PETER GÄRDENFORS AND NILS-ERIC SAHLIN*

UNRELIABLE PROBABILITIES, RISK TAKING, AND DECISION MAKING

1. THE LIMITATIONS OF STRICT BAYESIANISM

A central part of Bayesianism is the doctrine that the decision maker's knowledge in a given situation can be represented by a subjective probability measure defined over the possible states of the world. This measure can be used to determine the expected utility for the agent of the various alternatives open to him. The basic decision rule is then that the alternative which has the maximal expected utility should be chosen.

A fundamental assumption for this strict form of Bayesianism is that the decision maker's knowledge can be represented by a *unique* probability measure. The adherents of this assumption have produced a variety of arguments in favor of it, the most famous being the so-called Dutch book arguments. A consequence of the assumption, in connection with the rule of maximizing expected utility, is that in two decision situations which are identical with respect to the probabilities assigned to the relevant states and the utilities of the various outcomes the decisions should be the same.

It seems to us, however, that it is possible to find decision situations which are identical in all the respects relevant to the strict Bayesian, but which nevertheless motivate different decisions. As an example to illustrate this point, consider Miss Julie who is invited to bet on the outcome of three different tennis matches.[1] As regards match A, she is very well-informed about the two players – she knows everything about the results of their earlier matches, she has watched them play several times, she is familiar with their present physical condition and the setting of the match, etc. Given all this information, Miss Julie predicts that it will be a very even match and that a mere chance will determine the winner. In match B, she knows nothing whatsoever about the relative strength of the contestants (she has not even heard their names before) and she has no other information that

Synthese **53** (1982) 361–386. 0039–7857/82/0533–0361 $02.60

is relevant for predicting the winner of the match. Match C is similar to match B except that Miss Julie has happened to hear that one of the contestants is an excellent tennis player, although she does not know anything about which player it is, and that the second player is indeed an amateur so that everybody considers the outcome of the match a foregone conclusion.

If pressed to evaluate the probabilities of the various possible outcomes of the matches, Miss Julie would say that in all three matches, given the information she has, each of the players has a 50% chance of winning. In this situation a strict Bayesian would say that Miss Julie should be willing to bet at equal odds on one of the players winning in one of the matches if and only if she is willing to place a similar bet in the two other matches. It seems, however, perfectly rational if Miss Julie decides to bet on match A, but not on B or C, for the reason that a bet on match A is more *reliable* than a bet on the others. Furthermore she would be very suspicious of anyone offering her a bet at equal odds on match C, even if she could decide for herself which player to back.

The main point of this example is to show that the amount and quality of information which the decision maker has concerning the possible states and outcomes of the decision situation in many cases is an important factor when making the decision. In order to describe this aspect of the decision situation, we will say that the information available concerning the possible states and outcomes of a decision situation has different degrees of *epistemic reliability*. This concept will be further explicated later. We believe that the epistemic reliability of a decision situation is *one* important factor when assessing the *risk* of the decision. In our opinion, the major drawback of strict Bayesianism is that it does not account for the variations of the epistemic reliability in different decision situations.

The concept of epistemic reliability is useful also in other contexts than direct decision making. In the next section, after presenting the models of decision situations, we will apply this concept in a discussion of Popper's 'paradox of ideal evidence'.

In order to determine whether empirical support could be obtained for the thesis that the epistemic reliability of the decision situation affects the decision, Goldsmith and Sahlin [15], [16] and [17] performed a series of experiments. In one of these, test subjects were first presented with descriptions of a number of events and were asked to estimate for each event the probability of its occurrence. Some events were of the well-known parlor game type, e.g. that the next card drawn from an ordinary deck of cards will be a spade; while other events were ones about which the subjects presumably had very limited information, e.g. that there will be a bus strike in Verona, Italy next week. Directly after estimating the probability of an event, subjects were asked to show, on a scale from 0 to 1, the perceived reliability of their probability estimate. The experiment was constructed so that for each subject several sets of events were formed, such that all the events in a set had received the same probability estimate but the assessed reliability of the various estimates differed. For each set, the subject was then asked to choose between lottery tickets involving the same events, where a ticket was to be conceived as yielding a win of 100 SwKr if the event occurred but no monetary loss if it did not occur. One hypothesis that obtained support in this experiment was that for probabilities other than fairly low ones, lottery tickets involving more reliable probability estimates tend to be preferred. This, together with the results of similar experiments, suggested the reliability of probability estimates to be an important factor in decision making.

The aim of the present paper is to outline a decision theory which is essentially Bayesian in its approach but which takes epistemic reliability of decision situations into consideration. We first present models of the knowledge relevant in a decision situation. One deviation from strict Bayesianism is that we use a class of probability measures instead of only one to represent the knowledge of an agent in a given decision situation. Another deviation is that we add a new measure which ascribes to each of these probability measures a degree of epistemic reliability. The first step in a decision, according to the decision theory to be presented here, is to select a class of probability measures with acceptable degrees of reliability on which a decision is to be based. Relative to this class, one can then, for each decision alternative, compute the minimal expected utility of the alternative. In the second step the alternative with the largest minimal expected utility is chosen.[2] This decision theory is then compared to some other generalized Bayesian decision theories, in particular Levi's theory as presented in [26].

2. MODELS OF DECISION SITUATIONS

Our description of a decision situation will have many components in common with the traditional Bayesian way of describing decision problems. A decision is a choice of one of the *alternatives* available in a given situation. For simplicity, we will assume that in any decision situation there is a finite set $\mathscr{A} = \{a_1, a_2, \ldots, a_n\}$ of alternatives.

Though the decision maker presumably has some control over the factors which determine the outcome of the decision, he does not, in general, have complete control. The uncertainty as to what the outcome of a chosen alternative will be is described by referring to different *states of nature* (or just *states*, for brevity). We will assume that, in any given decision situation, only a finite number of states are relevant to the decision. These states will be denoted $s_1, s_2, \ldots, s_m$.

The result or *outcome* of choosing the alternative a_i, if the true state of nature is s_j will be denoted o_{ij}. An important factor when making a decision is that of the *values* the decision maker attaches to outcomes. We will make the standard assumption that this valuation can be represented by a utility measure u.[3] The utility of the outcome o_{ij} will be denoted u_{ij}. It is assumed that all information on how the decision maker values the outcomes is summarized by the utility measure.

A final factor in describing a decision situation is that of the beliefs the decision maker has concerning which of the possible states of nature is the true state. Within strict Bayesianism it is assumed that these beliefs can be represented by a single probability measure defined over the states of nature. This assumption is very strong since it amounts to the agent having *complete* information in the sense that he is *certain* of the probabilities of the possible states of nature. The assumption is unrealistic, since it is almost only in mathematical games with coins and dice that the agent has such complete information, while in most cases of practical interest the agent has only partial information about the states of nature.

In the strict form of Bayesianism which is advocated by de Finetti [10] and Savage [35] among others, it is assumed that the agent's subjective probability of a state of nature can be determined by his inclination to accept bets concerning the state.[4] The so-called Dutch book theorem states that if it is not possible to construct a bet where the agent will lose money independently of which state turns out to be the actual one, then the agent's degrees of beliefs satisfy Kolmogoroff's axioms, i.e., there is a unique probability measure that describes these degrees of belief.

