ISAAC LEVI

# IGNORANCE, PROBABILITY AND RATIONAL CHOICE

## 1. POSSIBLE VERSUS PERMISSIBLE PROBABILITIES

Strict Bayesians are legitimately challenged to tell us where they get their numbers. Typically, they concede that ordinary mortals with imperfect memories, computational capacity, and emotional balance are not able to specify what their credal probability judgments or their utilities are with perfect precision. Nonetheless, so they insist, such precision is an ideal to which we should be committed. We should regard our less definite appraisals of probability to be indicators of what our commitments concerning definite probabilities should be. Sometimes the indicators can be exploited to derive quite detailed specifications of the numerically precise credal probability distributions to which agents are often unwittingly committed; but even when they cannot, rational agents should be treated as if their minds were black boxes containing definite probability judgments whose contents are partially revealed by the explicit discernments they can make (Good, 1962).

But the challenge to strict Bayesians is not merely about the strong demands they make on our intellectual and emotional capacities. Any sophisticated theory of rational choice and belief will impose demands which cannot be fully met even by the most intelligent and mature individuals when the context of choice reaches a certain level of complexity. How we are to counsel real agents to cope with predicaments where their limitations prevent them from conforming to demands of ideal rationality is a very serious problem for any approach to rational choice; but to respond to it requires a conception of ideal rationality. The anxiety about where the numbers come from which I wish to consider addresses such ideals. Should we require that rational agents be committed to numerically definite judgments of credal probability – i.e., definite probability distributions to be used in the computation of expected utilities of feasible options?

Some critics of strict Bayesian doctrine answer this last question

*Synthese* **53** (1982) 387–417. 0039–7857/82/0533–0387 $03.10

with a resounding "no!" In their excellent paper (Gärdenfors and Sahlin, 1982), Gärdenfors and Sahlin have joined the nay-sayers. I am delighted to have them in our company. But there is a considerable difference between us concerning how one should proceed in dissenting from strict Bayesian doctrine. In this paper, I shall review some of the main points under dispute.

The deterrent to the construction of views rival to strict Bayesianism is not the absence of alternatives. To the contrary, there seem to be so many moves one can make in deviating from Bayesian scripture that straying from the straight and narrow path might appear to threaten decision theoretic anarchy. An examination of the dispute between Gärdenfors and Sahlin and myself may serve to show that there is more order in Babel than appears to be the case at first blush. Our disagreements, if I understand them correctly, concern discussable implications of treating refusal to commit oneself to a single expectation determining distribution as uniquely *permissible* or (b) as refusing to recognize a single such distribution as uniquely *possible*.

Gärdenfors and Sahlin take the position that a rational agent can allow two or more credal, expectation determining, or subjective probability distributions to be seriously or epistemically possible in the sense that "it does not contradict the decision maker's knowledge in the given decision situation" (p. 365). This characterization suggests that such epistemically or seriously possible probability distributions resemble truth value bearing hypotheses in that one can suspend judgment as to their truth values, can regard them to be more or less probable in their own right and can modify one's body of knowledge by rejecting those distributions whose probabilities or "reliabilities" are sufficiently low to render the risk incurred by rejecting them worthwhile.[1]

Gärdenfors and Sahlin appear prepared to allow for all these operations on possible probability distributions. According to their account, therefore, strict Bayesians who insist that rational agents restrict attention to one expectation determining credal distribution at a time insist that the credal distribution adopted is one of whose truth the agent is certain. Their objection to this view appears to be that rational agents are invited to be opinionated about credal probability distributions even though there are many contexts when an agent may lack the knowledge warranting such confidence.

Following De Finetti (1972, 189ff) and Savage (1954, p. 58), I contend that this view appears to suffer from contradictions which can be avoided only by desparate ad hoc repairs. Nonetheless, I reject the strict Bayesianism which De Finetti and Savage advocate. Rational agents can be in suspense between rival credal probability distributions. However, the suspense does not concern which of the rival distributions is true. There is no sense in which one such distribution may be regarded as more or less certain, reliable or probable than another. Error is not risked in rejecting such a distribution.

To explain what is to be meant by speaking of suspension of judgment concerning credal probability distributions, we should consider the function of such distributions in the assessment of expected utilities. If an agent is restricted to using exactly one distribution for such a purpose in a given context, he is committed to restricting his choice between options to one of those bearing maximum expected utility where expected utility is calculated according to that distribution. If more than one option bears maximum expected utility, the agent may appeal to other considerations to render a verdict.

Suppose, however, that two or more distributions are *permissible* to use in calculating expected utilities. Those options which bear maximum expected utility relative to at least one of those permissible distributions have not been prohibited from being chosen by considerations of expected utility. All such options are *E-admissible* (Levi, 1974, 1980). In a state of suspense between two or more credal probability distributions, options which are optimal with respect to expected utility relative to at least one of the distributions survive the test of expected utility.

Strict Bayesians require rational agents, at least ideally, to recognize exactly one distribution to be permissible in this sense at any given time. I reject this claim. I contend that a rational agent may embrace a *credal state* according to which more than one probability distribution over a given space of truth value bearing hypotheses) is permissible. But permissible distributions are neither true nor false. Hence, they are neither possibly true nor possibly false. One cannot compare such distributions with respect to probability; and although credal states may be modified by throwing out erstwhile permissible distributions, there is no risk of error involved in doing so.[2]

Thus, while I agree with Gärdenfors and Sahlin in rejecting strict Bayesianism, I part company with them and with I. J. Good (1962) concerning the treatment of credal probabilities as if they were truth value bearing and as if they could be assigned 'second order' probabilities.

Nonetheless, counter to De Finetti and Savage, I do think that there are truth value bearing hypotheses specifying probability distributions. I do not mean hypotheses of a biographical nature telling us what Miss Julie's state of belief happens to be at a given time. That there are such truth value bearing hypotheses and that we and Miss Julie can be in doubt about the truth of hypotheses of this sort is not a matter for dispute here. My reference is to statements of objective statistical probability or chance such as claims about the chance of a coin landing heads on a toss or of player 1 winning from player 2 in a tennis match between them on a sunny day in Lund, Sweden. We may be in doubt as to whether such a chance equals some value $p$ between 0 and 1 and, in our doubt, might assign credal probabilities to hypotheses as to what that chance $p$ is. And we can take risks of error by rejecting some of these hypotheses.[3]

However, such chance probabilities are not to be confused with credal probabilities. An agent may be in suspense as to what the true chance of coin $a$ landing heads on a toss is and yet have a definite credal probability assignment for the hypothesis that coin $a$ will land heads on the toss the agent knows will take place at some time $t$. In that case, there are many possibly true chance distributions but only one permissible distribution.

As is well known, De Finetti and Savage rejected the notion of objective statistical probability or chance as metaphysical nonsense. Hence, they denied that there is any good sense in which one can suspend judgment as to the truth of a probability distribution. But one does not have to saddle all strict Bayesians with the positivistic prejudices of De Finetti. If a strict Bayesian regards chance as intelligible, he can allow that rational agents may be in suspense as to the truth of statistical hypotheses while still insisting that to be rational, the agent must allow at most one credal distribution to be permissible at a time.

Thus, my complaint against Gärdenfors and Sahlin is not that they allow for possible probability distributions which vary in their probability or reliability but that they confuse such distributions with credal distributions. It is this confusion which leads to the inconsistencies which I allege are implicit in their view.

## 2. CREDAL PROBABILITIES OF CREDAL PROBABILITIES

Suppose that Miss Julie is in suspense between two rival credal probability distributions. One distribution assigns degree of credence .9 to the hypothesis $b$ that player 1 will win in match $C$ and the other assigns degree of credence .1 to the same hypothesis. This set of distributions is not convex; but Gärdenfors and Sahlin do not appear to require convexity in general; and their instincts are right insofar as they think of credal distributions as if they were truth value bearing hypotheses. Miss Julie might quite plausibly end up assuming that exactly one of these two distributions is true while remaining in suspense as to which one it is.

