RISK: PHILOSOPHICAL PERSPECTIVES

Edited by Tim Lewens

Routledge
Taylor & Francis Group
LONDON AND NEW YORK

7

ACTING UNDER RISK

D.H. Mellor

Ends and means

Although many actions, such as playing or listening to music, may be done just for their own sake, most are not. Most actions, like taking exercise to keep fit, are done partly or mainly as a means to an end, the end being an intended effect of the action. Many actions, indeed, are only done as means to ends, as when we pay money to buy goods. In these cases agents need to decide whether (they think) the end justifies the means, e.g. whether (they think) the goods are worth the price; and the job of decision theorists is to say on what grounds such decisions are or should be made. And so they do; but the grounds that most of them give are wrong, I shall argue, because they are *subjective*. Merely *thinking* that an end justifies a means, or fails to justify it, is not enough to justify our adopting, or declining to adopt, that means to that end. But to say this is easier than to say what more it takes to justify making such decisions, and what we should do when, as is often the case, we know too little to justify them. Those are the questions I set out to answer in this chapter.

Deciding whether to adopt an undesirable means to a desirable end, such as having painful dentistry to cure a toothache, is often easy enough in practice, even when, because 'the spirit is willing but the flesh is weak', mental or physical incapacity makes it hard or impossible to act on the decision. Here, however, we can mostly ignore complications raised by weakness of the will and other mental or bodily limitations. For present purposes, most questions of how we do or should act may be treated as questions of how we do or should *decide* to act.

Yet easy though it often is in practice to decide whether an end justifies a given means, these decisions are less easy to understand in theory. The first theoretical problem is to understand how to balance the utility of an end against the disutility of the means it needs to justify. This seems easy to understand in some cases, as when we pay money for goods, but not in all. How, for example, can we decide if the positive utility of keeping fit outweighs the negative utility of taking the necessary exercise? The usual answer is that these utilities are utilities *for us*, the agents, and can therefore be defined, or at

least constrained, by what we actually decide to do. Thus deciding to exercise to keep fit shows that, for us at the time, the utility of keeping fit exceeds the utility of not exercising, while deciding not to exercise shows the opposite.

This inference does not make decision theory vacuous; it merely follows the common practice of measuring causes by their effects, as when we infer temperatures from thermometer readings. What saves that inference from vacuity is the assumption that things have real properties (temperatures) whose effects include the thermometer readings we use to measure them. Similarly, what stops inferences from decisions to utilities being vacuous is the substantial assumption that we have desires whose degrees (utilities) affect the actual or hypothetical decisions we can then use to measure those utilities.

If decisions are theoretically problematic when we know that some available means certainly *will* achieve our ends, they are even more so when we don't, as when we gamble. Or, to take a few other obvious examples: we know that medicines don't always work; that giving up smoking doesn't always stop us getting cancer; that trains, planes and cars sometimes crash; and so on. The theoretical problem such cases pose is not that decision theory cannot cope with them: on the contrary, its ability to do so is what makes it practically useful and conceptually interesting. The problem is how to read the *probabilities* that the theory uses in these cases to say when we should adopt such fallible means. That is the problem whose orthodox solution I think is wrong and to which I shall now propose an alternative.

First, however, I need to say how and why I think the problem arises, given that probabilities do not figure in *all* principles that tell us how to decide whether to use means that we don't think are effective. For example, they figure in neither the *dominance* principle (Jeffrey 1983: ch. 1.5) nor the *maximin* principle (Luce *et al.* 1957: 278): principles which, if they always gave us credible answers, would enable us to dispense with probabilities altogether. But they don't, as the classic case of smoking and cancer shows.

Suppose I am wondering whether to quit smoking as a means M to the end E of escaping cancer. Here and hereafter I shall take M and E to be *states of affairs*, like my quitting smoking and escaping cancer, where E is the end and M the supposed means. Being states of affairs, M and E correspond to propositions, 'M' and 'E', that are true if and only if M and E 'obtain' or 'come about'. This lets me use the negation symbol '~' to represent the states of affairs (my going on smoking, and getting cancer) that will, in the circumstances, come about if M and E do not, as ~M and ~E. And in the simple two-option case, which is all we need consider, the four possible combinations of these states of affairs – M&E, M&~E, ~M&E, ~M&~E – are all we need take into account.