However, a presupposition of this theorem is that the agent be willing to take either side of a bet, i.e., if the agent is not willing to bet on the state s_j at odds of $a:b$, then he should be willing to bet on not-s_j at odds of $b:a$. But this assumption makes too heavy demands on people's willingness to make bets. One is often not willing to accept either of the two bets.[5] In our opinion, this is explained by the fact that the estimated probability of the different states of nature are unreliable and one is not willing to take the risk connected with this uncertainty.[6] This criticism is directed against the assumptions behind the Dutch book theorem, but similar criticism can be constructed against other arguments in favor of the assumption of representing beliefs by a unique probability measure.

In this paper we will relax this assumption and, as a first step in the description of the beliefs which are relevant in a decision situation, we instead assume that the beliefs about the states of nature can be represented by a set $\mathcal{P}$ of probability measures. The intended interpretation of the set $\mathcal{P}$ is that it consists of all epistemically possible probability measures over the states of nature, where we conceive of a probability measure as epistemically possible if it does not contradict the decision maker's knowledge in the given decision situation.[7] In this way, we associate with each state s_j a set of probability values $P(s_j)$, where $P \in \mathcal{P}$. The values may be called the epistemically possible probabilities of the state s_j. For simplicity, we will assume that the probabilities of the outcomes o_{ij} are independent of which alternative is chosen, so that $P(o_{ij}) = P(s_j)$, for all $P \in \mathcal{P}$ and all alternatives a_i. Since this assumption can be relaxed, the decision theory to be presented can be extended to the more general case.[8]

The idea of representing a state of belief by a class of probability measures is not new but has been suggested by various authors.[9] It has been most extensively discussed by Levi in [26] and [27], but, as will be seen in the sequel, he does not use the class of probability measures in the same way as we do.

Levi also assumes that the set of probability measures is *convex*, i.e., that if P and P' are two measures in the set, then the measure $\alpha \cdot P + (1-\alpha) \cdot P'$ is also in the set, for any α between 0 and 1.[10] The motivation for this assumption is that if P and P' both are possible probability distributions over the states, then any mixture of these distributions is also possible. We will discuss the requirement of convexity in section 5.

If $\mathcal{P}$ is assumed to be convex, then the set of epistemically possible probabilities associated with a state s_j by the elements of $\mathcal{P}$ will form an *interval* from the lowest probability assigned to s_j to the highest. Some authors have taken such intervals as basic when describing beliefs about the states of nature – to each state is assigned a probability interval and this assignment is governed by some consistency restrictions.[11] The representation by a convex set of probability measures is, however, more general, since from such a set one can always compute a unique set of associated intervals, but starting from an assignment of consistent probability intervals, there will in general be a large number of convex sets of probability measures that will generate the intervals.[12]

We believe that not all of an agent's beliefs about the states of nature relevant to a decision situation can be captured by a set $\mathcal{P}$ of probability measures. As a second element in describing the beliefs relevant to a decision situation, we introduce a (real-valued) measure ρ of the *epistemic reliability* of the probability measures in $\mathcal{P}$. Even if several probability distributions are epistemically possible, some distributions are more reliable – they are backed up by more information than other distributions.

The measure ρ is intended to represent these different degrees of reliability. In the introductory examples, Miss Julie ascribes a much greater epistemic reliability to the probability distribution where each player has an equal chance of winning in match A where she knows a lot about the players than in match B where she knows nothing relevant about the players. In match C, where she knows that one player is superior to the other, but not which, the epistemically most reliable distributions are the two distributions where one player is certain to win. Since there are only two relevant states of nature in these examples, viz. the first player wins (s_1) and the second player wins (s_2), a probability distribution can be described simply by the probability of one of the states. We can then illustrate the epistemic reliability of the various distributions in the three matches by diagrams as in Figure 1.

Even if examples such as these illustrate the use of the measure ρ, its properties should be specified in greater detail. Technically, the

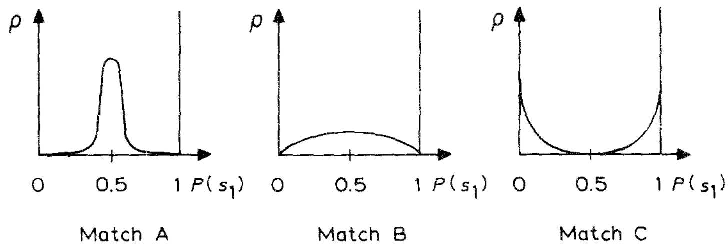

Fig. 1.

only property of ρ that will be needed in this paper is that the probability distributions in $\mathcal{P}$ can be *ordered* with respect to their degrees of epistemic reliability. However, it seems natural to postulate that ρ has an upper bound, representing the case when the agent has complete information about a probability distribution, and a lower bound, representing the case when the agent has no information at all about these distributions. However, we will not attempt a full description of the properties of the measure ρ, since we believe that this can be done only in a more comprehensive decision theory in which the relations between different decision situations are exploited.[13]

A fundamental feature of the epistemic reliability of the probability distributions possible in a decision situation, as we conceive of the measure, is that the less relevant information the agent has about the states of nature, the less epistemic reliability will be ascribed to the distributions in $\mathcal{P}$. Where little information is available, therefore, all distributions will, consequently, have about the same degree of epistemic reliability. Conversely, in a decision situation where the agent is well-informed about the possible states of nature, some distributions will tend to have a considerably higher degree of epistemic reliability than others.

A problem which strongly supports our thesis that the measure of epistemic reliability is a necessary ingredient in the description of a decision situation is Popper's paradox of ideal evidence. Popper asks us to consider the following example ([31], pp. 407–408):

Let z be a certain penny, and let a be the statement 'the nth (as yet unobserved) toss of z will yield heads'. Within the subjective theory, it may be assumed that the absolute (or prior) probability of the statement a is equal to 1/2, that is to say,

(1) $P(a) = 1/2$

Now let e be some *statistical evidence*; that is to say, a *statistical report*, based upon the observation of thousands or perhaps millions of tosses of z; and let this evidence e be *ideally favourable* to the hypothesis that z is strictly symmetrical.... We then have no other option concerning $P(a, e)$ than to assume that

(2) $\quad P(a, e) = 1/2$

This means that the probability of tossing heads remains unchanged in the light of the evidence e; for we now have

(3) $\quad P(a) = P(a, e).$

But, according to the subjective theory, (3) means that e is, on the whole (absolutely) *irrelevant* information with respect to a.

Now this is a little startling; for it means, more explicitly, that our so-called *'degree of rational belief' in the hypothesis, a, ought to be completely unaffected by the accumulated evidential knowledge, e*; that the absence of any statistical evidence concerning z justifies precisely the same 'degree of rational belief' as the weighty evidence of millions of observations which, *prima facie*, support or confirm or strengthen our belief.

The 'subjective theory' which Popper is referring to in this example is what we have here called strict Bayesianism.

