



OXFORD JOURNALS
OXFORD UNIVERSITY PRESS

Risk, Ambiguity, and the Savage Axioms

Author(s): Daniel Ellsberg

Source: *The Quarterly Journal of Economics*, Nov., 1961, Vol. 75, No. 4 (Nov., 1961), pp. 643-669

Published by: Oxford University Press

Stable URL: <https://www.jstor.org/stable/1884324>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



JSTOR

Oxford University Press is collaborating with JSTOR to digitize, preserve and extend access to *The Quarterly Journal of Economics*

RISK, AMBIGUITY, AND THE SAVAGE AXIOMS*

By DANIEL ELLSBERG

I. Are there uncertainties that are not risks? 643. — II. Uncertainties that are not risks, 647. — III. Why are some uncertainties not risks? — 656.

I. ARE THERE UNCERTAINTIES THAT ARE NOT RISKS?

There has always been a good deal of skepticism about the behavioral significance of Frank Knight's distinction between "measurable uncertainty" or "risk," which may be represented by numerical probabilities, and "unmeasurable uncertainty" which cannot. Knight maintained that the latter "uncertainty" prevailed — and hence that numerical probabilities were inapplicable — in situations when the decision-maker was ignorant of the statistical frequencies of events relevant to his decision; or when a priori calculations were impossible; or when the relevant events were in some sense unique; or when an important, once-and-for-all decision was concerned.¹

Yet the feeling has persisted that, even in these situations, people tend to *behave* "as though" they assigned numerical probabilities, or "degrees of belief," to the events impinging on their actions. However, it is hard either to confirm or to deny such a proposition in the absence of precisely-defined procedures for *measuring* these alleged "degrees of belief."

What might it mean operationally, in terms of refutable predictions about observable phenomena, to say that someone behaves "as if" he assigned quantitative likelihoods to events: or to say that he does not? An intuitive answer may emerge if we consider an example proposed by Shackle, who takes an extreme form of the Knightian

* Research for this paper was done as a member of the Society of Fellows, Harvard University, 1957. It was delivered in essentially its present form, except for Section III, at the December meetings of the Econometric Society, St. Louis, 1960. In the recent revision of Section III, I have been particularly stimulated by discussions with A. Madansky, T. Schelling, L. Shapley and S. Winter.

1. F. H. Knight, *Risk, Uncertainty and Profit* (Boston: Houghton Mifflin, 1921). But see Arrow's comment: "In brief, Knight's uncertainties seem to have surprisingly many of the properties of ordinary probabilities, and it is not clear how much is gained by the distinction. . . . Actually, his uncertainties produce about the same reactions in individuals as other writers ascribe to risks." K. J. Arrow, "Alternative Approaches to the Theory of Choice in Risk-taking Situations," *Econometrica*, Vol. 19 (Oct. 1951), pp. 417, 426.

position that statistical information on frequencies within a large, repetitive class of events is strictly irrelevant to a decision whose outcome depends on a single trial. Shackle not only rejects numerical probabilities for representing the uncertainty in this situation; he maintains that in situations where all the potential outcomes seem "perfectly possible" in the sense that they would not violate accepted laws and thus cause "surprise," it is impossible to distinguish meaningfully (i.e., in terms of a person's behavior, or any other observations) between the relative "likelihoods" of these outcomes. In throwing a die, for instance, it would not surprise us at all if an ace came up on a single trial, nor if, on the other hand, some other number came up. So Shackle concludes:

Suppose the captains in a Test Match have agreed that instead of tossing a coin for a choice of innings they will decide the matter by this next throw of a die, and that if it shows an ace Australia shall bat first, if any other number, then England shall bat first. Can we now give any meaningful answer whatever to the question, "Who will bat first?" except "We do not know?"²

Most of us might think we could give better answers than that. We could say, "England will bat first," or more cautiously: "I think England will probably bat first." And if Shackle challenges us as to what we "mean" by that statement, it is quite natural to reply: "We'll bet on England; and we'll give you good odds."

It so happens that in this case statistical information (on the behavior of dice) is available and does seem relevant even to a "single shot" decision, our bet; it will affect the odds we offer. As Damon Runyon once said, "The race is not always to the swift nor the battle to the strong, but that's the way to bet." However, it is our bet itself, and not the reasoning and evidence that lies behind it, that gives operational meaning to our statement that we find one outcome "more likely" than another. And we may be willing to place bets — thus revealing "degrees of belief" in a quantitative form — about events for which there is no statistical information at all, or regarding which statistical information seems in principle unobtainable. If our pattern of bets were suitably orderly — if it

2. G. L. S. Shackle, *Uncertainty in Economics* (London: Cambridge University Press, 1955), p. 8. If this example were not typical of a number of Shackle's works, it would seem almost unfair to cite it, since it appears so transparently inconsistent with commonly-observed behavior. Can Shackle really believe that an Australian captain who cared about batting first would be *indifferent* between staking this outcome on "heads" or on an ace?

satisfied certain postulated constraints — it would be possible to infer for ourselves numerical subjective probabilities for events, in terms of which some future decisions could be predicted or described. Thus a good deal — perhaps all — of Knight's class of "unmeasurable uncertainties" would have succumbed to measurement, and "risk" would prevail instead of "uncertainty."

A number of sets of constraints on choice-behavior under uncertainty have now been proposed, all more or less equivalent or closely similar in spirit, having the implication that — for a "rational" man — *all* uncertainties can be reduced to *risks*.³ Their flavor is suggested by Ramsay's early notions that, "The degree of a belief is . . . the extent to which we are prepared to act upon it," and "The probability of 1/3 is clearly related to the kind of belief which would lead to a bet of 2 to 1."⁴ Starting from the notion that gambling choices are influenced by, or "reflect," differing degrees of belief, this approach sets out to *infer* those beliefs from the actual choices. Of course, in general those choices reveal not only the person's relative expectations but his relative preferences for outcomes; there is a problem of distinguishing between these. But if one picks the right choices to observe, and if the Savage postulates or some equivalent set are found to be satisfied, this distinction can be made unambiguously, and either qualitative or, ideally, numerical probabilities can be determined. The propounders of these axioms tend to be hopeful that the rules *will* be commonly satisfied, at least roughly and most of the time, because they regard these postulates as normative maxims, widely-acceptable principles of rational behavior. In other words, people should tend to behave in the postulated fashion, because that is the way they would *want* to behave. At the least, these axioms

3. F. P. Ramsey, "Truth and Probability" (1926) in *The Foundations of Mathematics and Other Logical Essays*, ed. R. B. Braithwaite (New York: Harcourt Brace, 1931); L. J. Savage, *The Foundations of Statistics* (New York: Wiley, 1954); B. de Finetti, "Recent Suggestions for the Reconciliation of Theories of Probability," pp. 217-26 of *Proceedings of the Second (1950) Berkeley Symposium on Mathematical Statistics and Probability*, Berkeley, 1951; P. Suppes, D. Davidson, and S. Siegel, *Decision-Making* (Stanford University Press, 1957). Closely related approaches, in which individual choice behavior is presumed to be stochastic, have been developed by R. D. Luce, *Individual Choice Behavior* (New York: Wiley, 1959), and J. S. Chipman, "Stochastic Choice and Subjective Probability," in *Decisions, Values and Groups*, ed. D. Willner (New York: Pergamon Press, 1960). Although the argument in this paper applies equally well to these latter stochastic axiom systems, they will not be discussed explicitly.