Suppose, then, that, despite knowing that ~E may come about (I may get cancer) whether or not I 'do M', i.e. bring M about (by quitting smoking), I don't want my decision, to do M or not to do it, to depend on E's relative probabilities given M, $p(E|M)$, and given ~M, $p(E|\sim M)$. I may resist this dependency for various reasons – I may not know what $p(E|M)$ and $p(E|\sim M)$

are, I may not believe in such probabilities, or I may deny that they apply to individuals – it doesn't matter. What matters is that, for whatever reason, I want a principle that will tell me what to do given only the four relevant utilities, which are of the following states of affairs:

(a) smoking and avoiding cancer, ~M~E;
(b) quitting smoking and avoiding cancer, M~E;
(c) smoking and getting cancer, and ~M&~E; and
(d) quitting smoking and getting cancer, M&~E.

Then since I know that, while I would rather smoke whether I get cancer or not, getting cancer will, for me, be even worse than being deprived of tobacco, I prefer (a) to (b) to (c) to (d). In symbols, if the relevant utilities U(. . .) are governed by these preferences, then

(1) U(~M&E) > U(M&E) > U(~M&~E) > U(M&~E),

where ' > ' means 'greater than'. What, then, given (1), do the dominance and maximin principles tell me to do?

The simplest way to show what these principles say in cases of this simple kind is to display the relevant utilities in the matrix shown in Table 7.1.

The *dominance* principle, then, comes in two versions. The weak one says that one decision *dominates* if and only if, whatever the outcome (E or ~E), it has a greater utility than every alternative; the strong one only requires the dominant decision's utility to be no less than any alternative, and greater than at least one. In either version the dominant decision, if any, is the one the principle prescribes. In my case, therefore, both versions of the principle tell me to *go on smoking* (~M) because, since I will prefer to smoke whether I get cancer or not, going on dominates quitting. In general and in symbols, (1) implies that ~M dominates M because, in each column of Table 7.1, it makes the utility in the ~M row exceed the utility in the M row.

The *maximin* principle tells us to choose the course of action with the least bad worst outcome, i.e. to maximise the minimum utility (hence the name). In my case, therefore, maximin also tells me to *go on smoking* (~M) because, given (1), the lowest utility in the ~M row of Table 7.1, U(~M&~E), is greater than the lowest utility in the M row, U(M&~E). So in this case maximin agrees with dominance, as indeed it will whenever both principles apply. For

Table 7.1

	E	~E
M	U(M&E)	U(M&~E)
~M	U(~M&E)	U(~M&~E)

one decision can only have a greater utility than every alternative *whatever* the outcome if, in particular, its *minimum* utility exceeds that of every alternative.

Why, then, if these two principles, neither of which invokes probabilities, always agree in their prescriptions, should our decisions depend on probabilities at all? There are two main reasons. One is that neither dominance nor maximin applies to all credible combinations of utilities. If, for example, I know I'll prefer to smoke if I get cancer, but will prefer not to if I don't, so that

$$U(\sim M\&E) < U(M\&E) \text{ and } U(M\&\sim E) > U(\sim M\&\sim E),$$

then neither M nor ~M will dominate, and dominance will not apply. While if I know that cancer would leave me not caring whether I smoked or not, so that

$$U(\sim M\&\sim E) = U(M\&\sim E),$$

then the lowest utilities in the M and ~M rows of Table 7.1 will be equal, and maximin will also not tell me what to do.

The other, more serious, reason for rejecting dominance and maximin, even when they do apply, is their inability to distinguish relevantly different situations. Thus, in the smoking case, given (1), dominance tells me to go on smoking, and maximin tells me to quit, whatever the probabilities, $p(E|M)$ and $p(E|\sim M)$, of my avoiding cancer (E) if I quit (M) and if I don't. This implies that it neither would nor should affect my decision if I learned either that

$$p(E|M) = p(E|\sim M),$$

i.e. that I have the same probability of escaping cancer whether I quit smoking or not, or that

$$p(E|M) > p(E|\sim M),$$

i.e. that quitting smoking actually reduces my probability of escaping cancer. Yet it is clearly irrational to do something as a means to an end when we know that the end will be at least as probable if we don't do it. That is why the decision theories of Jeffrey (1983) and others take account of the relevant probabilities, which in our simple case they do as follows.

