Introduction: Making Decisions

R. Bradley

May 13, 2015

Abstract

1 Uncertainty

Uncertainty is a pervasive feature of human life and one of the fundamental parameters of decision making. Arguably indeed choice only has significance in the context of uncertainty about what will or could happen and what is worth doing. Yet it is was suprisingly late in human history, the 17th century, that a theory of uncertainty, and a way of measuring it, emerged in the work of Pascal and Fermat in the form of probability theory.

Yet philosophy has largely treated uncertainty as just lack of certainty—apparently the real interest—which is a bit like examining time through the concept of immortality. And for all the important work that has been done in other disciplines on decision making in conditions of uncertainty, very little attention has been paid to the nature of uncertainty itself.

[One of central achievements of the 20th century is the discovery fo probability as a measure of uncertainty the development on techniques for managing it.]

Standard presentations of decision theory recognises just two forms of uncertainty: objective uncertainty or *risk*, which derives from indeterminacy in the world, and subjective or *epistemic uncertainty*, which derives from lack of information. For each there is a dominant theory of how such uncertainty should be quantified (in both cases, by a probability measure) and of how, so quantified, it should weigh in the evaluation of actions. For situations of risk, decision theorists look to the version of expected utility theory originally due to John von Neumann and Oscar Morgenstern. For those characterised by subjective uncertainty they look to subjective expected utility theory, whose classic statement is to be found in the work of Leonard Savage. Finally, if the situation is one combining both risk and uncertainty, and states can be factored into elements of known objective probability and the rest, then the richer theory of Aumann and Anscombe, which combines the former two, can be employed.

The distinction between risk and epistemic uncertainty, useful though it may be, does not remotely do justice to the variety of forms and degrees of uncertainty. Firstly, our uncertainty concerns not just what is the case (empirical uncertainty), but also what should be the case (evaluative uncertainty), what could be the case (modal uncertainty) and what would be the case if we were to make an intervention of some kind (option uncertainty). Much of this section will be devoted to exploring this different forms of uncertainty and the question of how they should be measured.

Secondly, epistemic uncertainty comes in degrees of severity deriving from the quantity and quality of information that we hold. There is significant difference, for instance, between being unsure about when someone will arrive because one lacks precise information about their time of departure, traffic conditions, the route they have taken, and so on, and having absolutely no idea when they will arrive because you don't know when or whether they have left, whether they are walking or driving or indeed whether they even intend to come. In the former case, the information one holds is such as to make it possible to assign reasonable probabilities to the person arriving within various time intervals. In the latter, one has no basis at all for assigning probabilities, a situation of radical uncertainty or *ignorance*. It may be rare for us to be totally ignorant, but situations of partial ignorance, or *ambiguity*, in which the decision maker is unable to assign determinate probabilities to all relevant contingencies, are both common and important.

More generally, according to some critics, Bayesian theory fails to distinguish between the different levels of confidence we might have, or have reason to have, in our probability judgements. Compare a situation in which we are presented with a coin about which we know nothing and one in which we are allowed to conduct lengthy trials with it. In both situations we might ascribe probability one-half to it landing heads on the next toss: in the first case for reasons of symmetry, in the second because the frequency of heads in the trials was roughly 50%. It seems reasonable however to say that our probability ascriptions are more reliable in the second case than the first and hence that we should feel more confident in them. We will return to these issues in the second half of the book.

1.1 Evaluative Uncertainty

The distinctions between certainty, risk and uncertainty are standardly used only to characterise the agent's state of knowledge of the world. But it is equally important to distinguish cases in which consequences have known, or given, objective values and those in which these values are either unknown and the decision maker must rely on subjective evaluations of them, or do not exist and the decision maker must construct them. The possibility of evaluative uncertainty is typically ignored by decision theorists, because of their (often unconscious) attachment to Ethical Subjectivism, the aforementioned view that values are determined by the agent's subjective preferences, in the sense that what makes a consequence valuable or otherwise (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. If this view were correct, talk of evaluative uncertainty would be misleading as one is not normally uncertain about what one's own judgement on something is (just 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 to attach to a particular consequence 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 it is, how comfortable and so on. This is uncertainty of the ordinary kind and, if one wishes, 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 outcome of 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 whether these properties are valuable or as to how valuable they are. One may know all the specifications, technical or otherwise, 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.