Now, with the aid of the models of decision situations presented above, we can describe Popper's example in the following way. There is a set $\mathcal{P}$ of possible probability measures concerning the states of nature described by a and not-a. If one is forced to state the probability of a, before the evidence e is obtained, the most reliable answer seems to be 1/2. The degree of epistemic reliability of this estimate is, however, low, and there are many other answers which seem almost as reliable. After the evidence e is obtained, the most reasonable probability assessment concerning a is still 1/2, but now the distribution associated with this answer has a much higher degree of epistemic reliability than before, and the other distributions in $\mathcal{P}$ have correspondingly lower degrees of reliability. It should be noted that this distinction between the two cases cannot be formulated with the aid of the set $\mathcal{P}$ only, but the measure ρ of epistemic reliability is also necessary.[14]

We will conclude this section by briefly mentioning some related attempts to extend models of belief by some measure of 'reliability'.[15] An interesting concept was introduced by Keynes in [25], p. 71:

As the relevant evidence at our disposal increases, the magnitude of the probability of the argument may either decrease or increase, according as the new knowledge strengthens the unfavourable or favourable evidence; but *something* seems to have increased in either case – we have a more substantial basis upon which to rest our conclusion. I express this by saying that an accession of new evidence increases the *weight* of an argument. New evidence will sometimes decrease the probability of an argument, but it will always increase its "weight".

Here Keynes writes about the probability of an argument, while we are concerned with probability distributions over states of nature. Even if the intuitions behind Keynes's 'weight of evidence' and our 'epistemic reliability' are related, it is difficult to say how far this parallel can be drawn.

In [4], pp. 554–555, Carnap discusses 'the problem of the reliability of a value of degree of confirmation' which obviously is the same as Keynes's problem. Carnap remarks that Keynes's concept of 'the weight of evidence' was forestalled by Peirce who mentioned it in [30], p. 421, in the following way:

... to express the proper state of belief, not *one* number but *two* are requisite, the first depending on the inferred probability, the second on the amount of knowledge on which that probability is based.

The models of the agent's beliefs about the states of nature in a decision situation which have been presented here contain the two components $\mathcal{P}$ and ρ, i.e., the set of epistemically possible probability distributions, and the measure of epistemic reliability. These two components can be seen as an explication of the two numbers required by Peirce.[16]

3. A DECISION THEORY

The models of decision situations which were outlined in the previous section will be used now as a basis for a theory of decision. This theory can be seen as a generalization of the Bayesian rule of maximizing expected utility.

A decision, i.e. a choice of one of the alternatives in a decision situation, will be arrived at in two steps. The first step consists in restricting the set $\mathcal{P}$ to a set of probability measures with a 'satisfactory' degree of epistemic reliability. The intuition here is that in a given decision situation certain probability distributions over the states of nature, albeit epistemically possible, are not considered as serious possibilities. For example, people do not usually check whether there is too little brake fluid in the car or whether the wheels are loose before starting to drive, although, for all they know, such events are not impossible, and they realize that if any such event occurred they would be in danger.

Now examples of this kind seems to show that what the agent does is to disregard certain states of nature rather than probability distributions over such states. But if a certain state of nature is not considered as a serious possibility, then this means that all probability distributions which assign this state a positive probability are left out of consideration. And there may be cases when some probability distributions are left out of account even if all relevant states of nature are considered to be serious possibilities. So, restricting the set $\mathcal{P}$ is a more general way of modelling the process than restricting the set of states.

Deciding to consider some distributions in $\mathcal{P}$ as not being serious possibilities means that *one takes a risk.* The less inclined one is to take risks, the greater the number of distributions in $\mathcal{P}$ will be that are taken into account when making the decision.

A fundamental question is how the agent determines which probability distributions in $\mathcal{P}$ are 'satisfactorily reliable' and which are not. In our view, the answer is that the measure ρ of epistemic reliability should be used when selecting the appropriate subset $\mathcal{P}/\rho_0$ of $\mathcal{P}$. The agent selects a *desired level* ρ_0 *of epistemic reliability* and only those probability distributions in $\mathcal{P}$ which pass this ρ-level are included in $\mathcal{P}/\rho_0$, but not the others.[17] $\mathcal{P}/\rho_0$ can be regarded as the set of probability distributions that the agent takes into consideration in making the decision. An obvious requirement on the chosen level of reliability is, of course, that there be *some* distribution in $\mathcal{P}$ which passes the level.

Which 'desired level of epistemic reliability' the agent will choose depends on how large the risks are he is willing to take. The more risk aversive the agent is, the lower the chosen level of epistemic reliability will be. It is important to note that two agents in identical epistemic situations, here identified by a set $\mathcal{P}$ and a measure ρ, may indeed choose different values of ρ_0 depending on their different risk taking tendencies. This is the reason why it is assumed that ρ yields an ordering of $\mathcal{P}$ and not merely a dichotomy.

If the agent is willing to take a maximal risk as regards which probability distribution to consider as 'satisfactory' it may happen that there will be only one distribution which passes the desired level of reliability. After such a choice his relevant information about the states of nature will be of the same type as for the strict Bayesian, i.e. a unique probability measure, but this situation will arise for quite different reasons.[18]

The second step in the decision procedure starts with the restricted set $\mathcal{P}/\rho_0$ of probability distributions. For each alternative a_i and each probability distribution P in $\mathcal{P}/\rho_0$ the expected utility e_{ik} is computed in the ordinary way. The *minimal expected utility* of an alternative a_i, relative to a set $\mathcal{P}/\rho_0$, is then determined, this being defined as the lowest of these expected utilities e_{ik}. Finally the decision is made according to the following rule:[19]

The maximin criterion for expected utilities (*MMEU*): The alternative with the largest minimal expected utility ought to be chosen.

In order to illustrate how this decision procedure works, we will return to the introductory examples. For simplicity, let us, for all three tennis matches, denote by s_1 the event that the first player to serve wins the match, and by s_2 the event that the other player wins. Now assume that Miss Julie is offered the following bet for each of the three matches: She wins 30$ if s_1 occurs and loses 20$ if s_2 occurs. For each match, she must choose between the alternative a_1 of accepting the bet and the alternative a_2 of declining it. Let us furthermore assume that Miss Julie's utilities are mirrored by the monetary values of the outcomes.

In match *A*, where Miss Julie is very well informed about the players, she considers that the only probability distribution that she needs to take into consideration is the distribution P_1, where $P_1(s_1) = P_1(s_2) = 0.5$. She is willing to take the (small) risk of letting $\mathcal{P}/\rho_0$ consist of this distribution only. The only, and hence minimal, expected utility to compute is then $0.5 \cdot 30 + 0.5 \cdot (-20)$ for a_1 and $0.5 \cdot 0 + 0.5 \cdot 0$ for a_2. Hence, according to MMEU, she should accept the bet in match *A*.

