18 Letter from Robert Aumann to Leonard Savage and Letter from Leonard Savage to Robert Aumann

Letter from Robert Aumann to Leonard Savage, 8 January 1971

There is a conceptual question regarding subjective probabilities that has been puzzling me, about which I would like to consult you.

Consider the following two acts:

Act A : You get an umbrella if it rains, nothing if it does not rain.

Act B : You get an umbrella in either case.

Suppose your utility for an umbrella is 1, and for no umbrella is 0; suppose further that your subjective probability for rain is 1/2. Then acts A and B have utilities 1/2 and 1 respectively. On the other hand, I don't think it would be unreasonable for you to be indifferent between the two acts, since an umbrella is useless in fine weather.

Obviously, the answer is that your utility for umbrellas depends on the weather, i.e. on the state of the world. But that leads rather quickly to the conclusion that your postulate P3 is unreasonable; for example, all in all, I prefer an umbrella to a nickel, but if it does not rain, I prefer the nickel.

The only conclusion seems to be that it is improper to call an umbrella a "consequence." An umbrella is really an act; the consequence is getting wet or not. But that does not quite get us out of trouble. Most people prefer to get wet in fine weather; the fact is, most swimming pools are poorly attended when it rains. So even the utility of getting wet is state-dependent. The appropriate answer then seems to be: "Getting wet is also an act; the consequence is getting wet in the rain or getting wet in the sunshine." But you will agree that that isn't very satisfactory; we have now made the *description* of the consequence state-dependent. By the same token, sunshine and rain themselves could be called "consequences," and then one could construct nonsensical acts such as "You get sunshine if it rains, and rain otherwise."

It seems that the notions of "state," "act," and "consequence" have rather fuzzy interpretations; in particular, it is not always easy conceptually to distinguish between them. But to make sense of the axioms, it is essential to have a fairly sharp idea of what these notions mean.

The main question that is puzzling me is more basic, though; it does not concern your derivation of the subjective probability notion, but the very possibility of defining this notion—in any way—via preferences. Suppose Mr X loves his wife very much, he feels that if he should lose

This chapter originally appeared in *Essays on Economic Decisions under Uncertainty*, by J. H. Dreze, pp. 76–78 and 78–81, Cambridge University Press, Cambridge, 1987. Reprinted with permission.

her, life would be somehow less interesting, less attractive—less "worth living." His wife falls ill, and it is decided that if she is to survive, she must undergo an operation. This operation is well-known in medicine, one might even say routine; but it is very dangerous. In fact, 1/2 of the patients die on the operating table, whereas the remaining 1/2 survive the operation and then are entirely cured. (Of course the whole scenario is oversimplified, but I think it captures the essentials of some very important real-life situations.) Now the man is asked whether he would rather bet $100 on his wife's survival, or on heads in a coin-toss. I think he would not be unreasonable strongly to prefer the bet on his wife's survival. If she should not survive, the $100 is somehow worth less; and if he bets on the coin toss, he might get the $100 in a situation in which he would not be able to enjoy it. Nevertheless, Mr X might well agree with the medical information with which he is supplied, and estimate his "personal probability" for her survival—whatever that may mean—at 1/2. The point of the example, of course, is that in this situation there is *nothing* that one could truly call a 'consequence' in the sense in which I think you meant it, i.e. something whose 'value' is state-independent.

Let's change the scene. Professor Y is a fifty-year-old nineteenth-century physicist, whose life-work is strongly based on the notion of "ether" (again the scenario is oversimplified). He hears of the Michelson–Morley experiment, is very upset, and decides to repeat it. Now what about his a priori? I think it's fair to say that Y's feelings toward the ether are at least qualitatively comparable to X's feelings towards his wife. Can we at all *define* the a priori?

Lester Dubins was here for a few days, and we discussed this matter. He said he had heard this question before, and that you probably had an appropriate answer.