Let $Q_1(b) = .9$ and $Q_2(b) = .1$ represent these two distributions for the hypotheses $b$ that player 1 wins and $b'$ that player 2 wins. ($Q_i(b') = 1 - Q_i(b)$.) Given that Miss Julie is in suspense as to which is true, she should adopt some sort of credal state for these rival distributions just as she should for any rival hypotheses concerning whose truth she is in doubt. And if she is a strict Bayesian, she will assign a single distribution granting credal probability $x$ to $Q_1$ and $1-x$ to $Q_2$. Furthermore, it should be the case that if Miss Julie comes to be certain that $Q_1$ is the correct distribution, her degree of credence for $b$ should equal .9. Hence, while in suspense, she should embrace a probability function meeting the requirement that $Q(b; Q_1$ is true$) = .9$. Likewise, $Q(b; Q_2$ is true$) = .1$.

By the calculus of probabilities, it follows that $Q(b) = Q(b; Q_1$ is true$)Q(Q_1$ is true$) + Q(b; Q_2$ is true$)Q(Q_2$ is true$) = .9x + .1(1-x)$. The resulting value $Q(b)$ is the sole value which Miss Julie should use to compute expected utilities. Hence, Miss Julie should regard it as the sole possibly true distribution over $b$ and $b'$ – i.e., if the possibly true distributions are those the agent is free to use in computing expected utilities.

But if $x$ is different from 0 and 1, $Q(b)$ will differ from .1 and .9 and, hence, must by hypothesis be ruled out as a serious possibility. And if $x$ is 0 or 1, then exactly one of the two distributions, $Q_1$ or $Q_2$, is a serious possibility in contradiction to the claim that Miss Julie is not certain concerning either of them.

Matters do not improve if, instead of supposing that Miss Julie is in suspense between just two distributions, we think of her as in suspense between all distributions over $b$ and $b'$ assigning $b$ a degree of credence y and $b'$ $1-y$ where y is any value in the closed interval from 0 to 1. If we then suppose that Miss Julie adopts a single credal distribution over the range of values of y representable by a continuous density function $f(y)$ and assume that $Q(b; Q_y(b) = y$ is true$) = y$,[4] we have the result that $Q(b) = \int_0^1 yf(y)\,dy$. This takes a definite value between 0 and 1 and only this value is permissible. Miss Julie is not in suspense between it and the others in contradiction with the assumptions initially made about the situation.[5]

One might seek to postpone trouble by suggesting that Miss Julie be in suspense as to the truth of the second order probability distributions. But then Miss Julie must have a third order credal state concerning the possible second order distributions. Integration threatens to introduce contradiction once again. This can be avoided only by suspending judgment between all possible third order distributions over second order distributions, all fourth order distributions over third order distributions, etc. The evil day may then be postponed indefinitely.

Following this approach will not help Gärdenfors and Sahlin. It implies that the set of second order distributions possible must consist of all distributions consonant with the calculus of probabilities. Hence, it precludes a categorical weak ordering of the first order possible distributions with respect to second order probability. No two such distributions will be categorically comparable – i.e., orderable in the same way by all possible second order distributions. Gärdenfors and Sahlin, however, require that first order distributions be weakly ordered with respect to "epistemic reliability" (p. 367). This ordering seems, on their view, to exhibit the properties of a comparative probability.

There are other moves one might make. Perhaps the claim that $Q(b; Q_1(b) = .9) = .9$ might be rejected. But given a careful formulation of the principles licensing this constraint, it is difficult to see how Gärdenfors and Sahlin could wish to reject it. After all, on their own account, if one comes to be certain that a given credal distribution obtains, one should assign degrees of credence in accordance with that distribution.

Another (more promising) response would be to deny that second order distributions are to be treated like first order distributions. Whereas one can regard first order distributions to be like truth value bearing hypotheses which are possibly true or possibly false and which may be evaluated with respect to credal probability or reliability, second order distributions are permissible credal distributions along lines I have suggested. Thus, Miss Julie might suspend judgment as to the truth of $Q_1$ and $Q_2$ and assign a degree of credence $x$ to $Q_1$ and $1-x$ to $Q_2$. As before, $Q(b) = .9x + .1(1-x)$. But now the $Q$-function is not a truth value bearing distribution but is part of a second order credal or expectation determining credal distribution. Hence, in claiming that either $Q_1$ or $Q_2$ is true without declaring which one it is, one does not rule out the permissibility of using a second order $Q$-distribution to assign a value $Q(b) = .9x + .1(1-x)$ distinct from both .9 and .1.

To make this procedure work, one has to determine how to assess expected utility. If exactly one value of $x$ is permissible, it would seem that $Q(b) = .9x + .1(1-x)$ is uniquely permissible for that purpose.

The upshot is that in our effort to repair the inconsistencies in what appears to be the Gärdenfors–Sahlin position, we have been led to draw a distinction analogous (although perhaps not identical) to the distinction I have drawn between chance distributions and credal distributions. The erstwhile first order distributions have become objective statistical or chance probability distributions. The second order distributions have become credal distributions.

I cannot claim to have exhausted all the moves one might make in order to escape the inconsistencies which threaten the Gärdenfors–Sahlin approach; but enough has been said to suggest that there is some doubt as to whether a useful approach can be developed to an account of rational probability judgment and choice which treats credal probabilities as if they were truth value bearing hypotheses which can be possibly true and false and more or less reliable or probable in their own right.

The next task is to examine the scenarios involving Miss Julie presented for our consideration by Gärdenfors and Sahlin when the distinction between credal and chance probability is faithfully drawn.

## 3. CHANCE AND DIRECT INFERENCE

Miss Julie knows that the two players confronting each other in match $A$ have (more or less) equal skills, are in equally good physical and mental condition and are not handicapped by the weather, the court or other conditions of play. This information can be summed up as claiming that the chance of player 1 winning in a tennis match with player 2 under the given conditions of play is, at least for the present, equal to the chance of player 2 winning.

Such statements about chance or objective statistical probability characterize properties of experimental arrangements or systems. What is specified is a distribution of chance probabilities over a sample space of 'possible outcomes' of an experiment or trial of some kind on the experimental system. In our example, the experimental system is the pair of players together with the tennis court, the referees, tennis paraphernalia, etc. The kind of trial is the playing of a match by these players under conditions of play of some appropriate kind. The sample space consists of the two 'points': player 1 winning and player 2 winning. The chance distribution specifies a probability $x$ for the first point and $1-x$ for the second.

The experimental set up has the chance property (or lacks it) regardless of Miss Julie's beliefs. It also has the property, if it has it, regardless of whether the experiment is actually conducted – i.e., a match actually played. In this respect, it is like a coin which, if unbiased (thereby having an equal chance of landing heads up on a toss and landing tails up) is unbiased whether it is ever tossed or not. (Levi, 1967 ch. XIV and 1980 ch. 12.)

Miss Julie is supposed to know that a given simple chance hypothesis specifying a definite chance distribution over the given sample space relative to the given kind of trial of a definite experimental arrangement is true. Such knowledge should *not* be confused with Miss Julie's assigning a degree of credence of .5 to the hypothesis that player 1 will win.

Thus, Miss Julie might know that the simple chance hypothesis is true but not know that match $A$ will be held. Indeed, she might be certain that it will not be held. In that case, she will assign a degree of credence of 0 to the hypothesis that player 1 will win and do likewise for the hypothesis that player 2 will win. Even so, she still might assume (be certain, know) that the chance hypothesis $c_{.5}$ asserting that the chance of player 1 winning on a play of a match is .5 and that of player 2 winning is also .5.

