

Types of Uncertainty

Richard Bradley · Mareile Drechsler

Received: 19 June 2012 / Accepted: 26 July 2013 / Published online: 31 August 2013
© Springer Science+Business Media Dordrecht 2013

Abstract We distinguish three qualitatively different types of uncertainty—ethical, option and state space uncertainty—that are distinct from state uncertainty, the empirical uncertainty that is typically measured by a probability function on states of the world. Ethical uncertainty arises if the agent cannot assign precise utilities to consequences. Option uncertainty arises when the agent does not know what precise consequence an act has at every state. Finally, state space uncertainty exists when the agent is unsure how to construct an exhaustive state space. These types of uncertainty are characterised along three dimensions—nature, object and severity—and the relationship between them is examined. We conclude that these different forms of uncertainty cannot be reduced to empirical uncertainty about the state of the world without inducing an increase in its severity.

1 Introduction

It's decision time on war in the Middle East. A war which could follow an Israeli attack on Iran's nuclear facilities. Its aim would be to stop any plans that Iran might have to develop a nuclear bomb. America's defense secretary is reported to believe that there is a strong likelihood of such an attack within months. The consequences of any war are incalculable, but so too are the consequences of a nuclear armed Iran [...].

R. Bradley · M. Drechsler (✉)
Department of Philosophy, Logic and Scientific Method, London School of Economics and Political Science, Houghton Street, London WC2A 2AE, UK
e-mail: m.drechsler@lse.ac.uk

R. Bradley
e-mail: r.bradley@lse.ac.uk

So began the BBC Radio 4 programme “Decision Time”¹ on the question of “How to prevent a war in the Middle East over fears that Iran might build nuclear weapons”. The programme panel went on to consider the probability that Iran was building nuclear weapons, to examine the various actions Israel might take—including launching an attack on Iran, applying sanctions and doing nothing—and to assess the possible consequences of each of those actions in the event that Iran either was or was not building nuclear weapons. On all of these questions the panelists expressed great uncertainty: about what state Iran was in, about the actions that might be taken, about the consequences of doing so and about the desirability of these consequences.

The topic of this paper is the nature of these uncertainties, how they should be quantified and how they should be reflected in decision rules. The topic is not new and techniques for measuring and managing uncertainty have advanced considerably over the last century and a half, in tandem with the development of probability theory and modern decision theory. Indeed so close is this connection between them, that the concept of uncertainty has come to be inseparable from that of probability. In this paper we want to challenge a particular view on uncertainty associated with this development, namely that *all* uncertainty can be captured quantitatively by a *single probability* function on a suitably rich set of events or propositions. This view, although rarely articulated explicitly, can reasonably be regarded as the default in disciplines such as statistics, economics and philosophy and is testimony to the emergence of Bayesianism as a significant (and in some disciplines, dominant) intellectual current.²

What is wrong with the default view? One problem, now commonly recognised, is that it does not allow for differences in the *severity* of the uncertainty that we face. In particular it does not do justice to the difference between the situation of someone who does not know whether some event will occur or not, but knows the probability of its occurrence (i.e. who faces a known risk) and that of someone who does not have an adequate basis on which to judge how probable its occurrence is (i.e. who doesn't know what the risks are). Compare, for instance, the situation of someone who is about to toss a coin of unknown bias with someone who has had an opportunity to toss it 1,000 times and has been able to establish to their satisfaction that it is fair. While the latter can confidently assign a probability of one-half to the coin landing heads on the next toss, the former may reasonably feel that they lack the information required to settle on this assignment and choose instead to regard all probabilities for it landing heads as admissible.³

But this is not the only problem with the default view. It is equally important to recognise that we face qualitatively different *kinds* of uncertainty as well as different severities. In this paper we will attempt to enumerate and classify these different kinds and consider the degree to which they can be reduced to a single one. The main thesis of the paper is that while there *is* scope for reducing these different

¹ The program is available online at <http://www.bbc.co.uk/programmes/b01hxm1>.

² This is not of course to say that it is the *only* view being expressed. Inevitably such a broad characterisation of the state of thinking in a field will be a bit of a caricature and it is quite possible that nobody holds the default view in its completely unqualified form.

³ See Popper (1959) for an early statement of this argument.

kinds of uncertainty to what we call empirical-factual uncertainty about the state of the world, this can only be achieved at the cost of a large increase in its severity. So uncertainty can be transformed but not eliminated.

To examine and criticise the default view we will take the work of Leonard Savage as our point of departure. The choice of Savage is motivated by the fact that his version of Bayesian decision theory is the most widely known and used in economics and the other social sciences, but it is important to recognise that some of the details of our argument depend on this choice and that had we chosen to investigate uncertainty within the framework of, say, Richard Jeffrey's decision theory, there would have been some differences in emphasis (see, for instance, the discussion of Jeffrey in the section on option uncertainty). Nonetheless we do not think that anything very substantial depends on this choice of framework.

We proceed as follows. First, we present Savage's treatment of uncertainty and some of the basic problems with it. We then offer a taxonomy of uncertainty that allows us to classify the various types of uncertainty faced by decision makers along three dimensions: nature, object and severity. In subsequent sections we turn to a more detailed treatment of some different kinds of uncertainty, focusing on the question of whether they are reducible to a more fundamental kind susceptible to measurement by a single probability function.

2 Savage's Theory

From the perspective of a decision maker the most basic form of uncertainty concerns *what to do*. To decompose this basic uncertainty let us start with Savage's (1954) convenient representation of a decision problem by a matrix of the kind exhibited in Table 1, in which the A^i are the actions available to the agent, the S_j are the possible states of the world and each C_j^i is the consequence of performing action A^i when the state of the world is S_j .

Savage's way of presenting decision problems shows that in trying to decide what to do we can be uncertain about:

1. What states and consequences there are.
2. What actions are available/feasible.
3. Which state of the world is the actual one.
4. What the consequences are, in each state of the world, of performing an action.
5. What value to attach to each consequence.
6. How to evaluate acts (that is, what decision rule to use).

Table 1 Savage's decision problem

Options	States of the world		
	S_1	...	S_n
A^1	C_1^1	...	C_n^1
\vdots	\vdots	\ddots	\vdots
A^m	C_1^m	...	C_n^m

Savage's proposed resolution of all this uncertainty is both well-known and widely accepted. He argued that the decision maker should, when faced with a decision problem of the kind represented by Table 1, choose the action which maximises the subjective expectation of utility relative to a utility function on the set of consequences, measuring the degree to which she desires or values their realisation, and a subjective probability function on sets of states of the world (events), measuring the degree to which she believes the actual or true state to be contained in the set. The existence of such utility and probability measures is guaranteed, Savage showed, if the agent's preferences over actions satisfy a number of well-known conditions, including completeness, transitivity and separability ('the sure-thing principle').

The details of Savage's argument are not our main concern here. What does matter is that his treatment of decision problems seems to allow for a reduction of the decision maker's basic uncertainty about what to do to uncertainty about what the true state of the world is, i.e. to what we will call *state uncertainty*. Savage illustrates this kind of uncertainty with the example of someone who is cooking an omelet and has already broken five good eggs into a bowl, but is uncertain whether the sixth egg is good or rotten. In deciding whether to break the sixth egg into the bowl containing the first five eggs, to break it into a separate saucer, or to throw it away, the only question this agent has to grapple with is whether the last egg is good or rotten, for she knows both what the consequence of breaking the egg is in each eventuality and how desirable each consequence is. And in general it would seem that for Savage once the agent has settled the question of how probable each state of the world is, she can determine what to do simply by averaging the utilities of each action's consequences by the probabilities of the states of the world in which they are realised.

