



[Risk, Ambiguity, and the Savage Axioms]:Reply

Daniel Ellsberg

The Quarterly Journal of Economics, Vol. 77, No. 2. (May, 1963), pp. 336-342.

Stable URL:

<http://links.jstor.org/sici?sici=0033-5533%28196305%2977%3A2%3C336%3A%5BAATSA%3E2.0.CO%3B2-S>

The Quarterly Journal of Economics is currently published by The MIT Press.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/mitpress.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

Ellsberg's paradox serves as a useful example for showing that one's first intuitive judgments may be analyzed and illuminated by the very approach that they seem at first to threaten. I do not contend that all subjects would hold to the Savage axioms as a result of such analysis, but I believe, partly on the basis of informal experimentation, that many of them would.

In any event the remedy for vagueness is an honest attempt to recognize genuine vagueness, to deal with it directly if possible, or to bypass it skillfully by less formal and complete analyses. My personal opinion is that the problem of vagueness will be most successfully met in situations in which at least part of the information comes from sample data, that is, numbers generated by such processes as the Bernoulli, Poisson, or normal. When sample data are absent or when vagueness threatens our attempts to assess the sampling process, the role of formal analysis may have great conceptual value, for example, in disentangling the *probabilities* of events from the *utilities* of the consequences of events. The formal approach may hint at good informal ones, as when we graph data in ways that cast light on such assumptions as independence or normality. But the literal application of formal methods is likely to be much more restricted. Even so, the normative value of Bayesian decision theory can be great. If we cannot always eliminate vagueness about the answers, we never need be vague about the right questions to ask.

UNIVERSITY OF CHICAGO

REPLY

DANIEL ELLSBERG

There is so much in these matters on which Professor Roberts and I agree that, if a few summary sentences were slightly different, I would be tempted simply to thank him for supplementing my earlier discussion and underlining certain aspects of it. But that would be rather tactlessly to ignore Roberts' own interpretation of his remarks, which seems quite otherwise. If I am to take issue with him, it must be not with his treatment of my position — as criticism, his comments are admirably restrained, fair-minded and cautious — but with his understanding of the thrust of his own specific comments.

A major point on which there seems no disagreement is the fact and the general pattern of violations of the Savage axioms in connection with my hypothetical choice problems. Of course, in testing the acceptability of normative postulates, hasty, undeliberated choices are in no way conclusive and violation must be considered tentative¹ pending thorough analysis and reflection in the light of all implications of the theory.² "I do not contend," Roberts concludes, "that all subjects would hold to the Savage axioms as a result of such analysis, but I believe, partly on the basis of informal experimentation, that many of them would." That conclusion could summarize my own observations, though I might reverse the emphasis. As I reported in my article,³ responses do vary. Of those who tentatively violate the postulates, some conclude upon reflection or interrogation by unsympathetic critics (sometimes, themselves) that their initial choices were "mistakes" which they now wish to modify; others do not. For further reference, I shall call the former group "transient" violators and the latter, "deliberate" violators.⁴ Roberts' introspective experience with his own violations places him, so far, in the first category; I remain, so far, in the second. But though our own reactions fail to match, I gather that our observations of response varieties have been similar. Our notions of the relative frequency of "deliberate" violation may differ somewhat, and Roberts is more concerned to lower it, if possible. But he does not claim that the violations of his subjects have all so far, like his own, proved transient (I presume he would not fail to mention it if they had), nor does he conjecture that with adequate reflection by all concerned the class of deliberate violators will prove empty. It is only this residual group of deliberate violators who raise prob-

1. Likewise conformity, where the appropriateness of the theory is in question.

2. See "Risk, Ambiguity, and the Savage Axioms," this *Journal*, LXXV (Nov. 1961), 655-56, for comments and admonitions paralleling Roberts'. For a much fuller discussion of the nature, validation and functions of normative theories of choice, see "Risk, Ambiguity and Decision," The RAND Corporation, RM-3543 (forthcoming); again, I think we are in complete agreement on these questions.