4. Ramsey, *op. cit.*, p. 171.

are believed to predict certain choices that people will make when they take plenty of time to reflect over their decision, in the light of the postulates.

In considering only deliberate decisions, then, does this leave any room at all for "unmeasurable uncertainty": for uncertainties *not* reducible to "risks," to quantitative or qualitative probabilities?

A side effect of the axiomatic approach is that it supplies, at last (as Knight did not), a useful operational meaning to the proposition that people do *not* always assign, or act "as though" they assigned, probabilities to uncertain events. The meaning would be that with respect to certain events they did not obey, nor did they wish to obey — *even on reflection* — Savage's postulates or equivalent rules. One could emphasize here either that the postulates failed to be acceptable in those circumstances as normative rules, or that they failed to predict reflective choices; I tend to be more interested in the latter aspect, Savage no doubt in the former. (A third inference, which H. Raiffa favors, could be that people need more drill on the importance of conforming to the Savage axioms.) But from either point of view, it would follow that *there would be simply no way to infer meaningful probabilities for those events from their choices*, and theories which purported to describe their uncertainty in terms of probabilities would be quite inapplicable in that area (unless quite different operations for measuring probability were devised). Moreover, such people could not be described as maximizing the mathematical expectation of utility on the basis of numerical probabilities for those events derived on *any* basis. Nor would it be possible to derive numerical "von Neumann-Morgenstern" utilities from their choices among gambles involving those events.

I propose to indicate a class of choice-situations in which many otherwise reasonable people neither wish nor tend to conform to the Savage postulates, nor to the other axiom sets that have been devised. But the implications of such a finding, if true, are not wholly destructive. First, both the predictive and normative use of the Savage or equivalent postulates might be improved by avoiding attempts to apply them in certain, specifiable circumstances where they do not seem acceptable. Second, we might hope that it is precisely in such circumstances that certain proposals for alternative decision rules and nonprobabilistic descriptions of uncertainty (e.g., by Knight, Shackle, Hurwicz, and Hodges and Lehmann) might prove fruitful. I believe, in fact, that this is the case.

II. UNCERTAINTIES THAT ARE NOT RISKS

Which of two events, α , β , does an individual consider "more likely"? In the Ramsey-Savage approach, the basic test is: *On which event would he prefer to stake a prize, or to place a given bet?* By the phrase, "to offer a bet *on* α " we shall mean: to make available an action with consequence a if α occurs (or, as Savage puts it, if α "obtains") and b if α does not occur (i.e., if $\bar{\alpha}$, or "not- α " occurs), where a is preferable to b .

Suppose, then, that we offer a subject alternative bets "on" α and "on" β (α , β need not be either mutually exclusive or exhaustive, but for convenience we shall assume in all illustrations that they are mutually exclusive).

		Events		
		α	β	$\bar{\alpha} \cap \bar{\beta}$
Gambles	I	a	b	b
	II	b	a	b

The Ramsey-Savage proposal is to interpret the person's preference between I and II as revealing the relative likelihood he assigns to α and β . If he does not definitely prefer II to I, it is to be inferred that he regards α as "not less probable than" β , which we will write: $\alpha \geq \beta$.

For example, in the case of Shackle's illustration, we might be allowed to bet either that England will bat first or that Australia will (these two events being complementary), staking a \$10 prize in either case:

		England first	Australia first
	I	\$10	\$0
	II	\$0	\$10

If the event were to be determined by the toss of a die, England to bat first if any number but an ace turned up, I would strongly prefer gamble I (and if Shackle should really claim indifference between I and II, I would be anxious to make a side bet with him). If, on the other hand, the captains were to toss a coin, I would be indifferent between the two bets. In the first case an observer might infer, on the basis of the Ramsey-Savage axioms, that I regarded England as

more likely to bat first than Australia (or, an ace as less likely than not to come up); in the second case, that I regarded heads and tails as "equally likely."

That inference would, in fact, be a little hasty. My indifference in the second case would indeed indicate that I assigned equal probabilities to heads and tails, *if I assigned any probabilities at all to those events*; but the latter condition would remain to be proved, and it would take further choices to prove it. I might, for example, be a "minimaxer," whose indifference between the two bets merely reflected the fact that their respective "worst outcomes" were identical. To rule out such possibilities, it would be necessary to examine my pattern of preferences in a number of well-chosen cases, in the light of certain axiomatic constraints.

In order for any relationship \succsim among events to have the properties of a "qualitative probability relationship," it must be true that:

(a) \succsim is a complete ordering over events; for any two events α , β , either α is "not less probable than" β , or β is "not less probable than" α , and if $\alpha \succsim \beta$ and $\beta \succsim \gamma$, then $\alpha \succsim \gamma$.

(b) If α is more probable than β , then "not- α " (or, $\bar{\alpha}$) is less probable than not- β ($\bar{\beta}$); if α is equally probable to $\bar{\alpha}$, and β is equally probable to $\bar{\beta}$, then α is equally probable to β .

(c) If α and γ are mutually exclusive, and so are β and γ (i.e., if $\alpha \cap \gamma = \beta \cap \gamma = 0$), and if α is more probable than β , then the union $(\alpha \cup \gamma)$ is more probable than $(\beta \cup \gamma)$.

Savage proves that the relationship \succsim among events, inferred as above from choices among gambles, will have the above properties if the individual's pattern of choices obeys certain postulates. To indicate some of these briefly:

P1: Complete ordering of gambles, or "actions." In the example below either I is preferred to II, II is preferred to I, or I and II are indifferent. If I is preferred to II, and II is preferred or indifferent to III, then I is preferred to III (not shown).

	α	β	$\bar{\alpha} \cap \bar{\beta}$
I	a	b	b
II	b	a	b

P2: The choice between two actions must be unaffected by the value of pay-offs corresponding to events for which both actions have the *same* pay-off (i.e., by the value of pay-offs in a constant column).

Thus, if the subject preferred I to II in the example above, he should prefer III to IV, below, when a and b are unchanged and c takes any value:

	α	β	$\bar{\alpha} \cap \bar{\beta}$			
III	<table><tr><td>a</td><td>b</td><td>c</td></tr></table>	a	b	c		
a	b	c				
IV	<table><tr><td>b</td><td>a</td><td>c</td></tr></table>	b	a	c		
b	a	c				

This corresponds to Savage's Postulate 2, which he calls the "Sure-thing Principle" and which bears great weight in the analysis. One rationale for it amounts to the following: Suppose that a person would not prefer IV to III if he *knew* that the third column would not "obtain"; if, on the other hand, he knew that the third column *would* obtain, he would still not prefer IV to III, since the pay-offs (whatever they are) are equal. So, since he would not prefer IV to III "in either event," he should not prefer IV when he does not know *whether or not* the third column will obtain.