First, they define the so-called *expected* utilities of the supposed means M and its alternative. ~M. M's expected utility EU(M) is the sum of the utilities of M's two possible outcomes, M&E and M&~E, weighted by the probabilities, $p(E|M)$ and $p(\sim E|M)$, that they *are* the outcomes, where

$$(2a)\ p(\sim E|M) = 1 - p(E|M),$$

(2b) $p(\sim E|\sim M) = 1–p(E|\sim M)$;

and similarly for ~M. That is:

(3a) $EU(M) = U(M\&E).p(E|M) + U(M\&\sim E).p(\sim E|M)$;

(3b) $EU(\sim M) = U(\sim M\&E).p(E|\sim M) + U(\sim M\&\sim E).p(\sim E|\sim M)$.

Then the '*maximise expected utility*' principle, or MEUP for short, tells us to do M if M's expected utility EU(M) exceeds ~M's, to do ~M if EU(~M) exceeds EU(M), and tells us nothing if the two are equal.[1]

MEUP does not, however, imply that what makes M a means to E is that EU(M) > EU(~M). I have remarked already that an end E can *fail* to justify a means M if its utility fails to outweigh M's disutility, i.e. – according to MEUP – if EU(M) does *not* exceed EU(~M). That can happen because whether E justifies M depends on the relevant utilities, which M's being a means to E does not. Yet we can still use MEUP to tell us what makes M a means to E, as the following argument, taken from Mellor (1995: Ch. 7.4), shows.

First, if M's being a means to E is independent of the utilities of M&E, M&~E, ~M&E and ~M&~E, we can take these utilities to be whatever they need to be in order to let MEUP tell us what makes M a means to E. The utilities that do this are those that make M what I shall call a *pure* means to the end E. By this I mean that it makes no difference to E's utility, or to ~E's, whether we do M or ~M, so that

(4a) $U(M\&E) = U(\sim M\&E) = U(E)$ and

(4b) $U(M\&\sim E) = U(\sim M\&\sim E) = U(\sim E)$,

and, by definition, since E and not ~E is the end, U(E) > U(~E). Thus if, for example, quitting smoking *is* a means of escaping cancer, then it will be a *pure* means to that end if I know that, whether or not I escape cancer, I won't care whether I smoke or whether I don't.

The point of supposing M to be a pure means to E is that, by definition, it will then make no difference to E's utility whether M is brought about or not. This makes M's value to me purely instrumental, which means that MEUP should tell me to do M only if M *is* a means to E. In other words, MEUP should imply that M *is a means to E only if, were it to be a* pure *means to E, EU(M) would exceed EU(~M)*.

To see how this condition requires M and E to be related, we modify (3a) by replacing U(M&E) and U(~M&E) with U(E) – from (4a) – and $p(\sim E|M)$ with $1–p(E|M)$ – from (2a)–and similarly for (3b). This gives us

(5a) EU(M) = U(E).p(E|M) + U(~E).(1–p(E|M)) and

(5b) EU(~M) = U(E).p(E|~M) + U(~E).(1–p(E|~M));

which in turn, given that U(E) > U(~E), entails that

(6) EU(M) > EU(~M) if and only if p(E|M) > p(E|~M),

and hence that, according to MEUP,

(7) M is a means to E only if p(E|M) > p(E|~M).

In other words, MEUP implies the independently plausible condition that *M is a means to E only if E will be more probable if M is done than if it isn't.*

Risk and uncertainty

Note that the 'if and only if' in (6) has, in (7), been weakened to 'only if'. This is because at least three other conditions, besides M's raising E's probability, must be met if M is to be a means to E.

First, for M to be worth calling a means to E, it must raise E's probability by more than an infinitesimal amount. How much more is a moot but secondary question, like asking how hot is hot, or how much hair a man must lose before he becomes bald: questions whose answers raise Sorites questions, and depend on context in various ways, that for present purposes we may set aside.