Such is the lesson that has been drawn from Savage's work by mainstream economics (see, for instance, Mas-Colell et al. (1995) and Al-Najjar and Weinstein (2009)). The view is mistaken however. Firstly, Savage's formulation of decision problems does not in itself imply that the utilities of consequences are given or known. When they are not, agents may face what we will call *ethical uncertainty*, namely uncertainty as to the value they should attach to possible consequences of their actions. Secondly the representation of the decision problem that Savage starts with assumes what he calls a 'small world' in which all contingencies are foreseeable and in which each action determines a maximally specific consequence for each state of the world. But, as he puts it:

... what are often thought of as consequences (that is, sure experiences of the deciding person) in isolated decision problems typically are in reality highly uncertain—(Savage 1954, p. 84).

Savage was well aware that not all decision problems could be represented in a small world decision matrix. In Savage's words, you are in a *small world* if you can "look before you leap"; that is, it is feasible to enumerate all contingencies and you know what the consequences of actions are. You are in a *grand world* when you must "cross the bridge when you come to it", either because you are not sure what the possible states of the world, actions and/or consequences are—i.e. you face what

we will call *state space uncertainty*—or because you are not sure what the exact consequences of your actions are in each state of the world—i.e. because you face what we will call *option uncertainty*.

Most criticism of Savage has been directed not against his treatment of grand-world uncertainty however, but against the requirements of rationality that he postulates for small-world decision making. Two lines of criticism have predominated: one emanating from the Allais' paradox and directed primarily against his famous sure-thing principle and a second, emanating from the Ellsberg paradox, that is primarily directed against the implication of his postulates that rational agents act as if they have precise probabilities for all contingencies. We return to the latter issue in Sect. 5. But first we turn our attention to the problem of providing a taxonomy of uncertainty that accommodates the various forms identified thus far.

3 A Taxonomy of Uncertainty

What are the basic forms and properties of uncertainty? Most presentations of decision theory work from Luce and Raiffa's (1957) classic distinction between situations of certainty (when the consequences of actions are known), risk (when the probability of each possible consequence of an action is known, but not which will be the actual one) and uncertainty (when these probabilities are unknown). In this section, we propose a more wide-ranging classification and consider its implications.

Our basic suggestion is that there are three fundamental dimensions to uncertainty: its nature, object and severity.

3.1 Nature

The first dimension relates the kind of uncertainty to the *nature* of the judgement being made. We distinguish three basic forms of uncertainty—modal, empirical and normative—corresponding to the nature of the judgement that we can make about the prospects we face, or to the nature of the question we can ask about them.

1. *Modal* uncertainty is uncertainty about what is possible or about what *could* be the case. It arises in connection with our possibility judgements: those concerned with what is conceivable, logically possible, feasible, and so on. For instance, in thinking about how to represent a decision problem we might be unsure as to what the possible states of the world are or what possible consequences could follow from the choice of an action. This uncertainty thus concerns the make-up of the space of states and consequences, and hence what actions are possible. (In the most severe case of modal uncertainty, the agent is *unaware* of certain states and/or consequences).
2. *Empirical* uncertainty is uncertainty about what *is* the case (or has been or would be the case). It arises in connection with our descriptive judgements.

Such uncertainty can be present even if all modal uncertainty is resolved, since we may be sure about what the relevant possible states are, but unsure as to which is the one that actually holds. (The opposite is true as well: we may be sure what the actual state of the world is, while being unsure about what it could have been).

3. *Normative* uncertainty is uncertainty about what is desirable or what *should* be the case. It arises in connection with our evaluative judgements. Normative uncertainty can be present even if all modal and empirical uncertainty is resolved: we may be sure what the state of the world is or could have been, but unsure what value to attach to either the state, or to the consequences that follow from performing an action when that state is the prevailing one. (Once again the opposite is true as well: we can be sure what value to attach to an outcome without knowing whether we are in a state in which it will come about).

3.2 Object

A second dimension relates to the *objects* of the judgements that agents make; the features of reality that their judgements are directed at. Here we distinguish two fundamental classes of object—facts and counterfactuals—and associated forms of uncertainty:

1. *Factual* uncertainty is uncertainty about the actual world; about the way things are—the facts.
2. *Counterfactual* uncertainty is uncertainty about non-actual worlds; about the way things could or would be if things were other than the way they are—the counterfactuals.

The distinction between factual and counterfactual uncertainty is orthogonal to that between the various natures of uncertainty, for there can be modal, empirical and normative uncertainty concerning the counterfactuals as well as concerning the facts. For instance I can be uncertain whether, if someone were to break into my house, the alarm would sound, whether it could fail to do so and whether it is desirable that it would do so. If in fact no-one will break into the house then my uncertainty about these questions is counterfactual. On the other hand if it's true that someone will break in then my uncertainty is factual.

The role of factual uncertainty in decision making is obvious. But counterfactual uncertainty is equally important because it conditions the agent's deliberations about what to do. Someone who finds themselves at a fork in the road in an unfamiliar part of the country might ask themselves what would happen if they were to take the left fork and what would happen if they were to take the right instead. Believing that were they to take the left fork they would come to a dead-end would give them reason to take the right fork and hence to take an action which makes this belief (about what would happen if they were to go left) one that concerns a counterfactual possibility.

3.3 Severity

The third dimension relates to the *difficulty* the agent has in making a judgement about the prospects they face, a feature that depends on the amount of judgement-relevant information that is available to them, how coherent this information is, and what inferential and judgemental skills they possess. The dimension of severity is orthogonal to the other two introduced for one can just as much face different severities of uncertainty in making normative judgements as in making empirical ones, or in making judgements about the counterfactuals as in making ones about the facts.

If we focus on the informational element of agent's uncertainty, then we can build on Luce and Raiffa (1957) to classify levels of severity as follows. In order of decreasing severity:

1. *Ignorance*: When the agent has no judgement-relevant information.
2. *Severe uncertainty*: When they only have enough information to make a partial or imprecise judgement. (This is termed ambiguity when the context is that of empirical-factual judgement).
3. *Mild uncertainty*: When they have sufficient information to make a precise judgement.
4. *Certainty*: When the value of the judgement is given or known.

This classification applies most naturally to empirical uncertainty, but even in this case there are some subtleties. For instance, when agents make decisions in a situation conventionally described as one of risk, the probabilities of the states of the world are considered part of the information the agent holds. So on our classification they are certain about the probabilities of the states but mildly uncertain about what state is the true one. In other circumstances probabilities may not be given in this way, but the agent nonetheless holds enough information to make precise probabilistic judgements. In this situation her probability judgements reflect her mild uncertainty about the state of the world, without implying any probabilistic certainty. In other words she may or may not be sure what the true probabilities are or even acknowledge that there are any such things.