When we assume that values are given, we take this uncertainty to have been resolved in some way. This could be because we assume that there is a fact of the matter as to how good a consequence is or as to whether one outcome is better than another, a fact that would detailed by the true axiology. But it could also be because the description of the decision problem itself comes with values 'built-in'. For instance, in a problem involving a decision between two courses of medical treatment, it may be that a limited number of value considerations apply in the assessment of these treatments: number of patients saved, amount of discomfort caused, and so on. The decision theorist will be expected in such circumstances to apply only the relevant values to the assessment of the options, and to set aside any other considerations that he or she might 'subjectively' consider to be of importance. A large number of applications of expected utility theory take place in this sort of environment, when the issue of what values to be apply have been settled by prior public policy debate.

In many situations, however, values are not given in any of these ways and the agent may be uncertain as to the value she should attach to the relevant prospects. She may, for this reason, also be willing revise to revise her evaluations in the face of new considerations or the criticism of others. These facts would seem to render Ethical Subjectivism unsustainable. But this is too quick. The Ethical Subjectivist will insist that an agent cannot be wrong about what value to attach to fully specified consequences, but will point out that it suffices that there be factual uncertainty for us to be unsure about the desirability of any less than fully specified prospect, including actions. Since in practice complete specification is impossible, this means that evaluative uncertainty of the kind that derives from factual uncertainty will be ubiquitous. Furthermore, since we may hold false beliefs our ethical judgements on incompletely specified prospects are criticisable. All that is ruled out is pure value uncertainty.

Other views take the possibility of value uncertainty more seriously. There are three that I will discuss here. The first draws on the cognitivist rather than expressivist form of subjectivism, according to which the decision maker's probabilities and utilities are her degrees of belief and preference, understood to be factual properties of her. On this view evaluative uncertainty is just ordinary uncertainty about the agent's tastes or more generally about the features of agents that are relevant to the utility of the option. The thought is this. What value an agent will assign to a commodity depends not just on features of the commodity itself (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 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 affected will judge it because we are uncertain about their tastes, needs and other features determining their judgements.[1] 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.[2]

This view allows for uncertainty about what we and others in fact prefer, uncertainty that derives not from our values but from the subjective features of the agent upon which the application of these values depends. On this view, for instance, there is some fact of the matter as to whether one prefers speed to safety and to what extent, but it takes some experimentation to work out what this is. This position has some plausibility in cases such as the car one, where preferences are largely based on fairly static tastes. 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 I doubt that in such cases there really is anything like a set of pre-given preferences waiting to be discovered. An intermediate case would be trying to decide whether to take up the violin or fencing. Can the problem be described as trying to work

[1] See Hansson (ref).

[2] See for instance Sugden (reference?)

out what one's tastes are? I think not. 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.

On a second cognitivist view, what I am calling value uncertainty is just ordinary uncertainty about normative facts. The uncertainty I experience about whether to help my friend 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 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 this 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. David Lewis [**?**][**?**] takes this as an example of what he calls the 'Desire-as-Belief thesis' and shows 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 belief about what is good or best). Lewis' argument has attracted a good deal of controversy with some authors (e.g. Oddie **??** and Weintraub [**?**]) seeing it as a refutation of cognitivism and others (in particular Broome [**?**]) arguing that it establishes no more than the non-existence of the kind of goodness propositions that Lewis postulates. If Broome is right, as I believe he is, then other ways of spelling out a cognitivist interpretation of evaluative uncertainty may viable.

Both these cognitivist views treat value 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 of the views I want to consider holds that evaluative 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 probabilities judgements about normative facts. Making this 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 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 uncertainty is removed and the value 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.