"Except possibly for the assumption of simple ordering," Savage asserts, "I know of no other extralogical principle governing decisions that finds such ready acceptance."⁵

P4: The choice in the above example must be independent of the values of a and b , given their ordering. Thus, preferring I to II, the subject should prefer V to VI below, when $d > e$:

	α	β	$\bar{\alpha} \cap \bar{\beta}$			
V	<table><tr><td>d</td><td>e</td><td>e</td></tr></table>	d	e	e		
d	e	e				
VI	<table><tr><td>e</td><td>d</td><td>e</td></tr></table>	e	d	e		
e	d	e				

This is Savage's Postulate 4, the independence of probabilities and pay-offs. Roughly, it specifies that the choice of event on which a

5. *Op. cit.*, p. 21. Savage notes that the principle, in the form of the rationale above, "cannot appropriately be accepted as a postulate in the sense that P1 is, because it would introduce new undefined technical terms referring to knowledge and possibility that would render it mathematically useless without still more postulates governing these terms." He substitutes for it a postulate corresponding to P2 above as expressing the same intuitive constraint. Savage's P2 corresponds closely to "Rubin's Postulate" (Luce and Raiffa, *Games and Decisions*; New York: Wiley, 1957, p. 290) or Milnor's "Column Linearity" postulate, *ibid.*, p. 297, which implies that adding a constant to a column of pay-offs should not change the preference ordering among acts.

If numerical probabilities were assumed known, so that the subject were dealing explicitly with known "risks," these postulates would amount to Samuelson's "Special Independence Assumption" ("Probability, Utility, and the Independence Axiom," *Econometrica*, Vol. 20 (Oct. 1952), pp. 670-78), on which Samuelson relies heavily in his derivation of "von Neumann-Morgenstern utilities."

person prefers to stake a prize should not be affected by the size of the prize.

In combination with a "noncontroversial" Postulate P3 (corresponding to "admissibility," the rejection of dominated actions), these four postulates, if generally satisfied by the individual's choices, imply that his preference for I over II (or III over IV, or V over VI) may safely be interpreted as sufficient evidence that he regards α as "not less probable than" β ; the relationship "not less probable than" *thus operationally defined*, will have all the properties of a "qualitative probability relationship." (Other postulates, which will not be considered here, are necessary in order to establish *numerical* probabilities.) In general, as one ponders these postulates and tests them introspectively in a variety of hypothetical situations, they do indeed appear plausible. That is to say that they do seem to have wide validity as normative criteria (for me, as well as for Savage); they are probably⁶ roughly accurate in predicting certain aspects of actual choice behavior in many situations and better yet in predicting reflective behavior in those situations. To the extent this is true, it should be possible to infer from certain gambling choices in those situations at least a qualitative probability relationship over events, corresponding to a given person's "degrees of belief."

Let us now consider some situations in which the Savage axioms do *not* seem so plausible: circumstances in which none of the above conclusions may appear valid.

Consider the following hypothetical experiment. Let us suppose that you confront two urns containing red and black balls, from one of which a ball will be drawn at random. To "bet on Red_I" will mean that you choose to draw from Urn I; and that you will receive a prize a (say \$100) if you draw a red ball ("if Red_I occurs") and a smaller amount b (say, \$0) if you draw a black ("if not-Red_I occurs").

You have the following information. Urn I contains 100 red and black balls, but in a ratio entirely unknown to you; there may be from 0 to 100 red balls. In Urn II, you confirm that there are exactly 50 red and 50 black balls. An observer — who, let us say, is ignorant of the state of your information about the urns — sets out to measure your subjective probabilities by interrogating you as to your preferences in the following pairs of gambles:

1. "Which do you prefer to bet on, Red_I or Black_I: or are you indifferent?" That is, drawing a ball from Urn I, on which "event" do you prefer the \$100 stake, red or black: or do you care?

6. I bet.

2. "Which would you prefer to bet on, Red_{II} or Black_{II}?"
3. "Which do you prefer to bet on, Red_I or Red_{II}?"⁷
4. "Which do you prefer to bet on, Black_I or Black_{II}?"⁸

Let us suppose that in both the first case and the second case, you are indifferent (the typical response).⁹ Judging from a large number of responses, under absolutely nonexperimental conditions, your answers to the last two questions are likely to fall into one of three groups. You may still be indifferent within each pair of options. (If so, you may sit back now and watch for awhile.) But if you are in the majority, you will report that you prefer to bet on Red_{II} rather than Red_I, and Black_{II} rather than Black_I. The preferences of a small minority run the other way, preferring bets on Red_I to Red_{II}, and Black_I to Black_{II}.

If you are in either of these latter groups, you are now in trouble with the Savage axioms.

Suppose that, betting on red, you preferred to draw out of Urn II. An observer, applying the basic rule of the Ramsey-Savage approach, would infer tentatively that you regarded Red_{II} as "more probable than" Red_I. He then observes that you also prefer to bet on Black_{II} rather than Black_I. Since he cannot conclude that you regard Red_{II} as more probable than Red_I and, at the same time, not-Red_{II} as more probable than not-Red_I — this being inconsistent with the essential properties of probability relationships — he must conclude that your choices are not revealing judgments of "probability" at all. So far as these events are concerned, it is *impossible* to infer probabilities from your choices; you must inevitably be violating some of the Savage axioms (specifically, P1 and P2, complete ordering of actions or the Sure-thing Principle).¹

7. Note that in no case are you invited to choose both a color and an urn freely; nor are you given any indication beforehand as to the full set of gambles that will be offered. If these conditions were altered (as in some of H. Raiffa's experiments with students), you could employ randomized strategies, such as flipping a coin to determine what color to bet on in Urn I, which might affect your choices.

8. See immediately preceding note.

9. Here we see the advantages of purely hypothetical experiments. In "real life," you would probably turn out to have a profound color preference that would invalidate the whole first set of trials, and various other biases that would show up one by one as the experimentation progressed inconclusively.

However, the results in Chipman's almost identical experiment (*op. cit.*, pp. 87–88) do give strong support to this finding; Chipman's explanatory hypothesis differs from that proposed below.

1. In order to relate these choices clearly to the postulates, let us change the experimental setting slightly. Let us assume that the balls in Urn I are each

marked with a I, and the balls in Urn II with a II; the contents of both urns are then dumped into a single urn, which then contains 50 Red_{II} balls, 50 Black_{II} balls, and 100 Red_I and Black_I balls in unknown proportion (or in a proportion indicated only by a small random sample, say, one red and one black). The following actions are to be considered:

	100		50	50
	R _I	B _I	R _{II}	B _{II}
I	a	b	b	b
II	b	a	b	b
III	b	b	a	b
IV	b	b	b	a
V	a	a	b	b
VI	b	b	a	a

Let us assume that a person is indifferent between I and II (between betting on R_I or B_I), between III and IV and between V and VI. It would then follow from Postulates 1 and 2, the assumption of a complete ordering of actions and the Sure-thing Principle, that I, II, III and IV are all indifferent to each other.

To indicate the nature of the proof, suppose that I is preferred to III (the person prefers to bet on R_I rather than R_{II}). Postulates 1 and 2 imply that certain transformations can be performed on this pair of actions *without affecting their preference ordering*; specifically, one action can be replaced by an action indifferent to it (P1 — complete ordering) and the value of a constant column can be changed (P2 — Sure-thing Principle).

Thus starting with I and III and performing such “admissible transformations” it would follow from P1 and P2 that the *first* action in *each* of the following pairs should be preferred:

	R _I	B _{II}	R _{II}	B _{II}	
I	a	b	b	b	
III	b	b	a	b	
I'	a	b	b	a	P2
III'	b	b	a	a	
I''	a	b	b	a	P1
III''	a	a	b	b	
I'''	b	b	b	a	P2
III'''	b	a	b	b	
I''''	b	b	a	b	P1
III''''	a	b	b	b	

Contradiction: I preferred to III, and I'''' (equivalent to III) preferred to III'''' (equivalent to I).