4 Classification and Reduction

The sole form of uncertainty recognised by the default view is mild *state uncertainty*, i.e. factual-empirical uncertainty as to the state of the world. State uncertainty is factual because the uncertainty concerns the actual state of the world; it is empirical because it pertains to descriptive judgements of this world. We have no quarrel with the claim that uncertainty of this kind is adequately represented by a single probability function on the states of the world. But state uncertainty need not be mild: when agents lack the information or skills necessary to assign a precise probability to each state of the world – a situation typically termed *ambiguity*—then the empirical uncertainty they face is severe. Nor is empirical-factual uncertainty the only form relevant to decision-making: an agent can also face *state space*

uncertainty, when she doesn't know what the possible states of the world are; *ethical uncertainty*, when she does not know how to value the consequences of her actions; and *option uncertainty*, when she does not know what the consequences of her action are.

These forms of uncertainty occupy positions in our three-dimensional system of classification different from that of state uncertainty. State-space uncertainty is a form of modal uncertainty; ethical uncertainty a form of normative uncertainty. Option uncertainty, on the other hand, *is* a kind of empirical uncertainty, but it is of a counterfactual type, since it pertains to the question of what would be the case if a particular action were performed, rather than what is the case. This classification is summarised in Table 2 which displays two of the three proposed dimensions. The table also highlights the possibility of two other forms of uncertainty—normative-counterfactual and modal-counterfactual—and tentative examples of each have been entered into it, corresponding to uncertainty about how to value acts (act-value) and about what acts, qua functions from states to consequences, are possible (act-space).

This three dimensional taxonomy raises the question as to whether these various forms of uncertainty are independent of one another, or whether it is possible to reduce some or even all of them to some basic form of uncertainty. In particular, given the current state of the literature on the topic, it is natural to ask:


1. Can we reduce ambiguity to mild empirical uncertainty?
2. Can we reduce normative uncertainty to empirical uncertainty?
3. Can we reduce counterfactual uncertainty to factual uncertainty?
4. Can we reduce modal uncertainty to empirical uncertainty?

Each of these questions will be addressed in the subsequent sections where we examine the different forms of uncertainty individually. Although the reduction issues differ to some extent in these detailed treatments, it is possible to draw some general conclusions. Our central conclusion is that a partial reduction of nature and object *is* often possible, but only at the expense of an increase in severity. To draw an analogy with a domestic problem: just as it is possible to sweep the dirt lying under one part of the carpet to another part, but only at the cost of creating a bigger mound of dirt there, so too can uncertainty of one nature and object be converted into that of another nature and object only by increasing the severity of the uncertainty. So, in a very rough sense, total uncertainty is conserved.

This thesis can be represented diagrammatically with the help of the uncertainty simplex in Fig. 1. Every point in the simplex represents a combination of nature, object and severity. At point A, for instance, we face counterfactual, empirical

Table 2 Classification of uncertainty

Object	Nature		
	Empirical	Normative	Modal
Factual	State	Ethical	State-space
Counterfactual	Option	[Act-value]	[Act-space]

Fig. 1 Uncertainty simplex

uncertainty of moderate severity. At point *B* in the diagram we face severe empirical and factual uncertainty. At point *C* we face normative, factual uncertainty of only mild severity. Now our hypothesis is that it is only possible to travel on lines inside the simplex, so that for instance in attempting to eliminate the counterfactual uncertainty present at *A* by moving to *B* we are forced to take on uncertainty of a greater severity as the price for the nature change. Similarly if we are at *C* and wish to eliminate our normative uncertainty we can convert it into empirical uncertainty by moving to *B*, but only at the price of an increase in severity, or by moving to *A* but only by taking on counterfactual uncertainty as well.

5 Ambiguity

By “uncertain” knowledge, let me explain, I do not mean merely to distinguish what is known for certain from what is only probable. The game of roulette is not subject, in this sense, to uncertainty; . . . Even the weather is only moderately uncertain. The sense in which I am using the term is that in which the prospect of a European war is uncertain, or the price of copper . . . About these matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know—Keynes (1937, pp. 213–214).

The view that we can face uncertainty of such severity with regard to certain classes of events that we cannot ascribe a numerical probability to them goes back to at least Knight (1921) and was defended by luminaries such as Keynes. But it fell out of favour as the Bayesian view gained ascendancy in the latter half of the twentieth century, as did the complementary literature on decision making under complete ignorance (see Binmore 2009 for a discussion). Savage’s *Foundations of Statistics* can indeed be read as an argument that considerations of rational preference imply that all empirical uncertainty is ‘mild’ in the sense that a rational agent will act as if she maximises expected utility relative to a unique probability

Table 3 The Ellsberg paradox

	Red	Black	Yellow
L_1	\$100	\$0	\$0
L_2	\$0	\$100	\$0
L_3	\$100	\$0	\$100
L_4	\$0	\$100	\$100

function on the states of the world. To put it slightly differently: an agent in a situation that is objectively one of ignorance or ambiguity (because the objective probabilities of the relevant events are either non-existent or unknown) must, on pain of inconsistency, reduce it to one of mild uncertainty by assigning a subjective probability to each state of the world in accordance with the degree to which she believes that it is the actual one.⁴

Early doubts about whether in situations of ambiguity rationality does require conformity with Savage's theory were expressed by Ellsberg (1961), who conducted a set of now very famous experiments showing that agents do not in fact choose in accordance with the dictates of subjective expected utility theory in these conditions. The predominant concern of the literature on decision making under ambiguity that followed in his wake is to trace out the implications of these empirical observations for Savage's theory and more generally for Bayesian theories of decision making, an endeavour that is typically justified via an appeal to the robustness of the empirical findings made in Ellsberg's experiments (for instance, Slovic and Tversky 1974 show that if given the opportunity to reconsider the preference expressed in Ellsberg's experiment, subjects choose not to reverse their preferences).

Let us start by recalling Ellsberg's three colour experiment. An urn contains 90 balls, 30 of which are red, and the remaining 60 are black or yellow in an unknown proportion (see Table 3). Subjects are asked to choose between two bets. The first, L_1 , pays off \$100 if in a random draw from the urn a red ball is drawn. The second, L_2 , pays off \$100 if a black ball is drawn. Most subjects express a preference of L_1 over L_2 . In a second choice problem, subjects are asked to choose between L_3 and L_4 , which pay out \$100 in the events "red or yellow" and "black or yellow" respectively. Here, most subjects express a preference of L_4 over L_3 . As can easily be verified, the choice of L_1 and L_4 is inconsistent both with the sure-thing principle and with the existence of a unique probability distribution over the states: Whilst the preference of L_1 over L_2 (L_2 over L_1) implies that the agent ranks event "red" ("black") as subjectively more likely than "black" ("red"), the preference of L_4 over L_3 (L_3 over L_4) entails that the agent ranks the event "black" ("red") as subjectively more likely than "red" ("black").

Ellsberg's own explanation for these preferences is that subjects are *averse* to the ambiguity about the precise probability distribution over the state space. In the first choice situation subjects are given information which makes it reasonable for them

⁴ Though, as Binmore (2009) emphasises, Savage only held that this applied in small-world decision making.

to put the probability of drawing a red ball at one-third, but with regard to the probability of a black ball they know only that it is no more than two-thirds. In view of this many subjects, Ellsberg conjectured, would ‘play it safe’ and opt for the lottery which pays out with a known probability over the one in which there is a good deal of uncertainty about the probability of it paying out. Similar reasoning would lead them, in the second choice problem, to pick lottery L_4 which has a ‘known’ probability of two-thirds of a win over L_3 with its unknown probability of a win⁵.