In match *B*, where she has no relevant information at all, the epistemic reliability of the epistemically possible distributions is more evenly spread out. Consequently, Miss Julie is not so willing to leave distributions out of account when forming the subset $\mathcal{P}/\rho_0$ as in the previous case. For simplicity, let us assume that $\mathcal{P}/\rho_0 = \{P_1, P_2, P_3\}$, where P_1 is as before, P_2 is defined by $P_2(s_1) = 0.25$ and $P_2(s_2) = 0.75$, and P_3 is defined by $P_3(s_1) = 0.75$ and $P_3(s_2) = 0.25$. The expected utilities for the alternative a_1 of accepting the bet, determined by these three distributions, are 5, −7.5 and 17.5 respectively; while the expected utilities for the alternative a_2 of not accepting the bet are 0 in all three cases. Since the minimal expected utility of a_1 is less than that of a_2, the MMEU criterion demands that a_2 be the chosen alternative, i.e. Miss Julie should decline the bet offered.

In match *C*, it is reasonable to assume that $\mathcal{P}/\rho_0$ contains some probability distribution which assigns s_1 a very high probability and some distribution which assigns it a very low probability. A similar analysis as above then shows that Miss Julie should not accept the bet in this case either.

This example can be heuristically illustrated as in Figure 2. In this figure the broken horizontal line indicates the desired level of epistemic reliability.

To give a further illustration of the decision theory, it should be noted that the hypothesis from the Goldsmith–Sahlin experiments, mentioned in the introduction, is well explained by the MMEU criterion. When an agent is asked to choose between tickets in two lotteries which are estimated to have the same primary probability of winning, he should choose the ticket from the lottery with the epistemically most reliable probability estimate, since this alternative will have the highest minimal expected utility.[20] Still other applications of the decision theory will be presented in the next two sections.

A limiting case of information about the states of nature in a decision situation is to have *no* information at all. In the decision models presented here, this would mean that all probability distributions over the states are epistemically possible and that they have equal epistemic reliability. In such a case, the minimal expected

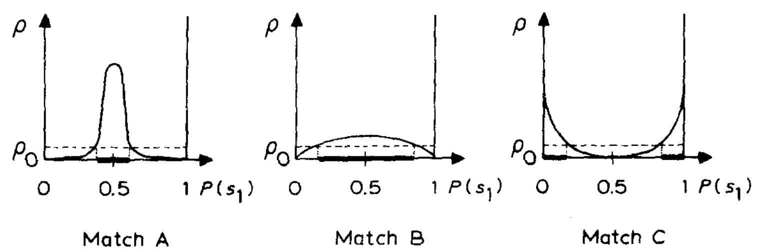

Fig. 2.

utility of an alternative is obtained from the distribution which assigns the probability 1 to the worst outcome of the alternative. This is, however, just another way of formulating the classical maximin rule, which has been applied in what traditionally has been called 'decision making under uncertainty' (a more appropriate name would be 'decision making under ignorance'). Hence, the classical maximin turns out to be a special case of the MMEU criterion.

At the other extreme, having *full* information about the states of nature implies that only one probability distribution is epistemically possible.[21] In this case the MMEU criterion collapses into the ordinary rule within strict Bayesianism, i.e. the rule of maximizing expected utility, which has been applied to what traditionally, but somewhat misleadingly, has been called 'decision making under risk'.

The decision theory which has been presented here thus covers the area between the traditional theories of 'decision making under uncertainty' and 'decision making under risk' and it has these theories as limiting cases.

4. RELATION TO EARLIER THEORIES

Several authors have proposed decision theories which are based on more general ways of representing the decision maker's beliefs about the states than what is allowed by strict Bayesianism. The most detailed among these is Levi's theory ([26] and [27]), which will be discussed in a separate section. In this section we will compare the present theory with some earlier statistical decision theories.

In [38], Wald formulates a theory of 'statistical decision functions' where he considers a set Ω of probability measures and a 'risk' function. He says that 'the class Ω is to be regarded as a datum of the decision problem' (p. 1). He also notes that the class Ω will generally vary with the decision problem at hand and that in most cases 'will be a proper subset of the class of all possible distribution functions' (p. 1). In the examples, Ω is often taken to be a parametric family of known functional form. A risk function is a function which determines the 'cost' of a wrong decision. Such a function can be seen as an inverted utility function restricted to negative outcomes. On this interpretation it is easier to compare Wald's theory to the present decision theory.

Wald suggests two alternative decision rules. The first is the traditional Bayesian method when an 'a priori' probability measure in Ω can be selected, or, as Wald puts it, when 'it exists and is known to the experimenter' (p. 16). The second case is when the entire set Ω is employed to determine the decision. Wald suggests that in such cases one should 'minimize the maximum risk'. In our terminology, using utility functions instead of risk functions, this is the same as maximizing the minimal expected utility with respect to the set Ω.

If Wald's theory is interpreted as above, the difference between his and our theory is mainly of epistemological character. Since Wald does not say anything about how Ω is to be determined it is difficult to tell whether it corresponds to our set $\mathcal{P}$ of epistemically possible probability functions or to the set $\mathcal{P}/\rho_0$ of 'reliable' functions. In particular, he does not introduce any factor corresponding to the measure ρ of epistemic reliability, nor does he associate the choice of Ω with any form of risk taking.

Hurwicz [22] apparently interprets Wald's set Ω as corresponding to our set $\mathcal{P}$. He notes that sometimes some of the distributions in Ω are more 'likely' than others (p. 343). For example, assume that Ω consists of all normal distributions with mean zero and standard deviation σ. In a particular decision situation evidence at hand may support the assumption that σ is considerably small. It thus seems reasonable to select a proper subset Ω_0 of Ω which is restricted to those distributions with standard deviation less than or equal to some value σ_0.

Hurwicz assumes that such a subset Ω_0 ($\Xi^{(0)}_{\mathcal{F}}$ in his terminology) of Ω can be selected in any decision situation. He then suggests a 'generalized Bayes-minimax porinciple' which amounts to using Ω_0 as the base when maximizing the minimal expected utility (minimizing the maximal risk). Obviously, the set Ω_0 corresponds closely to our set $\mathcal{P}/\rho_0$. The main difference between Hurwicz' theory and the present one is that he does not give any account of how the set Ω_0 is to be determined. In particular he, as Wald, does not introduce any factor corresponding to the measure ρ.

In [21], Hodges and Lehmann suggest an alternative to Wald's minimax solution which they call a 'restricted Bayes solution'. It is of interest here since it is adopted by Ellsberg [8] as a solution to his 'paradox'. Let us start by considering Ellsberg's problem.[22]

Ellsberg ([8], pp. 653–654) asks us to consider the following decision problem. Imagine an urn known to contain 30 red balls and

60 black and yellow balls, the latter in unknown proportion. One ball is to be drawn at random from the urn. In the first situation you are asked to choose between two alternatives a_1 and a_2. If you choose a_1 you will receive \$100 if a red ball is drawn and nothing if a black or yellow ball is drawn. If you choose a_2 you will receive \$100 if a black ball is drawn, otherwise nothing. In the second situation you are asked to choose, under the same circumstances, between the two alternatives a_3 and a_4. If you choose a_3 you will receive \$100 if a red or a yellow ball is drawn, otherwise nothing and if you choose a_4 you will receive \$100 if a black or yellow ball is drawn, otherwise nothing. This decision problem is shown in the following decision matrix.

