Milne, Frank; Kelsey, David

Working Paper

Induced Preferences and Decision-Making Under Risk and Uncertainty

Queen's Economics Department Working Paper, No. 897

Provided in Cooperation with:
Queen’s University, Department of Economics (QED)

Queen's Economics Department Working Paper No. 897

Induced Preferences and Decision-Making Under Risk and Uncertainty

Frank Milne David Kelsey

Department of Economics
Queen's University
94 University Avenue
Kingston, Ontario, Canada
K7L 3N6

1-1994

Discussion Paper #897

Induced Preferences and Decision-Making Under Risk and Uncertainty

by

Frank Milne
Queen's University

and

David Kelsey
University of Birmingham

January 1994

INDUCED PREFERENCES AND DECISION-MAKING UNDER RISK AND UNCERTAINTY*

by

David Kelsey
Department of Economics,
The University of Birmingham.

and

Frank Milne
Department of Economics,
Queens University,
Canada.

23 January 1994

Abstract In this paper we suggest a new interpretation of non-additive probabilities. We study a decision-maker who follows the Savage axioms. We show the if (s)he is able to take unobservable actions which influence the probabilities of outcomes then it can appear to an outsider as if the his/her subjective probabilities are non-additive. We make a related analysis of models with objective probabilities and show that the induced preferences can have the rank dependent expected utility form. Implications for multi-period decisions are explored. We show that such preferences are not vulnerable to "Dutch books".

Keywords Uncertainty aversion, induced preferences, Dutch Book, rank dependent expected utility.

JEL Classification D81.

Address for Correspondence David Kelsey, Department of Economics, The University of Birmingham, Birmingham B15 2TT, UK.

* We would like to acknowledge financial assistance from SSHRC (Canada) grant No. 410-92-0190 and the Department of Economics, The University of Birmingham. We would like to thank Richard Cornes, Jurgen Eichberger, Simon Grant and John Quiggin for helpful comments.

INTRODUCTION

1.1 Background

The dominant theory of decision-making under uncertainty is subjective expected utility theory SEU (Savage 1954). This has the following implications;

(1) Beliefs may be represented as a unique finitely additive probability distribution;

(2) Preferences are linear in these probabilities and hence may be represented as maximising the expected value of a utility function.

Together these imply that the decision-maker is locally risk-neutral, (where the utility function is differentiable).[1] Local risk-neutrality has some implausible implications, realist risk premia for small gambles imply excessively large risk premia for large gambles (Epstein 1992 p.11), individuals will never buy full insurance at an unfair premium (Segal and Spivak 1990) and investors will either buy or sell short *every* asset (Dow and Werlang 1992a). There is also experimental evidence which fails to support SEU. The Ellsberg Paradox, (see Ellsberg 1961) contradicts implication (1).

Recent years have seen the development of some alternative decision theories, which do not imply local risk neutrality and are more compatible with the experimental evidence. Despite this, the theory which has the strongest normative justification is SEU. In the present paper we suggest a way to resolve this tension. We show that the decision-maker may at a fundamental level have SEU preferences. However to an outsider who is not able to observe all the relevant variables the induced preferences (Milne 1981) may appear to satisfy the axioms of one of the alternative theories. This theory of decision-making embodies the normative properties of SEU and the superior descriptive performance of the alternatives, at the same time.

Following the seminal work of Ellsberg (1961) there is a strong argument both on theoretical and empirical grounds that individuals often do not use conventional subjective probabilities, when making decisions under uncertainty. One of the more prominent alternatives is that individuals have subjective probabilities which fail to be additive across disjoint events, Schmeidler (1989). The expected value of utility with respect to a non-additive probability distribution may be defined to be a Choquet integral, (Choquet 1953-4). The theory of maximising a Choquet integral of utility with respect to a non-additive probability will henceforth be referred to as *Choquet expected utility* (CEU). CEU is a generalisation of SEU.

CEU preferences may, under some assumptions, be given an alternative interpretation in terms of additive, but non-unique subjective probabilities, Gilboa and Schmeidler (1989), Kelsey (1994). In this case preferences may be represented as the minimum over a *set* of additive probabilities of the expected value of utility. This is known as *Maxmin Expected Utility theory (MMEU)*. It embodies the intuition that when probabilities are not clear, individuals may simultaneously consider a number of probability distributions to be feasible. MMEU has the advantage that it can be expressed in terms of conventional expected values rather than the less familiar Choquet integral.