3. *Op. cit.*

4. This distinction can be made hollow, to whatever end that serves, by taking the period of "adequate" reflection to be sufficiently short or sufficiently long. For practical purposes, those violators who conclude that they wish to persist in their violating choices after conscientious consideration of such critiques as those of Roberts and of H. Raiffa ("Risk, Ambiguity, and the Savage Axioms: Comment," this *Journal*, LXXV, Nov. 1961, 690-94) may be classed as "deliberate." Although such persons may always change their minds eventually, the Savage axioms do not constitute an appropriate, usable model of their actual, deliberated preferences *now*; if they are to benefit from a normative logic of choice, it must be somewhat different.

lems of normative theory, and my "diagnosis" (that "ambiguity" was a major contributing factor) concerned them exclusively.

Roberts defines "vagueness" in terms that correspond to my notions of "ambiguity" well enough for this particular discussion:

A person is vague about the probability assigned to a single trial if he cannot obtain from himself a clear answer as to what probability to assign to it. He is vague about a probability distribution if introspection fails to reveal clearly what the distribution is . . . we are always more or less vague.⁵

I presume that Roberts agrees that my hypothetical examples tend to induce considerable vagueness of opinion concerning certain alternatives (his subjects would tell him, if his own introspection did not).⁶ Moreover, Roberts emphasizes that the problem of vagueness is "real and important." What he denies is that the evident vagueness of opinion in these instances contributes in any important way to the violations of the Savage postulates that admittedly do occur.⁷

5 Though Roberts furnishes no formal model for this state of mind, his definition immediately suggests one: the whole *set* of probability distributions that are not *ruled out* by those probability comparisons of which the individual feels relatively "sure." This assumes that introspection does reveal a relatively clear answer with respect to *some* probability comparisons; typically, these judgments are in the form of inequalities, statements that the probability of one event is greater (or, not less) than that of another. To say that "we are always more or less vague" may be interpreted to mean that the set of probability distributions compatible with (not excluded by) all our definite probability judgments at a given moment typically contains more than one member. In general, this model would associate a particular event not with a precise probability but with an *interval* of probability-numbers circumscribed by inequalities, and the relation "not less probable than" is regarded as providing only a *partial ordering* among events.

The distribution-set is denoted Y° in "Risk, Ambiguity, and the Savage Axioms" (p 661). At the time of writing that article, I was unaware that this model of opinion, encompassing "vague" opinion and "ambiguity," had been elaborately developed in earlier, important works by B O Koopman ("The Axioms and Algebra of Intuitive Probability," *Annals of Mathematics*, Series 2, Vol. 41, 1940, pp. 269-92) and I J Good ("Rational Decisions," *Journal of the Royal Statistical Society*, Series B, Vol. 14, 1952, pp. 107-14; more recently, "Subjective Probability as the Measure of a Non-Measurable Set," *Proceedings of the International Congress for Logic, Methodology and Philosophy of Science*, Stanford University Press, 1962, pp 319-29). These works are further discussed in "Risk, Ambiguity and Decision."

6. In the specific example he discusses, the subjects are given no explicit information on the ratio of red to black balls in Urn I; in practice, subjects readily report great vagueness of opinion concerning their prospects of winning bets on the color of a ball to be drawn from Urn I (bets on Red₁ or Black₁), in contrast to their precise opinions on the probabilities of winning corresponding bets on Urn II, in which the proportion of red to black is known to be 50 50. In terms of the model in the preceding footnote, the opinion of such a subject concerning a drawing from Urn I (*or*, concerning the *ratio* of red to black in Urn I: an uncertain fact on which the prize does not depend directly but which underlies the subject's evaluation of the bet) must be represented by a large set of distributions, while his opinion concerning a drawing from Urn II may be represented by a single, precise distribution.

7. Why, then, is it "important"? How does vagueness affect decision? I

He concludes: "It is hard to see any important role for vagueness in Ellsberg's paradox, at least for a person who makes definite choices for all four questions." The challenge seems clear enough. But on a closer look, the conflict blurs.

The heart of Roberts' specific critique is a list of considerations that could induce a subject to make choices in the "two-urn example" cited above in a pattern that evidently violates the Savage postulates: specifically, Postulate 2, the Sure-thing Principle. These considerations include possible mistaken beliefs as to the explicit conditions of the bets: for example, the facts that the game is to be played only once, that a single drawing is to be made from the selected urn, and that the money payoff depends only on the color of the ball drawn, not on the proportion of colors in the urn.