Furthermore, even if Miss Julie knows not only $c_{.5}$ but also that match $A$ will be played, her degree of credence that $b$ (player 1 will win) need not be .5. She knows, for example, that the match is to be played in Lund and unless she can also assume that this information is stochastically irrelevant to the outcome, she cannot assign that degree of credence. (If the chance of a result of kind *R* on a trial of kind *S* is known to be *p* and it is known that the trial on a particular occasion is of kind *S* and also of kind *T*, the information that the trial is of kind *T* is stochastically irrelevant for the agent if the agent knows that the chance of an *R* on a trial of kind *S* & *T* is equal to the chance of an *R* on a trial of kind *S*.)

Specifying the precise conditions under which knowledge of chances warrants a judgment of credal probability concerning the outcome of some specific trial is afforded by a principle of direct inference or what Peirce called 'statistical deduction'. The formulation of such a principle is a matter of some controversy. The issues need not detain us here. (See Levi, 1980 ch. 12 and 16 and 1982 for discussion and references.) But whatever the final resolution of the controversies involved might be, the status of principles of direct inference appears to be similar to principles of coherence mandating the use of measures obeying the requirements of the calculus of probabilities as expectation determining probability measures. They are principles of probabilistic or inductive logic which impose constraints on how ideally rational agents should assign credal probabilities to hypotheses relative to knowledge of chances or statistical probabilities.

B. De Finetti and those who follow him deny the need for any principles of probabilistic logic other than coherence requirements; and it is easy to see why. De Finetti denies the intelligibility of the concept of statistical probability. (De Finetti, 1964, pp. 140–142.) If there can be no knowledge of statistical probability, principles of direct inference are devoid of applicability.[6]

However, if there are truth value bearing statements of chance which are knowable, we must have some way of understanding how such knowledge ramifies for the behavior of experimental arrangements to which statistical properties are attributed. Principles of direct inference elaborate upon these ramifications in a manner bearing an attenuated analogy to the way in which Carnap held that reduction sentences explicate the conditions for disposition statements.

Thus, Miss Julie can be said to have information about chances and the conditions of playing the match which obligate her via direct inference to assign the hypothesis $b$ that player 1 will win the degree of credence of .5.

We can say that Miss Julie is 'sure' of the truth of a chance distribution and uses this information to ground her assessment of the credal probability that $b$.

On this view, it is misleading to say that Miss Julie is certain of the truth of her subjective or credal probability assignment to $b$. What she is certain of is the truth of propositions warranting such an assignment to $b$. In particular, she is certain that $c_{.5}$ is true and that the match $A$ will be played under appropriate conditions.

In the case of match $B$, Miss Julie knows that the two players are about to play a match under the given conditions but knows nothing about their abilities and, hence, knows nothing about the chances of player 1 winning in a match between the two. Thus, Miss Julie knows that one and only one hypothesis of the form $c_p$ specifying a chance between 0 and 1 of player 1 winning in a match with player 2 is true. In virtue of direct inference and the other knowledge available, Miss Julie judges for each value of $p$ (where this value is designated by a standard designator for a real number) that the conditional credal probability that $b$ given $c_p$ is equal to $p$. $Q(b; c_p) = p$. The calculus of probabilities implies that $Q(b) = \int_0^1 Q(b; c_p)f(p)\,\mathrm{d}p = \int_0^1 pf(p)\,\mathrm{d}p$ where $f(p)$ is a density for $p$.

To determine a value for $Q(b)$, therefore, it is necessary to have a credal distribution $f(p)$ over the values of $p$. (In cases where the distribution is not continuous or fails to range over all real values in the interval from 0 to 1, the credal distribution over the hypotheses $c_p$ would, of course, not be representable simply by a density as I am doing here.)

In one special situation, it is possible, even in match $B$, to obligate Miss Julie to assign $Q(b) = .5$ by the calculus of probabilities and the principle of direct inference just as in case $A$.

Even though Miss Julie does not know which value of $p$ represents the chance of player 1 winning on a play of a match like match $B$, she might believe that tennis ability is genetically determined and, indeed, that the relative abilities of the two players is itself the outcome of some stochastic process. To stretch our imaginations still farther, she might assume that the chance distribution over outcomes of this process (which allocates values of $p$ to pairs of players) is such that the chance of the value equalling $p^*$ is equal to the chance of its equalling $1 - p^*$ for every value of $p^*$ in the interval from 0 to 1.

If Miss Julie knows this second level chance distribution, then she also knows (via the use of the calculus of probabilities applied to chance probabilities) that the chance of player 1 winning in a match where the abilities of the two players is selected by a stochastic process corresponding to such a chance distribution is equal to .5. Given that she knows that the particular match *B* she is about to witness is of this kind, direct inference mandates a degree of credence for *b* equal to .5 – just as in case *A*.

As in case *A*, Miss Julie is certain of information about chances warranting this degree of credence; but it is a different knowledge of chances. There is a difference between $c_{.5}$ which asserts that the chance of player 1 winning in a match with player 2 is .5 and the chance claim made here which asserts that the chance of player 1 winning in a match with player 2 where the abilities of the two players are determined by the given stochastic process is .5.

The predicament can be redescribed in a slightly different way. Direct inference can be said to require assigning $c_p$ an equal credal probability as the credal probability assigned $c_{1-p}$. That is to say, $f(p) = f(1-p)$. Then $\int_0^1 pf(p)\,dp$ must equal $.5 = Q(b)$.

No matter how the situation is described, Miss Julie is ignorant as to which of the many rival conjectures as to the relative abilities of the two players is true. Yet, it is a situation where Miss Julie is as "sure" of her credal judgment as in case *A*. Strictly speaking, of course, to speak of her being sure of her credal probabilities threatens us with incoherence. What is common between this variant of case *B* and match *A* is that Miss Julie has knowledge in both cases which justifies the assignment of $Q(b) = .5$ as uniquely permissible on the basis of inductive or probabilistic logic alone – i.e., principles of credal coherence requiring the use of the calculus of probabilities and principles of direct inference.

In this version of *B*, the mere fact that Miss Julie does not know which of the rival chance hypotheses of type $c_p$ is true cannot be sufficient for abandoning strict Bayesian doctrine. Indeed, Miss Julie should behave like a strict Bayesian in this case. The principles of inductive logic require it. Only one credal distribution for *b* is permissible even though there are many possible chance distributions specifying chances of player 1 winning in a match with player 2.

Needless to say, Gärdenfors and Sahlin did not have this version of match *B* in mind (or the corresponding version of match *C*) when they sought to contrast $B$ with match $A$. But their description of match $B$ in English does not preclude this possibility unless their interpretation of the sense in which Miss Julie is ignorant of the abilities of the two players is taken in the strong sense which rules out knowledge of a stochastic process of the sort I just described. And it is important to keep this in mind. In any case, let us now turn to the predicament vis-à-vis match $B$ in a manner closer to the intent of Gärdenfors and Sahlin.

## 4. ENTROPIC VERSUS CREDAL IGNORANCE

To get at the intent of Gärdenfors and Sahlin we should not only suppose that Miss Julie is modally ignorant concerning the $c_p$'s (so that the truth and the falsity of $c_p$ are both serious possibilities for each value of $p$); but she cannot justify a credal state for the $c_p$'s via direct inference from knowledge of chances.

Thomas Bayes (1963) suggested a long time ago that such cases should be treated as if Miss Julie did have knowledge of chances of a certain kind – namely of the sort which obligated assigning each of the $c_p$'s equal probability (i.e., using a uniform density $f(p) = 1$ to represent the credal distribution). The idea that in the face of modal ignorance one should be *entropically ignorant* when one cannot derive a credal state via direct inference from knowledge of chances gained additional authority from the efforts of Laplace; and many other distinguished writers have undertaken efforts to polish and immunize the principle of insufficient reason which thus emerged from the threat of contradiction.