Dreze (1987) has modelled situations where an individual is able to influence the probabilities of the states by means of an (unmodelled) hidden action. He deduces that the individual has a convex set C of subjective probabilities. Any action is evaluated by the maximum value over the set C of expected utility and (s)he chooses the action which maximises the maximal value of expected utility. This can be motivated by assuming that the individual influences the probabilities to make them more favourable. Although the motivation is different, Dreze's axioms are similar to Gilboa and Schmeidler's (1989) axioms for (MMEU), which contains CEU as a special case, provided uncertainty aversion is assumed.

In this paper we explicitly model the hidden actions. We assume that the state space consists of both observable and unobservable states. We show that SEU preferences over the full state space will induce preferences of the CEU form over the observable states. Two hypothesis are used. First that the decision-maker is able to take an unobservable action, which makes the probability distribution more favourable. Second that nature may also take an unseen move which is unfavourable to the decision-maker. Of these, the first seems more plausible, however only the second is able to generate uncertainty-averse preferences. This is unfortunate since uncertainty-aversion is required to explain the Ellsberg paradox. Uncertainty-aversion appears to have more reasonable implications in some theoretical models, (see for instance Dow and Werlang 1992a, Eichberger and

Kelsey 1994). However as Kelsey and Quiggin (1992) note, preferences which are uncertainty loving (at least in some range) appear to be able to explain gambling behaviour better than SEU.

It could be that there are no actual moves by nature, but the perception of them causes individuals to act in an uncertainty-averse manner. It may be the case that the hidden actions are perceived to be available by the decision-maker but not by an outside observer. For instance, participants in a lottery may believe that choosing their "lucky number" will favourably influence the probability distribution.

Some support from our theory may be obtained from attitudes towards danger. The concern which individuals express about certain risks is often not directly proportionate to their probabilities, eg individuals are often more concerned about the dangers of air travel than those of car travel, even though car travel is statistically more likely to result in an accident. This may be connected with individuals inability to influence probabilities in the case of air travel.

1.2 Updating Uncertainty Averse Preferences

At present there is no generally accepted method of updating CEU preferences. This is a major barrier to applying CEU, since most economic problems, involve decisions taken at a number of points in time. The theory of induced preferences suggests a procedure for updating CEU preferences. It is generally agreed that SEU preferences should be updated by using Bayes' rule to revise the subjective probabilities. Hence when updating CEU preferences, which are induced by SEU preferences over some underlying space, it seems natural to take the Bayesian update of the underlying preferences. The new observable preferences are those induced by the updates of the underlying preferences. We show that under certain assumptions this procedure yields the same results as the commonly used Dempster-Shafer rule (Shafer 1976) for updating CEU preferences.

It has often been argued that individuals who do not follow SEU will present an opportunity to an informed outsider to make a profit at their expense, see for instance de Finetti (1974), Freedman and Purves (1969). This is know informally as the "Dutch book" argument. Related problems have been observed with CEU preferences, Dow and Werlang (1992b), Epstein and le Breton (1993) and Kelsey (1993). The theory of induced preferences suggests a possible resolution of these problems. Apparent dynamic inconsistencies in the induced preferences arise, because the decision-maker can change the hidden actions, when updating. Likewise the apparent opportunities to make Dutch books are not real because they do not take account of payments in the hidden states.

Organisation of Paper In the next section we introduce the CEU model. Section 3 shows how CEU preferences can arise as induced preferences when an individual is able to take a hidden action. In section 4 we study updating rules for induced preferences and show that they are dynamically consistent. Section 5 contains our conclusions.

2 DEFINITIONS

In this section we present our definitions.

Notation 2.1 There is a finite set S of states of nature. A subset of S will be referred to as an *event*. The set of possible outcomes is denoted by X. An *action* is a function from S to X. The space of actions is denoted by $A(S)$. The decision-maker's preferences over $A(S)$ are denoted by $\succeq$.