There are three aspects of Ellsberg’s account that need to be separated: firstly, the question of which of Savage’s axioms are normatively compelling in ambiguous decision problems, secondly, how agents should represent their uncertainty in intuitively ambiguous situations and finally, the question of how they should make choices, given their uncertainty.

While Savage’s postulates are, as many decision theorists are agreed, convincing for small world decision problems, they are less so in ambiguous decision problems (much less decision problems under ignorance). For large world decision situations, Savage’s argument has an obvious weakness: the postulate of complete preferences. For in situations of empirical ambiguity it is implausible that rationality requires agents to be able to judge of any two prospects that one is better or preferable than the other or that they are equally so. Even if we set aside the forms of normative uncertainty that we discuss in the next section, so that it is reasonable to assume that agents have complete preferences over *final* outcomes, it remains the case that if they are unable to judge how probable the various contingencies upon which the realisation of these outcomes depend, given a choice of act, they will simply be unable to assess the various acts amongst which they must choose.⁶ But if this is the case why should they conform to Savage’s other postulates? The question of which postulates can be seen as requirements of rationality from a normative point of view remains, to a great extent, an open one.

Ambiguity not only impacts on the formation of preference, but also on the agent’s representation of uncertainty. For can we require that in a situation of ambiguity a rational agent forms a unique and additive probability distribution over the state space? One well-known argument against the reduction of ambiguity to mild uncertainty is given in the following example: suppose a coin of unknown bias is to be tossed and a prize will be awarded depending on whether it lands heads or tails and on what act I choose. Suppose that I must choose between an act which wins the prize if the coin lands heads, one which wins the prize if the coin lands tails, and one which gives me a 50 % chance of the prize whether the coin lands heads or tails. Savage’s theory requires that if I am indifferent between the first two acts then I must be indifferent between them and the third. But it does not seem irrational for me to choose the third on the grounds that in doing so I am able to fix my chances of a prize at 50:50. If this is so then I am not rationally required to be a

⁵ This epistemic reading of ambiguity aversion is not the only one to be found in the literature. See, for instance, Fox and Tversky (1995).

⁶ An argument of this kind has been made for the case of Knightian uncertainty by Bewley (1986).

subjective expected utility maximiser; hence not required to quantify my uncertainty with a single probability function.

Based on this intuition, a number of economists and philosophers, notably Gilboa and Schmeidler (1989), Levi (1974, 1985) and Joyce (2010), have argued that the most natural way of understanding the epistemic situation of the agent in Ellsberg's set-up is that they are unable to determine which of a set of possible probability distributions is the true one. Indeed they should not form beliefs that are any more precise than is permitted by the information they hold.

Imprecise beliefs are naturally represented by *sets* of probability functions, rather than a singleton probability, the intuitive idea being that each member of the set is a candidate for being the true probability (if there is one) or a probabilistic belief that is admissible in the light of the evidence. The severity of the uncertainty the agent faces will be reflected in the size of the set of permissible probabilities; in the limiting case of ignorance the set will contain all probability distributions, in the other limiting case of mild uncertainty it will just contain one. (Following Schmeidler (1989), some economists prefer to represent imprecise beliefs by capacities or non-additive belief functions rather than sets of probabilities, but the differences between these representations will not matter to our discussion).

The view that agents can, and sometimes should, have imprecise beliefs is now quite widely held.⁷ There is far less of a consensus on how choices should be made by agents with imprecise beliefs. One common view is that they should calculate the expected utilities of acts relative to each of the probability distributions they regard as permissible, identify the minimum expected utility (MEU) of each act and choose the one which has maximum MEU. For instance suppose that the state of uncertainty of an agent facing the Ellsberg problem is given by a family of probability measures each assigning some value p in the interval $[0, \frac{2}{3}]$ to the probability of Black and corresponding value $1 - p$ to Yellow. Then while lottery L_1 has expected utility $\frac{1}{3} \times U(\$100) + \frac{2}{3} \times U(\$0)$, lottery L_2 has expected utility in the range $[U(\$0), \frac{2}{3} \times U(\$100) + \frac{1}{3} \times U(\$0)]$. Assuming that utility is a positive function of money, the minimum value here is $U(\$0)$, so lottery L_1 is better according to the MEU criterion. On the other hand lottery L_4 is better than lottery L_3 since it has expected utility of $\frac{2}{3} \times U(\$100) + \frac{1}{3} \times U(\$0)$ while L_3 has a minimum expected utility value of $\frac{1}{3} \times U(\$100) + \frac{2}{3} \times U(\$0)$.

To give axiomatic foundations to maximisation of MEU, Gilboa and Schmeidler (1989) use the Anscombe-Aumann (1963) framework (which is a variation of Savage's theory that uses objective probabilities to figure in the objects of choice), but restrict separability to convex combinations of any act with a constant act. Their main additional postulate, the axiom of *uncertainty aversion*, requires agents to weakly prefer convex combinations of acts that are equally preferred to one another to either of the individual acts. Agents that are uncertainty averse in this sense choose as if they had imprecise beliefs and maximise MEU in the light of them. So the normative status of this axiom is of central importance in establishing the permissibility of behaviour contrary to Savage's prescriptions.

⁷ See for instance Walley (1991), Joyce (2010), Bradley (2009) and Levi (1985).

Maximising minimum expected utility is arguably too conservative a decision rule, focusing as it does on the worst case scenario. But the sets-of-probabilities representation of the agent's subjective uncertainty can support a number of other, more plausible rules that permit hedging against uncertainty; for instance, that agents should maximise the average expected utility of acts relative to some set of weights on admissible probability functions, or a weighted sum of the upper and lower expected utilities. For a detailed survey of the now extensive literature on decision making under ambiguity we refer the reader to Gilboa and Marinacci (2011).

6 Ethical Uncertainty

Ethical or value uncertainty arises when the values to be used in assessing the desirability of decision-relevant prospects are either unknown, so that the decision maker must rely on subjective evaluations of them, or do not exist, so that the decision maker must construct them.⁸ Ethical uncertainty is typically ignored by decision theorists, because of their (often unconscious) attachment to the view that values are determined by the agent's subjective preferences, in the sense that what makes a consequence valuable to the agent is just that she desires it to some degree, or that she prefers it to a greater or lesser extent to other consequences. Let us call this view *Ethical Subjectivism*. If it were correct, ethical uncertainty would be a minor phenomenon, as one is not normally uncertain about what one's own judgement on something is (just about what it should be). Indeed questions such as 'What utility should I attach to this outcome' seem barely intelligible on this view. If a prospect's value for an agent is determined by her preferences, she cannot be right or wrong about what value to attach to them; nor can her preferences be criticised on grounds of their failure to adequately reflect one value or another.

There are, however, at least two ways in which one can be uncertain about the value of consequences or, more generally, whether one consequence is preferable to another. Firstly one may be uncertain about the factual properties of the consequence in question. If possession of the latest Porsche model is the prize in a lottery one is considering entering, one may be unsure as to how fast it goes, how safe and how comfortable it is, and so on. This is mild factual-empirical uncertainty and it can be 'transferred' (subject to some qualifications discussed in the next section) from the consequence to the state of the world by making the description of the consequence more detailed. For example, the lottery may be regarded as having one of several possible consequences, each an instantiation of the schema 'Win a car with such and such speed, such and such safety features and of such and such comfort', with the actual consequence of winning depending on the uncertain state of the world.