To follow the principle in the case of match $B$ will, of course, lead to assigning $Q(b)$ the value .5. Moreover, advocates of insufficient reason contend that the assignment of this value is grounded as strongly in principles of probabilistic logic as is the assignment of the same value on the basis of knowledge of chances derived via direct inference.

Principles of insufficient reason have been the source of much controversy and I assume that Gärdenfors and Sahlin, like myself, reject them. (See Seidenfeld, 1979 for some quite decisive reasons for doing so.) Indeed, part of the force of their use of a contrast between match $A$ and match $B$ is to undermine the presystematic cogency of insufficient reason. In the case of match $A$ (or match $B$ where relative abilities are known to be the outcome of a stochastic process), Miss Julie's expectation determining or credal probability that $b$ can be derived from her knowledge of chances via direct inference. Hence, she is obliged to assign $b$ a degree of credence of .5 and to act accordingly.

For many agents, no sense of obligation is present in the case of match $B$ (at least when knowledge about the stochastic process for selecting a value of $p$ is absent). If such agents are reasonable in their responses (and I, along with Gärdenfors and Sahlin, think they are), the dictates of insufficient reason cannot be obligatory on all rational agents and, hence, cannot be accorded the status of a principle of probabilistic logic.

To say this, however, does not imply that rational agents are prohibited from assigning equal probabilities to all $c_p$'s. They cannot claim a warrant from insufficient reason for adopting a state of entropic ignorance; but perhaps they do not need to have their credal state certified by principles of probabilistic logic and their total knowledge. That, at any rate, is the view of those strict Bayesians who are, in the terminology of L. J. Savage 'personalists'.

Consequently, although there is no obligation to be entropically ignorant in the context of match $B$ as far as personalist strict Bayesians are concerned, there is no prohibition against it either; and those strict Bayesians who are entropically ignorant will behave the same when confronted with gambles on $b$ and $b'$ in match $B$ as they would in the case of match $A$.

Of course, a personalist of strict Bayesian persuasion will concede that Miss Julie could have assigned the $c_p$'s credal probability in a nonuniform manner and ended up with a different value than .5 for $Q(b)$. In that event, Miss Julie's evaluations of her gambles would have been different in the case of match $B$ than it is in the case of match $A$. Nonetheless, Miss Julie would still be behaving like a strict Bayesian.

Thus, even if we concede (as we should) that in the case of match $B$, Miss Julie is not obligated to assign $Q(b)$ the value .5, it does not follow that she is obligated to deviate from the dictates of strict Bayesianism. She might continue to adopt one value for $Q(b)$ as uniquely permissible.

Yet, there is more to be said. Once it is acknowledged that no warrant exists in probabilistic logic for mandating $Q(b) = .5$ in the case of match $B$ on the basis of Miss Julie's knowledge, not only is it open to Miss Julie to adopt another numerically definite expectation determining credal probability as uniquely permissible but she may rather prefer to remain in suspense between rival distributions – not in the sense that two or more are possibly true but in the sense in which two or more are permissible to use in the evaluation of expected utility. In that event, it seems reasonable to require that all real values in some subinterval of the unit interval be permissible credal probabilities for $b$ and, more generally, that the set of permissible distributions over the $c_p$'s form a convex set.[7]

The ignorance which emerges here is not modal ignorance of the truth values of the $c_p$'s for that is already built into the match $B$ predicament. Nor is it modal ignorance of the credal distributions over the $c_p$'s; for these lack truth values. There is no assessment of the relative reliability of the distributions over the $c_p$'s. Some distributions are permissible according to Miss Julie's credal state and others are not. That is all that can be said. I call a credal state for the $c_p$'s of this sort a state of partial or total credal ignorance.[8] It is total when all distributions obeying the calculus of probabilities over the $c_p$'s are permissible. It is partial in varying ways when the set of permissible distributions is a convex subset of this.[9]

In any case, whether the credal state is a state of partial or total ignorance concerning the $c_p$'s, the credal state for $b$ and $b'$ will, in general, fail to mandate assignment of a definite credal probability.[10]

My contention is that there are no canons of rational probability judgment mandating the assignment of a numerically definite distribution for the $c_p$'s or for $b$ and $b'$ as uniquely permissible in match $B$ as Gärdenfors and Sahlin understand that match.

Gärdenfors and Sahlin attempt to capture the same point (wrongly I think) by denying that there is a single *possible* distribution for $b$ and $b'$ in that case. Of course, in a sense, there are many possible distributions in case $B$. These are the possibly true statistical probability distributions represented by the $c_p$'s. But as we have seen, this does not entail deviation from the strict Bayesian requirement that exactly one credal distribution be permissible.

In section 2, I argued for the incoherence of the Gärdenfors and Sahlin approach to characterizing probabilistic ignorance as a form of modal ignorance. What I have sought to illustrate here is how the features of the contrast between matches $A$ and $B$ they sought to capture with their approach can be captured by an approach which utilises the notion of modal ignorance of statistical probability and notion of credal ignorance of expectation determining probability without lapsing into such incoherence.

It is now time to consider the import of these considerations for rational choice.

## 5. EXPECTATION AND SECURITY

In the case of match $B$ where the credal state for $b$ and $b'$ allows many different distributions to be permissible, Miss Julie will not and should not evaluate gambles in the manner strict Bayesians require.

Consider gambles on $b$ with payoff $S-P$ when $b$ is true and $-P$ when $b$ is false. Strict Bayesians require that the following conditions be satisfied: (a) If $S$ is positive, there is a lub $p_*$ to the set of ratios $P/S$ such that Miss Julie should accept the gambles with those values for the 'stakes' $S$ and 'price' $P$. (b) If $S$ is negative, there is a glb $p^*$ to the set of ratios $P/S$ for which Miss Julie should accept such gambles. (c) $p^* = p_*$.

Gärdenfors, Sahlin and I agree that conditions (a) and (b) should be met but, following C. A. B. Smith (1961), reject (c).

In spite of this agreement, there are differences in the decision theories we employ to enforce this view and these differences have ramifications for more complex contexts of decision making which ought to be brought out. I shall first explain the fundamentals of my approach. In the next section, I shall contrast it with the ideas of Gärdenfors and Sahlin.

According to the approach to rational choice I favor (Levi, 1974, 1980), rational agents should restrict their choice to those feasible options which are admissible. Admissibility, in the special sense in which I use that term, is determined by a sequence of lexicographically ordered tests of admissibility. In previous discussions, I have considered a sequence of three kinds of tests: tests for $E$-admissibility invoking considerations of expected utility, tests for $P$-admissibility to determine when if at all it is feasible to make a choice which expresses suspense between the $E$-admissible options and tests for $S$-admissibility (or lex admissibility) which identify those $P$-admissible options which promote security. In the context of gambling and other noncognitive decision making, consideration of $P$-

admissibility is rarely salient and, hence, may be ignored for the purpose of the present discussion. Thus, we may simplify the discussion by requiring that rational agents restrict their choices to $E$-admissible options and then identify those $E$-admissible options which are $S$-admissible (or lex admissible).

Given a set of of feasible options, the set of $E$-admissible options is that subset of feasible options which are optimal in expected utility when expected utility is computed by some permissible probability distribution in the credal state. The set of $E$-admissible options are those feasible options between which the agent is incapable of rendering a verdict insofar as consideration of expected utility is taken into account. Thus, if the agent endorses exactly one credal distribution as permissible, the $E$-admissible options will coincide with those that maximize expected utility. If two or more options are thus optimal in expected utility, considerations of expected utility can no longer render a verdict. Some other test for admissibility can then be used to reach a decision.

If the credal state contains more than one probability distribution as permissible, it can happen that an option is optimal in expected utility relative to one distribution and a different option is optimal relative to the other. In this case also, both options are $E$-admissible. But we can no longer regard either option as optimal in expected utility in any categorical sense. As far as the agent's categorical evaluations of expected utility are concerned, they are noncomparable. They cannot be used to render a verdict between the two options with respect to expected utility any more than when the two options tie for categorical optimality as in the previous case. Hence, in spite of the difference between being equal in expected utility and being noncomparable with respect to expected utility which will prove crucial subsequently, we may resort to another criterion to render a verdict.

