August 18, 2003

Aanund Hylland:#

DECISIONS UNDER UNCERTAINTY

Standard theory and alternatives

1. Introduction

Individual decision making under uncertainty can be characterized as follows:

– The decision maker has to choose one *act* from a given set of possible acts.
– A set of potential *states* is given, representing the circumstances about which there is uncertainty. It is logically necessary that one and only one of the states will occur; otherwise, the states have been erroneously specified. Eventually, the uncertainty will be resolved and one state will be realized.

When the act has been chosen and the state has been realized, the outcome is determined, as can be illustrated by the following table:

	State				
Act	s_1	...	s_j	...	s_n
a_1	$O_{1,1}$		$O_{1,j}$		$O_{1,n}$
...					
a_i	$O_{i,1}$		$O_{i,j}$		$O_{i,n}$
...					
a_m	$O_{m,1}$		$O_{m,j}$		$O_{m,n}$

For simplicity, it is assumed that there are a finite number of possible acts, denoted a_1, a_2, ..., a_i, ..., a_m, and a finite number of potential states, denoted s_1, s_2, ..., s_j, ..., s_n. In many applications, both act and state contain continuous variables, making the sets of acts and states infinite.

Acts, states and outcomes can be complex objectss, but the nature of these objects are not discussed further.

#. Professor of economics and decision theory, University of Oslo. Address: Department of Economics, P.O. Box 1095, Blindern, N-0317 OSLO, Norway. Telephone: (+47) 22 85 42 71 Telefax: (+47) 22 85 50 35 Electronic mail: aanund.hylland@econ.uio.no. Home telephone: (+47) 22 60 47 00

Everything that can affect the outcome and about which there is uncertainty, shall be made part of the state. When the act (a_i) has been chosen and the state (s_j) has become known, the outcome ($O_{i,j}$) shall follow, mechanically and with certainty.

Everything the decision maker cares about, shall be a part of the outcome. That is, if certain acts or states in themselves are valued positively or negatively, the relevant aspects of the act or state shall be incorporated into the outcome. In other words, it is assumed that the decision maker has preferences over outcomes and only over outcomes.

Every combination of act and state must in principle be possible, although some combination may lead to very bad outcomes. The act is chosen without knowledge of which state will be realized, and this is inconsistent with some combinations being impossible.

2. Degrees of uncertainty

The set-up of Section 1 presupposes that the decision maker knows the set of states. In other words, there is not total ignorance. It is difficult to imagine what a theory of decision making under total ignorance should look like.

Concerning degree of uncertainty, two alternatives will be discussed:

(1) The decision maker knows the set of potential states, but nothing more.

(2) The decision maker is able to assign to each potential state a precise probability.

There is uncertainty in both cases, but the decision maker knows a lot more in case (2) than in case (1).[1]

This is by no means an exhaustive list of possible degrees of uncertainty. On the one hand, there is the case of complete ignorance, that is, less knowledge than assumed in (1). On the other hand, full certainty is a possibility, amounting to more knowledge than assumed in (2). Moreover, there is a range of possibilities between (1) and (2). In such intermediate cases the decision maker knows the set of potential states and has some opinion as to which of them are more or less likely, in absolute terms or relatively to one another, without being able to express this knowledge as precise probabilities. There is an infinity of different cases falling between (1) and (2), but they are not discussed further here.[2]

---

1. Some authors reserve the phrase decisions under *uncertainty* for case (1), referring to case (2) as decisions under *risk*, but this terminology is not adopted here.

2. It can be claimed that almost all practically important instances of decision making under uncertainty fall between (1) and (2). Therefore, it is unsatisfactory that no good theory exists for these cases, but this is the sad state of affairs.

3. Only the set of states is known

In standard microeconomic theory, case (1) is usually not discussed. In other professions, such as philosophy, it has received more attention. Here it will be discussed briefly.

When only the set of states is known, it does not make sense to put more weight on the outcomes of one state than another. If one, for example, puts greater weight on the outcome of s_1 than that of s_2, one would in a sense say that s_1 is considered more likely or probable than s_2. By assumption, this type of information is not available.

It is, however, possible to put greater weight on outcomes that are particularly bad or good.

The following decision rule is often recommended in case (1):

> For each possible act, that is, for each line in the table in Section 1, find the worst outcome. Judge the act as if this outcome will occur and choose the act that, on this basis, is the best one.

This is often called the *maximin-rule*. One maximizes – over the set of possible acts – the minimal (worst) outcome.[3] In order to apply the rule, the decision maker must be able to rank the outcomes, that is, ordinal preferences over outcomes must exist. More demanding requirements need not be imposed on the preferences.

In a sense, the maximin-rule is based on an extremely pessimistic view of the world. It may appear that the decision maker argues as follows:

> «No matter which act I choose, the world will turn against me and see to it that the outcome is as bad as possible, given the act I chose.»

It is, however, possible to give a normative justification for the rule without invoking this type of reasoning, but the issue is not discussed further here.