Second, doing M must be *feasible*, since nothing can be a useful means to an end E that is harder to bring about than E itself. If, for example, we could simply *will* our recovery from illness as easily as we can take medicine, there would be no point in taking medicine as a pure means to recovery even if the expected utility of doing so exceeded that of not doing so.

These two conditions may be vague, and perhaps debatable, but they are certainly less contentious that the third, which is the one I shall now argue for, namely that *M is a means to E only if it raises E's* objective *probability.*

What makes this condition contentious is the unorthodox reading of decision theory it requires, a reading I must therefore now explain and defend. To do so I start with the distinction between *risk* and *uncertainty* drawn by Luce and Raiffa (1957: 13)in their classic introduction to the theory of games and decision-making:

> As to the certainty–risk–uncertainty classification, let us suppose that a choice must be made between two actions. We shall say that we are in the realm of decision making under:

> (a) *Certainty* if each action is known to lead invariably to a specific outcome . . .
> Risk if each action leads to one of a set of possible specific outcomes, each outcome occurring with a known probability. The probabilities are assumed to be known to the decision maker. For example, an action might lead to this risky outcome: a reward of $10 if a 'fair' coin comes up heads, and a loss of $5 if it comes up tails. Of course, certainty is a degenerate case of risk where the probabilities are 0 and 1.
> (c) *Uncertainty* if either action or both has as its consequence a set of possible specific outcomes, but where the probabilities of these outcomes are completely unknown or are not even meaningful.

In this quotation 'probability' clearly means a probability, e.g. of a coin toss landing heads, or of my escaping cancer, that is both objective and empirical. In what follows I shall, for brevity, call such probabilities 'chances' and write them '*ch*(. . .)' but without, however, committing myself to any particular analysis of them.

Thus, for example, if E is my escaping cancer, *ch*(E) is my objective chance of doing so, whether this chance be identified with the relative frequency with which people relevantly like me (in age, sex, diet, etc.) escape cancer, or with a propensity to escape it, i.e. with a probabilistic disposition to escape cancer that I share with other such people (Mellor 2005a: Chs 3–4).

But however they are analysed, objective chances must be distinguished from probabilities of at least one other kind, i.e. from a different application of the mathematical calculus of probability (Mellor 2005a: Chs 1, 5–6). These are the so-called *subjective* probabilities, which for short I shall call *credences* and write '*cr*(. . .)', that measure our degrees of belief in possible states of affairs, e.g. that a coin toss will land heads or that I will escape cancer.

The difference between my credence *cr*(E) that I will escape cancer and the chance *ch*(E) of my doing so is, therefore, that while my *cr*(E) measures how strongly I *believe* I will escape cancer, *ch*(E) measures how likely I am to escape it *in fact*, whatever I or anyone else believes. This means that my *cr*(E) may well differ from *ch*(E), perhaps being lower if I am a confirmed hypochondriac, or higher if I smoke but am, like all too many smokers, convinced that the statistics on smoking and cancer don't apply to me.

I take this difference between chances and credences to correspond to Luce and Raiffa's distinction between risk and uncertainty. That is, I take the *risk* of a means M failing to bring about an end E to be the *chance* *ch*(~E|M) that it will fail to do so. *Acting under risk* I shall therefore take to be basing a decision, on whether to do M as a means to E, on the known value of *ch*(~E|M) and hence – since *ch*(~E|M) = 1–*ch*(E|M) – of *ch*(E|M), the chance that M *will* bring E about. In other words, my decision will be made using what I shall call the *objective* reading of MEUP, on which (3a) becomes

$$EU(M) = U(M\&E).ch(E|M) + U(M\&\sim E).(1–ch(E|M)),$$

and similarly for (3b).