The test I favor using is one which identifies for each feasible option a security level and enjoins picking that $E$-admissible option (or one of those $E$-admissible options) bearing maximum security level. In a more sophisticated version, each option is associated with a security level (worst possible case), secondary security level (second worst possible case), etc. until all possible cases are exhausted and a lexicographic version of the injunction to maximize security levels among the $E$-admissible options. And a still more general version considers permissible evaluations of security. For the sake of simplicity, I shall consider only the simpler criterion of $S$-admissible where there is a uniquely permissible assessment of security.

Consider a variant of Miss Julie's predicament when confronting match $B$ where she is in a state of partial credal ignorance concerning the $c_p$'s. In particular, her credal state for the $c_p$'s restricts the density $f(p)$ to the interval from 1/2 to 3/2. In that event, the lowest credal probability for $b$ permissible must be 3/8. That will obtain when $f(p)$ is 3/2 for the interval from 0 to .5 and is 1/2 for the interval from .5 to 1. The highest permissible credal probability for $b$ is 5/8 which obtains when $f(p) = 1/2$ from 0 to .5 and is 3/2 from .5 to 1.

If Miss Julie is offered a gamble with positive stake $S = 1$ utile and pays a price $P$ less than 3/8 of a utile, her expected utility for the gamble will be positive no matter what permissible value of $Q(b)$ she uses to compute expected utility. So accepting the gamble is uniquely $E$-admissible in a pairwise choice where the options are accepting a gamble for price $P$ less than 3/8 utiles or rejecting the gamble with neither gain nor loss. If $P$ is greater than 5/8, rejecting the gamble is uniquely $E$-admissible; for the expectation of the gamble is negative for all permissible distributions over $b$ and $b'$. When $P$ is greater than or equal to 3/8 and less than or equal to 5/8, both options are $E$-admissible.

So let us explore the situation where $P$ falls between 3/8 and 5/8. We urge Miss Julie to consider the worst possible case or payoff from accepting the gamble and from rejecting it and to favor the option where the worst possible consequence is best.

But how is a worst possible case or security level to be identified? That depends upon how we partition the space of possible consequences for each option.

One method is to consider the payoff of accepting the gamble when $b$ is true and when $b$ is false. Clearly the payoff when $b$ is false is the worst possible case. So the security level will be $-P$ for some value of $P$ in the interval from 3/8 to 5/8. The security level for refusing the gamble will be 0. Hence, refusing the gamble is uniquely $S$-admissible and should be chosen.

But there are other methods of partitioning the space of possible outcomes for each option. We can consider as possible consequences of accepting the gamble, accepting when $c_p$ is true for each value of $p$. The 'payoff' for each such possible case is the conditional expected utility of accepting the gamble given that $c_p$ is true. Because of the principle of direct inference, we are given that $Q(b; c_p) = p$ so that this conditional expectation has a definite value for each value of $p$. Hence, we obtain a unique payoff for each possible case and a definite security level. The security level for accepting the gamble will be the expected utility conditional on $c_0$ which is $-P$. That is, of course, the same numerical value as according to the preceding method of determining security levels. But in many situations, the result would be different. For example, had the values of $p$ been restricted to the range between .1 and .9, the security level would have been $.1 - P$ by the method just now described whereas it would have remained $-P$ by the previous method.

The interesting point to notice for our future discussion is that the security level determined in this fashion is the smallest expected utility conditional on a possible objective statistical probability distribution.

A third method of fixing security levels emerges if the possible consequences of accepting the gamble are reduced to one – the sure consequence that either $b$ or $b'$ is true. In that case, the security level goes indeterminate consisting of the range of permissible expected utilities computed by the permissible probability assignments to $b$.

Each of these three methods of fixing security levels can be applied to the option of refusing the gamble as well. For each of these methods, the security level is unequivocally 0.

Thus, according to the first two methods, refusing the gamble is uniquely $S$-admissible. According to the third method both options are $S$-admissible.

Clearly when the options become more complex, the opportunities for introducing new ways for individuating possible consequences of feasible options become even greater and, hence, the methods for fixing security levels multiply and ramify. In my judgment, it would be foolish to insist that there is a uniquely rational way to determine security levels. A theory of rational choice does not legislate how the rational agent should assign utilities to the possible consequences. That is a matter of the agent's goals, values, moral commitments, professional obligations and the like. Similarly, the choice of a method for fixing security levels should be understood as manifesting an aspect of the agent's values, moral commitments, etc.

No doubt agents are prone to favor some methods of individuating possible consequences rather than others. Thus, Miss Julie, I surmise, will not fix security levels for the gamble individuating possible consequences according to the third method described above. However, if she did, she would not be irrational – only atypical in the way she assesses security.

On the other hand, if Miss Julie were to fix security levels according to the third method, for some permissible security distributions, accepting the gamble would rank over rejecting it and for some refusing the gamble would rank over accepting it. Hence, both options would be $S$-admissible. Appeal to security would be futile. Hence, Miss Julie might be tempted (although not obliged) to invoke another way of thinking of security to render a verdict between the options.

Fixing security according to one of the two other procedures guarantees that accepting the gamble has a negative security whereas refusing the gamble has 0 security level. Refusing the gamble is uniquely $S$-admissible. This will be so regardless of the value of the price $P$ as long as it takes a value in the closed interval from 3/8 to 5/8. Thus, when security is fixed according to one of these two procedures, the lub of prices at which Miss Julie will accept a gamble on $b$ with a stake of 1 utile is 3/8 utile.

By similar reasoning, the glb of such prices where the stake is $-1$ utile is $-5/8$. Hence, for positive stakes the betting quotient threshold for acceptance is 3/8. For positive stakes, it is 5/8. Thus, condition (c) cited above requiring that $p^* = p_*$ is violated.

As another example of how my proposals work and how they relativize considerations of security to value judgments involved in specifying partitions of possible consequences, consider the paradox of Ellsberg discussed by Gärdenfors and Sahlin (pp. 374–77 and p. 379).

Ellsberg had claimed that the fact that some, indeed many, agents choose $a_1$ over $a_2$ while choosing $a_4$ over $a_3$ reveals a violation of Savage's sure thing principle.

This is true, however, only on assumptions which I have rejected. The Savage principle would be violated were it the case that an agent prefers $a_1$ strictly to $a_2$ while strictly preferring $a_4$ to $a_3$. Ellsberg tacitly assumes that if the agent regards $a_1$ as uniquely admissible in a pairwise choice between $a_1$ and $a_2$, he reveals his strict preference for $a_1$ over $a_2$. Similarly, if he chooses $a_4$ over $a_3$ and regards $a_4$ as uniquely admissible, he reveals his strict preference for $a_4$ over $a_3$.

But it should be apparent by now that I reject that view. Indeed, in the Ellsberg example, many individuals would refuse to endorse a single distribution as uniquely permissible and would, indeed, adopt a credal state sufficiently indeterminate so as to guarantee that both $a_1$ and $a_2$ are $E$-admissible in a pairwise choice and that both $a_3$ and $a_4$ are also both $E$-admissible in a pairwise choice.

According to my proposal, resort would have to be made to considerations of security. Among the many methods for partitioning the options into possible consequences, two stand out:

According to method 1, the possible consequences of option $a_i$ are the payoffs when a red ball, a black ball and a yellow ball is picked. The security level for $a_1$ according to this method is 0 as is the security level when $a_2$ is chosen. Similarly, the security levels for both options in the second pair is 0. Thus, in both decision problems, both options are $S$-admissible. On this account, the decisions of those respondents who definitely favor $a_1$ as uniquely admissible over $a_2$ and $a_4$ as uniquely admissible over $a_3$ cannot be accommodated.