The same applies if you preferred to bet on Red_I and Black_I rather than Red_{II} or Black_{II}. Moreover, harking back to your earlier (hypothetical) replies, any *one* of these preferences involves you in conflict with the axioms. For if one is to interpret from your answers to the first two questions that Red_I is "equally likely" to not-Red_I, and Red_{II} is equally likely to not-Red_{II}, then Red_I (or Black_I) should be equally likely to Red_{II} (or to Black_{II}), and any preference for drawing from one urn over the other leads to a contradiction.²

It might be objected that the assumed total ignorance of the ratio of red and black balls in Urn I is an unrealistic condition, leading to erratic decisions. Let us suppose instead that you have been allowed to draw a random sample of two balls from Urn I, and that you have drawn one red and one black. Or a sample of four: two red and two black. Such conditions do not seem to change the observed pattern of choices appreciably (although the reluctance to draw from Urn I goes down somewhat, as shown for example, by the amount a subject will pay to draw from Urn I; this still remains well below what he will pay for Urn II). The same conflicts with the axioms appear.

Long after beginning these observations, I discovered recently that Knight had postulated an identical comparison, between a man who knows that there are red and black balls in an urn but is ignorant of the numbers of each, and another who knows their exact proportion. The results indicated above directly contradict Knight's own intuition about the situation: "It must be admitted that practically, if any decision as to conduct is involved, such as a wager, the first man would have to act on the supposition that the chances are equal."³ If indeed people were compelled to act on the basis of some Principle of Insufficient Reason when they lacked statistical information, there would be little interest in Knight's own distinctions between risk and uncertainty so far as conduct was involved. But as many people predict their own conduct in such hypothetical situations, they do *not* feel obliged to act "as if" they assigned probabilities at all, equal or not, in this state of ignorance.

Another example yields a direct test of one of the Savage postulates. Imagine an urn known to contain 30 red balls and 60 black and yellow balls, the latter in unknown proportion. (Alternatively, imagine that a sample of two drawn from the 60 black and yellow

2. See immediately preceding note.

3. Knight, *op. cit.*, p. 219.

balls has resulted in one black and one yellow.) One ball is to be drawn at random from the urn; the following actions are considered:

	30	60	
	Red	Black	Yellow
I	\$100	\$0	\$0
II	\$0	\$100	\$0

Action I is "a bet on red," II is "a bet on black." Which do you prefer?

Now consider the following two actions, under the same circumstances:

	30	60	
	Red	Black	Yellow
III	\$100	\$0	\$100
IV	\$0	\$100	\$100

Action III is a "bet on red or yellow"; IV is a "bet on black or yellow." Which of these do you prefer? Take your time!

A very frequent pattern of response is: action I preferred to II, and IV preferred to III. Less frequent is: II preferred to I, and III preferred to IV. Both of these, of course, violate the Sure-thing Principle, which requires the ordering of I to II to be preserved in III and IV (since the two pairs differ only in their third column, constant for each pair).⁴ The first pattern, for example, implies that the subject prefers to bet "on" red rather than "on" black; and he also prefers to bet "against" red rather than "against" black. A relationship "more likely than" inferred from his choices would fail condition (b) above of a "qualitative probability relationship," since it would indicate that he regarded red as more likely than black, but

4. Kenneth Arrow has suggested the following example, in the spirit of the above one:

	100		50	50
	R _I	B _I	R _{II}	B _{II}
I	<i>a</i>	<i>a</i>	<i>b</i>	<i>b</i>
II	<i>a</i>	<i>b</i>	<i>a</i>	<i>b</i>
III	<i>b</i>	<i>a</i>	<i>b</i>	<i>a</i>
IV	<i>b</i>	<i>b</i>	<i>a</i>	<i>a</i>

Assume that I is indifferent to IV, II is indifferent to III. Suppose that I is preferred to II; what is the ordering of III and IV? If III is not preferred to IV, P2, the Sure-thing Principle is violated. If IV is not preferred to III, P1, complete ordering of actions, is violated. (If III is indifferent to IV, both P1 and P2 are violated.)

also "not-red" as more likely than "not-black." Moreover, he would be acting "as though" he regarded "red or yellow" as less likely than "black or yellow," although red were more likely than black, and red, yellow and black were mutually exclusive, thus violating condition (c) above.

Once again, it is impossible, on the basis of such choices, to infer even qualitative probabilities for the events in question (specifically, for events that include yellow or black, but not both). Moreover, for any values of the pay-offs, it is impossible to find probability numbers in terms of which these choices could be described — even roughly or approximately — as maximizing the mathematical expectation of utility.⁵

You might now pause to reconsider your replies. If you should repent of your violations — if you should decide that your choices implying conflicts with the axioms were "mistakes" and that your "real" preferences, upon reflection, involve no such inconsistencies — you confirm that the Savage postulates are, if not descriptive rules for you, your *normative criteria* in these situations. But this is by no means a universal reaction; on the contrary, it would be exceptional.

Responses do vary. There are those who do *not* violate the axioms, or say they won't, even in these situations (e.g., G. Debreu, R. Schlaiffer, P. Samuelson); such subjects tend to apply the axioms rather than their intuition, and when in doubt, to apply some form of the Principle of Insufficient Reason. Some violate the axioms cheerfully, even with gusto (J. Marschak, N. Dalkey); others sadly but persistently, having looked into their hearts, found conflicts with the axioms and decided, in Samuelson's phrase,⁶ to satisfy their

5. Let the utility pay-offs corresponding to \$100 and \$0 be 1, 0; let P_1 , P_2 , P_3 be the probabilities corresponding to red, yellow, black. The expected value to action I is then P_1 ; to II, P_2 ; to III, $P_1 + P_3$; to IV, $P_2 + P_3$. But there are no P_i s, $P_i \geq 0$, $\sum P_i = 1$, such that $P_1 > P_2$ and $P_1 + P_3 < P_2 + P_3$.

6. P. Samuelson, "Probability and the Attempts to Measure Utility," *The Economic Review* (Tokyo, Japan), July 1950, pp. 169-70.

To test the predictive effectiveness of the axioms (or of the alternate decision rule to be proposed in the next section) in these situations, controlled experimentation is in order. (See Chipman's ingenious experiment, *op. cit.*) But, as Savage remarks (*op. cit.*, p. 28), the mode of interrogation implied here and in Savage's book, asking "the person not how he feels, but what he would do in such and such a situation" and giving him ample opportunity to ponder the implications of his replies, seems quite appropriate in weighing "the theory's more important normative interpretation." Moreover, these nonexperimental observations can have at least negative empirical implications, since there is a presumption that people whose instinctive choices violate the Savage axioms, and who claim upon further reflection that they do not *want* to obey them, do not *tend* to obey them normally in such situations.

preferences and let the axioms satisfy themselves. Still others (H. Raiffa) tend, intuitively, to violate the axioms but feel guilty about it and go back into further analysis.