	Red	Black	Yellow
a_1	\$100	\$0	\$0
a_2	\$0	\$100	\$0
a_3	\$100	\$0	\$100
a_4	\$0	\$100	\$100

The most frequent pattern of response to these two decision situations is that a_1 is preferred to a_2 and a_4 is preferred to a_3. As Ellsberg notes, this preference pattern violates Savage's 'sure thing principle' (postulate $P2$ in [35]), which requires that the preference ordering between a_1 and a_2 be the same as the ordering between a_3 and a_4.

When applying the present decision theory to this problem, the main step is to determine the set $\mathcal{P}/\rho_0$. The set $\mathcal{P}$ should be the same in the two decision situations, since they do not differ with respect to the information about the states. Now, unless the decision maker believes that he is being cheated about the content of the urn, $\mathcal{P}$ is most naturally taken as the class of distributions $(1/3,\ x,\ 2/3-x)$, where x varies from 0 to 2/3.

If the decision maker chooses a low ρ_0, $\mathcal{P}/\rho_0$ will contain most of the distributions in this class. For simplicity, let us for the moment assume that $\mathcal{P}/\rho_0 = \mathcal{P}$. With this choice, the minimal expected utilities of the alternatives are $1/3 \cdot u(\$100)$, $1 \cdot u(\$0)$, $1/3 \cdot u(\$100)$ and $2/3 \cdot u(\$100)$, for a_1, a_2, a_3 and a_4, respectively. Assuming that $1/3 \cdot u(\$100)$ is greater than $u(\$0)$, the MMEU criterion requires that a_1

be preferred to a_2 and a_4 to a_3, which accords with Ellsberg's findings. Intuitively, a_2 and a_3 involve greater 'epistemic risks' than a_1 and a_4 – thus a_2 and a_3 are avoided by most subjects. This feature is well captured by the present decision theory.

If the decision maker is willing to take an epistemic risk and chooses a higher ρ_0, fewer distributions will be included in $\mathcal{P}/\rho_0$. If the decision maker has no further information about the distribution of black and yellow balls, then, because of symmetry, it is likely that he judges the distribution (1/3, 1/3, 1/3) as being the highest in the ρ-ordering. If this is the only distribution included in $\mathcal{P}/\rho_0$, he should be indifferent between a_1 and a_2 and between a_3 and a_4 according to the MMEU criterion. In this case Savage's sure thing principle is not violated.[23]

In order to explain the 'paradox' that decision makers do not act according to Savage's sure thing principle, Ellsberg ([8], p. 661) first introduces for a decision maker in a given situation a set Y^0 of probability distributions

> that still seem "reasonable", reflecting judgements 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.

Ellsberg also considers a particular probability distribution y^0, the "estimated" probability distribution, which can be viewed as the distribution a strict Bayesian decision maker would have adopted in the decision situation. Ellsberg also ascribes to y^0 a degree ρ_e (e for Ellsberg) of the decision maker's 'confidence' in the estimate.

The decision rule suggested by Ellsberg, which is the restricted Bayesian solution developed by Hodges and Lehmann, can now be described as follows: Compute for each action a_i the expected utility according to the distribution y^0 and the minimal expected utility relative to the set Y^0. Associate with each action an index based on a weighed average of these two factors, where ρ_e is the weight ascribed to the first factor and $1 - \rho_e$ is the weight ascribed to the latter factor. Finally, choose that act with the highest index.

When comparing this decision theory with the theory presented in this paper, one notes that there are differences both of epistemological and formal character. Firstly, since Ellsberg, like his predecessors, does not say anything about how the set Y^0 is to be determined, and it is therefore difficult to say whether it corresponds to our set $\mathcal{P}/\rho_0$.

Secondly, even if we identify Y^0 with $\mathcal{P}/\rho_0$, Ellsberg exploits the degree ρ_e of 'confidence', which is defined for only one distribution y^0, in a way that differs considerably from our use of the measure ρ, which is assumed to be defined for all distributions in $\mathcal{P}$. In particular we need not assume that ρ gives a numerical value, only that it orders the distributions in $\mathcal{P}$. Thirdly, since the decision rules are different, the theories will recommend different decisions in many situations. The most important disagreement here is that we reject the need of an estimated distribution y^0. We believe that once the set $\mathcal{P}/\rho_0$ has been selected, the distribution with the highest degree of reliability (corresponding to Ellsberg's y^0) does not play any outstanding role in the decision making.

This difference between the two theories is in principle testable, assuming that ρ_e is not always close to zero. The experiment performed by Becker and Brownson [2] is relevant here. They offered subjects ambiguous bets differing in the range of probabilities possible but of equal expected value and found that subjects were willing to pay money for obtaining bets with narrower ranges of probability. This finding seems to highlight the importance of a measure of reliability. But they did not, however, obtain support for Ellsberg's hypothesis that the distribution y^0 is relevant for the decision making.

5. A COMPARISON WITH LEVI'S THEORY

In this section we compare our theory with Levi's theory which is presented in [26] and elaborated on in [27]. Levi starts out from a description of the decision maker X's information at the time t about the states of nature. This information is contained in a convex set $B_{X,t}$ of probability distributions. The distributions in $B_{X,t}$ are, according to Levi, the 'permissible' distributions. As to the meaning of "permissible", he offers only indirect clarification by indicating the connections between permissibility and rational choice. In order to compare the theories, we will here assume that the set $B_{X,t}$ corresponds to the set $\mathcal{P}/\rho_0$ (or its convex hull) as presented in section 3.[24]

Levi also generalizes the traditional way of representing the utilities of the outcomes by introducing a class G of 'permissible' utility measures, such that not all of these measures need be linear transformations of one another.

An alternative a_i is said to be *E-admissible* if and only if there is some probability distribution P in $B_{X,t}$ and some utility function u in G such that the expected utility of a_i relative to P and u is maximal among all the available alternatives. A first requirement on the alternative to be chosen in a given decision situation is then that it should be E-admissible.

The second step in Levi's decision procedure concerns the opportunity to defer decision between two or more E-admissible alternatives. He argues that a rational agent should "keep his options open" whenever possible. An alternative is said to be *P-admissible* if it is E-admissible and it is 'best' with respect to E-admissible option preservation. Levi does not, however, explicate what he means by 'best' here, and we will ignore the effects of this requirement here, since we have not imposed any structure on the set of alternatives.[25]

Let us say that a P-admissible alternative a_i is *security optimal relative to* a utility function u if and only if the minimum u-value assigned to some possible outcome o_{ij} of a_i is at least as great as the minimal u-value assigned to any other P-admissible alternative. Levi then, finally, calls an alternative *S-admissible* if it is P-admissible and security optimal relative to *some* utility function in G.

Levi states (in [27], p. 412) that he "cannot think of additional criteria for admissibility which seem adequate" so he, tentatively, assumes that all S-admissible alternatives are 'admissible' for the final choice, which then, supposedly, is determined by some random device.