Definition 2.1 A *non-additive probability (or capacity)* on S is a real-valued function ν on the subsets of S which satisfies the following properties

a. $A \subseteq B \Rightarrow \nu(A) \leq \nu(B)$ b. $\nu(\emptyset) = 0, \quad \nu(S) = 1.$

An expected value with respect to a capacity can be defined as a Choquet integral (Choquet 1953-4). We explain the Choquet integral below.

Notation 2.2 If $\phi: S \to \mathbb{R}$, let $\phi_{(i)}$ be the ith highest consequence of ϕ and $s_{(i)} = \phi^{-1}\phi_{(i)}$ be the state in which consequence $\phi_{(i)}$ occurs.

Definition 2.2 Choquet Integral If $\phi: S \to \mathbb{R}$, the Choquet integral of ϕ with respect to the capacity ν is defined by the following formula,

$$\int \phi(s)\, d\nu(s) = \phi_{(1)}\nu(s_{(1)}) + \sum_{i=2}^{n} \phi_{(i)}[\nu(s_{(1)} \ldots s_{(i)}) - \nu(s_{(1)} \ldots s_{(i-1)})].$$

Note that if ν is additive then this coincides with the usual definition of an expected value (for a discrete distribution).

Choquet Expected Utility (CEU) allows the decision-maker's subjective probability ν to be non-additive across disjoint events, ie. we do not require that $\nu(A \cup B) = \nu(A) + \nu(B)$, for $A, B \subseteq S$, $A \cup B = \emptyset$. CEU postulates that the decision maker has a utility function u and maximises the expected value of u with respect to the non-additive probability distribution ν, ie. $\int u(a(s))\,d\nu(s)$.

3 INDUCED PREFERENCES

3.1 The CEU Model

In this section we model situations where the decision-maker can influence the probability of the states. Dreze (1987) has developed an axiomatic theory of decision-making under such circumstances. He shows that there exists a utility function $u: X \to \mathbb{R}$ and a convex set C of subjective probability distributions on S such that,

$$a \succeq b \Leftrightarrow \max_{p \in C} \mathbf{E}_p\, u(a) \geq \max_{p \in C} \mathbf{E}_p\, u(b),$$

where $\mathbf{E}_p\, u(a)$ denotes the expected utility of action a with respect to the probability distribution p. This is a special case of CEU.

Dreze assumes that the decision-maker is able to influence the probabilities by some process which is not modelled explicitly. We shall represent this by assuming that there are in addition unobservable states and actions. We show that if the individual has SEU preferences over pairs of observable and hidden actions, the induced preferences over the observable actions will be of the CEU form.

Let S_1 = first set of unobservable states,

S_2 = second set of unobservable states,

$$\Sigma = S \cup S_1 \cup S_2 = \text{set of all states.}$$

We model the decision process by assuming that, first the decision-maker chooses an observable action $a \in A(S)$. Let $r(a) \subseteq X$ denote the range of action $a \in A(S)$. Let $H_i(a)$ be the set of all functions from S_i to $r(a)$, for $i = 1,2$. (S)he is assumed to make a subsequent choice of a hidden action, $h_1 \in H_1(a)$. Nature then chooses a hidden action, $h_2 \in H_2(a)$. The restrictions on the choice of the hidden actions, ensure that their effect is to modify the probabilities of the consequences of the observable action. No new consequences are introduced by the hidden actions. In our opinion this represents the situation intended by Dreze. (A possible generalisation would be to restrict h_i to lie in a subset of H_i.)

We shall assume that the decision maker has SEU preferences over the set of all triples $\langle a,h_1,h_2\rangle$, $a \in A(S)$, $h_1 \in H_1$, $h_2 \in H_2$, with subjective probability π on Σ and utility function $u:X \to \mathbb{R}$. Such preferences could be given an axiomatic justification, (see, for instance Savage 1954).

Notation 3.1 If $a \in A(S)$ then the minimum of a, $m(a)$ (resp. the maximum of a, $M(a)$) will be defined by $m(a) = \min_{s \in S}\{u(a(s))\}$ (resp. $M(a) = \max_{s \in S}\{u(a(s))\}$).

It is clear that the decision-maker will choose his/her own hidden action to be a constant action which yields utility $M(a)$.