Just as I take the chance $ch(\sim E|M)$ to measure the risk of M's failing to bring E about, so I take my credence $cr(\sim E|M)$ to measure my degree of belief that M will fail to bring E about, and thus how *uncertain* I am that it *will* bring E about. *Acting under uncertainty* I shall therefore take to be basing a decision, on whether to do M as a means to E, on $cr(\sim E|M)$ and hence on $cr(E|M)$, my credence that M will bring E about. In other words, my decision will be made using what I shall call the *subjective* reading of MEUP, or 'subjective MEUP', on which (3a) becomes

$$EU(M) = U(M\&E).cr(E|M) + U(M\&\sim E).(1–cr(E|M)),$$

and similarly for (3b).

In short, I take Luce and Raiffa's two concepts, of acting under risk and under uncertainty, to make MEUP provide two quite different decision theories, theories whose prescriptions may well differ when, and because, our credences differ from the corresponding chances. When, they, do so differ, which should we follow?

Description and prescription

The question just posed has a deceptively simple two-part answer. The first part says that when there is no such thing as the objective chance $ch(E|M)$, or there is but we don't know what it is, we cannot use an objective MEUP and must therefore use a subjective one. The second part says that, when $ch(E|M)$ does exist and we do know what it is, then since our credence $cr(E|M)$ should generally equal it (Lewis 1980), the subjective and objective theories will generally prescribe the same decision anyway. So either way the question of which theory to choose need not arise: since we always can, and often must, use a subjective MEUP, which is therefore the one to follow.

The error here lies in the assumption, usually taken for granted, that a subjective decision theory gives us a good enough reason to do what it tells us to do. To see the error, we must now distinguish two readings of MEUP, whether they be objective or subjective: a *prescriptive* reading and a *descriptive* one. Read prescriptively, MEUP says that we *should* do M if EU(M) exceeds EU(~M), and *should* do ~M if EU(~M) exceeds EU(M). (If EU(M) equals EU(~M), it says nothing either way.) Read descriptively, MEUP says that we *will* follow these prescriptions if EU(M) and EU(~M) differ (and it says nothing if they don't).

This distinction gives us four possible MEUPs: a prescriptive one, using either chances or credences; and a descriptive one, again using either chances or credences. Which of the four should we use?

The answer is that we can use two of them: a prescriptive MEUP to tell us what we *should* do, and a descriptive one to tell us what determines what we

will do. The latter will be a causal theory, i.e. one that tells us, rightly or wrongly, how relevant utilities and probabilities cause us to (decide to) act as we do. And this at once rules out a descriptive reading of an objective MEUP. For the direct causes of, say, my decision to quit smoking, which are what a descriptive MEUP offers to tell us, can hardly include the objective *chance* *ch*(E|M) of my escaping cancer if I quit. That chance may indeed be an indirect cause of my quitting, if my knowing its value causes me to have the corresponding credence *cr*(E|M). But that is only because the chance affects the credence which, together with my subjective utilities, is what directly causes me to quit. And similarly in all other cases: chances can only affect our decisions indirectly, via their effects our knowledge of them has on our credences. So it is the latter, not the former, that must figure in a descriptive MEUP; which is why that theory can only be subjective.

But how, then, can a subjective MEUP be prescriptive? How can a theory that claims to tell us what we *will* do also tell us what we *should* do? The answer, all parties agree, is that it can't; which is why most advocates of a subjective MEUP (generally called Bayesians, because they trace their views back to the work of Thomas Bayes (1763)), deny that it tells us what we will do:

> . . . the subjectivistic theory of probability is not an empirical psychological theory of degrees of belief . . . the object [of tests of it] is not to find out if the theory accurately describes the behaviour of people, but to find out whether people are rational according to the prescriptions of the theory.
>
> (Kyburg and Smokler 1964: 6)

That is the orthodox Bayesian view: a subjective MEUP is prescriptive, and only tells us what we will do *if we are rational*, which it is taken for granted that we should be. I think this is wrong, and that a prescriptive reading of MEUP can only be justified if it is objective. A merely subjective MEUP I follow Ramsey (1926) in taking to be a descriptive theory of the causes of our decisions. That claim I have argued elsewhere (Mellor 2005b); here I make the case for requiring a prescriptive MEUP to be objective. But as part of that case is the prescriptive inadequacy of a subjective MEUP, I must restate briefly why I think it is inadequate.