In order to illustrate the differences between the decision theory presented in the previous section and Levi's theory, we will consider the following example which contains two states and three alternatives:

	s_1	s_2
a_1	−10	12
a_2	11	−9
a_3	0	0

In this matrix the numbers denote the utilities of the outcomes. Assume that the set $\mathcal{P}/\rho_0$, which here is identified with Levi's set $B_{X,t}$, consists of the two probability distributions P, defined by $P(s_1) = 0.4$ and $P(s_2) = 0.6$, and P', defined by $P'(s_1) = 0.6$ and $P'(s_2) = 0.4$, together with all convex combinations of P and P'.

The minimal expected utility of a_1 is -1.2 and the minimal expected utility of a_2 is -1.0. The minimal expected utility of a_3 is of course 0, so MMEU requires that a_3 be chosen.

In contrast to this, only a_1 and a_2 are E-admissible. P-admissibility has no effect here, but a_2 is security optimal relative to the utility measure given in the matrix, so a_2 is the only S-admissible alternative. Thus, according to Levi's theory, a_2 should be chosen.

When the uncertainty about the states of nature in this decision situation, represented by $\mathcal{P}/\rho_0$, is considered, we believe that a_3 is intuitively the best alternative. Against this it may be argued that using a maximin principle is unnecessarily risk aversive. It should be remembered, however, that when restricting $\mathcal{P}$ to $\mathcal{P}/\rho_0$ the agent is already taking a risk and his choice of $\mathcal{P}/\rho_0$ indicates that he is not willing to take any further epistemic risks. On the other hand, Levi's requirement of E-admissibility has the consequence that, in many cases, the choices made by his theory seem unrealistically optimistic.

A strange feature of Levi's theory is that if the previous decision situation is restricted to a choice between a_2 and a_3, then his theory recommends choosing a_3 instead of a_2! In their chapter on individual decision making under uncertainty, Luce and Raiffa ([29], pp. 288–290) introduces the condition of *independence of irrelevant alternatives* which in its simplest form demands that if an alternative is not optimal in a decision situation it cannot be made optimal by adding new alternatives to the situation. The example presented here shows that Levi's theory does not satisfy this condition since in the decision situation where a_2 and a_3 are the only available alternatives and where a_3 is optimal according to Levi's theory, a_2 can be made optimal by adding a_1.[26] It is easy to show, however, that the MMEU criterion which has been presented here satisfies the condition of independence of irrelevant alternatives.[27]

Levi's theory also seems to have problems in explaining some of the experimental results considered in this paper. If Levi's theory is applied to Ellsberg's two decision situations as presented earlier, it gives the result that both alternatives in the two situations are S-admissible, and hence that a_1 is equally good as a_2 and a_3 is equally good as a_4. This contrasts with Ellsberg's findings which are in accordance with the recommendations of the present theory. Similar considerations apply to the experiments presented in Becker and Brownson [2].

Decision theories of the kind presented in this paper are based on several idealizations and they will unavoidably be exposed to some refractory empirical material. We believe, however, that the considerations of this section show that the decision theory presented in section 3 is a more realistic theory than Levi's.

6. CONCLUSION

The starting-point of this paper is that in many decision situation the assumption, made in the strict form of Bayesianism, that the beliefs of an agent can be represented by a single probability distribution is unrealistic. We have here presented models of belief which contain firstly, a class of probability distributions, and, secondly, a measure of the epistemic reliability of these probability distributions. Several authors before us have suggested that a class of probability distributions should be exploited when describing the beliefs of the agent. A main thesis of this paper is that this is not sufficient, but an assessment of the information on which the class of probability distributions is based is also necessary. We have here tried to capture this assessment by the measure ρ of epistemic reliability. With the aid of this measure we can account for one form of risk taking in decision situations.

On the basis of the models of the beliefs which are relevant in a decision situation we have formulated a decision theory. We have argued that this theory has more desirable properties and is better supported than other decision theories which also generalize the traditional Bayesian theory.

The MMEU criterion, which has been suggested here as the main rule of the decision theory, is generally applicable to decision situations where the possible outcomes are non-negative from the point of view of the decision maker. However, there are situations where the MMEU criterion seems to be too cautious. For example, the 'reflection effect' and the 'isolation effect' suggested by Kahneman and Tversky [24] cannot be explained directly with the decision theory of this paper. We believe that in order to cover these phenomena a more general and comprehensive decision theory is needed which includes references to the decision maker's 'levels of aspiration'. A special case of the effects of levels of aspiration would be the 'shifts of reference point' discussed by Kahneman and Tversky. Introducing 'levels of aspiration' means that the part of traditional Bayesian theory which refers to utilities has to be considerably extended and modified.

Lund University, Sweden

NOTES

*The order of the authors' names is based only upon age and (or?) wisdom. The authors wish to thank Robert Goldsmith, Sören Halldén, Bengt Hansson and Isaac Levi for helpful criticism and comments.

[1] This example was chosen in order to simplify the exposition. We believe, however, that similar examples can be found within many areas of decision making, e.g. medical diagnosis and portfolio selection.

[2] This decision rule is a generalization of the rule suggested by Gärdenfors in [13], p. 169.

[3] This measure is assumed to be unique up to a positive linear transformation. In [26] and [27], Levi has generalized another dimension of the traditional Bayesian decision theory by allowing sets of utility functions which are not linear transformations of each other. We believe that this generalization is beneficiary in some contexts, but we will not discuss it further in the present paper. Cf. note 20.

[4] It is interesting to note that de Finetti in [10], p. 62, note (a), has recognized some problem in such an operational definition as a way of representing the agent's beliefs.

[5] The central axiom is the so-called coherence criterion which assumes that if the agent is willing to bet on state s_j at the least odds of $a:b$, then he should be willing to bet on not-s_j at odds of $b:a$. The first of these betting ratios will thus be equal to one minus the second betting ratio, i.e. $a/(a+b) = 1 - b/(a+b)$. Smith [37], among others, points out that this assumption need not be satisfied. An agent may very well be willing to bet on least odds of $a:b$ for s_j, but at the same time bet on least odds of $c:d$ against s_j, where $a/(a+b) \neq 1 - c/(c+d)$, which contradicts the coherence criterion.

[6] This aspect of risk taking will be further discussed in the next section.

[7] In the present paper we do not aim at an elaborate analysis of the concept of knowledge, but we take this as a primitive notion.

[8] In this paper we use a decision theory similar to Savage's [35]. We thus deliberately exclude problems connected with conditional probabilities and probabilities of conditionals. One reason for this is that it is rather straightforward to generalize a decision theory based on such probabilities in the same way as we have generalized Savage's theory. The second reason is that Luce and Krantz [28] have shown that in decision situations with only finitely many states and outcomes it is possible to translate a decision situation containing conditional probabilities into a Savage type situation (and vice versa). For a discussion of this result, cf. Jeffrey [23]. As is easily seen, this result also holds for sets $\mathcal{P}$ of probability measures.

[9] See, e.g., Dempster [5], Good [18], Smith [36] and [37].

[10] Levi [27], p. 402, requires that the set of 'permissible' probability measures be convex. The interpretation of 'permissible' is discussed in section 4. In this connection it is interesting to note that Savage [35], p. 58, note (+), mentions that "one tempting

representation of the unsure is to replace the person's single probability measure P by a set of such measures, especially a convex set".