Finally, I'd like to mention that there is no particular difficulty in extending utility theory to this kind of situation. For each state, one can define the conditional utility of a consequence (or an act) *given the state*, and one can also define numbers that behave mathematically like subjective probabilities of the states. The trouble is that utilities that are conditional on disjoint states can be normalized independently, as long as one adjusts appropriately the so-called "subjective probabilities." Thus mathematically, the subjective probabilities and the utilities get all mixed up, and cannot be separated from each other. Conceptually, of course, that is exactly what happens also.

Since Frank Anscombe and I once did some joint work on subjective probabilities, I am taking the liberty of sending him a copy of this letter. I hope very much that you will be able to clear me up on this.

Letter from Leonard Savage to Robert Aumann, 27 January 1971

Thank you for your letter of 8 January. It is the sort of letter that one is tempted to postpone answering until there is time to reflect and prepare a thorough answer. For me, this temptation often results in letters altogether unanswered, so let me say something promptly and perhaps return to the theme later, especially if you raise new questions.

The difficulties that you mention are all there; I have known about them in a confused way for a long time; I believe they are serious but am prepared to live with them until something better comes along. The theory of personal probability and utility is, as I see it, a sort of framework into which I hope to fit a large class of decision problems. In this process, a certain amount of pushing, pulling, and departure from common sense may be acceptable and even advisable. There are minds that think it absurd to accept zero as a number and the null set as a set; it seems idle to say that these minds are objectively wrong, but you and I prefer the other way and have come to think it natural. To some—perhaps to you—it will seem grotesque if I say that I should not mind being hung so long as it be done without damage to my health or reputation, but I think it desirable to adopt such language so that the danger of being hung can be contemplated in the framework of F. of S. An extremely able discussion of such pushing and pulling is Quine's famous essay, "Two Dogmas of Empiricism," *Philosophical Review*, vol. 60, pp. 20–43 (reprinted in *From a Logical Point of View*, Harvard University Press, Cambridge, 1953).

Let me point out some passages in F. of S. that speak a little to the problems raised in your letter, though of course they cannot really put the problems to rest. Perhaps the first complete paragraph on page 15 is relevant. On page 14, the sentence, "Consequences might appropriately be called states of the person, as opposed to states of the world," seems suggestive. The lower half of page 25 describes an example much like some of yours. Finally, Section 5.5 of F. of S., "Small Worlds," is about philosophical problems that seem to be close to, or at any rate entangled with, yours.

Now let us see what specific comments the examples in your letter may suggest. On the first page, you take an example, at first it is carelessly described, and subjected to increasingly more careful descriptions, much as I would do myself. The very last line of the page is telling. I would regard it as fanciful but not as nonsense to say, "You experience sunshine if it rains, and rain otherwise." In this, I have changed your "get" into "experience" to emphasize my notion that a consequence is in the last

analysis an experience. The insistence that consequences are experiences or sensations does not of course sweep all problems away. These terms may be suggestive, but they might defy definition. Also, the appreciation of uncertainty is itself an important aspect of sensation, and that seems to contradict the notion of a sure consequence.

This returns to what was hinted at early in this letter. I believe, and examples have confirmed, that decision situations can be usefully structured in terms of consequences, states, and acts in such a way that the postulates of F. of S. are satisfied. Just how to do that seems to be an art for which I can give no prescription and for which it is perhaps unreasonable to expect one—as we know from other postulate systems for application. Thus to paraphrase the middle of your first page, I would be glad to pay a nickel to rent an umbrella for a fall football match but given that it will not rain, I would prefer the nickel. I analyze this in terms of several consequences: the status quo, being miserably drenched, and being undrenched but out a nickel. These "consequences" seem to enable me to describe the situation in terms of, and consistent with, the postulates in F. of S. Of course, they are not ultimate. A nickel is itself a lottery ticket, and one objection to getting miserably drenched is that it seems conducive to illness. If the problem were concerned with illness or the possibility of accidentally buying poisoned food, then of course the notion of consequence would have to be further analyzed. An ultimate analysis might seem desirable, but probably it does not exist and certainly threatens to be cumbersome.