Broome [**?**] showed that the Rationality Hypothesis does not rule out the possibility of objective values. The converse is true as well: the objectivity of values doesn't imply ethical cognitivism. Or to put it slightly differently the expressivist view sketched above is neutral on the question of whether there are normative facts or not. Suppose I can give an exhaustive specification of all the possible facts concerning both the ways things are and how good it is or would be that they are so. Such a specification determines my judgements in the sense that, conditional on these being the facts, I must judge what is the case and what would be good in accordance with the specification. But does such a specification determine its own probability (unconditionally)? Of course not. That judgement as to which specification of the facts corresponds to the actual state of affairs is not determined by those facts alone. It is something that I can be uncertain about, an uncertainty that I encode in my probability assignment to these specifications. What if we also exhaustively specify the probabilistic facts regarding what is the case? In this case a specification does determine the probability of the basic facts, but it still does not determine the probability of the probabilistic facts. Of these I can remain uncertain, an uncertainty encoded in a probability assignment to possibe combinations of probabilistic facts. Similarly, although the evaluative facts, as specified, determine how desirable the non-evaluative facts are, they don't determine their own desirability. Of these I can remain uncertain, an uncertainty that I encode in my utility assignment to each exhaustive specification of the facts.

1.2 Modal Uncertainty

Empirical uncertainty arises when we are unsure as to what *is* the case; evaluative uncertainty when we are unsure as to what *should* be the case. The third fundamental kind of uncertainty—modal uncertainty— arises when we are unsure as to what is possible or about what *could* be the case: what contingencies might arise, what consequences might follow from our actions and what actions are feasible. In the Savage framework, no modal uncertainty can arise as both the state space and the set of possible consequences are exogenously given. In real decision problems, however, a state space may be hard to construct and agents must grapple with the possibility of *unforeseen contingencies*: eventualities that even the educated decision maker will fail to take account of. For instance, the 2011 tsunami and subsequent nuclear meltdown in Japan were 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 [**?**] call 'conscious unawareness'—then they face modal 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 [**?**] 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, $\mathcal{S} = \{S_1, ..., S_n\}$ 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 very 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 [**?**], 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." [**?**, 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 [**?**] and [**?**]). 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 credited with suitability to grand world decision 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, Lipman and Rustichini [**?**] 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.

1.3 Option uncertainty

In the state-consequence representation of a decision problem that we started with, actions were associated with definite consequences, one for each state of the world. 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 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. Whether we describe the problem in these terms–as uncertainty about what options we have–or as uncertainty about the consequences, in each state of the world, of exercising any of the options we know we have, is of little substance, and I shall use the same term–option uncertainty–to denote both.

Option uncertainty arises when we are unsure about what would happen if we were to act in some way, or perform some kind of intervention in, or manipulation of, our environment. When an agent faces option uncertainty, she cannot represent her decision problem with the simple state-consequence matrix represented by 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 1 in which each act γ is represented as a function from each state S_i to a set of associated possible consequences $\{c_1^1, ..., c_1^j\}$. The larger the sets of possible consequences the greater the option uncertainty facing the agent.

There are three strategies that can be pursued in handling option uncertainty. The first is to try to reduce or transform option uncertainty into empirical uncertainty about the state of the world. The second is to reduce it to

	States of the world		
Options	S_1	...	S_n
α	$\{a_1^1, ..., a_1^j\}$	...	$\{a_n^1, ..., a_n^j\}$
...	...	...	...
γ	$\{c_1^1, ..., c_1^j\}$	...	$\{c_n^1, ..., c_n^j\}$

Table 1: State-Consequence Correspondence

evaluative uncertainty. And the third is to treat it as a *sui generis* form of uncertainty. Let's consider each turn.

Reduction to Empirical Uncertainty: Decision theorists 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 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. 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 me dry.