I start with a presupposition shared by both sides to the dispute. This is that doing M as a means to E may be justified even if it is neither certain, nor thought to be certain, that M will in fact bring E about. The greater probability, subjective or objective, of my escaping cancer if I quit smoking may justify my quitting even if I still get cancer in the end. But that is just bad luck: it does not show that my decision to quit was wrong. But what, then, if my decision *was* right, made it so: how can doing M's raising E's probability make it right to do M? We all agree that doing M will be right if it raises E's probability enough

to make M's expected utility exceed ~M's. What we disagree about is the *kind* of probability of E that doing M must raise if M is to be the right thing to do. Bayesianism implies that M need only raise the agent's *credence* in E; I say it must raise E's *chance*.[2]

To see why I say this, suppose I think that, whether I escape cancer or not, the utility for me of smoking will always exceed that of not smoking. Suppose too that, being 'in denial' about the relevant statistics, my credence that I will escape cancer (E) will be the same whether I quit smoking (M) or not, so that my $cr(E|{\sim}M) = cr(E|M)$. Then, however much E's utilities for me with and without M exceed ~E's, a subjective MEUP will always tell me to go on smoking.

To me and many others this prescription seems absurd, given the evidence for an objective if indeterministic causal link between smoking and cancer. (See the references to 'Fisher's problem' in Jeffrey 1983: 15, 25.) This is why I think that MEUP can only tell me whether to quit smoking if it is *objective*, i.e. if, for given utilities, EU(M) and EU(~M) are fixed not by my credences, $cr(E|M)$ and $cr(E|{\sim}M)$, but by the objective chances, $ch(E|M)$ and $ch(E|{\sim}M)$, of my escaping cancer if I quit smoking and if I don't. All a subjective MEUP can do is tell us what would make me *think* that quitting smoking is the right thing for me to do, whether it is or not.

Another way of putting the same point is to ask what makes my quitting smoking, M, a *means* to my end E of escaping cancer. I argued in Sect. 1 that a prescriptive MEUP will credit M with being a means to E only if it would tell us to do M when

(4a) $U(M\&E) = U({\sim}M\&E) = U(E)$ and

(4b) $U(M\&{\sim}E) = U({\sim}M\&{\sim}E) = U({\sim}E)$,

i.e. when, if M *is* a means to E, it is a *pure* means. This test, as we saw, makes a prescriptive MEUP say that if E is an end, i.e. if $U(E) > U({\sim}E)$, then whatever the relevant utilities,

(7) M is a means to E only if $p(E|M) > p(E|{\sim}M)$.

So here again the question is: should the probabilities in (7) be objective or subjective; should (7) require doing M to raise E's *chance*, or my *credence* in E? And again the answer seems obvious: my quitting smoking is only a means of escaping cancer if it raises my *chance* of escaping cancer, not merely my *credence* that I will escape it. My $cr(E|M)$ exceeding my $cr(E|{\sim}M)$ may be what makes me *think* that M is a means to E. But thinking that will only make it so if raising my credence in E also raises E's chance – as indeed the so-called 'placebo effect' in medicine shows that it may.

The placebo effect is the benefit referred to in the medical definition of a placebo as 'a chemically inert substance given instead of a drug. Benefit may be gained from a placebo because the person taking it believes it will have a positive effect.' (British Medical Association 2002: 451). This definition implies that when the placebo effect occurs, i.e. when taking a placebo *is* in fact a means to the medical end E, what makes it so is the fact that, by raising a patient's credence in E, it thereby raises E's chance. Only by accepting that this is what a means to E must do that we can explain how placebos that work (i.e. produce the placebo effect) differ both from placebos that don't work and from 'non-placebo' medicines, i.e. medicines that need not raise the patient's credence in E in order to raise E's chance. (This is not to deny that patients may only take non-placebo medicines if they believe that doing so *will* raise their credence in E, merely that this effect on patients' credences has to be what makes the medicines work.)