The first paragraph on the second page of your letter seems excellent to me. The terms are indeed "fuzzy" and it is indeed "essential to have a fairly sharp idea of what these notions mean." My own notion seems to be fairly sharp in that I seem to be able to couch decision problems in terms of them, and while these formulations involve various choices, tempting choices do not seem to lead to different practical conclusions. It seems something like the following familiar phenomenon. We usually couch probability problems in terms of the Kolmogorov theory and in particular in terms of atomic, or unsubdividable, events; these are the points of the probability space. But in practice, any event can be further subdivided by flipping still another coin. Yet we feel, and find, that there is no harm in this ambiguity. I do not mean to insist that the ambiguities of "state," "act," and "consequence" are that innocuous but only to remind you of a certain kind of floating flexibility that we expect in the formulation of applied problems.

Let us try to reflect on the medical example in the middle of your page 2. It is quite usual in this theory to contemplate acts that are not actually available. These serve something like construction lines in geometry. A

typical decision theoretic argument runs, "If B were available, I would clearly prefer A to B and B to C, therefore, my momentary impression that C is more attractive than A will not bear inspection."

In particular, I can contemplate the possibility that the lady dies medically and yet is restored in good health to her husband. Put a little differently, I can ask Mr. Smith how he would bet on the operation if the continuance of his family life were not dependent on its outcome. Make believe is certainly involved, and indeed it is extremely difficult to make believe to the required extent. Yet, it does seem to be a helpful goal. Incidentally, it would not be nonsensical, though unmannerly, for the experimenter to guarantee to execute Mrs. Smith if she recovers from the operation. And I see no real objection to Mr. Smith imagining this cruel situation if it helps him appraise his own probabilities.

Another line of thought to which it might be well to return more thoroughly is this. By betting both on and off the recovery, Mr. Smith's hedging could be detected and perhaps measured and corrected for. It would be good to know how much can and cannot be done thus in principle.

Yes, a bet on the survival of one's favorite theory does seem somewhat like a bet on the survival of one's wife. If Professor Y wants to know his own personal probability for the event he has a very severe problem in detaching himself. Whether it is in principle different from any other such introspection such as telling himself the price at which he would sell his car is hard for me to decide. Any statement of the form, "I would do this if that," is somewhat mysterious philosophically; are the ones you emphasize so much worse than the others? When you underline the word "define" I understand that you want to try to be operational in eliciting this chap's probability that his favorite theory is true and you find it difficult to imagine any bet on the issue that would not encourage hedging. I can imagine some procedures that have the defect of being extremely expensive and possibly extremely cruel but that might yet have the merit of showing certain things to be possible in principle, as in disposing of Mrs. Smith. We might be able to find some way to so blacken Mr. Y's reputation, painting him as a plagiarist and a fool, that he no longer cares whether his favorite theory is true or not and can bet on it dispassionately.

One of the problems raised about a Professor Y is that of his optimism or pessimism and what it actually means for him to combat those tendencies in himself. I have some tendency to confuse that with your riddle which is really a different one, namely that of insulating Mr. Y from consequences normally associated with the events about which we want his opinion.

I suppose that the mixing up of things that you mention on page 3 is the same as, or closely related to, a mixing up of things mentioned in my section on small worlds but have not checked closely.

As promised this reply is prompt, I am sorry that it is not more satisfactory. There is certainly much in what you say. What is not clear to me is what, if anything, had best be done about it. A person who has published much in the spirit of your example is Jacques Drèze. One reference of his of which I have record is: "Fondements logiques de la probabilité subjective et de l'utilité," pp. 73–87 in *La Décision*, 1961, Paris, Centre National de la Recherche Scientifique.

I shall give Frank a copy of this letter and shall be as interested as you to hear his reactions to these matters.

Under separate cover are a few reprints and preprints of possible interest including a little reading note about the Michelson–Morley experiment. Michelson himself, incidentally, fully believed in the ether–drift effect even after his most refined and successful experiment.