To get around these difficulties the reductionist can make a two-pronged attack. To handle objective indeterminacy, she can allow consequences to be objective probability distributions (lotteries) over outcomes, and apply von Neumann and Morgestern's [**?**] theory to give a measure of their utility. And to handle enumerability-of-states problems, she can draw on 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, she 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 2, in which each $S_j(A^i, C^i_j)$ denotes the state that maps action A^i to consequence C^i_j.

This decriptive strategy has some notable advocates. Lewis [**?**], for instance, treats states as 'dependency hypotheses', these being maximally specific propo-

	States of the world		
Options	$S_1(A_i, C_{1j})$	...	$S_n(A_i, C_{mj})$
A_1	C_{11}	...	C_{1n}
...	...	...	...
A_m	C_{m1}	...	C_{mn}

Table 2: State Functions

sitions about how consequences depend causally on acts. Similarly, Stalnaker [**?**] suggests 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 way 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 dependency hypothesis or particular conjunction of conditionals.

One might well wonder whether this isn't just sweeping the problem under the carpet. For how is the agent to assign probabilities to states so described? What probabilities should be attached to conditionals and conjunctions of them? These are questions to which we will return later on. But it is important not to confuse logical and epistemological worries. Talk of state-functions and/or counterfactuals helps deal with the problem of how to represent option uncertainty, even if it doesn't solve the problem of what beliefs to adopt.

Reduction to Evaluative Uncertainty: A second 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" [**?**, p. 84]. Formally this amounts to treating the sets of possible consequences associated with an action occuring in Table 1 as single coarse-grained consequences and giving it a utility value. Pursuit of this strategy converts option uncertainty, not into ordinary uncertainty about the state of the world, but into uncertainty about the desirability of the consequence as described. We may be sure that the act of taking an umbrella will have the consequence in a rainy state of being able to shield ourselves 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. Whatever utility value we assign to the coarse-grained consequence of having an umbrella as rain-protection will embody this uncertainty.

	Consequences		
Options	c_1	...	c_n
α	$\Pr_\alpha(c_1), U_\alpha(c_1)$	...	$\Pr_\alpha(c_n), U_\alpha(c_n)$
...	...	...	...
γ	$\Pr_\gamma(c_1), U_\gamma(c_1)$	...	$\Pr_\gamma(c_n), U_\gamma(c_n)$

Table 3: Act-Dependent Matrix

Non-Reduction: The last strategy to consider is to accept the presence of option uncertainty and try and develop both a measure of it. The form of Bayesian decision theory espoused by Richard Jeffrey [**?**] offers a way of doing this. Jeffrey dispenses with the formal distinction between states and consequences entirely and assigns a probability to consequences that depends on the action performed. So instead of trying to enumerate the features of the state of the world that will ensure that I stay dry if I take an umbrella, I simply assess the probability that I will stay dry if I take the umbrella and the probability that I will get wet anyhow (even if I take it). In making these probability judgements, I may well try and conceive of the various contingencies under which staying dry will be a consequence of my action, but I need not be able to conceive of all of them in order to do so. Having done so I may directly represent the decision problem in terms of the probabilities and utilities of the various possible consequences induced by each option. This yields a new kind of probability-utility matrix as in Table 3 in which $\Pr_\gamma(C)$ and $\Pr_\gamma(C)$ are the probability and desirability of consequence C induced by action γ.

Whereas the other strategies led to alternatives to our initial state-consequence matrix representation of a decision problem, this last one offers an alternative to initial quantitative representation of a decision problem. It differs in another important respect: it is no longer assumed that it possible to identify states that are probabilistically independent of the available actions.

In recent years the debate between evidential and causal decision theorists has raged as to the nature of the act-dependent quantities. Evidentialists such as Richard Jeffrey [**?**] regard the conditional probability, given that an action is performed, as giving the correct measure of the uncertainty associated with acting (i.e. $\Pr_\gamma(C) = \Pr(C|\gamma)$) , while causal decision theorists such as James Joyce [**?**] argue that what is required is a measure of the probability under the counterfactual supposition that the action is performed. If evidentialism is correct then a single probability function suffices 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 suggests 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.