Secondly one can be unsure as to the value of a consequence, not because of uncertainty about its factual properties, but because of uncertainty about the extent

⁸ The term 'ethical' is used here in the same way that it is used by Ramsey, to denote that which has to do with what matters to the agent. It is not meant to be read as having only to do with morality.

to which these properties are valuable. One may know all the specifications of the latest Porsche and Ferrari models, so that they can be compared on every dimension, but be unsure whether speed matters more than safety or comfort. Once all factual uncertainty has been stripped from a consequence by detailed description of its features, one is left with pure value uncertainty of this kind.

The Ethical Subjectivist may draw on the first point to elaborate her position. What she rules out is *pure* value uncertainty. But if there is factual uncertainty then an agent may well be unsure about the desirability of any less than fully specified prospect. Likewise her judgements about them may be criticised if they are based on false beliefs and revised by the agent in the face of evidence. In a nutshell, what might look like value uncertainty is in fact just factual uncertainty in disguise.

Ethical Subjectivism, as we have characterised it, is a species of non-cognitivism. It involves two claims: firstly, that desirability or utility judgements don't express beliefs and secondly, that they don't track any kind of objective value facts. The 'opposite' view is *Ethical Cognitivism*: the view that utility judgements do express beliefs about objective normative facts. On this view, what we are calling value uncertainty is just factual-empirical uncertainty about these normative facts. The uncertainty one experiences about whether to help a friend, for instance, is uncertainty about whether it is in fact good to help one's friend or whether it is true that it is better to help one's friend than to further one's own interests. So, on this view the difference between uncertainty about whether it will rain and about whether it is good that it rains is to be located in the type of proposition about which one is uncertain, not in the nature of uncertainty.

Ethical Cognitivism can be made more precise in a number of different ways. One (rather simplistic) way would be to say that the desirability or utility of any prospect just is the probability that the prospect in question is desirable or good. Lewis (1988) takes this as an example of what he calls the 'Desire-as-Belief thesis' and argues that it is inconsistent with decision theory together with some mild assumptions about how agents revise their beliefs. From this he draws the conclusion that no desire or preference is fully determined by a belief (and so, in particular, by a normative belief). There has been considerable debate about the significance of Lewis' result, with some authors (e.g. Oddie 1994 and Weintraub 2007) seeing it as a refutation of Cognitivism and others (e.g. Broome 1991) arguing that it is not. The debate has been conducted in the framework of Jeffrey's (1965) decision theory, but we can nonetheless give a flavour of the problem by transcribing Lewis' argument into Savage's framework.

For simplicity suppose that there are just two utility values assigned to consequences: 0 for 'bad' and 1 for 'good'. Let A be any action and let the event \mathring{A} be the set of states of the worlds in which the action A has a consequence with utility 1. Intuitively, \mathring{A} is the event of the action A being a good one. Let the agent's degrees of belief be given by probability measure Pr . It follows that $EU(A) = 1 \times \text{Pr}(\mathring{A}) + 0 \times (1 - \text{Pr}(\mathring{A})) = \text{Pr}(\mathring{A})$. Suppose the agent comes to believe that they will in fact perform action A , so that their degrees of belief are now given by a probability measure Pr_A . Does this change of belief state have any effect on the expected utility of A ? No, because in Savage's framework the probabilities of states

of the world are independent of the action performed. It follows that $\Pr_A(\overset{\circ}{A}) = \Pr(\overset{\circ}{A})$. But this must be true irrespective of what the agent believes to be the case. In particular it must be true even if the agent learns or believes that Q : that either not $\overset{\circ}{A}$ or not A . But this leads to a contradiction, for $\Pr_A(\overset{\circ}{A}|Q) = 0$ but $\Pr(\overset{\circ}{A}|Q) > 0$.

There are a host of objections one can make to this argument, some of which will apply to Lewis' version too and some of which will be peculiar to Savage's framework (for instance, one might object that this just goes to show that states cannot be probabilistically independent of the action performed). But the point that we want to make is that although the Desire-as-Belief thesis in the form presented here is *implied* by Savage's framework, given the various assumptions made in the course of the argument, this does not in itself commit Savage to the existence of normative facts, qua properties of the world. Or to put it a bit more carefully, our argument drew on normative facts, but not necessarily ones that are independent of the agent's degrees of desire. Consequently, our Lewis-style argument doesn't offer any support for or against either Ethical Subjectivism or Ethical Cognitivism. What is at stake between these accounts of value uncertainty is not therefore commitment or otherwise to the Desire-as-Belief thesis, but the interpretation of the utilities assigned to consequences: in particular, whether these utilities are objective (i.e. they are features of the consequences) or subjective (i.e. they are features of the agent's judgements).⁹

It is worth mentioning a third view, intermediate between Ethical Subjectivism and Ethical Cognitivism, namely that ethical uncertainty is uncertainty about the agent's tastes or fundamental preferences. This view accepts the subjectivist line that there are no preference-independent values, but treats the agent's preferences, or those features of her that determine her preferences, as factual properties of the agent. The thought is, what value an agent will assign to a commodity or good depends not just on features of the commodity itself (for instance, the speed, safety and comfort of the cars) but also on features of the consumer: their likes and dislikes, their capacities (for instance, their driving skills) and their needs. One can be just as uncertain about the latter class of facts as the former. Consider, for instance, a policy decision which has ramifications for a large number of people and which we want to evaluate in terms of the attitudes those affected will take to its consequences. Even if we are certain about what the consequences of the adoption of the policy will be, we may be uncertain about how those who are concerned will judge it because we are uncertain about their preferences. The same is true for actions which have consequences for ourselves that lie well into the future, when our tastes, skills or needs might have changed in ways that we cannot predict with certainty.

On this view ethical uncertainty is just factual uncertainty on the part of the decision maker about what the true preferences are of those affected by their decision (including themselves). The view has some plausibility in cases like the car purchasing one: there could be some fact of the matter as to whether one prefers speed to safety and to what extent, even if it takes some experimentation to work out

⁹ This is broadly the same conclusion as is reached by Broome (1991).

what this is. But many cases are not like this. When we are uncertain about whether it is more important to help a friend or to further one's own interests, the difficulty that we have in deciding the question stems not from the fact that we don't know what we in fact prefer but that we don't know what we *should* prefer. Indeed we doubt that in such cases there really is anything like a set of pre-given preferences waiting to be discovered. Or to take a different type of example, consider trying to decide whether to take up the violin or fencing. Can the problem be described as trying to work out what one's tastes are? This seems implausible. One's tastes are likely to be shaped by the decision itself, for in pursuing the violin one will learn to appreciate one set of skills, in taking up fencing one will learn to appreciate another.