[11] See, e.g., Dempster [5], Edman [6], Ekelöf [7], Gärdenfors [13], Good [18], Halldén [19], Smith [36], and [37].

[12] We say that a set of probability intervals associated with the states of a decision situation is *consistent* if and only if, for any state s_i and for any number x within the interval associated with s_i, there is a combination of numbers, which lie within the intervals associated with the remaining states, such that the sum of x and these numbers equal 1. Levi [26], p. 416–417, gives an example which shows that there may be two decision situations with the same alternatives, states and outcomes, but with different sets of 'permissible' probability measures, which give different decisions when his decision theory is used, although the intervals that can be associated with the states are identical.

[13] An interesting possibility is to take ρ to be a second order probability measure, i.e. to let ρ be a probability measure defined over the set $\mathcal{P}$ of epistemically possible probability measures. If $\mathcal{P}$ is finite there seem to be no problems connected with such a measure. But if $\mathcal{P}$ is taken to be a convex set of probability measures and we at the same time want all measures in $\mathcal{P}$ to have a non-zero second order probability, we run into problems. However, nothing that we have said excludes the possibility of taking ρ to be a non-standard probability measure. For a discussion of such measures see Bernstein and Wattenberg [3].

[14] We can compare this example with the difference between match A and match B in the introductory example. The reliability measure connected with match A which is depicted in Figure 1, can be seen as corresponding to the reliability measure after the evidence e has been obtained, and the measure connected with match B as corresponding to the reliability measure before e is obtained. Ideas similar to those presented here have been discussed by Bar-Hillel [1] and Rosenkrantz [32].

[15] Models of belief similar to ours have been presented in terms of fuzzy set theory by [39] and [11]. However, we will not consider this theory in the present paper.

[16] In the quotation above it is obvious that Popper uses the traditional definition of relevance, i.e., e is *relevant* to a if and only if $P(a) \neq P(a, e)$. We believe that this definition is too narrow. Instead we propose the definition that e is relevant to a iff $P(a) \neq P(a, e)$ or the evidence e changes the degree of epistemic reliability of P. Keynes is also dissatisfied with the traditional definition of 'relevance'. He wants to treat 'weight of evidence' and 'relevance' as correlative terms so that "to say that a new piece of evidence is 'relevant' is the same thing as to say that it increases the 'weight' of the argument" ([25], p. 72). For a proof that Keynes's definition of 'relevance' leads to a trivialization result, and for a discussion of some general requirements on a definition of 'relevance' the reader is referred to Gärdenfors [12].

[17] The 'desired level of epistemic reliability' can be interpreted in terms of levels of aspiration so that the ρ_0 chosen by the decision maker is his level of aspiration as regards epistemic reliability.

[18] We have mentioned that an agent is taking a risk by not taking all epistemically possible measures under consideration. 'Risk' is a notion of great complexity and the literature is flooded with papers trying to capture all aspects of risk in *one* measure. However, we do not believe that this is possible and we will give a brief explanation

why. Let us, as an example, return to Miss Julie. It seems reasonable that she perceives a greater risk in betting on match B (and an even greater one in match C) than in match A. She may very well, if forced to do so, estimate that each player has a 50% chance of winning in both match A and B, but still regard match B as riskier. This is due to the fact that her state of knowledge is very different in the two decision situations and it shows that the degree of epistemic reliability is an important factor when determining the risk involved in a decision situation. This aspect of risk taking has not been considered in the traditional theories of risk.

This dimension of risk taking is represented in our theory by the selection of a subset $\mathcal{P}/\rho_0$ of $\mathcal{P}$. An agent who takes all epistemically possible measures into consideration takes no 'epistemic' risk at all. If it is assumed that ρ is a second order probability distribution (cf. note 13 and [33]), we suggest the following measure of the epistemic risk R taken by an agent in a decision situation:

$$R(\mathcal{P}/\rho_0) = 1 - \rho(\mathcal{P}/\rho_0)/\rho(\mathcal{P}),$$

where $\rho(\mathcal{P})$ is equal to $\Sigma_{P\in\mathcal{P}}\rho(P)$ and similar for $\rho(\mathcal{P}/\rho_0)$. As is easily seen, Miss Julie will take a rather great epistemic risk if she acts as a strict Bayesian in match B. But whether she will do so or not is dependent on her risk preferences.

For a criticism of the traditional risk concept and for a discussion of other solutions see Sahlin [34] and Hansson [20].

19 Cf. Gärdenfors [13], p. 169.

20 The results of the majority of subjects of the Goldsmith–Sahlin experiments support the thesis that the degree of epistemic reliability of the probability estimates is an important factor when choosing lottery tickets, but, it should be admitted that not all of these results can be explained by the MMEU criterion. The main reason for this is, in our opinion, that the agent's *values* are not completely described by a utility measure. In this paper we have concentrated on the epistemic aspects of decision making and used the traditional way, i.e. utility measures, to represent the values of the decision maker. We believe that this part of the traditional Bayesian decision theory should be modified as well, perhaps by including a 'level of aspiration', but such an extension lies beyond the scope of this paper.

21 Such a distribution may, in many cases, assign the probability 1 to one of the states, but we do not assume that it always will. We interpret 'full information' in a pragmatic way, meaning something like 'having as much information as is practically possible', so, even if the world is deterministic, having full information does not entail that one knows which is the true state of nature.

22 Fellner [9] considers problems similar to Ellsberg's paradox.

23 However, if the agent has some information that he judges relevant for the distribution of the black and yellow balls, then the epistemic reliability of the distributions in $\mathcal{P}$ may be quite different. He may, for example, believe that the distributions in $\mathcal{P}/\rho_0$ cluster around e.g. (1/3, 1/2, 1/6). Then the MMEU criterion will recommend that a_2 be preferred to a_1 and that a_4 be preferred to a_3. This recommendation does not conflict with Savage's sure-thing principle either (cf. [14]).

24 It should be noted that even if $\mathcal{P}$, the set of all epistemically possible probability measures, is convex, the set $\mathcal{P}/\rho_0$ selected with the aid of the desired level of epistemic reliability need not be convex. For example, in the representation of the epistemic

reliability connected with match C as depicted in Figure 2, the corresponding set $\mathcal{P}/\rho_0$ will consist of two disconnected intervals at the end points 0 and 1. If Levi's $B_{X,t}$ is identified with $\mathcal{P}/\rho_0$, this example shows that his requirement of convexity is not always realistic.

[25] Further information on the notion of P-admissibility is to be found in Ch. 6 of Levi [27].

[26] Levi discusses this problem in [27], pp. 208–210.

[27] It is also easy to see that according to Levi's theory there may be a change in the set of optimal alternatives if two (or more) states are conjoined. This is the case if it is assumed that if s_i and s_j are the two states to be conjoined, then, for any probability measure, P, the probability of the conjoined state is equal to $P(s_i) + P(s_j)$ and the utility of the outcome in the new state, if alternative a_k is chosen, is equal to $(P(s_i) \cdot u_{ki} + P(s_j) \cdot u_{kj})/(P(s_i) + P(s_j))$. Such a contraction of the decision problem is reasonable, if, for example, it is realized that the partitioning adopted initially was unnecessarily refined. It is easy to verify that the set of alternatives which are optimal according to MMEU is not altered by such a conjoining of states.