Consider, however, the partitioning into possible consequences where a possible consequence of accepting $a_i$ is accepting it when one of the 61 possible constitutions of the urn is the correct one. (Remember that there are 60 balls which are black or yellow in unknown proportion while there are 30 red balls. There are 61 distinct compositions possible meeting this requirement.) Direct inference mandates a conditional probability distribution over red, black and yellow for each specific hypothesis about the composition of the urn and, hence, a definite conditional expectation for $a_i$. For $a_1$, the security level on this account is 100/3. For $a_2$ it is 0. Hence, $a_1$ is uniquely $S$-admissible. The security level for $a_3$ is 100/3 and for $a_4$ it is 200/3. Hence, $a_4$ is uniquely $S$-admissible.

Thus, if security levels are fixed according to method 2, $a_1$ is uniquely admissible in the pairwise choice with $a_2$ whereas $a_4$ is uniquely admissible in a pairwise choice with $a_3$.

Clearly the result conforms with the reaction Ellsberg alleges violates the sure thing principle. But it does not represent such a violation. $a_1$ is uniquely admissible in a pairwise choice over $a_2$. But it is not strictly preferred in any categorical sense over $a_2$. Both $a_1$ and $a_2$ are $E$-admissible. They are noncomparable with respect to expected utility. Neither is preferred to or indifferent to the other with respect to expected utility. I fail to see, therefore, how the Savage sure thing principle is violated in any sense in which it is obvious that it should not be violated.

In footnote 14 of my (1974), I suggested a way of determining security levels for mixed options which favors the use of method 1 for fixing security levels and explicitly contrasted it with the alternative method advocated by Von Neumann and Morgenstern and by Wald in formulating their minimax criteria. When this paper was reissued in 1978, I wrote an addendum in which I explicitly withdrew the implication that one method of fixing a security level is the rational one and instead suggested that "the partitioning of possible consequences used to fix security levels depends on the agent's goals and values just as the set of permissible u-functions (utility functions) does." This position was further explained in Levi, 1980, 7.4, 7.5, 17.3 and 17.4.

If one were to take the position suggested by footnote 14 of Levi (1974), then it would be correct to say, as Gärdenfors and Sahlin allege, that my approach fails to account for the Ellsberg phenomenon. But the attitude I expressed in 1978 and later in 1980 quite clearly allows for behavior in conformity with the Ellsberg response.

It is true that my decision theory does not obligate all rational agents to choose in the manner which Ellsberg considered. But Ellsberg himself did not assume that rational agents ought always to choose in that way. That depends on the exogenous variables of his decision theory just as my proposal depends on another set of exogenous variables. All I need do to accommodate as reasonable the allegedly paradoxical Ellsberg response is to provide an account of rational choice which does not condemn those who choose $a_1$ over $a_2$ and $a_4$ over $a_3$ is irrational. In this respect, my theory is all that Ellsberg could ask for.

Gärdenfors and Sahlin allege (p. 379) that my theory "seems to have problems in explaining some of the experimental results" they cite in their paper including results pertaining to Ellsberg-like phenomena because my theory claims that $a_1$ is equally as good as $a_2$, and $a_3$ is equally as good as $a_4$. In the light of the preceding discussion, we can identify two errors in this contention.

First, on my view, to claim that $a_1$ is "equally as good as" $a_2$ is not equivalent to saying that $a_1$ and $a_2$ are both admissible. The two options might be both admissible even though neither is better than the other or as good as the other. Hence, they cannot correctly charge that according to my theory $a_1$ is equally as good as $a_2$ or that $a_3$ is equally as good as $a_4$.

Second, if they mean to allege that on my view $a_1$ and $a_2$ are both admissible according to the first pairwise choice and $a_3$ and $a_4$ are both admissible according to the second pairwise choice, this is true only if security levels are fixed according to method 1; but not if they are fixed according to method 2. As I have explained, my theory refuses to mandate use of one method over others as rational.

## 6. THE DECISION THEORY OF GÄRDENFORS AND SAHLIN

According to Gärdenfors and Sahlin, the decision maker begins with a set of possible credal probability distributions, assesses them with respect to epistemic reliability, rejects those which are sufficiently unreliable and then uses the remainder to evaluate the feasible options. I have already registered my difficulties with the initial steps of this approach; but nothing has been said thus far about the final steps where the agent uses the surviving possible distributions to assess the feasible options.

The Gärdenfors–Sahlin theory contains one rule rather than a lexicographically ordered series of rules. For each feasible option, one is supposed to compute the expected utility relative to every possible probability distribution and determine the minimal expected utility for that option. The maximum criterion for expected utilities (MMEU) recommends maximizing the minimal expected utility.

In their essay, Gärdenfors and Sahlin sought to interpret my theory within their framework. I shall now repay their kindness by doing the same to their theory in reverse. That is to say, I shall explain their theory and assess its merits on the assumption that the distinction I have insisted upon between truth value bearing statistical probability distributions which are possibly true or false and permissible expectation determining credal probability distributions which are neither true nor false is preserved.

This can be done by construing their possible credal distributions as possible statistical probability distributions and not as credal distributions at all. By direct inference, one can compute the expected utility of each feasible option conditional on each possible chance distribution. The minimal expected utility associated with a given feasible option according to their theory then translates into a security level when security is determined by using what is, in effect, the negative of A. Wald's "risk function" (Wald, 1950, p. 12). In the case of the Ellsberg example, this is using method 2 for determining security.

On this construal, their recommendation to maximize the minimal expected utility reduces to Wald's injunction to minimize the maximum risk (or to maximize the minimum negative risk).

As I have just explained, there is nothing objectionable on my view in fixing security levels in this manner. What is objectionable is to insist that agents are obliged to fix security levels in this way as a matter of rational principle. To do so is to confuse questions concerning value commitments which go beyond issues of coherence with questions of rational or coherent decision making.

However, assuming for the sake of the argument that security levels are to be fixed with the aid of risk (or negative risk) functions, my theory still conflicts with the Gärdenfors–Sahlin approach. In my view, we should restrict choice to those options which maximize security among *E-admissible* options and not necessarily among all *feasible* options. Thus, according to my theory, prior to assessing security, the credal state of permissible distributions is to be used to evaluate expected utility in order to weed out the $E$-inadmissible options. This initial step is left entirely out of account in the Gärdenfors–Sahlin theory. All that matters is the use of possible statistical probability distributions to determine the risk functions to be employed in evaluating security levels.

This difference makes no difference in those cases where the set of $E$-admissible options coincides with the set of feasible options or, indeed, in those cases where the options bearing maximum security level among the feasible options are themselves a subset of the $E$-admissible options. This will not, however, be the case in general when the credal state for the rival chance hypotheses expresses only partial credal ignorance or becomes so definite as to allow exactly one distribution to be permissible.

To illustrate, consider the example of match $B$ under the conditions where Miss Julie is modally ignorant concerning which $c_p$ is true for values of $p$ between 0 and 1 but where she is only partially credally ignorant so that the permissible range of values for $b$ consists of real values in the interval from 3/8 to 5/8.

If Miss Julie is offered a gamble on $b$ with a prize of 7/8 utiles if $b$ is true and a loss of 1/8 utile if $b'$ is true, accepting the gamble is uniquely $E$-admissible in a pairwise choice between accepting the gamble and refusing it with neither gain nor loss. Considerations of security need not be invoked at all and the gamble should be accepted.

But if the Gärdenfors–Sahlin possible probability distributions are the possibly true distributions of statistical probability as I have been assuming, the chance of player 1 winning ranges from $p = 0$ to $p = 1$. Since the set of permissible credal distributions is ignored along with the assessment of $E$-admissibility, we must compare only security levels as determined by the risk function. The security level for the gamble is $-1/8$ utile and for refusal is 0. The gamble, according to the Gärdenfors–Sahlin theory, must be rejected.