One can also imagine a "maximax"-rule, where each act is judge by its *best* outcome, but it is difficult to justify for that rule normatively.

4. The states have known probabilities

Given the assumption that there is a finite number of potential states, case (2) amounts to postulating the existence of numbers $p_1, p_2, ..., p_j, ..., p_n$, where $p_j \geq 0$ for all j and $p_1 + p_2 + ... + p_j + ... + p_n = 1$. Here p_j is the probability that state s_j will occur.[4] The case of infinitely many states is not discussed.

3. Hopefully, this sentence will not be misunderstood, although it is slightly misleading. One shall not look for the worst of all outcomes, but the worst outcome for each act. Of course, the outcomes that are worst for two different acts, may belong to different states, that is, they may be found in different columns in the table.

4. In the finite case it does not matter whether we assume $p_j \geq 0$ or $p_j > 0$. If $p_j = 0$, s_j is impossible and can be

From where do probabilities come? In some cases, it can reasonably be claimed that there exist objectively given probabilities. This is true for controlled games of chance, where the rules of the game specify the probabilities of the outcomes. In other cases, it may be possible to base (objective) probabilities on experience, that is, probability is identified with observed frequency in (a large number of) previous cases of the same type. Often, however, the probabilities will contain an element of subjective judgment on part of the decision maker. Therefore, two persons who face what objectively might appear to be the same choice situation, can specify the probabilities differently.

Much can be said about probabilities and their intrinsic nature. It would lead too far to start that discussion here. For the present purpose, it suffices to assume that the decision maker is able to assign definite probabilities to the states.

When probabilities have been introduced, an act can be identified with a lottery over outcomes. In the table of Section 1 the act a_i will lead to the outcome $O_{i,1}$ with probability p_1, $O_{i,2}$ with probability p_2, ..., $O_{i,j}$ with probability p_j, ... and $O_{i,n}$ with probability p_n. Such a lottery is usually called a *prospect*. In order to be able to make rational choices, the decisions maker must have preferences prospects.

As mentioned above, the outcomes can in principle be very complex objects. For the subsequent discussion it does not really matter what the outcomes are, but it may be easier to follow the argument if attention is restricted to a simple example. Therefore, it is assumed that an outcome simply is an amount of money. Previous assumptions then imply that the decision maker is concerned only with this amount. No other aspect of the outcome is of importance, and acts and states have no (positive or negative) value in themselves.

Under these assumptions, an act can be identified with a (finite) list of possible amounts of money, with a probability assigned to each amount. Such a prospect can be denoted as follows:

$$A = \{p_1{:}x_1;\ p_2{:}x_2;\ ...\ p_j{:}x_j;\ ...\ p_n{:}x_n\}.$$

The decision maker receives the amount x_j with probability p_j for $j = 1, 2, ... n$. It is assumed that $p_j \geq 0$ for all j and $p_1 + p_2 + ... + p_j + ... + p_n = 1$.

5. Expected utility

The decision maker is assumed to have preferences over monetary prospects, that is, objects of the form A.

What kind of requirements can reasonably be imposed on these preferences? Usually, more stringent conditions are imposed than those normally used in standard consumer theory under certainty.

(..continued)
deleted from the set of potential states.

For one thing, it is implied by the formalism that only the amounts of money and their probabilities shall play a role. The process producing the probabilities shall be irrelevant. If one receives NOK 100 with probability 0.25 and 0 with probability 0.75, it shall not matter if a coin is tossed twice and one wins the NOK 100 if both tosses give tails, or if a roulette wheel is used, on which one quarter of the numbers give a favorable outcome, or if some other procedure giving the same probability is used.

Let A be the prospect given above and let B be another prospect. If there are amounts that have positive probabilities in B but do not occur in A, we can add these to A with probability 0, and vice versa. Therefore, with no loss of generality we can write:

$$A = \{p_1{:}x_1;\ p_2{:}x_2;\ \ldots\ p_j{:}x_j;\ \ldots\ p_n{:}x_n\},$$

$$B = \{q_1{:}x_1;\ q_2{:}x_2;\ \ldots\ q_j{:}x_j;\ \ldots\ q_n{:}x_n\}.$$

Let λ be a number satisfying $0 \leq \lambda \leq 1$. Then we can consider the prospect:

$$C = \{\lambda{:}A;\ 1\text{-}\lambda{:}B\}.$$

First, the lot shall be drawn between A and B, with probabilities λ og $1\text{-}\lambda$, respectively. Then one conducts the lottery specified by either A or B. All in all, the amount x_j is chosen with probability $\lambda p_j + (1\text{-}\lambda)q_j$, for all $j = 1, 2, \ldots n$. It shall make no difference if one conducts a direct lottery with these probabilities, rather than going through the two-stage lottery involving A or B.