Levi discusses these matters on pp. 161–162 in [27]. He notes that one can obtain different classes of S-admissible alternatives by using different partitions of the states, but he contends that the adoption of a method for fixing security levels in determining S-admissibility is a moral or political value judgement distinct from a principle of rational choice. We do not have any such problems, since we do not obtain different results when conjoining states as above.

REFERENCES

[1] Bar-Hillel, M.: 1979, 'The Paradox of Ideal Evidence and the Concept of Relevance', a translation of a paper appearing in *IYYUN*, a Hebrew Philosophical Quarterly, **27**, 1976–1977.

[2] Becker, S. W. and Brownson, F. O.: 1964, 'What Price Ambiguity? Or the Role of Ambiguity in Decision Making', *Journal of Political Economy* **72**, 62–73.

[3] Bernstein, A. R. and Wattenberg, F.: 1969, 'Non-Standard Measure Theory', in W.A.J. Luxemburg, ed., *Applications of Model Theory of Algebra, Analysis and Probability*, Holt, Rinerhart & Winston, New York, 171–185.

[4] Carnap, R.: 1950, *Logical Foundations of Probability*, University of Chicago Press, Chicago.

[5] Dempster, A. P.: 1967, 'Upper and Lower Probabilities Induced by a Multivalued Mapping', *Ann. Math. Stat.* **38**, 325–339.

[6] Edman, M.: 1973, 'Adding Independent Pieces of Evidence', in *Modality, Morality and Other Problems of Sense and Nonsense; Essays Dedicated to Sören Halldén*, C.W.K. Gleerup, Lund, 180–191.

[7] Ekelöf, P.O.: 1977, *Rättegång*, P. A. Norstedt & Söner, Stockholm.

[8] Ellsberg, D.: 1961, 'Risk, Ambiguity, and the Savage Axioms', *Quarterly Journal of Economics* **75**, 643–669.

[9] Fellner, W.: 1961, 'Distortion of Subjective Probabilities as a Reaction to Uncertainty', *Quarterly Journal of Economics* **75**, 670–689.

[10] de Finetti, B.: 1980, 'Foresight: Its Logical Laws, Its Subjective Sources', in H. E.

Kyberg and H. E. Smokler, eds., *Studies in Subjective Probability*, R. E. Krieger, New York, 53–118.
[11] Freeling, A.: 1980, 'Fuzzy Sets and Decision Analysis', *IEEE Transactions on Systems, Man, and Cybernetics* SMC-10, no. 1, 341–354.
[12] Gärdenfors, P.: 1978, 'On the Logic of Relevance', *Synthese* **37**, 351–367.
[13] Gärdenfors, P.: 1979, 'Forecasts, Decisions and Uncertain Probabilities', *Erkenntnis* **14**, 159–181.
[14] Gärdenfors, P. and Sahlin, N-E.: 1981, 'Decision Making with Unreliable Probabilities'. Paper presented at Eighth Research Conference on Subjective Probability, Utility and Decision Making, Budapest.
[15] Goldsmith, R. W. and Sahlin, N-E.: 1982, 'The Role of Second-order Probabilities in Decision Making', in Humphreys, P. C., Svenson, O. and Vari, A., eds., *Analysing and Aiding Decision Processes*, North-Holland, Amsterdam.
[16] Goldsmith, R. W. and Sahlin, N-E.: 1981, 'Do Second-order Probabilities Affect Decisions?', forthcoming.
[17] Goldsmith, R. W. and Sahlin, N-E.: 1981 'Second-order Probabilities and Risk in Decision Making', forthcoming.
[18] Good, I. J.: 1962, 'Subjective Probability as a Measure of a Non-Measurable Set', in E. Nagel, P. Suppes and A. Tarski, eds., *Logic, Methodology and Philosophy of Science; Proc. of the* 1960 *Int. Congress*, Stanford University Press, Stanford.
[19] Halldén, S.: 1973, 'Indiciemekanismer', *Tidskrift for Rettsvitenskap* **86**, 55–64.
[20] Hansson, B.: 1981, 'The Decision Game – the Conceptualisation of Risk and Utility', in E. Morscher and R. Stranzinger, eds., *Ethics – Foundations, Problems, and Applications*, Wien.
[21] Hodges, J. L. and Lehmann, E. L.: 1952, 'The Uses of Previous Experience in Reaching Statistical Decisions', *Annals of Mathematical Statistics* **23**, 396–407.
[22] Hurwicz, L: 1951, 'Some Specification Problems and Applications to Econometric Models', *Econometrica* **19**, 343–344 (abstract).
[23] Jeffrey, R. C.: 1977, 'Savage's Omelet', in F. Suppe and P. D. Asquith, eds., *PSA 1976*, volume 2, East Lansing, Michigan, 361–371.
[24] Kahneman, D. and Tversky, A.: 1979, 'Prospect Theory', *Econometrica* **47**, 263–292.
[25] Keynes, J. M.: 1921, *A Treatise on Probability*, Macmillan, London.
[26] Levi, I.: 1974, 'On Indeterminate Probabilities', *Journal of Philosophy* **71**, 391–418.
[27] Levi, I.: 1980, *The Enterprise of Knowledge*, MIT Press, Cambridge, Mass.
[28] Luce, D. and Krantz, D. H.: 1971, 'Conditional Expected Utility', *Econometrica* **39**, 253–271.
[29] Luce, D. and Raiffa, H.: 1957, *Games and Decisions*, Wiley, New York, 288–290.
[30] Peirce, C. S.: 1932, *Collected Papers*, ed. by Hartshorne and Weiss, Belknap Press, Cambridge, Mass.
[31] Popper, K. R.: 1974, *The Logic of Scientific Discovery*, Hutchinson, London.
[32] Rosenkrantz, R. D.: 1973, 'Probability Magic Unmasked', *Philosophy of Science* **40**, 227–233.
[33] Sahlin, N-E.: 1980, 'Generalized Bayesian Decision Models', *Working Paper*.
[34] Sahlin, N-E.: 1980, 'Level of Aspiration and Risk', *Working Paper No.* 21, *The Mattias Fremling Society*.

[35] Savage, L. J.: 1972, *The Foundations of Statistics*, Dover, New York.
[36] Smith, C. A. B.: 1961, 'Consistency in Statistical Inference and Decision', *Journal of the Royal Statistical Society Ser. B* **23**, 1–25.
[37] Smith, C. A. B.: 1965, 'Personal Probabilities and Statistical Analysis', *Journal of the Royal Statistical Society Ser. A* **128**, 469–499.
[38] Wald, A.: 1950, *Statistical Decision Functions*, J. Wiley & Sons, New York.
[39] Watson, S. R., Weiss, J. and Donell, M.: 1979, 'Fuzzy Decision Analysis', *IEEE Transactions on Systems, Man, and Cybernetics* SMC-9, no. 1, 1–9.