All the views discussed so far treat ethical uncertainty as a kind of factual uncertainty, differing only with regard to the kinds of facts that they countenance and consider relevant. They are, in that sense, reductive views. The last view to consider holds that ethical uncertainty is different in kind from factual uncertainty and is directly expressed in utility judgements, rather than in second-order probability judgements about tastes or first-order probability judgements about normative facts. Making this claim precise requires some care. Utility judgements are like probability judgements in that they are judgements about the world (and not just expressions of the agent's mental state), but they are nonetheless a different kind of judgement. While we can say that one's probability for rain tomorrow, say, reflects the degree to which one is uncertain as to whether it will rain then, it is not the case that one's utility for rain expresses the degree to which one is uncertain as to whether it is good that it rains. Rather it expresses one's uncertainty as to how good it would be if it rained. On the reductive views, once we know all the facts—about what will happen when it rains, how much people like getting wet, and so on—all such uncertainty is removed and the desirability of rain is fully determined by either the relevant normative facts or by the agent's subjective degrees of desire for rain, given the facts. On the non-reductive view, even when we know all the facts we can be unsure as to how desirable rain is, given the facts. There can, as it were, be value uncertainty all the way down.

It is not our intention to adjudicate on these competing views, but rather to point out the trade-off we face in the choice of which to adopt. Suppose for instance we adopt Ethical Cognitivism. Then we must introduce evaluative prospects into the domain of the probability function measuring decision-relevant factual uncertainty. The effect of this will normally be to increase both the amount and the severity of the uncertainty that an agent must handle. For now agents must attribute probabilities not only to 'material' features of their environment but also to the evaluative features determining the value of the consequences of their action. Since the truth or falsity of evaluative propositions are much more difficult to settle, the severity of the uncertainty is also higher. Indeed we would claim that the normal condition we are in with regard to evaluative propositions is neither that of mild uncertainty, nor that of ignorance, but of severe uncertainty.

If we adopt a non-reductive view, on the other hand, then we are required to acknowledge the existence of irreducible ethical uncertainty and to develop techniques for measuring and managing it. In this regard the current situation is mixed. For situations of mild uncertainty, they already exist. For adoption of the

non-reductive view requires no revision to formal decision theory in this case. Utility, on one of its standard interpretations, measures the agent's degree of preference for a prospect. When she is uncertain as to how valuable she should regard a prospect then her preferences will embody this uncertainty. For instance, if she is unsure which of two goods is better she may prefer a few of each to all of either of them i.e. the utility she assigns to receiving bundles of the two goods will reflect her uncertainty as to which is best. However, from the point of view of rational decision making nothing changes: an agent should still maximise her subjective expectation of subjective utility.

The situation is different if her ethical uncertainty is more severe. In this case, our recommendation would be that the agent represents her value uncertainty by working with *sets* of utility functions, rather than a single one, analogously to the use of multiple probabilities in decision making under ambiguity. Each utility function in the set represents a possible resolution of her normative uncertainty, to be discarded if it proves untenable in the light of experience or deliberation. An agent facing severe normative uncertainty cannot of course make her choices simply on the basis of maximisation of expected utility, but she could choose in such a way as to maximise the subjective expectation of *imprecise* utility, relative to a probability on states of the world and an imprecise utility function that assigns real numbers to sets of utilities. This leaves open the question of the form of imprecise utility measure should take. There is a small literature on multiple utility representations—see for instance Levi (1986), Schervish et al. (1995), Bradley (2009) and Karni (2013)—but it offers sparse help on this question. Several natural candidates present themselves: the average utility, the minimum utility and a weighted average of the maximum and minimum utilities in the set. But proper appraisal of these alternatives is beyond the scope of this paper.

7 Option Uncertainty

In Savage's representation of a decision problem actions are associated with definite consequences, one for each state of the world. These consequences are, in Savage's words, "sure experiences of the deciding person" and the description of them should leave no decision-relevant aspect out of the model. But in real decision problems we are often unsure about the relationship between actions, worlds and consequences, either because we do not know what consequence follows in each possible state of the world from a choice of action, or because we don't know what state of the world is sufficient for a given action to bring about that consequence. For instance, we may be uncertain as to whether taking an umbrella will certainly have the consequence of keeping us dry in the event of rain. Perhaps the umbrella has holes, or the wind will blow it inside out or the rain will be blown in from the sides.

We can put this difficulty in slightly different terms. A possible action may be *defined*, as in Savage's theory, by a particular mapping from states to consequences. Then no uncertainty about the mapping itself can arise. But what we will then be unsure about is which actions are actually available to us, i.e. which of the various hypothetical actions are real options. So while on the first way of looking at the

problem the agent faces a form of empirical-counterfactual uncertainty that we earlier termed option uncertainty, on the second what she faces is a form of modal uncertainty, regarding what the space of possible actions looks like (which we called act-space uncertainty earlier on). Our focus here will be on the former.

Option uncertainty is an endemic feature of decision making, for it is rarely the case that we can predict consequences of our actions in every detail (alternatively, be sure what our options are). And although in many decision situations, it won't matter too much what the precise consequence of each action is, in some the details will matter very much. For example, when choosing between various package holidays my satisfaction may depend on such details as whether I will get food poisoning from the hotel food or a noisy construction project will begin outside my hotel room window. So understanding how to handle option uncertainty is important.

When an agent faces option uncertainty, she cannot represent her decision problem in the manner of Table 1. But she can do something quite similar by replacing the fine-grained consequences that play the role of Savage's 'sure experiences of the deciding person' with *sets* of such fine-grained consequences—intuitively, the set of consequences the agent regards as possible given the act and state in question. This is exhibited schematically in Table 4 in which each act A^i is represented as a function from each state S_j to a set of associated possible consequences $C_j^i = \{c_{j1}^i, c_{j2}^i, \dots, c_{jm}^i\}$, i.e. by what Ghirardato (2001) terms a correspondence from states into sets of consequences. The larger the sets of possible consequences the greater the option uncertainty facing the agent.

As before there are two basic strategies that can be pursued in handling option uncertainty. The first is to accept the presence of option uncertainty and try and develop both a measure of it and decision rules appropriate to this measure. Although it is not usually understood in these terms, one can regard the debate between evidential and causal decision theorists as being about how this should be done. Evidentialists such as Jeffrey (1965) regard the conditional probability, given that an action is performed, as giving the correct measure of the uncertainty associated with acting, while causal decision theorists such as Joyce (1999) argue that what is required is a measure of the probability under the counterfactual supposition that the action is performed. If evidentialism were correct then it would follow that a single probability function would suffice not only to measure state uncertainty but option uncertainty as well. But the difficulty that evidential decision theory faces in dealing with Newcomb's paradox and other more homely cases in which probabilistic correlation fails to provide a good guide to causal efficacy

Table 4 Decision making under option uncertainty

Options	States of the world		
	S_1	...	S_n
A^1	$C_{11}^1, C_{12}^1, \dots, C_{1y}^1$...	$C_{n1}^1, C_{n2}^1, \dots, C_{nx}^1$
\vdots	\vdots	\ddots	\vdots
A^m	$C_{11}^m, C_{12}^m, \dots, C_{1y}^m$...	$C_{n1}^m, C_{n2}^m, \dots, C_{nz}^m$

suggest that it is not. So it is reasonable to conclude that the upshot of this debate is that a different probability measure is required for option uncertainty from that used to measure state uncertainty.