Moreover, one usually introduces a so-called *independence axiom*. Let A, B and C be prosepcts, and assume that the decision maker prefers A to B. La λ be a number satisfying $0 < \lambda \leq 1$. Then the axion requires that the decision maker prefer

$$D = \{\lambda{:}A;\ 1\text{-}\lambda{:}C\}$$

to

$$E = \{\lambda{:}B;\ 1\text{-}\lambda{:}C\}.$$

This can be justified as follows: If the last alternative is realized in the lottery D or E, the outcome will be C in any case, so it does not matter whether D or E was chosen. If the first alternative is realized, which is possible since $\lambda > 0$, the outcome will be A or B, and it has been assumed that the decision maker prefers A to B. In one case D and E give equally good results; in another D is better than E. Therefore, it is being claimed, a rational decision maker must prefer D to E.

If A, B and C had been different goods, and λ and $1\text{-}\lambda$ had been amounts of these goods, there would have been nothing irrational in preferring A to B, but at the same time preferring E to D. This can be explained by effects occurring when A or B is being consumed together with C. In the case studied here, however, it is not a question of "consuming" A or B together with C. Either the last alternative is realized in the lottery D or E, and then A and B are definitely out of the picture, or the first alternative is realized, and then C is out.

When these assumptions, together with some others that are not controversial, are satisfied, the so-called *expected utility theorem* will hold. There exists a function u, known as the *utility function*, defined over amounts of money, with the property that the decision maker ranks prospect on the basis of their *expected utility*, that is

$$p_1u(x_1) + p_2u(x_2) + ... + p_ju(x_j) + ... + p_nu(x_n)$$

The prospect with the highest expected utility is chosen.

It is not assumed that the function u exists as *a* mental reality for the decision maker. It is only claimed that given the conditions imposed on the preferences, the decision maker will act *as if* expected utility is maximized.

6. Allais' paradox

The following example has been used as an argument against the independence axiom presented in Section 5. It also provides an argument against the expected utility theorem, which is based on this axiom.

The table below defines four acts or prospects, *A*, *B*, *C* and *D*. There are three states, with probabilities given at the top of the table. Inside the table is given the outcome for every combination of act and state, expressed as mounts of money measured in some appropriate unit.[5]

	0,10	0,01	0,89
A	5	0	0
B	1	1	0
C	5	0	1
D	1	1	1

The decision maker shall either chose between *A* and *B*, or chose between *C* and *D*.

If the expected utility theorem holds, *A* will be preferred to *B* if and only if *C* is preferred to *D*, and similarly for the opposite preferences and for indifference. Without loss of generality, $u(0) = 0$ og $u(5) = 1$ can be assumed. It is then easy to see that *A* has higher expected utility than *B* if and only if $u(1) < 10/11$. Exactly the same condition on $u(1)$ is equivalent to *C* having higher expected utility than *D*.

The same conclusion can be reached applying the argument behind the independence axiom: If the last of the three states is realized, the outcome is the same for *A* and *B*. Therefore, this state should be irrelevant for the choice between *A* and *B*. If the last state is realized, the outcome is also the same for *C* and *D*. Therefore, this state should be irrelevant for the choice between *C* and *D*. If the last state is ignored, *A* = *C* and *B* = *D*. The conclusion is that *A* should be preferred to *B* if and only if *C* is preferred to *D*.

5. The unit must be rather large in order that the statements made below about common intuitions shall hold. Perhaps a unit of NOK 10,000 would be appropriate.

Many people will intuitively choose *A* when facing the choice between *A* and *B*, but choose *D* when facing the choice between *C* and *D*. This contradicts the argument above.

Both in *A* and *B* there is a considerable danger of not winning anything. Measured by expected monetary value, *A* is a lot better than *B*. Therefore, it seems natural to prefer *A*. In the choice between *C* and *D*, certainty of winning something can be achieved by choosing *D*. Although *C* has a higher monetary expectation than *D*, the possibility of a certain positive outcome may weigh more heavily and lead to the choice of *D*.

Empirical studies, of this example and similar ones, convincingly show that many people make choices that are inconsistent with the expected utility theorem. The deviations from the predictions of the theorem are not random and arbitrary, bur systematic. Almost nobody chooses *B* and *C* in the examples presented here, while the combination *A* and *D* occurs frequently.

In spite of these empirical observations, it can be claimed that choices can only be *rational* if they are consistent with the expected utility theorem. What many people intuitively find reasonable, is then characterized as irrational.

On the other hand, it can be claimed that the reasoning given above to justify the choice of *A* and *D* is rational and consistent. If so, the arguments in support of the independence axiom, presented in Section 5, cannot be convincing.

There is an extensive discussion of this issue in the literature.

7. Concluding remarks

The purpose of this note has been to discuss the logic behind the theory of decision making under uncertainty. In particular, the independence axiom and the expected utility theorem have been presented, as well as objections to this line of thought, as represented by Allais' paradox.

It is not my purpose to draw conclusion concerning the validity of the standard theory or the objections to it.