In short, put in terms of the distinction between risk and uncertainty drawn in the second section of this chapter, M can only be a means to E if doing M reduces the *risk* of E's not coming about. It is not enough for doing M to reduce our *uncertainty* that E will come about, thereby making us *think* that M is a means to E. That thought cannot be self-verifying, as the contingency of the placebo effect shows. Yet that is what a prescriptive MUEP that is merely subjective must take it to be, by insisting that in reality uncertainty is all there to risk, and therefore that no one who thinks that M is – or is not – a means to E can be mistaken.

Subjective prescriptions

Bayesians may reply in three ways to the arguments of the third section of this chapter. First, those who, like de Finetti (1937), deny the very existence of objective chances, may therefore deny that in reality there *is* any more to risk than uncertainty. The objection to this is that the probabilities that all sciences, from microphysics to epidemiology, take to explain observed frequencies, can only do so if they are as objective as the frequencies they explain, i.e. are chances (Mellor 2005a: Ch. 10). My high credence in non-smokers escaping cancer may explain why I *expect* most of them to escape it: it cannot explain why most of them *do* escape it.

Second, most Bayesians who do admit chances will also admit that, as we noted on page 00, the values of $ch(E|M)$ and $ch(E|{\sim}M)$, when we know them, should usually also be the values of our $cr(E|M)$ and $cr(E|{\sim}M)$, and hence the values we should use in deciding whether to do M. Thus if all I know about a coin toss is that its chance of landing heads is p, then p is the credence I should use in deciding whether to accept a bet at any given odds that it *will* land heads. And this principle, which I shall here follow Lewis (1980: 87) in calling the 'principal principle', arguably exhausts chance's prescriptive role in MEUP. This lets Bayesians claim that a prescriptive MEUP, since it need not use

chances, but only credences derived from them by the principal principle, can itself be purely subjective. That claim, however, begs the question against an objective MEUP, since the case for the principal principle presupposes, rather than proves, that we should base our decisions on known chances when we can: since that is *why* our credences – which, like our subjective utilities, are *measured* by the decisions they cause us to make (Ramsey 1926) – should equal those chances in the first place.

The third and strongest Bayesian reply is to ask how we can apply an objective MEUP when the relevant chances are unknown. It seems obvious that we can't, which is why, in the quotation on page 00, Luce and Raiffa limit 'the realm of decisions under . . . risk' to situations where the relevant objective 'probabilities are assumed to be known to the decision maker'. If so, then when we must make decisions without knowing the relevant chances, as we often must, why should we not follow the prescriptions of a subjective MEUP, if only *faute de mieux*? How else can MEUP tell us how to act under uncertainty?

An objectivist response to this rhetorical question needs to make several points. First, it can and should insist that, as a descriptive theory of how our credences and subjective utilities make us act, a subjective MEUP is, if flawed, still

> a useful approximation to the truth particularly in the case of our self-conscious or professional life, and it is presupposed in a great deal of our thought. It is a simple theory and one which many psychologists would like to preserve by introducing unconscious desires and unconscious opinions in order to bring it more into harmony with the facts.
>
> (Ramsey 1926: 69)

This is what enables a subjective MEUP to explain many of our decisions, by giving our reasons for making them: as when it says that my reasons for quitting smoking are that I want to escape cancer and believe that quitting smoking is a means to that end.

However, the fact that a decision theory gives mental reasons (desires and beliefs) for our decisions makes some physicalists and others, who think all causes are physical, insist that since the theory cannot explain our decisions causally, it can only *rationalise* them, i.e. say how our beliefs and desires make them *rational*. Hence the temptation to think that since subjective MEUP tells us 'what we will do *if we are rational*, which . . . we should be' (page 00), it tells us not how we will but how we should act under uncertainty.

The temptation should be resisted. For however our reasons for a decision are related to its physical causes (see e.g. Crane *et al.* 1990, Davidson 1963: Ch. 1), they can only make the decision rational if *they* are rational, and they may not be. For example, as I noted on page 00, it would be as irrational on present evidence for my quitting smoking not to raise my credence that I will escape cancer as it would for me positively to desire cancer. Whatever such credences

and subjective utilities made me decide to do, they could not make that decision rational, and a prescriptive MEUP should not, as a subjective MEUP does, imply that they could.