Roberts' discussion both before and after this list is presented tends to suggest that *all* the considerations to be mentioned are of the character of the three above: "mistakes, misinterpretations, misconceptions" that may "confuse" a subject but which, once brought into consciousness and made explicit, will not induce him to persist in violations when the problem is properly understood. Moreover, his conclusion implies that vagueness plays no part in these considerations.

But neither of these characterizations of the proposed rationales seems accurate; and even if they were, Roberts' argument would not sustain his conclusion.

The first characterization, if valid, would imply that he has simply compiled a list of various sources of *transient* violations: an believe that Roberts would answer: It affects the difficulty of decision-making, the time and effort required, the pleasantness of the task and one's confidence or uneasiness in the results, the frequency of random errors and "brief" transient violations; but *not*, given enough time for reflection and analysis, the answers one will ultimately choose, hence *not* the acceptability of Savage's normative postulates. In short, vagueness of opinions affects "feelings" more than decisions, promotes indecisiveness and vacillation, and affects decision-making in the same sorts of ways as does *complexity*.

My own view is that in addition to these effects, and partly because of them, vagueness of opinion can affect the choices that seem preferable on thorough reflection, and for some people its influence does lead to deliberate conflicts with the Savage postulates. Vagueness need not be synonymous with indecision, as Roberts implies; *some choices are easy, just because certain opinions are vague* (for many people this is true for the choice between a bet on Red_I and a bet on Red_{II}, i.e., for questions 3 and 4 cited by Roberts). Nor must such people turn to informal analysis, if a more appropriate formal theory can be made available. I think it can.

I proposed *one* candidate in my article and discuss several others, including some I would favor (omitted, regrettably, from the article for reasons of space) in "Risk, Ambiguity and Decision." Most of this "therapy" was antedated by I. J. Good in his remarkable paper, "Rational Decisions" (cited above). For a recent parallel, see C. A. B. Smith, "Consistency in Statistical Inference and Decision," *Journal of the Royal Statistical Society, Series B*, Vol. 23 (1961), pp. 1-25.

interesting and useful effort, but not addressed to *my* problem. Used in self-interrogation, such a list (in which vagueness would surely be but one factor among many) would be helpful in sifting out the transient from the deliberate violators more efficiently. But unless one conjectured that such a sieve would, in practice, show the latter, residual class to be null — and Roberts does not press his argument this far — it would seem to provide, at best, imperfect insight into the “paradox” (in Roberts’ eyes) of persistent, deliberate violation.

What *are* the considerations, in Roberts’ view, that influence the deliberate violator, he whose transgressions of the Savage postulates must be adjudged neither lighthearted, irresponsible, nor unwise? Roberts does not dispute his existence, yet he ventures no explanation; nor does he directly attack the one I propose. In fact, he seems to be silent on the matter. But this is the precise point at issue, if my “diagnosis” — which concerned *only* these residual subjects — is in question!

The assertion that the diagnosis of vagueness has been shown to be “wrong” or irrelevant appears even more puzzling when one considers the actual substance of the points Roberts raises in his critique. In the six numbered passages and two of the footnotes (footnote 9 on p. 333 and footnote 4 on p. 336) toward the end of Section III, I count nine distinct “resolutions” of the pattern of violations in question. In *six* of these, implicitly or explicitly, *vagueness seems to play a critique role!* I shall comment on the more important of these (the reader is referred to Roberts’ text for the exact import of the propositions cited):

(a) “The subject may have had a much harder time in answering question 1 than question 2.” Really? Why? except that he finds his opinions more *vague*, with respect to question 1! (Question 1 concerns a drawing from Urn I, for which the ratio of red to black has not been specified.) Why else would he “certainly not have an easy mind” about any given resolution of the question; why else might he expect himself to be “wildly volatile in his choices” if he had to make the assessment repeatedly? (Either of these expectations could serve as a fairly adequate working definition of “vagueness” of opinion; but here and below, I rely on Roberts’ definition cited earlier.) As Roberts suggests, “it is easy to carry over this feeling” into a preference for bets based on Urn II (in which the ratio is known to be precisely 50:50). But this is a response to relative *vagueness!* And by the way . . . in what sense does it reflect “mistakes,” confusions, misinterpretation of the prob-

lem? Can, or should, an application of the Savage postulates abstract from such "difficulties" that issue from the very quality of the uncertainties in question?