The important finding is that, after rethinking all their "offending" decisions in the light of the axioms, a number of people who are not only sophisticated but reasonable decide that they wish to persist in their choices. This includes people who previously felt a "first-order commitment" to the axioms, many of them surprised and some dismayed to find that they wished, in these situations, to violate the Sure-thing Principle. Since this group included L. J. Savage, when last tested by me (I have been reluctant to try him again), it seems to deserve respectful consideration.

III. WHY ARE SOME UNCERTAINTIES NOT RISKS?

Individuals who would choose I over II and IV over III in the example above (or, II over I and III over IV) are simply not acting "as though" they assigned numerical or even qualitative probabilities to the events in question. There are, it turns out, other ways for them to act. But what *are* they doing?

Even with so few observations, it is possible to say some other things they are *not* doing. They are not "minimaxing"; nor are they applying a "Hurwicz criterion," maximizing a weighted average of minimum pay-off and maximum for each strategy. If they were following any such rules they would have been indifferent between each pair of gambles, since all have identical minima and maxima. Moreover, they are not "minimaxing regret," since in terms of "regrets" the pairs I-II and III-IV are identical.⁷

Thus, *none* of the familiar criteria for predicting or prescribing decision-making under uncertainty corresponds to this pattern of choices. Yet the choices themselves do not appear to be careless or random. They are persistent, reportedly deliberate, and they seem to predominate empirically; many of the people who take them are eminently reasonable, and they insist that they *want* to behave this way, even though they may be generally respectful of the Savage axioms. There are strong indications, in other words, not merely of the existence of reliable patterns of blind behavior but of the opera-

7. No one whose decisions were based on "regrets" could violate the Sure-thing Principle, since all constant columns of pay-offs would transform to a column of 0's in terms of "regret"; on the other hand, such a person would violate P1, complete ordering of strategies.

tion of definite normative criteria, different from and conflicting with the familiar ones, to which these people are trying to conform. If we are talking about you, among others, we might call on your introspection once again. What did *you* think you were doing? What were you trying to do?

One thing to be explained is the fact that you probably would *not* violate the axioms in certain other situations. In the urn example, although a person's choices may not allow us to infer a probability for yellow, or for (red or black), we may be able to deduce quite definitely that he regards (yellow or black) as "more likely than" red; in fact, we might be able to arrive at quite precise numerical estimates for his probabilities, approximating $2/3$, $1/3$. What is the difference between these uncertainties, that leads to such different behavior?

Responses from confessed violators indicate that the difference is not to be found in terms of the two factors commonly used to determine a choice situation, the relative desirability of the possible pay-offs and the relative likelihood of the events affecting them, but in a third dimension of the problem of choice: the nature of one's information concerning the relative likelihood of events. What is at issue might be called the *ambiguity* of this information, a quality depending on the amount, type, reliability and "unanimity" of information, and giving rise to one's degree of "confidence" in an estimate of relative likelihoods.

Such rules as minimaxing, maximaxing, Hurwicz criteria or minimaxing regret are usually prescribed for situations of "complete ignorance," in which a decision-maker lacks any information whatever on relative likelihoods. This would be the case in our urn example if a subject had no basis for considering any of the possible probability distributions over red, yellow, black — such as (1,0,0), (0,1,0), (0,0,1) — as a better estimate, or basis for decision, than any other. On the other hand, the Savage axioms, and the general "Bayesian" approach, are unquestionably appropriate when a subject is willing to base his decisions on a definite and precise choice of a particular distribution: his uncertainty in such a situation is unequivocally in the form of "risk."

But the state of information in our urn example can be characterized neither as "ignorance" nor "risk" in these senses. Each subject does know enough about the problem to *rule out* a number of possible distributions, including all three mentioned above. He *knows* (by the terms of the experiment) that there are red balls in the urn;

in fact, he knows that exactly $1/3$ of the balls are red. Thus, in his "choice" of a subjective probability distribution over red, yellow, black — if he wanted such an estimate as a basis for decision — he is limited to the set of potential distributions between $(1/3, 2/3, 0)$ and $(1/3, 0, 2/3)$: i.e., to the infinite set $(1/3, \lambda, 2/3 - \lambda)$, $0 \leq \lambda \leq 2/3$. Lacking any observations on the number of yellow or black balls, he may have little or no information indicating that one of the remaining, infinite set of distributions is more "likely," more worthy of attention than any other. If he should accumulate some observations, in the form of small sample distributions, this set of "reasonable" distributions would diminish, and a particular distribution might gather increasing strength as a candidate; but so long as the samples remain small, he may be far from able to select one from a number of distributions, or one composite distribution, as a unique basis for decision.

In some situations where two or more probability distributions over the states of nature seem reasonable, or possible, it may still be possible to draw on different sorts of evidence, establishing probability weights in turn to these different distributions to arrive at a final, composite distribution. Even in our examples, it would be misleading to place much emphasis on the notion that a subject has *no* information about the contents of an urn on which no observations have been made. The subject can always ask himself: "What is the likelihood that the experimenter has rigged this urn? Assuming that he has, what proportion of red balls did he probably set? If he is trying to trick me, how is he going about it? What other bets is he going to offer me? What sort of results is he after?" If he has had a lot of experience with psychological tests before, he may be able to bring to bear a good deal of information and intuition that seems relevant to the problem of weighting the different hypotheses, the alternative reasonable probability distributions. In the end, these weights, and the resulting composite probabilities, may or may not be equal for the different possibilities. In our examples, actual subjects do tend to be indifferent between betting on red or black in the unobserved urn, in the first case, or between betting on yellow or black in the second. This need not at all mean that they felt "completely ignorant" or that they could think of no reason to favor one or the other; it does indicate that the reasons, if any, to favor one or the other balanced out subjectively so that the possibilities entered into their final decisions weighted equivalently.

Let us assume, for purposes of discussion, that an individual can always assign relative weights to alternative probability distributions reflecting the relative support given by his information, experience and intuition to these rival hypotheses. This implies that he can always assign relative likelihoods to the states of nature. But how does he *act* in the presence of his uncertainty? The answer to that may depend on another sort of judgment, about the reliability, credibility, or adequacy of his information (including his relevant experience, advice and intuition) as a whole: not about the relative support it may give to one hypothesis as opposed to another, but about its ability to lend support to any hypothesis at all.

If all the information about the events in a set of gambles were in the form of sample-distributions, then ambiguity might be closely related, inversely, to the size of the sample.⁸ But sample-size is not a universally useful index of this factor. Information about many events cannot be conveniently described in terms of a sample distribution; moreover, sample-size seems to focus mainly on the quantity of information. "Ambiguity" may be high (and the confidence in any particular estimate of probabilities low) even where there is ample quantity of information, when there are questions of reliability and relevance of information, and particularly where there is *conflicting* opinion and evidence.

This judgment of the ambiguity of one's information, of the over-all credibility of one's composite estimates, of one's confidence in them, cannot be expressed in terms of relative likelihoods or events (if it could, it would simply affect the final, compound probabilities). Any scrap of evidence bearing on relative likelihood should already be represented in those estimates. But having exploited knowledge, guess, rumor, assumption, advice, to arrive at a final judgment that

8. See Chipman, *op. cit.*, pp. 75, 93. Chipman's important work in this area, done independently and largely prior to mine, is not discussed here since it embodies a stochastic theory of choice; its spirit is otherwise closely similar to that of the present approach, and his experimental results are both pertinent and favorable to the hypotheses below (though Chipman's inferences are somewhat different).