Axiomatic models of decision-making when the payoff to the individual depends on a move by nature have been studied by Arrow and Hurwicz (1972), Barbera and Jackson (1988) and Barrett and

Pattanaik (1994), among others. In these papers the decision-maker is assumed to have no beliefs over nature's move. Most of them conclude that the decision-maker must follow a "maximin-type" decision rule, which gives lexicographic priority to the worst outcome. We shall use these results to justify modelling a decision-maker who perceives nature's move to be the constant action which yields $m(a)$. This could also arise if the second hidden action were controlled by a "malevolent" individual who had interests diametrically opposed to those of the decision-maker. This would arise naturally in a two-person zero-sum game. These choices of hidden actions will yield induced preferences over observable actions which can be represented by the function,

$$V(a) = \sum_{s \in S} \pi_s u(a_s) + \delta_1 M(a) + \delta_2 m(a), \tag{3.1}$$

where $\delta_1 = \pi(S_1)$, $\delta_2 = \pi(S_2)$, π_s denotes the probability of state s, and a_s denotes the outcome which action a yields in state s. As the following Proposition states the induced preferences over observable actions have the CEU form.

Proposition 3.1 Suppose that the underlying preferences are SEU with subjective probability π, then the induced preferences over $A(S)$ have the CEU form with the same utility function and a capacity ν given by $\nu(A) = \pi(A) + \delta_1$, $(A \neq S)$, $\nu(S) = 1$.

Proof Let a be an action. Without loss of generality we may assume, $u(a(s_1)) > \ldots > u(a(s_n))$. If ν is defined as above

$$E_\nu\, u(a) = (\pi_{s_1} + \delta_1) u(a_{s_1}) + \pi_{s_2} u(a_{s_2}) + \ldots + \pi_{s_{n-1}} u(a_{s_{n-1}}) + (\pi_{s_n} + \delta_2) u(a_{s_n})$$

$$= \sum_{s \in S} \pi_s u(a_s) + \delta_1 M(a) + \delta_2 m(a).$$

This is the same as equation 3.1, hence the Choquet integral does indeed represent the induced preferences.

■

3.2 Rank Dependent Expected Utility

So far we have interpreted π as a subjective probability on Σ. However our analysis would remain valid if π were an objective probability distribution. In this case, our arguments would show that an individual whose underlying preferences satisfy the von-Neumann Morgenstern axioms would appear to an outsider who could not see the hidden states, to be maximising the expected value of utility with respect to a distortion of the true probability distribution. Such preferences have been axiomatised by Quiggin (1982) and Yaari (1987) and are known as *rank dependent expected utility* (RDEU). We define RDEU preferences below.

Notation 3.2 If $a{:}S \to X$, let $a_{(i)}$ be the ith most preferred consequence of a and $p_{(i)}$ be the probability of consequence $a_{(i)}$.

RDEU implies that the decision-maker maximises,

$$u(a_{(n)})g(p_{(n)}) + \sum_{i=1}^{n-1} u(a_{(i)})\,[g(p_{(i)}+\ldots+p_{(n)}) - g(p_{(i+1)}+\ldots+p_{(n)})],$$

where g is a distortion of the probability distribution.

Suppose that the decision-maker maximises the expected value of the utility function u, with respect to the objective probability distribution π over the underlying state space Σ. To an outsider who can only observe the state space S, it will appear as if the probability is the distribution p on S, defined by $p_s = \pi_s / (\sum_{\sigma \in S} \pi_\sigma)$, where p_s is the probability of state s, as it appears to the outsider. The decision-maker will appear as if (s)he has RDEU preferences with distortion function given by $g(0) = 0$; $g(p) = \delta_2 + (1-\delta_1-\delta_2)p$, $0 < p < 1$; $g(1) = 1$. Apart from the discontinuities at 0 and 1, this function satisfies the properties postulated by Quiggin (1982), section 4. He argues that these preferences are able to explain both the Allais paradox and simultaneous gambling and insurance. Chew, Karni and Safra (1987) show that RDEU preferences will be risk-averse provided both u and g are concave. This implies that if u is concave the induced preferences will be risk-averse.

4 UPDATING

4.1 Pseudo-Bayesian Updates

The CEU model has been successfully applied to a number of atemporal problems. (For a survey of applications see Kelsey and Quiggin 1992.) A major barrier to further applications is the lack of a generally accepted multi-period extension. For instance models of pricing derivative securities typically require trading in more than one time period.