(b) "The subject may realize he is vague, in Savage's sense, about the distribution" of the *ratio* of Red to Black in Urn I, and fail "to realize that this distribution is irrelevant to either question 3 or 4." But vagueness about the distribution is not conducive to precision or confidence about its mathematical expectation, which is relevant to questions 3 and 4. Roberts mentions earlier that "the answer given by the subject to question 1 suggests that, for him," this expectation = .5; but the *other* answers postulated for him, to questions 3 and 4, are *inconsistent* with this interpretation, as they are with the inference of a single, definite probability distribution over the possible ratios. As Roberts says, the subject "could" claim a precise opinion on the expectation despite vagueness as to the distribution (somewhat implausibly, unless he avows compelling intuitions of symmetry that are not obviously appropriate here); but he need not do so, merely to justify his answer to question 1.⁸ In any case, the role of vagueness is explicit here, whether or not a "mistake" is involved.

(c) *Why* might the subject "feel that his choice of Red_I could lead to unpleasant second guesses by someone who observed the experiment"? Why might he be criticized "for not taking an apparently 'safe' course of action (Red_{II}); is he lost by taking an 'unsafe' one (Red_I)"; on what grounds can bets on Red_I and Red_{II} be discriminated by a potential critic save relative vagueness of accompanying opinions? Indeed, on what other basis can the terms "safe" and "unsafe," as applied here, be interpreted meaningfully? (The very fact that these terms do seem apt in this context deserves some serious thought from Roberts; these examples were partly constructed just to elicit such notions without offering a basis for them in terms of the range, minimum payoff, variance or expected value of a specified distribution.) Even if, as Roberts postulates, the subject's own opinions happen to be precise, it is vagueness — in this case, the anticipated perception and evaluation of vagueness by *others* — that determines his hypothesized response.

(d) "In the same vein, the subject may fear that Urn I might contribute to an ulcer." Quite: but why Urn I? Obviously, because he sees the vagueness that we see, and that he might expect others

8. The *set* of distributions representing his vague opinions concerning Urn I may have certain symmetry features (by containing matched pairs of *asymmetric* distributions) that account for his indifference between bets on Red_I and Black_I; no precise "expectation" need attach to the set.

to see. Even if we abstracted, experimentally, from the problem of anticipated second guesses by others by keeping his choices private, we could not protect him from no less unpleasant second-guessing by himself; from self-reproaches, from *regret* (in a familiar, not a technical, sense that would repay analysis) evoked by losses on Urn I, for reasons intimately associated with vagueness.

Once again, are most of these considerations based on mistakes, confusions, misconceptions? Is the subject wrong to expect the epithet "unsafe" to be attached to Red_I and not to Red_{II} ; and is he foolish to take that into account? Is the ulcer-prone individual mistaken, or arbitrary, to see Urn I as the more threatening to him? And is a subject likely to change his choices when such considerations underlying them are made fully explicit? Somewhat ironically, it is just because these factors do *not* reflect misinterpretations that they may, after all, help explain some *deliberate* violations.

This is not to say that vagueness, as defined, is typically the sole factor underlying deliberate choices in conflict with the Savage postulates, even in the situations I described, or that such choices reflect mainly a simple aversion to vagueness (though my article may have given those impressions). My own thinking has moved recently toward recognizing the influence of various dimensions of the decision problem under uncertainty that are strongly associated with vagueness but distinct from it; several of Roberts' remarks are highly pertinent and stimulating along these lines.

Nevertheless, as I indicated at the outset, the careful reader of Roberts' catalogue of rationales for violation of the postulates may well find his appreciation of vagueness as a *contributing* factor enhanced rather than diminished. In fact, Roberts' summary remark, "It is hard to see any important role for vagueness . . .," seems to me to make a distinctly odd impression following immediately, as it does, his discussion in Section III. That section ends with a warning which I second, but which also seems cogent in slightly paraphrased form: Conformity to the Savage axioms in spite of vagueness should not be contemplated lightheartedly if the decisions or inferences involved are taken seriously.