See also the comments by N. Georgescu-Roegen on notion of "credibility," a concept identical to "ambiguity" in this paper: "The Nature of Expectation and Uncertainty," in *Expectations, Uncertainty, and Business Behavior*, ed. Mary Bowman, Social Science Research Council (New York, 1958), pp. 24-26; and "Choice, Expectations and Measurability," this *Journal*, LXVIII (Nov. 1954), 527-30. These highly pertinent articles came to my attention only after this paper had gone to the printer, allowing no space for comment here.

one event is more likely than another or that they are equally likely, one can still stand back from this process and ask: "How much, in the end, is all this worth? How much do I really know about the problem? How firm a basis for choice, for appropriate decision and action, do I have?" The answer, "I don't know very much, and I can't rely on that," may sound rather familiar, even in connection with markedly unequal estimates of relative likelihood. If "complete ignorance" is rare or nonexistent, "considerable" ignorance is surely not.

Savage himself alludes to this sort of judgment and notes as a difficulty with his approach that no recognition is given to it:

... there seem to be some probability relations about which we feel relatively "sure" as compared with others. . . The notion of "sure" and "unsure" introduced here is vague, and my complaint is precisely that neither the theory of personal probability, as it is developed in this book, nor any other device known to me renders the notion less vague. . . A second difficulty, perhaps closely associated with the first one, stems from the vagueness associated with judgments of the magnitude of personal probability.⁹

Knight asserts what Savage's approach tacitly denies, that such over-all judgments may influence decision:

The action which follows upon an opinion depends as much upon the amount of confidence in that opinion as it does upon the favorableness of the opinion itself . . . Fidelity to the actual psychology of the situation requires, we must insist, recognition of these two separate exercises of judgment, the formation of an estimate and the estimation of its value.¹

Let us imagine a situation in which so many of the probability judgments an individual can bring to bear upon a particular problem are either "vague" or "unsure" that his confidence in a particular assignment of probabilities, as opposed to some other of a set of "reasonable" distributions, is very low. We may define this as a situation of high ambiguity. The general proposition to be explored below is that it is precisely in situations of this sort that self-consistent behavior violating the Savage axioms may commonly occur.

Ambiguity is a subjective variable, but it should be possible to identify "objectively" some situations likely to present high ambiguity, by noting situations where available information is scanty or

9. Savage, *op. cit.*, pp. 57-58, 59. Savage later goes so far as to suggest (*op. cit.*, pp. 168-69) that the "aura of vagueness" attached to many judgments of personal probability might lead to systematic violations of his axioms although the decision rule he discusses as alternative — minimaxing regret — cannot, as mentioned in footnote 7 on p. 656 above, account for the behavior in our examples.

1. Knight, *op. cit.*, p. 227.

obviously unreliable or highly conflicting; or where expressed expectations of different individuals differ widely; or where expressed confidence in estimates tends to be low. Thus, as compared with the effects of familiar production decisions or well-known random processes (like coin flipping or roulette), the results of Research and Development, or the performance of a new President, or the tactics of an unfamiliar opponent are all likely to appear ambiguous. This would suggest a broad field of application for the proposition above.

In terms of Shackle's cricket example: Imagine an American observer who had never heard of cricket, knew none of the rules or the method of scoring, and had no clue as to the past record or present prospects of England or Australia. If he were confronted with a set of side bets as to whether England would bat first — this to depend on the throw of a die or a coin — I expect (unlike Shackle) that he would be found to obey Savage's axioms pretty closely, or at least, to want to obey them if any discrepancies were pointed out. Yet I should not be surprised by quite different behavior, at odds with the axioms, if that particular observer were forced to gamble heavily on the proposition that England would *win the match*.

Let us suppose that an individual must choose among a certain set of actions, to whose possible consequences we can assign "von Neumann-Morgenstern utilities" (reflecting the fact that in choosing among *some* set of "unambiguous" gambles involving other events and these same outcomes, he obeys the Savage axioms). We shall suppose that by compounding various probability judgments of varying degrees of reliability he can eliminate certain probability distributions over the states of nature as "unreasonable," assign weights to others and arrive at a composite "estimated" distribution y^0 that represents all his available information on relative likelihoods. But let us further suppose that the situation is ambiguous for him. Out of the set Y of all possible distributions there remains a set Y^0 of distributions that still seem "reasonable," reflecting judgments that he "might almost as well" have made, or that his information — perceived as scanty, unreliable, *ambiguous* — does not permit him confidently to rule out.

In choosing between two actions, I and II, he can compute their expected utilities in terms of their pay-offs and the "estimated" probability distribution y^0 . If the likelihoods of the events in question were as unambiguous as those in the situations in which his von Neumann-Morgenstern utilities were originally measured, this would

be the end of the matter; these pay-offs embody all his attitudes toward "risk," and expected values will correspond to his actual preferences among "risky" gambles. But in this case, where his final assignment of probabilities is less confident, that calculation may leave him uneasy. "So I has a lower expectation than II, on the basis of *these* estimates of probabilities," he may reflect; "How much does that tell me? That's not much of a reason to choose II."

In this state of mind, searching for additional grounds for choice, he may try new criteria, ask new questions. For any of the probability distributions in the "reasonably possible" set Y^o , he can compute an expected value for each of his actions. It might now occur to him to ask: "What might happen to me if my best estimates of likelihood don't apply? What is the *worst* of the reasonable distributions of pay-off that I might associate with action I? With action II?" He might find that he could answer this question about the lower limit of the reasonable expectations for a given action much more confidently than he could arrive at a single, "best guess" expectation; the latter estimate, he might suspect, might vary almost hourly with his mood, whereas the former might look much more solid, almost a "fact," a piece of evidence definitely worth considering in making his choice. In almost no cases (excluding "complete ignorance" as unrealistic) will the *only* fact worth noting about a prospective action be its "security level": the "worst" of the expectations associated with reasonably possible probability distributions. To choose on a "maximin" criterion alone would be to ignore entirely those probability judgments for which there is evidence. But in situations of high ambiguity, such a criterion may appeal to a conservative person as deserving *some* weight, when interrogation of his own subjective estimates of likelihood has failed to disclose a set of estimates that compel exclusive attention in his decision-making.

If, in the end, such a person chooses action I, he may explain:

In terms of my best estimates of probabilities, action I has almost as high an expectation as action II. But if my best guesses should be rotten, which wouldn't surprise me, action I gives me better protection; the worst expectation that looks reasonably possible isn't much worse than the "best guess" expectation, whereas with action II it looks possible that my expectation could really be terrible.

An advocate of the Savage axioms as normative criteria, foreseeing where such reasoning will lead, may interject in exasperation: Why are you double-counting the "worst" possibilities? They're already taken into account in your over-all estimates of likelihoods, weighted in a reasoned, realistic way that represents — by your own claim — your best judgment. Once you've arrived at a probability distribution that reflects everything you know

that's relevant, don't fiddle around with it, use it. Stop asking irrelevant questions and whining about how little you really know.

But this may evoke the calm reply:

It's no use bullying me into taking action II by flattering my "best judgment." I know how little that's based on; I'd back it if we were betting with pennies, but I want to know some other things if the stakes are important, and "How much might I expect to lose, without being unreasonable?" just strikes me as one of those things. As for the reasonableness of giving extra weight to the "bad" likelihoods, my test for that is pragmatic; in situations where I really can't judge confidently among a whole range of possible distributions, this rule steers me toward actions whose expected values are relatively *insensitive* to the particular distribution in that range, without giving up too much in terms of the "best guess" distribution. That strikes me as a sensible, conservative rule to follow. What's wrong with it?