The second strategy one can pursue is to try to reduce or transform option uncertainty into some other kind of uncertainty so that existing decision rules can be applied. Decision theorists who subscribe to the default view, for instance, typically attempt to reduce option uncertainty to uncertainty about the state of the world by refining their description of the states until all contingencies are taken care of. They will regard a state of the world as insufficiently described by the absence or presence of rain, for instance, and argue that one needs to specify the speed and direction of the wind, the quality of the umbrella, and so forth. There are at least two reasons why this reductive strategy will not work on all occasions. Firstly because, according to our best scientific theories, the world is not purely deterministic. When the conditions under which a coin is tossed do not determine whether a coin will land heads or tails, for instance, the act of tossing the coin does not have a predictable consequence in each state of the world. Secondly, even if we are in a purely deterministic set-up, it may be subjectively impossible for the decision maker to conceive of and then weigh up all the relevant contingencies or to provide descriptions of the states of the worlds that are sufficiently fine-grained as to ensure that a particular consequence is certain to follow, in each state, from the choice of any of the options open to them. It is, as Savage (1954, p. 16) noted, “utterly beyond our power to plan a picnic or to play a game of chance in accordance with the principle [of considering all relevant contingencies], even when the world of states and set of available acts to be envisaged are artificially reduced to the narrowest reasonable limits”. And even if one could envisage all the possibilities, one may simply not know what state of the world is sufficient for the act of taking an umbrella to keep one dry.

To get around these difficulties the reductionist can employ two different strategies. The first is to use descriptions of the states of the world that identify the set of conditions sufficient for the determination of the consequence, given the performance of the action, without actually listing the conditions. For instance, one can turn Savage’s theory around and take actions and consequences as the primitives and then define states of the world as consequence-valued functions ranging over actions. This would lead to a decision matrix of the kind exhibited in Table 5, in which each $S_j(A^i, C_j^i)$ denotes the state that maps action A^i to consequence C_j^i .

Table 5 States as functions of acts and consequences

Options	State		
	$S_1(A^i, C_1^i)$...	$S_n(A^i, C_n^i)$
A^1	C_1^1	...	C_n^1
\vdots	\vdots	\ddots	\vdots
A^m	C_1^m	...	C_n^m

Table 6 Coarsened consequences

Options	State		
	S_1	...	S_n
A^1	$C(A^1, S_1)$...	$C(A^1, S_n)$
\vdots	\vdots	\vdots	\vdots
A^m	$C(A^m, S_1)$...	$C(A^m, S_n)$

Similar moves are advocated in the philosophical literature. Lewis (1981), for instance, treats states as ‘dependency hypotheses’, which are just maximally specific propositions about how consequences depend causally on acts. Many causal decision theorists follow Stalnaker’s (1981) suggestion that a state of the world be denoted by a conjunction of conditional sentences of the form ‘If action A were performed then consequence C would follow; if action A' were performed then consequence C' would follow; if ...’. In this case we can replace the state function with the conjunction of conditional sentences that describe it.

By pursuit of any version of this reductive strategy, option uncertainty is transformed into a particular kind of state uncertainty, namely uncertainty as to the true mapping from actions to consequences, or as to the truth of a conjunction of conditionals. This strategy is likely to lead to an increase in the severity of the uncertainty that the agent faces. For how is she to assign probabilities to states so described? What probabilities should be attached to counterfactual conditionals and conjunctions of them? If she is unable to resolve these questions completely then she will find herself in a situation of ambiguity rather than mild uncertainty.

A second reductionist strategy for dealing with option uncertainty is to coarsen the description of the consequences to the degree necessary to ensure that we can be certain it will follow from the exercise of an option in a particular state. In his treatment of what he calls ‘small worlds’, for instance, Savage acknowledges the need for “acts with actually uncertain consequences to play the role of sure consequences in typical isolated decision situations” (Savage 1954, p. 84). Following this advice leads to a decision matrix of the kind displayed in Table 6 in which $C(A^i, S_j)$ denotes the consequence (whatever it is) of performing the action A^i in state S_j .

Pursuit of this strategy converts option uncertainty, not into empirical uncertainty about the state of the world, but into uncertainty about the desirability of the consequence as described—a version of what we previously called ethical uncertainty. We may be sure that the act of taking an umbrella in a rainy state will have the consequence of being able to protect us against the rain by opening the umbrella. But whether this is a good thing or not depends on contingencies that by assumption we are unable to enumerate or identify. How bad it is to get soaked, for instance, depends on how cold the rainwater is and rain temperature may be a variable about whose determinants we know very little. And in this case our ethical uncertainty is likely to be severe for the consequences are normatively ambiguous.

In conclusion, it would seem that option uncertainty can be ‘swept under the carpet’ by reducing it to either state uncertainty or ethical uncertainty. But the cost

in both cases is an overall increase in the severity of the uncertainty experienced in the dimension to which option uncertainty is reduced.

8 State Space Uncertainty

In the Savage framework, the state space is a primitive of the theory and is exogenously given. In real decision problems, however, a state space the elements of which are mutually exclusive and collectively exhaustive may not be given and may be hard to construct. There are, in the real world, *unforeseen contingencies*: eventualities that even the educated decision maker will fail to foresee. For instance, the recent tsunami and subsequent nuclear meltdown in Japan are events that most agents would have omitted from their decision models. If a decision maker is aware of the possibility that they may not be aware of all relevant contingencies—a state that Walker and Dietz (2011) call ‘conscious unawareness’—then they face state space uncertainty.

There are two variants of the problem of unforeseen contingencies that should be distinguished. The first concerns cases in which the state space is not sufficiently detailed to capture all decision-relevant factors; in this case, the state space is excessively coarse-grained. For instance, my decision as to whether to go for a walk on the beach might not only hinge on whether it is sunny, but also on whether the beach is littered, or whether a nearby bar is playing music too loudly. Omission of these details regarding the state of the world may make it impossible to be sure what the exact consequence of performing an action in that state would be, i.e. it gives rise to option uncertainty (indeed, Walker and Dietz (2011) analyse the problem of unforeseen contingencies in precisely this way). Alternatively, the use of a coarse partition may reflect the ambiguity of the fully refined partition. The agent is simply unable to assign probabilities to the fine partition and hence is forced to work with the coarser one. So in this case the agent can eliminate her state space uncertainty in favour of either option uncertainty or ambiguity.

A distinct, second type of state space uncertainty arises when an agent is aware of the possibility that she has entirely omitted a state of the world from the state space. Again it is important to distinguish the case in which she nonetheless knows what the range of possible consequences are and the case in which there may be unforeseen consequences. The first case too can be recast as a combination of option uncertainty and ambiguity. Suppose the agent starts with a state space, $\mathcal{S} = \{S_1, \dots, S_n\}$ which enumerates all contingencies she can think of. Aware that there might be further contingencies that she is unaware of, she adds a catch-all state S_{n+1} —the ‘any other contingencies not yet enumerated’ state—thereby eliminating her state space uncertainty. But this catch-all state has a completely unknown consequence, so option uncertainty now arises. Furthermore, she has no way of assigning a probability to this state and so she faces empirical ambiguity.

The second case, where the agent is aware that there might be unforeseen consequences of actions as well as unforeseen states, is one which *prima facie* can’t be reduced to any other type of uncertainty. Irreducible state space uncertainty is the most difficult kind of grand-world uncertainty to handle with the resources of

Savage's framework, for the notion of an optimal choice is defined relative to a given state-consequence matrix. But if the agent is unsure whether the matrix is exhaustive then they can't rely on calculations of expected utility relative to that matrix to prescribe the best course of action. Indeed the rationale for consistency itself is considerably weakened. This latter point has been made forcefully by Binmore (2009), who argues:

In a large world, the possibility of an unpleasant surprise that reveals some consideration overlooked in [the] model can't be discounted. [...] Someone who insists on acting consistently come what may is just someone who obstinately refuses to admit the possibility of error. (Binmore, 2009, p. 117).