The theory of induced preferences suggests a solution to this problem. There is a strong case for using the Bayesian update of the underlying preferences, which will induce a set of updated preferences over the observable actions.

One difficulty with implementing this approach is that it is not clear how to apply Bayes' rule in this context. Suppose an event $E \subset S$, is observed. Then any subset $E' \subset \Sigma$ of the underlying state space is compatible with this observation provided, $E \subseteq E' \subseteq E \cup S_1 \cup S_2$. In other words, it is not clear after the observation of a subset E of the observable states, how many of the unobservable states will remain possible.

According to the most usual interpretation of CEU preferences, the non-observable states represent ambiguity in the decision-makers' beliefs. It would seem likely that as observations are made there would be a reduction in ambiguity. We shall model this in a simple way by assuming that after an observation fractions θ_1, θ_2 of S_1 and S_2 respectively, remain possible. (This is a special case of a proposal made by Gilboa and Schmeidler 1992). We formalise this in the following definition.

Definition 4.1 Suppose that the decision-maker has SEU preferences over actions on the underlying state space Σ with subjective probability π on Σ. Let θ denote the pair (θ_1, θ_2). The θ-Bayesian update π_E on Σ conditional on the observable event $E \subseteq S$ is given by,

$$\pi_E(s) = \frac{\pi_s}{\sum_{t \in E} \pi_t + \theta_1\delta_1 + \theta_2\delta_2}, \quad \pi_E(S_1) = \frac{\theta_1\delta_1}{\sum_{t \in E} \pi_t + \theta_1\delta_1 + \theta_2\delta_2}$$

and $\pi_E(S_2) = \theta_2\delta_2 / (\sum_{t \in E} \pi_t + \theta_1\delta_1 + \theta_2\delta_2)$.

Note that we do not require the same θ to apply each time preferences are updated.

Proposition 4.1 Under assumption 4.1 the induced preferences conditional on the observable event $E \subset S$ have the CEU form, with respect to the updated capacity ν_E defined by

$$\nu_E(A) = \frac{\pi(A \cap E) + \theta_1\delta_1}{\pi(E) + \theta_1\delta_1 + \theta_2\delta_2}, \qquad A \cap E \neq E, \quad \nu(E) = 1.$$

Proof Note that the underlying preferences have the same general structure, before and after updating, hence it follows from Proposition 3.1 that the induced preferences have the CEU form. The form of the updated capacity ν_E follows from applying Proposition 3.1 to the updated probabilities π_E on the underlying state space.
■

Remark In the special case where $\theta_1 = 0$ and $\theta_2 = 1$,

$$\nu_E(A) = \frac{\pi(A \cap E)}{\pi(E) + \delta_1} = \frac{\nu((A \cap E) \cup \neg E) - \nu(\neg E)}{1 - \nu(\neg E)}.$$

This is the Dempster-Shafer rule, which has been axiomatised by Gilboa and Schmeidler (1993).

4.2 Dynamic Inconsistency and Dutch Books

It is often alleged that individuals who use non-expected utility theory are liable to violate dynamic consistency. A particular form of dynamic inconsistency is what is known as a "Dutch book". This consists of a sequence of trades each of which on its own appears acceptable, however the effect of the sequence as a whole is that the decision-maker is certain to lose money. These arguments can be applied to CEU decision-makers as can be seen from the example below. In this section we shall argue that these apparent dynamic inconsistencies are not real since they fail to take account of the hidden actions. The following example is typical of the structure of alleged dynamic inconsistencies.