"What's wrong with it" is that it will lead to violations of Savage's Postulate 2, and will make it impossible for an observer to describe the subject's choices as though he were maximizing a linear combination of pay-offs and probabilities over events. Neither of these considerations, even on reflection, may pose to our conservative subject overwhelming imperatives to change his behavior. It will *not* be true that this behavior is erratic or unpredictable (we shall formalize it in terms of a decision rule below), or exhibits intransitivities, or amounts to "throwing away utility" (as would be true, for example, if it led him occasionally to choose strategies that were strongly "dominated" by others). There is, in fact, no obvious basis for asserting that it will lead him in the long run to worse outcomes than he could expect if he reversed some of his preferences to conform to the Savage axioms.

Another person, or this same person in a different situation, might have turned instead or in addition to some other criteria for guidance. One might ask, in an ambiguous situation: "What is the *best* expectation I might associate with this action, without being unreasonable?" Or: "What is its average expectation, giving all the reasonably possible distributions equal weight?" The latter consideration would not, as it happens, lead to behavior violating the Savage axioms. The former would, in the same fashion though in the opposite direction as the "maximin" criterion discussed above; indeed, this "maximaxing" consideration could generate the minority behavior of those who, in our urn example, prefer II to I and III to IV. Both these patterns of behavior could be described by a decision rule similar to the one below, and their respective rationales might be similar to that given above. But let us continue to focus on the particular pattern discussed above, because it seems to predominate

empirically (at least, with respect to our examples) and because it most frequently corresponds to *advice* to be found on decision-making in ambiguous situations.

In reaching his decision, the relative weight that a conservative person will give to the question, "What is the worst expectation that might appear reasonable?" will depend on his confidence in the judgments that go into his estimated probability distribution. The less confident he is, the more he will sacrifice in terms of estimated expected pay-off to achieve a given increase in "security level"; the more confident, the greater increase in "security level" he would demand to compensate for a given drop in estimated expectation. This implies that "trades" are possible between security level and estimated expectation in his preferences, and that does seem to correspond to observed responses. Many subjects will still prefer to bet on R_{II} than R_I in our first example even when the proportion of red to black in Urn II is lowered to 49:51, or will prefer to bet on red than on yellow in the second example even when one red ball is removed from the urn. But at *some* point, as the "unambiguous" likelihood becomes increasingly unfavorable, their choices will switch.²

Assuming, purely for simplicity, that these factors enter into his decision rule in linear combination, we can denote by ρ his degree of confidence, in a given state of information or ambiguity, in the estimated distribution y^o , which in turn reflects all of his judgments on the relative likelihood of distributions, including judgments of equal likelihood. Let \min_x be the minimum expected pay-off to an act x as the probability distribution ranges over the set Y^o ; let est_x be the expected pay-off to the act x corresponding to the estimated distribution y^o .

The simplest decision rule reflecting the above considerations would be:³ *Associate with each x the index:*

$$\rho \cdot \text{est}_x + (1 - \rho) \cdot \min_x$$

Choose that act with the highest index.

2. This contradicts the assertions by Chipman (*op. cit.*, p. 88) and Georgescu-Roegen ("Choice, Expectations and Measurability," pp. 527-30), and "The Nature of Expectation and Uncertainty," p. 25) that individuals order uncertainty-situations lexicographically in terms of estimated expectation and "credibility" (ambiguity); ambiguity appears to influence choice even when estimated expectations are not equivalent.

3. This rule is based upon the concept of a "restricted Bayes solution" developed by J. L. Hodges, Jr., and E. L. Lehmann ("The Uses of Previous Experience in Reaching Statistical Decision," *Annals of Mathematical Statistics*,

An equivalent formulation would be the following, where y^o is the estimated probability vector, y_x^{\min} the probability vector in Y^o corresponding to \min_x for action x and (X) is the vector of payoffs for action x : Associate with each x the index:

$$[\rho \cdot y^o + (1 - \rho) y_x^{\min}](x)$$

Choose that act with the highest index.

In the case of the red, yellow and black balls, supposing no samples and no explicit information except that $1/3$ of the balls are red, many subjects might lean toward an estimated distribution of $(1/3, 1/3, 1/3)$: if not from "ignorance," then from counterbalancing considerations. But many of these would find the situation ambiguous; for them the "reasonable" distributions Y^o might be all those between $(1/3, 2/3, 0)$ and $(1/3, 0, 2/3)$. Assuming for purposes of illustration $\rho = 1/4$ (Y^o , y^o , X and ρ are all subjective data to be inferred by an observer or supplied by the individual, depending on whether the criterion is being used descriptively or for convenient decision-making), the formula for the index would be:

$$1/4 \cdot \text{est}_x + 3/4 \min_x.$$

The relevant data (assigning arbitrary utility values of 6 and 0 to the money outcomes \$100 and \$0) would be:

	Red	Yellow	Black	Min _x	Est _x	Index
I	6	0	0	2	2	2
II	0	6	0	0	2	.5
III	6	0	6	2	4	2.5
IV	0	6	6	4	4	4

A person conforming to this rule with these values would prefer I to II and IV to III, in violation of the Sure-thing Principle: as do most people queried. In justifying this pattern of behavior he might reproduce the rationale quoted above (*q.v.*); but most verbal explanations, somewhat less articulately, tend to be along these lines:

Vol. 23 (Sept. 1952), pp. 396–407. The discussion throughout Section III of this paper derives heavily from the Hodges and Lehmann argument, although their approach is motivated and rationalized somewhat differently.

See also, L. Hurwicz, "Some Specification Problems and Applications to Econometric Models," *Econometrica*, Vol. 19 (July 1951), pp. 343–44 (abstract). This deals with the same sort of problem and presents a "generalized Bayes-minimax principle" equivalent, in more general form, to the decision rule I proposed in an earlier presentation of this paper (December, 1960), but both of these lacked the crucial notions developed in the Hodges and Lehmann approach of a "best estimate" distribution y^o and a "confidence" parameter ρ .

The expected pay-off for action I is definite: 2. The risks under action II may be no greater, but I know what the risk is under action I and I don't under action II. The expectation for action II is ambiguous, it might be better or it might be worse, anything from 0 to 4. To be on the safe side, I'll assume that it's closer to 0; so action I looks better. By the same token, IV looks better than III; I *know* that my expected pay-off with IV is 4, whereas with III it might be as low as 2 (which isn't compensated by the chance that it could be 6). In fact, I know the whole *probability distribution of payoffs* (though not the distribution over *events*) for I and IV, but I don't for II and III. I know that a payoff of 6 is *twice as likely* as 0 under IV, whereas 6 *may* be only half as likely as 0 under III.

Leaving the advocate of the Savage axioms, if he is still around to hear this, to renew his complaints about the silliness and irrelevance of such considerations, let us note a practical consequence of the decision rule which the above comment brings into focus. It has already been mentioned that the rule will favor — other things (such as the estimated expectation) being roughly equal — actions whose expected value is less sensitive to variation of the probability distribution within the range of ambiguity. Such actions may frequently be those definable as “status quo” or “present behavior” strategies. For these, ρ may be high, the range of Y_0 small.