The problem of how to handle the possibility of unforeseen contingencies has received little attention in philosophical decision theory—a notable exception being Hansson who gives careful consideration to the question of rational responses to the problem (see Hansson 1994; 2013). In economics, Gilboa and Schmeidler's (1995) theory of case-based decisions, where decision makers evaluate prospects according to both their similarity to previously encountered problems and their expected utility, is often cited with suitability to large world decision making because it does not require that agents have a complete state space in order to make decisions. Similar considerations might seem to motivate attention to theories of bounded rationality, but while the descriptive promise of such a move is clear enough, it is far less clear how this would help in answering the normative problem. Going in a more conservative direction, Kreps (1992) introduces unforeseen contingencies in a Savage model through the idea of a subjective state space, arguing that an agent's anticipation of unforeseen contingencies is expressed in their preference for flexibility. The idea of a subjective state space has been developed further by Dekel et al. (2001) and Epstein and Seo (2009). But none of these theories directly engage with the normative question of how an agent *should* deal with the possibility of unforeseen contingencies. So this remains an open question to a great extent.

9 Conclusion

Contrary to the default view, uncertainty has different natures, objects and severities, a fact which precludes capture of all uncertainty by a single probability function. It follows that the scope of standard subjective expected utility theory is restricted to situations in which a number of forms of uncertainty is resolved. What rationality requires of decision makers in situations in which this is not so is still to a large extent an open question. We have attempted to suggest some answers in this paper, drawing on recent work in economics and philosophy. But our main aim has been to establish the importance of this work for a full understanding of decision making under uncertainty.

Acknowledgements We would like to acknowledge the support of the Arts and Humanities Research Council (Grant Reference: AH/I003118/1), The Netherlands Organisation for Scientific Research (Grant

Number 236-20-005), and the Fonds National de la Recherche Luxembourg (Grant Number 09-194). We are grateful to Hykel Hosni and Casey Helgeson for their comments on an earlier draft.

References

- Al-Najjar, N. I., & Weinstein, J. (2009). The ambiguity aversion literature: A critical assessment. *Economics and Philosophy*, 25, 249–284.
- Anscombe, F. J., & Aumann, R. J. (1963). A definition of subjective probability. *The Annals of Mathematical Statistics*, 34(1), 199–205.
- Bewley, T. F. (1986/2002). Knightian decision theory. Part I. *Decisions in Economics and Finance*, 25, 79–110.
- Binmore, K. (2009). *Rational decisions*. Princeton, NJ: Princeton University Press.
- Bradley, R. (2009). Revising incomplete attitudes. *Synthese*, 171(2), 235–256.
- Broome, J. (1991). Desire, belief and expectation. *Mind*, 100(2), 265–267.
- Dekel, E., Lipman, B. L., & Rustichini, A. (2001). Representing preferences with a unique subjective state space. *Econometrica*, 69(4), 891–934.
- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics*, 75(4), 643–669.
- Epstein, L. G., & Seo, K. (2009). Subjective states: A more robust model. *Games and Economic Behavior*, 67(2), 408–427.
- Fox, C. R., & Tversky, A. (1995). Ambiguity aversion and comparative ignorance. *The Quarterly Journal of Economics*, 110(3), 585–603.
- Ghirardato, P. (2001). Coping with ignorance: Unforeseen contingencies and non-additive uncertainty. *Economic Theory*, 17, 247–276.
- Gilboa, I., & Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18, 141–153.
- Gilboa, I., & Schmeidler, D. (1995). Case-based decision theory. *The Quarterly Journal of Economics*, 110(3), 605–639.
- Gilboa, I., & Marinacci, M. (2011). Ambiguity and the Bayesian paradigm. *Mimeo*.
- Hansson, S.-O. (1994). Decision making under great uncertainty. *Philosophy of the Social Sciences*, 26(3), 369–386.
- Hansson, S.-O. (2013). *Decision theory. A brief introduction*. <http://home.abe.kth.se/soh/decisiontheory.pdf>.
- Jeffrey, R. C. (1965). *The logic of decision*. New York: McGraw Hill.
- Joyce, J. (1999). *Foundations of causal decision theory*. Cambridge: Cambridge University Press.
- Joyce, J. (2010). A defense of imprecise credences in inference and decision making. *Philosophical Perspectives*, 24(1), 281–323.
- Karni, E. (2013). Subjective expected utility with incomplete preferences. *Econometrica*, 81(1), 255–284.
- Keynes, J. M. (1937). The general theory of employment. *The Quarterly Journal of Economics*, 51(2), 2009–2023.
- Knight, F. H. (1921). *Risk, uncertainty and profit*. Boston, MA: Hart, Schaffner and Marx.
- Kreps, D. M. (1992). Static choice in the presence of unforeseen contingencies. In P. Dasgupta, D. Gale, O. Hart, E. S. Maskin (Eds.), *Economic analysis of markets and games* (pp. 258–281). Cambridge, MA: MIT Press.
- Levi, I. (1974). On indeterminate probabilities. *The Journal of Philosophy*, 71(13), 391–418.
- Levi, I. (1985). Imprecision and indeterminacy in probability judgment. *Philosophy of Science*, 25(3), 390–409.
- Levi, I. (1986). *Hard choices: Decision making under unresolved conflict*. Cambridge: Cambridge University Press.
- Lewis, D. (1981). Causal decision theory. *Australasian Journal of Philosophy*, 59, 5–30.
- Lewis, D. (1988). Desire as belief. *Mind*, 97, 323–332.
- Luce, R. D., & Raiffa, H. (1957/1989). *Games and decisions. Introduction and critical survey*. New York: Dover Publications.
- Mas-Colell, A., Whinston, M. D., & Green, J. R. (1995). *Microeconomic theory*. New York: Oxford Press.
- Oddie, G. (1994). Harmony, purity, truth. *Mind*, 103(412), 451–472.

- Popper, K. (1959). *The logic of scientific discovery*. London: Hutchinson.
- Savage, L. J. (1954). *The foundations of statistics*. New York: Wiley.
- Schervish, M. J., Seidenfeld, T., & Kadane, J. B. (1995). A representation of partially ordered preferences. *The Annals of Statistics*, 23(6), 2168–2217.
- Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica*, 57(3), 571–587.
- Slovic, P., & Tversky, A. (1974). Who accepts Savage's axioms?. *Behavioral Science*, 14, 368–373.
- Stalnaker, R. (1981). Letter to David Lewis. In W. Harper, Stalnaker & G. Pearce & (Eds.), *IFS: Conditionals, belief, decision, chance, and time*. Dordrecht, NL: Reidel.
- Walker, O., & Dietz, S. (2011). A representation result for choice under conscious unawareness. Working Paper, Grantham Research Institute, London, UK.
- Walley, P. (1991). *Statistical reasoning with imprecise probabilities*. London: Chapman and Hall.
- Weintraub, R. (2007). Desire as belief, Lewis notwithstanding. *Logic and Analysis*, 67(294), 116–122.