Example 4.1 Suppose that there are three states of nature, $S = \{s_1, s_2, s_3\}$. Let E denote the event $\{s_1, s_2\}$. Suppose that a decision-maker's prior beliefs are given by the capacity $\nu(s_1) = \nu(s_2) = \nu(s_3) = \frac{1}{4}$, $\nu(s_1, s_2) = \nu(s_1, s_3) = \nu(s_2, s_3) = \frac{1}{2}$. According to the Dempster-Shafer rule the updated beliefs conditional on event E are given by the capacity, $\nu_E(s_1) = \nu_E(s_2) = \frac{1}{3}$. Consider the following two actions $a = \langle 10,8,2 \rangle$ and $b = \langle 24,0,1 \rangle$. According to the initial preferences $V(a) = 5\frac{1}{2}$ and $V(b) = 6\frac{1}{4}$, where V denotes the expected value of utility with respect to capacity ν. Hence initially b is preferred to a. However according to the updated preferences, $V_E(a) = 8\frac{2}{3}$, $V_E(b) = 8$, where V_E denotes the expected value of utility with respect to capacity ν_E. Although b is initially preferred to a, a is preferred to b after observation of either element of the partition E, $\{s_3\}$.

This could be used to construct a Dutch book. Suppose the decision-maker's endowment corresponded to action a. Then an outsider could offer him/her action b. After the observation of either element of the partition the outsider could then offer to exchange b for $a-\epsilon$, provided $\epsilon \leq \frac{2}{3}$, making a certain profit of £ϵ in the process.

We shall now show that such inconsistencies are not a problem for induced preferences. Consider a simple model, in which the preferences specified above could arise as induced preferences. There are two individuals a punter (he) and a bookmaker (she). There is an urn which contains 75 balls which are equally red R, green G and yellow Y. The punter holds endowment a which pays £10 if a red ball is drawn, £8 if a green ball is drawn and £2 if a yellow ball is drawn. The payment to the punter is made by the bookmaker. The punter may trade with the bookmaker. Then bookmaker may add 25 balls of a colour of her choice. Finally a ball is drawn from the urn which determines the payoff to the punter.

The apparent opportunity for profit presented in Example 4.1 does not exist when we take account of the hidden actions. Suppose that the bookmaker does not offer to trade with the punter, adds 25 yellow balls to the urn and allows the draw to proceed. Then her expected payment to the punter is £5½. Suppose now that the bookmaker naively tries to construct the Dutch book of Example 4.1. Then initially she will be able to charge the agent £¾ to swap the endowment for option *b*.

Suppose now, event *E* is observed. We shall model this by assuming the 25 yellow balls are removed from the urn. The bookmaker is still able to add 25 balls which may be either red or green. Then the punter's beliefs may be represented by the Dempster-Shafer update of the original capacity conditional on event E. As shown above the punter will now be prepared to pay £⅔ to exchange *b* for *a*.

Finally the bookmaker adds 25 green balls to the urn and the draw is made. The bookmaker's expected payout is £8⅔. Thus the net expected payout from the "Dutch book" strategy is £7¼, which is higher than the expected payout with no trade.

Some intuition for why the bookmaker fails to make a profit from a "Dutch book" may be gained from the following argument. Let us describe an action by a four-component vector where the first three indicate the punter's monetary payoff when the original red green and yellow balls respectively are drawn and the fourth describes the punter's payoff if one of the balls chosen by the bookmaker is drawn. Then the endowment may be represented by $a = \langle 10,8,2,2 \rangle$. However the action which the punter pays to recover after the observation of event E has the representation $a' = \langle 10,8,2,8 \rangle$. It is clear that a' is more desirable to the punter and less profitable to the bookmaker than a. Put another way, with no trade the bookmaker can ensure that her 25 balls are assigned to the overall worst outcome of a, while with the "Dutch book" strategy these balls are assigned to the worst outcome of a in event E. The key point is that the bookmaker's hidden action has changed, hence the punter is not really back where he started, as it might appear to somebody who cannot observe the hidden actions.

While the above example is very abstract, we believe that it contains, albeit in a very simplified form, some essential features of decision-making when one party can influence the probabilities. More generally, induced preferences which arise from SEU preferences over some underlying space will have the same dynamic consistency as SEU preferences themselves. Clearly it does not make sense to regard the induced preferences as less rational then the underlying SEU preferences.

In fact there is not one Dutch book argument but many, depending on the strategies allowed to the punter and the bookmaker and the information they are assumed to have. Of course it is not necessarily the case that a decision theory which is liable to one of these Dutch books will be vulnerable to all of them. It should be emphasised that no decision theory is likely to be immune to all possible Dutch books. The strongest of these

Dutch book arguments even apply to SEU (see Yaari 1985, Wakker 1993). Despite this it seems that SEU is the decision theory, which satisfies the highest standards of consistency. We believe CEU (or other) preferences which are induced by SEU preferences also satisfy the highest standards of consistency.