A familiar, ongoing pattern of activity may be subject to considerable uncertainty, but this uncertainty is more apt to appear in the form of “risk”; the relation between given states of nature is known precisely, and although the random variation in the state of nature which “obtains” may be considerable, its stochastic properties are often known confidently and in detail. (Actually, this confidence may be self-deceptive, based on ignoring some treacherous possibilities; nevertheless, it commonly exists.) In contrast, the ambiguities surrounding the outcome of a proposed *innovation*, a departure from current strategy, may be much more noticeable. Different sorts of events are relevant to its outcome, and their likelihoods must now be estimated, often with little evidence or prior expertise; and the effect of a given state of nature upon the outcome of the new action may itself be in question. Its variance may not appear any higher than that of the familiar action when computed on the basis of “best estimates” of the probabilities involved, yet the meaningfulness of this calculation may be subject to doubt. The decision rule discussed will not preclude choosing such an act, but it will definitely bias the choice away from such ambiguous ventures and toward the strategy with “*known* risks.” Thus the rule is “conservative” in a sense more familiar to everyday conversation than to statistical decision theory; it may often favor traditional or current strategies, even perhaps at high risk, over innovations whose

consequences are undeniably ambiguous. This property may recommend it to some, discredit it with others (some of whom might prefer to reverse the rule, to emphasize the more hopeful possibilities in ambiguous situations); it does not seem irrelevant to one's attitude toward the behavior.

In the equivalent formulation in terms of y_x^{\min} and y^o , the subject above could be described "as though" he were assigning weights to the respective pay-offs of actions II and III, whose expected values are ambiguous, as follows (assuming $y^o = (1/3, 1/3, 1/3)$ in each case):

	y_x^{\min}	$\rho \cdot y^o + (1 - \rho) y_x^{\min}$
II	$\left(\frac{1}{3}, 0, \frac{2}{3}\right)$	$\left(\frac{1}{3}, \frac{1}{12}, \frac{7}{12}\right)$
III	$\left(\frac{1}{3}, \frac{2}{3}, 0\right)$	$\left(\frac{1}{3}, \frac{7}{12}, \frac{1}{12}\right)$

Although the final set of weights for each set of pay-offs resemble probabilities (they are positive, sum to unity, and represent a linear combination of two probability distributions), they differ for each action, since y_x^{\min} will depend on the pay-offs for x and will vary for different actions. If these weights were interpreted as "probabilities," we would have to regard the subject's subjective probabilities as being dependent upon his pay-offs, his evaluation of the outcomes. Thus, this model would be appropriate to represent cases of true pessimism, or optimism or wishfulness (with y_x^{\max} substituting for y_x^{\min}). However, in this case we are assuming *conservatism*, not pessimism; our subject does not actually expect the worst, but he chooses to act "*as though*" *the worst were somewhat more likely than his best estimates of likelihood would indicate*. In either case, he violates the Savage axioms; it is impossible to infer from the resulting behavior a set of probabilities *for events* independent of his pay-offs. In effect, he "distorts" his best estimates of likelihood, in the direction of increased emphasis on the less favorable outcomes and to a degree depending on ρ , his confidence in his best estimate.⁴

4. This interpretation of the behavior-pattern contrasts to the hypothesis or decision rule advanced by Fellner in the accompanying article in this symposium. Fellner seems unmistakably to be dealing with the same phenomena discussed here, and his proposed technique of measuring a person's subjective probabilities and utilities in relatively "unambiguous" situations and then using these measurements to calibrate his uncertainty in more ambiguous environments seems to me a most valuable source of new data and hypotheses. Moreover, his

Not only does this decision model account for "deviant" behavior in a particular, ambiguous situation, but it covers the observed shift in a subject's behavior as ambiguity decreases. Suppose that a sample is drawn from the urn, strengthening the confidence in the best estimates of likelihood, so that ρ increases, say, to $\frac{3}{4}$. The weights for the pay-offs to actions II and III would now be:

$$\begin{aligned} & \rho \cdot y^o + (1-\rho) y_x^{\min} \\ \text{II} & \left(\frac{1}{3}, \frac{1}{4}, \frac{5}{12} \right) \\ \text{III} & \left(\frac{1}{3}, \frac{5}{12}, \frac{1}{4} \right) \end{aligned}$$

and the over-all index would be: *Index*

I	2
II	1.5
III	3.5
IV	4

In other words, the relative influence of the consideration, "What is the worst to be expected?" upon the comparison of actions is lessened. The final weights approach closer to the "best estimate" values, and I and II approach closer to indifference, as do III and IV. This latter aspect might show up behaviorally in the amount a subject is willing to pay for a given bet on yellow, or on (red or black), in the two situations.

In the limit, as ambiguity diminishes for one reason or another and ρ approaches 1, the estimated distribution will come increasingly to dominate decision. With confidence in the best estimates high, behavior on the basis of the proposed decision rule will roughly conform to the Savage axioms, and it would be possible to infer the estimated probabilities from observed choices. But prior to this, a large number of information states, distinguishable from each other and all far removed from "complete ignorance," might all be suffi-

descriptive data and intuitive conjectures lend encouraging support to the findings reported here. However, his solution to the problem supposes a single set of weights determined independently of pay-offs (presumably corresponding to the "best estimates" here) and a "correction factor," reflecting the degree of ambiguity or confidence, which operates on these weights in a manner independent of the structure of pay-offs. I am not entirely clear on the behavioral implications of Fellner's model or the decision rule it implies, but in view of these properties I am doubtful whether it can account adequately for all the behavior discussed above.

ciently ambiguous as to lead many decision-makers to conform to the above decision rule with $p < 1$, in clear violation of the axioms.

Are they foolish? It is not the object of this paper to judge that. I have been concerned rather to advance the testable propositions: (1) certain information states can be meaningfully identified as highly ambiguous; (2) in these states, many reasonable people tend to violate the Savage axioms with respect to certain choices; (3) their behavior is deliberate and not readily reversed upon reflection; (4) certain patterns of "violating" behavior can be distinguished and described in terms of a specified decision rule.

If these propositions should prove valid, the question of the optimality of this behavior would gain more interest. The mere fact that it conflicts with certain axioms of choice that at first glance appear reasonable does not seem to me to foreclose this question; empirical research, and even preliminary speculation, about the nature of actual or "successful" decision-making under uncertainty is still too young to give us confidence that these axioms are not abstracting away from vital considerations. It would seem incautious to rule peremptorily that the people in question should not allow their perception of ambiguity, their unease with their best estimates of probability, to influence their decision: or to assert that the manner in which they respond to it is against their long-run interest and that they would be in some sense better off if they should go against their deep-felt preferences. If their rationale for their decision behavior is not uniquely compelling (and recent discussions with T. Schelling have raised questions in my mind about it), neither, it seems to me, are the counterarguments. Indeed, it seems out of the question summarily to judge their behavior as irrational: I am included among them.

In any case, it follows from the propositions above that for their behavior in the situations in question, the Bayesian or Savage approach gives wrong predictions and, by their lights, bad advice. They act in conflict with the axioms deliberately, without apology, because it seems to them the sensible way to behave. Are they clearly mistaken?

THE RAND CORPORATION
SANTA MONICA, CALIFORNIA