Of course, if the possible values of $p$ had ranged between 3/8 and 5/8 and there had been complete credal ignorance concerning these values (i.e., all credal distributions over this range allowed by the calculus of probabilities were permissible), the Gärdenfors–Sahlin theory would recommend accepting the gamble – just as my approach would do.

The divergence of the Gärdenfors–Sahlin theory from my own is even more dramatic in cases where there are three or more feasible options.

Suppose that in the case of match $B$, Miss Julie regards all values of $p$ as serious possibilities and is maximally credally ignorant concerning these possibilities. Let her then be offered a choice between the following three options:

$d_1$ Receive 55 utiles if $b$ is true and lose 45 if $b'$ is true.
$d_2$ Lose 46 utiles if $b$ is true and receive 54 if $b'$ is true.
$d_3$ Refuse $d_1$ and $d_2$ and receive 0 utiles no matter what happens.

The security levels for the three options are $-45$, $-46$ and 0 respectively. *But $d_3$ is not E-admissible.* There is no permissible probability distribution over the $c_p$'s which determines a probability distribution for $b$ and $b'$ such that $d_3$ is optimal in expected utility among the three options. This is immediately apparent once it is recognized that for no value of $p$ between 0 and 1 is the expected utility of $d_3$ on $c_p$ at least as great as the expected utilities of both the other alternatives. To be sure, $d_3$ is not dominated by either of the other options. Nonetheless, if Miss Julie knew for sure which of the $c_p$'s is true, she would refuse $d_3$.

In spite of these considerations, Gärdenfors and Sahlin recommend that Miss Julie should choose $d_3$. My theory recommends choosing $d_1$ on the grounds that it and $d_2$ are $E$-admissible and that $d_1$ maximizes security among the $E$-admissible options.

In discussing an analogous example (pp. 378–79), Gärdenfors and Sahlin allege that choosing the analogue of $d_3$ is "intuitively the best option". To claim that $d_3$ is intuitively best even though it is never optimal given the truth of $c_p$ is, in my view, a very difficult claim to sustain.

Gärdenfors and Sahlin worry about the charge that they are excessively "risk aversive" (p. 379). They respond by claiming that some distributions are already rejected because they are unreliable and that this already incurs a risk. Whatever the merits of this argument might otherwise be, in the example now under consideration it cannot apply because no possible distributions have been rejected. (The example discussed by Gärdenfors and Sahlin differs from the one discussed here in this respect.)

In any case, the issue under dispute cannot be concerning whether $d_3$ is the 'best' option or not. $d_3$ is best in security when security is measured by risk. What is under dispute is whether security should be maximized or whether, prior to invoking considerations of security, one should first rule out those options failing to maximize expected utility against any permissible credal distribution.

When the three options are compared with respect to expected utility, for each $i$ and $j$, $d_i$ and $d_j$ are noncomparable with respect to expected utility. There is no best option in this respect. $d_1$ and $d_2$ are $E$-admissible and $d_3$ is not; but $d_3$ is not inferior to the other two options in expected utility – just noncomparable.

To repeat what I have stated before, the dispute between Gärdenfors–Sahlin and myself does not concern whether $d_3$ is best but whether it is uniquely admissible.

Gärdenfors and Sahlin contend that my approach seems "unrealistically optimistic" (p. 379). The contention is without foundation as is the countercharge that the Gärdenfors–Sahlin view is excessively pessimistic.

A. Wald showed that under certain conditions the options which are minimizers of maximum risk also maximize expected utility relative to a probability distribution which he called the "least favorable" distribution (Wald, 1950, pp. 18 and 91). Wald's rhetoric has been taken to imply that someone who minimizes maximum risk (or maximizes minimum negative risk) is pessimistic or paranoid.

The metaphor could be taken seriously if we were to suppose that the appeal to considerations of security is a way of selecting a credal distribution to be used to compute expected utilities; for in that case, it might be said that the injunction of maximize minimum negative risk is pessimistic – i.e., is to adopt a numerically definite and pessimistic credal state.

If one understands the use of minimax risk in this way, rejection or mitigation of the use of that principle might seem "unrealistic optimistic". But such a charge cannot apply to my theory. There are two reasons why it fails:

(1) Wald's theorem establishing that minimax risk solutions are Bayes solutions applies to situations where the set of feasible options is "convex" (Wald, 1950, Assumption 3.6(i), p. 68). This means, in effect, that all mixtures of feasible options are also feasible. But that need not be true in real life and is surely not true in any of the examples we are now considering.

Thus, there is no probability distribution over $b$ and $b'$ relative to which $d_3$ maximizes expected utility. This does not violate Wald's theorem; for we have not considered all mixtures of the three options to be options. If we did, there would be a different minimax risk solution – namely the option of choosing $d_1$ or $d_2$ depending on the outcome of a toss of an unbiased coin. But there is no reason to suppose that such an option is always feasible or that rational agents will be prepared to use the risk function for a mixed option to determine the security level – even though statisticians and decisions have often followed Wald in doing so. The lack of realism applies to those who assume that the set of feasible options should be convex.

If this is granted, one cannot say that Gärdenfors and Sahlin are unduly pessimistic in recommending the choice of $d_3$. They cannot be pessimistic at all – even in Wald's sense; for there is no probability distribution relative to which $d_3$ is a Bayes' solution maximizing expected utility. Hence, it cannot be the Bayes' solution relative to a 'least favorable solution'.

But even if we overlook this point, there is yet another more fundamental one. Suppose we follow my approach which recommends minimizing maximum risk not among the feasible options but among the $E$-admissible options. In that case, there will, indeed, always be a probability distribution relative to which the admissible (S-admissible) option chosen is optimal in expected utility. It does not follow from this that one is thereby committed to adopting as one's credal state those credal distributions which render the S-admissible options optimal in expected utility. In my view, one invokes considerations of security when one's beliefs and expectations cannot be used to render a verdict. For me, using considerations of security is not an indirect method of picking out a credal probability distribution. To the contrary, security comes into its own when considerations of expected utility become useless. Thus, on my approach, the use of considerations of security reveals no predilection in favor of either optimism or pessimism.

Gärdenfors and Sahlin argue in favor of their approach as compared to mine by considering a case analogous to the three way choice between $d_1$, $d_2$ and $d_3$ to show that my theory leads to violation of what A. K. Sen called "property $\alpha$" (Sen, 1970, p. 17) and a condition called "independence of irrelevant alternatives" by Luce and Raiffa (1958, p. 288) – not to be confused with the Arrovian principle of the same name.

According to this principle, if $d_3$ is inadmissible in the three way choice, it cannot be admissible in a pairwise choice between $d_3$ and one of the other options. Yet, in a pairwise choice between $d_1$ and $d_3$, both options are $E$-admissible and $d_3$ uniquely $S$-admissible even though it is $E$-inadmissible and, hence, inadmissible in the three way choice.

Gärdenfors and Sahlin think this implication of my theory to be a strong objection against it. In presenting their case, they follow Luce and Raiffa in using 'optimal' where I use 'admissible'. And if my theory did indeed imply that when $d_3$ is optimal in the three way choice, it is not optimal in a pairwise choice between $d_3$ and $d_1$, the objection would be telling. But what does "optimal" mean here? $d_3$ is not optimal in expected utility in either the three way or the two way choice. It is optimal in security both in the three way and the two way choice. But even though it is optimal in security in the three way choice it is not $E$-admissible and, hence, not even $S$-admissible in the three way choice. Since optimality with respect to expected utility cannot be legislative in either the pairwise or three way choice because there is no weak ordering of the options with respect to expected utility, that sort of optimality is irrelevant to the problem. And since optimality with respect to security only applies to those options surviving the test of expected utility, the fact that $d_3$ is optimal in security among the three options does not make it admissible in the three way choice. Thus, to equate admissibility with optimality and then argue that property $\alpha$ should be obeyed is clearly to beg the question against approaches such as my own which reject the equation of admissibility with optimality either with respect to expected utility or security.