5 CONCLUSION

If it is assumed that the actions contain complete descriptions of all the relevant variables then we would agree that Example 4.1 is an instance of dynamic inconsistency. However in many economic problems we believe that it would be impractical to use complete descriptions. Complete descriptions would include factors such as regret and disappointment, other options available when a particular action was chosen etc. Many of these variables would be difficult to include in standard economic models. It may be more practical to use induced preferences over incompletely specified actions, these may take the CEU form or another non-expected utility form. This does not necessarily imply dynamic inconsistencies, since we cannot be sure that the non-described components of the actions are held constant.

Although for some purposes complete descriptions may be desirable, on other occasions, it may be more economical to apply theories such as CEU over truncated descriptions of the options.

Related Literature

Machina (1989) has suggested a method of updating non-expected utility preferences which is not vulnerable to Dutch books. After observation of event E the decision-maker should use the original preferences contingent upon E. This is roughly equivalent to the decision-maker formulating a plan of action for all possible contingencies at the beginning of time and not deviating from it as uncertainty is resolved. It should be noted that this procedure is equivalent to Bayesian updating for SEU preferences. This approach is not very suitable for CEU preferences since as Eichberger and Kelsey (1993), show the conditional preferences cannot be guaranteed to have the CEU form. Since the original preferences would typically be an update of some earlier preferences, it is difficult to sustain the joint hypotheses of CEU preferences and Machina's updating rule. We believe that the theory of induced preferences provides the best way to use the CEU model in a dynamically consistent way. Our proposal is compatible with Machina updating of the underlying preferences.

There is some similarity between our results and those of Gilboa and Schmeidler (1992), however the motivation is different. They are presenting new representations of CEU preferences rather than finding functional forms for induced preferences. Like the present paper they show that CEU preferences are compatible with SEU preferences over an enlarged state space. However they do not explicitly model the hidden actions, nor do they consider the Dutch book argument.

Laboratory Evidence

One objection to using induced preferences to explain deviations from SEU is that this approach has difficulties explaining the relevant laboratory evidence. If the probabilities are controlled in experiments, there would be no scope for hidden actions. There are two defences of our approach. Firstly it might be the case that subjects form habits in the real world to cope with the possibility of hidden actions. By force of habit they continue to behave in a similar way in the laboratory. Secondly they may perceive the possibility that the experimenter will take hidden actions. There are a number of psychological experiments where the experimenter has deviated from the procedure originally explained to the subjects. In addition subjects may feel that they are at a disadvantage relative to the experimenter since they will typically have a poorer understanding of both the experiment itself and the relevant theory. The very fact that somebody has decided that it is worthwhile to conduct an experiment on a phenomena, may suggest to subjects that there are hidden subtleties in the questions being asked and induce them to behave cautiously. These problems are relevant in the theory of uncertainty-aversion since, the ambiguity in the Ellsberg experiment may create the impression that the experimenter might perform a hidden action (viz change the ratio of balls in the urn).

Notes on the text

1. Note that standard economic assumptions such as monotonicity or concavity imply differentiability almost everywhere.

REFERENCES

Arrow, K.J. and L. Hurwicz, (1972), in, "An Optimality Criterion for Decision-Making Under Ignorance." in C.F. Carter and J.L. Ford (eds.), *Uncertainty and Expectations in Economics*. Oxford, Blackwell.

Barbera, S. and M. Jackson, (1988), "Maximin, Leximin and the Protective Criterion." *Journal of Economic Theory,* vol. 46, pp. 34-44.

Barrett, C.R. and P.K. Pattanaik, (1994), "Decision Making Under Complete Uncertainty: The Ordinal Framework." forthcoming, in D. Dickinson and S. Sen (eds), *Uncertainty and Risk in Economic Life*, Edward Elgar, Aldershot, UK.

Chew, S.H., E. Karni and Z. Safra, (1987), "Risk Aversion and the Theory of Expected Utility with Rank Dependent Probabilities." *Journal of Economic Theory,* vol. 42, pp. 370-381.

Choquet, G., (1953-4), "Theory of Capacities", *Annales Institut Fourier*, vol. 5, pp. 131-295.

