainty. *Oper. Res.* **30** 961–981.
- PRATT, J. W., RAIFFA, H. and SCHLAIFER, R. (1964). The foundations of decision under uncertainty: An elementary exposition. *J. Amer. Statist. Assoc.* **59** 353–375.

## Rejoinder

Glenn Shafer

The main thesis of my article was that Savage did not establish the unique normativeness of subjective expected utility. It appears that three of the commentators, Robin Dawes, Phil Dawid, and Peter Fishburn, agree, while two, Dennis Lindley and John Pratt, disagree. In my rejoinder, I will concentrate on this central issue of normativeness. I will also respond, briefly, to the question about alternatives to subjective expected utility.

Fishburn gently notes that aspects of my constructive viewpoint are not altogether new. He adds that the idea of using subjective expected utility constructively was not altogether absent from Savage's own thinking. The points could be put more strongly. My viewpoint has, I hope, all the triteness of common sense. Common sense and historical perspective also tell us that Savage, like everyone else, expected to use

subjective expected utility in the constructive direction, from probabilities and utilities to preferences between acts.

One aspect of my constructive viewpoint is the idea that one deliberately compares a problem to a scale of canonical examples involving chance. This aspect is scarcely new. It can be found in Bertrand (1907, page 26) and in Ramsey (1931, page 256). Pratt, Raiffa, and Schlaifer (1964) very effectively incorporated it into their alternative axiomatization of subjective expected utility.

I did not venture, in my article, to survey the many alternative axiomatizations of subjective expected utility that have followed Savage's. Had I done so, I would have had an opportunity to agree with the widespread opinion that Pratt, Raiffa, and Schlaifer's is the most attractive of these. Making explicit the

comparison with games of chance (or with reference probabilities) is clearly the right thing to do. As Dawid points out, the worst absurdities involved in Savage's version of the problem of small worlds are a result of his trying to avoid this explicit comparison.

At the risk of being tendentious, I would like to underline a point Pratt seems to ignore. When we give up the attempt to avoid reference probabilities, we do weaken the plausibility of the argument for normativeness. Savage avoided reference probabilities because he wanted to avoid the issue of adequacy of evidence. He wanted to avoid the traditional objection that evidence for probabilities may be inadequate or altogether missing. So he tried to make all his comparisons look like comparisons of value. When we ask people to compare their evidence to knowledge of chances, we bring the traditional objection back in the form of the retort that there is no evidence to compare.

Pratt's attitude is only partially constructive. A fully constructive attitude involves the admission, already made by Bertrand, that the comparison to games of chance may be infeasible or undesirable. It may be infeasible because we lack the evidence needed to make our knowledge comparable to knowledge of chances. It may be undesirable because the commitments of value we want to make do not take the form of the needed utilities.

### COHERENCE AND SMALL WORLDS

Robin Dawes, Phil Dawid, and Dennis Lindley emphasize the virtues of coherence. We can often learn more about our own values, or create more satisfying values, if we consider a variety of situations and adopt preferences that are consistent across these situations. This message is valid and important. But are there limits? Is more coherence always better? Is it always useful to reach for more coherence, or does the reach become counterproductive when it goes beyond our evidence or our capacity for value?

Common sense tells us that there are limits, but it is difficult to incorporate this common sense into the normative viewpoint. As soon as we admit that the usefulness of a particular comparison depends on our evidence and our commitments, the assertion that it is normative to make the comparison becomes meaningless. This is the problem of small worlds that remains even after we give up Savage's avoidance of reference probabilities.

Lindley struggles manfully with the problem, but the only solution he can find is to posit "normative" probabilities and utilities and to suppose that our actual assessments are measurements with error of these normative quantities. What are these normative quantities? Hidden properties of the person and his evidence? The normative doctrine takes refuge again in fable.

Dawes has read far too much into my article. Neither I nor Wolfowitz would deny the value of thinking in terms of policy. We would deny only that more comprehensive policies are always better. I very much agree with Dawes that the choice of a broader context in which to embed a problem should depend on empirical facts about people and about their evidence.

### ALTERNATIVES

What alternatives to subjective expected utility does a constructive framework permit? There are many. I would include Wald's decision theory and the allied frequentist methods of inference among them. The more sterile aspects of the Bayesian versus frequentist controversy might be dispelled if we recognized that so-called frequentist methods also involve subjective comparisons to canonical examples.

I am also interested, as Lindley suspects, in a decision theory based on belief functions (Shafer, 1976). Such a decision theory would generalize subjective expected utility. A belief function is more general than a probability measure; it attaches basic probability masses to subsets (called focal elements) rather than to points. We can similarly generalize a utility function by attaching numerical utilities to sets (called goals) rather than to points. If the belief function represents our evidence about the result of an act, then we can calculate a generalized expected utility for the act by summing the products of those probability/utility pairs for which the focal element falls inside the goal, indicating evidence that the goal will be achieved. We can also calculate a generalized disutility by summing the products for which the focal element falls outside the goal, indicating evidence that the goal will not be achieved. In symbols, the generalized expected utility is

$$\sum \{m(A)v(B) \mid A \subseteq B\},$$

and the generalized expected disutility is

$$\sum \{m(A)v(B) \mid A \subseteq B^c\},$$

where the  $m(A)$  are the probability masses (non-negative numbers adding to one;  $m(A)$  is zero unless  $A$  is a focal element), and the  $v(B)$  are the utilities (non-negative numbers;  $v(B)$  is zero unless  $B$  is a goal).

One act will dominate another if its generalized expected utility is greater and its generalized expected disutility is less. This is only a partial ordering. We will have to say of some pairs of acts that our utilities and our evidence are insufficient to determine a choice.

Both frequentist methods and belief-function decision theory differ from subjective expected utility by failing to make some comparisons. In general, they give only a partial ordering of acts. They combine

limited commitments of value with probability judgments based on limited evidence, and consequently they draw only limited conclusions.

Advocates of the unique normativeness of subjective expected utility consider this failure to totally order acts an inexcusable fault; but, if we take a constructive viewpoint, then the problem of small worlds shows us that an ordering by subjective expected utility is not necessarily more complete. A subjective expected utility analysis does not determine preferences among all acts—it only determines preferences among acts at a given level of description. An analysis using an alternative decision theory may consider deeper levels of description instead of ironing out a complete ranking of acts at a given level; but, it is hard to see why this is bad. It is hard to see any justification for always insisting on a breadth first study of a problem.

### CONCLUSION

In conclusion, I would like to express my appreciation of the thoughtfulness and depth of the comments.

When an apostate inflicts a critique of scripture on the faithful, he can expect two reactions. Some will defend every implausibility and contradiction he has criticized. Others will scold him for being so rude and obtuse, assuring him that no one nowadays takes the passages literally. There are elements of both these reactions in the comments, but there is also genuine life and thought. Savage's work is not yet dead scripture.

### ADDITIONAL REFERENCES

- BERTRAND, J. (1907). *Calcul des Probabilités*, 2nd ed. Gauthier-Villars, Paris.
- PRATT, J. W., RAIFFA, H. and SCHLAIFER, R. (1964). The foundations of decision under uncertainty: An elementary exposition. *J. Amer. Statist. Assoc.* **59** 353-375.
- RAMSEY, F. P. (1931). *The Foundations of Mathematics and Other Logical Essays*. Routledge and Kegan Paul, London.
- SHAFER, G. (1976). *A Mathematical Theory of Evidence*. Princeton Univ. Press.