Gärdenfors and Sahlin appear to be in the thrall of the behavioralist myth that options which are judged admissible are thereby revealed to be optimal. Those who subscribe to this prejudice are, indeed, driven to requiring that criteria of admissibility satisfy property $\alpha$; for only then can transitivity and asymmetry of strict preference be preserved.

The same motivation appears to lurk behind the inclination of Gärdenfors and Sahlin to endorse a solution to the Ellsberg problem which violates the sure thing principle. My approach, which recommends the same choices for the Ellsberg problem as theirs when security is determined by risk functions, avoids violating the sure thing principle because it avoids implying that $a_1$ is preferred to $a_2$ or that $a_4$ is preferred to $a_1$.

In this discussion of the Gärdenfors–Sahlin theory, I have sought to show (1) that they have confused modal ignorance concerning truth value bearing statistical probabilities which may be possibly true and possibly false and credal ignorance concerning permissible but non-truth value bearing credal or expectation determining probabilities, (2) that when this confusion is removed, their theory reduces to the Wald minimax risk theory, (3) that when the set of minimax risk solutions among the feasible options is a subset of the set of $E$-admissible options (Bayes solutions) among the feasible options, their theory is a special case of mine and (4) that the fact that my theory violates property $\alpha$ in those limited contexts where it does is, if anything, a mark in its favor and a mark against theories like the Gärdenfors–Sahlin proposal which fail to take note of the legitimacy of such violations.

My critique of the Gärdenfors–Sahlin theory is predicated on interpreting their decision theory in a manner which distinguishes between possibly true statistical probability distributions and permissible credal distributions. Gärdenfors and Sahlin proceed as if this were an untenable dualism and, hence, might wish to reject my criticisms on the grounds that they are predicated on a misrepresentation of their proposals.

If so, I refer back to my original objections to conflating possible with permissible probabilities. I have no doubt that there are ways to reply to the objections I raised which I have failed to consider; although I am inclined to doubt that there are any promising responses.

In spite of the several points under dispute between us, it is important in closing to emphasize the shared perception common to Gärdenfors, Sahlin and myself that there is something deeply dogmatic in the strict Bayesian insistence on the notion that ideal rationality requires numerically definite probability judgments. Of course, a cast of thousands has shared this perception with us; but most critics of strict Bayesianism have tended to be eclectic suggesting one set of prescriptions for one class of cases and another for a different set of cases without attempting to formulate a coherent system of underlying principles unifying their prescriptions. Gärdenfors and Sahlin agree with me that such eclecticism is not enough to meet the strict Bayesian challenge and they have embarked on a program designed to show that one can provide a unified approach to probability judgment and rational choice which is more comprehensive than the strict Bayesian approach. For them, as for me, strict Bayesian doctrine covers a special limiting case not typical of most practical situations.

Given our shared perception of the problematic facing critics of strict Bayesianism, our disagreements ought to be understood to be results of our efforts to explore different solutions to a common problem. What matters at the present stage of the discussion is that the *problematic* be taken seriously by careful study of potential solutions of the difficulties faced. The contribution of Gärdenfors and Sahlin is, in spite of the many reservations I have expressed, a most welcome addition to the current discussion of this problematic.

*Columbia University*

NOTES

[1] See Levi (1980), 9.1–9.4. The term "cognitive ignorance" which appears in 9.1 corresponds to what I here call "modal ignorance" here. What I now call "entropic ignorance" was called "Bayesian ignorance" there. In (1977), what I now call "modal

ignorance" was called by the same name. Entropic ignorance was called "probabilistic ignorance" and credal ignorance was called "probabilistic ignorance in the extreme sense".

[2] An account of the revision of credal states involving the removal of erstwhile permissible credal distributions is given in Levi (1980), ch. 13.

[3] I first considered the topic of rejecting statistical probability distributions in Levi (1962); but the first version of the approach which evolved into my current position is in Levi (1967), especially chapter XVI. The latest representation of my view is given in Levi (1980), ch. 13.

[4] I assume that the designators to be substituted for 'y' are standard designators for real numbers.

[5] See Levi (1980), 9.3. As I understand it, the objection I am raising is an elaboration of the point atrributed by L. J. Savage to Max Woodbury in Savage (1954), p. 58. It is interesting to note that in the second edition of 1972, Savage added a footnote called the notion of representing "the unsure" by convex sets of distributions "tempting". I was unaware of this until Nils-Eric Sahlin drew it to my attention via the good offices of Peter Gärdenfors.

[6] This is not true without qualification. H. E. Kyburg has proposed an account of direct inference which substitutes knowledge of relative frequencies for knowledge of statistical probabilities; but this possibility depends critically on the adoption of principles of direct inference which allow for violation of conditionalizing arguments – indeed, mandate such violation – in ways which Bayesians would find unacceptable. See Levi (1980), ch. 12.

[7] For some further discussion of convexity, see Levi (1980), 9.5 and the related discussion in 8.3 concerning utility functions.

[8] See footnote 1 for reference to other terminology I have used elsewhere.

[9] As Gärdenfors and Sahlin correctly observe, the convex set representing a credal state cannot be uniquely picked out by the interval valued probability function which envelops it. This point is emphasized in Levi (1974) and in (1980), 9.8–9.9.

[10] There are some exceptional cases. Thus, the set of credal distributions such that $\int_0^1 pf(p)\,dp = r$ is convex. However, in my judgment, the importance of such sets in representing credal states is extremely limited.

## REFERENCES

De Finetti, B.: (1964) 'Foresight: Its Logical Laws, Its Subjective Sources', translated from the French original by H. E. Kyburg, *Studies in Subjective Probability*, 1st edition, edited by H. E. Kyburg and H. E. Smokler, New York: Wiley, pp. 95–158.

De Finetti, B.: (1972) *Probability, Induction and Statistics*, London: Wiley.

Gärdenfors, P. and Sahlin, N.-E.: (1982) 'Unreliable Probabilities, Risk Taking and Decision Making,' *Synthese*, this issue, pp. 361–86.

Good, I. J.: (1960) 'Subjective Probability as the Measure of a Non-measureable Set,' E. Nagel, P. Suppes and A. Tarski (eds.) *Logic, Methodology and Philosophy of Science*, Stanford: Stanford University Press (1962), pp. 319–329.

Levi, I.: (1962) 'On the Seriousness of Mistakes,' *Philosophy of Science* **29**, pp. 47–65.

Levi, I.: (1967) *Gambling with Truth*, New York: Knopf, reissued by MIT Press in 1973.

Levi, I.: (1974) 'On Indeterminate Probabilities,' *Journal of Philosophy* **71**, pp. 391–417. Reissued with some added commentary in C. A. Hooker, J. J. Leach and E. F. McClennen (eds.), *Foundations and Applications of Decision Theory*, v. I, Dordrecht: D. Reidel, 1978, pp. 233–261.

Levi, I.: (1977) 'Four Types of Ignorance,' *Social Research* **44**, pp. 745–756.

Levi, I.: (1980) *The Enterprise of Knowledge*, Cambridge, Mass.: MIT Press.

Levi, I.: (1982) 'Direct Inference and Confirmational Conditionalization,' *Philosophy of Science* **48**, 532–552.

Luce, R. D. and Raiffa, H.: (1958) *Games and Decisions*, New York: Wiley.

Savage, L. J.: (1954) *The Foundations of Statistics*, New York: Wiley revised slightly and reissued by Dover in 1972.

Seidenfeld, T.: 'Why I am not an Objective Bayesian,' *Theory and Decision* **11**, pp. 413–440.

Sen, A. K.: (1970) *Collective Choice and Social Welfare*, Holden Day.

Wald, A.: (1950) *Statistical Decision Functions*, New York: Wiley.