Dow, J. and S.R.C. Werlang, (1992a), "Risk Aversion, Uncertainty Aversion and the Optimal Choice of Portfolio." *Econometrica,* vol. 60, pp. 197-204.

Dow, J. and S.R.C. Werlang, (1992b), "The ex ante non-Optimality of the Dempster-Shafer Updating Rule for Non-additive Probabilities." *Proceedings of the Fourth Conference on Theoretical Aspects of Reasoning About Knowledge,* San Mateo CA.

Dreze, J.H., (1987), "Decision Theory with Moral Hazard and State-Dependent Preferences." in J.H. Dreze, *Essays on Economic Decisions Under Uncertainty,* CUP, Cambridge.

Eichberger J. and D. Kelsey, (1993), "Uncertainty Aversion and Dynamic Consistency." University of Birmingham discussion paper no. 93-08

Eichberger J. and D. Kelsey, (1994), "Non-Additive Beliefs and Game Theory." University of Birmingham discussion paper no. 94-01.

Ellsberg, D., (1961), "Risk, Ambiguity and the Savage Axioms." *Quarterly Journal of Economics,* vol. 75, pp. 643-669.

Epstein, L., (1992), "Behaviour Under Risk: Recent Developments in Theory and Applications." in J.J.Laffont (ed), *Advances in Economic Theory,* vol. II, CUP, Cambridge.

Epstein, L. G. and M. le Breton (1993), "Dynamically Consistent Beliefs must be Bayesian." *Journal of Economic Theory*, vol. 61, pp. 1-22.

de Finetti, B. (1974), *Theory of Probability*, Wiley, New York.

Freedman, D. A. and Purves, R. A. (1969), "Bayes' Method for Bookies." *Annals of Mathematical Statistics*, vol. 40, pp. 1177-1186.

Gilboa, I. and D. Schmeidler, (1989), "Maxmin Expected Utility with a Non-Unique Prior." *Journal of Mathematical Economics,* vol. 18, pp. 141-153.

Gilboa, I. and D. Schmeidler, (1992), "Additive Representations of Non-Additive Measures and the Choquet Integral." working paper, Northwestern University.

Gilboa, I. and D. Schmeidler, (1993), "Updating Ambiguous Beliefs." *Journal of Economic Theory,* vol. 59, pp. 33-49.

Kelsey, D., (1993), "Dutch Books and Learning in a Non-Expected Utility Framework." University of Birmingham discussion paper no. 93-03.

Kelsey, D., (1994), "Maxmin Expected Utility and Weight of Evidence." *Oxford Economic Papers,* forthcoming.

Kelsey, D. and J. Quiggin, (1992), "Theories of Choice Under Ignorance and Uncertainty." *Journal of Economic Surveys,* vol. 6, pp. 133-153.

Machina, M., (1989), "Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty", *Journal of Economic Literature*, vol. 27, pp. 1622-1668.

Milne, F., (1981), "Induced Preferences and the Theory of the Consumer." *Journal of Economic Theory,* vol. 24, pp. 205-213.

Quiggin, J., (1982), "A Theory of Anticipated Utility." *Journal of Economic Behaviour and Organisation,* vol. 3, pp. 323-43.

Savage, L.J., (1954), *Foundations of Statistics,* Wiley, New York.

Schmeidler, D., (1989), "Subjective Probability and Expected Utility without Additivity." *Econometrica*, vol. 57, pp. 571-587.

Segal, U. and A. Spivak, (1990), "First-order versus Second-Order Risk Aversion." *Journal of Economic Theory*, vol. 51, pp. 111-125.

Shafer, G. (1976), *A Mathematical Theory of Evidence*, Princeton University Press, New Jersey.

Wakker, P., (1993), "Savage's Axioms Usually Imply Violation of Strict Stochastic Dominance." *Review of Economic Studies*, vol. 60, pp. 487-493.

Yaari, M., (1985), "On the Role of 'Dutch Books' in the Theory of Choice Under Risk" Nancy L. Schwartz Memorial Lecture, Northwestern University.

Yaari, M., (1987), "The Dual Theory of Choice Under Risk." *Econometrica,* vol. 55, pp. 95-115.