Bayesians may reply to this by invoking an analogy between a rational decision and a deductively valid inference, i.e. one whose conclusion follows logically from its premises, and so cannot be false if they are true. For since an inference's validity does not depend on its premises' truth, false conclusions may well be validly inferred from false premises, as when, thinking on Sunday that today is Monday, I infer validly but falsely that today is a weekday. Yet the falsity of this conclusion is clearly no reflection on the inference itself, whose validity makes it equally rational whether or not its premises are true.

Similarly with inductive inferences (e.g. from 'X is copper' to 'X conducts electricity') that depend on laws of nature, which cannot be known *a priori*, to make their conclusions true when their premises are. Our knowledge of the laws that make these inferences truth-preserving may then make them as rational as deductive inferences, even when their conclusions ('This conducts electricity') are false because their premises ('This is copper') are.

And as with these deductive or inductive 'theoretical' inferences, so with the 'practical inferences', from credences and subjective utilities (the 'premises') to decisions (the 'conclusions'), that a subjective MEUP provides. Here too the inference may be rational even when the credences in its premises are so far from the relevant chances that its conclusion (e.g. to go on smoking) is objectively wrong, i.e. differs from that of an objective MEUP. Hence, Bayesians may say, the prescriptive force of a subjective MEUP when the relevant chances are unknown. Using it to derive decisions from credences and subjective utilities that may or may not match their objective counterparts is just as rational as using truth-preserving theoretical inferences to draw conclusions from beliefs that may or may not be true.

Some of this argument we may accept. No objectivist need deny that always using a subjective MEUP to make decisions is as rational as always using truth-preserving inferences to derive new beliefs from old, and for the same reason. For just as the latter is the only way to ensure that *if* the premises are true, they will yield true conclusions, so the former, objectivists will say, is the only way to ensure that, *if* our credences and subjective utilities match their objective counterparts, they will yield objectively right decisions. Moreover, if the conditions stated by these if-clauses are *not* met, no alternative principles of theoretical or practical inference can be relied on to do any better, if only because no principles of inference can then be relied on to do any good at all. So the weak dominance principle mentioned in the first section of this chapter, if nothing else, tells us to use the present principles of inference always, whether or not we know that their if-clauses are true. All objectivists can therefore admit that requiring practical inferences to conform to a subjective MEUP is as rational as requiring theoretical inferences to be truth-preserving.

So far so good for a prescriptive reading of a subjective MEUP. But that is as far as it goes, and it is not far enough for Bayesians. The trouble is that an inference's rationality, so understood, is not inherited by its conclusion when Bayesians need it to be. That is because, as we have seen, the merit of this kind of rationality is *conditional*: the condition in a theoretical inference being that its premises be true; and in a practical one that its credences and subjective utilities yield the same decision as their objective counterparts. In both cases, therefore, an inference's rationality only makes it rational to accept its conclusion when these conditions are met. When any of a theoretical inference's premises are false, the fact that the inference is truth-preserving is no reason to believe its conclusion. Similarly, objectivists will say, with decisions derived from our credences and subjective utilities. When those credences differ from the corresponding chances, the fact that the decision is prescribed by a subjective MEUP is no reason to think that it is right, i.e. that an objective MEUP would prescribe it.

This distinction, between the rationality of inferences and that of their conclusions, is unfortunately obscured in English by our calling the latter 'inferences' too, as when we call an inference false, meaning that its conclusion is. The ambiguity is especially unfortunate here because it encourages the false belief that an inference's being rational is enough to make its conclusion rational, which we have just seen that it isn't. That error is then reinforced by another ambiguity in claiming that inferences (and hence their conclusions) are justified *given their premises*. For this to be true, 'given the premises' needs to mean 'given that the premises are *justified*', whereas all it usually means is 'given that these *are* the premises', which is not enough to justify the conclusion of any inference, theoretical or practical. My believing that quitting smoking will lower rather than raise my chance of escaping cancer may partly *explain* my decision to remain a smoker; but given that this belief is neither true nor justified, it cannot *justify* that decision. That is why a merely subjective MEUP cannot tell us what we should do; only an objective MEUP can do that.

Acting under uncertainty

What, then, can tell us how to act under uncertainty as opposed to risk, i.e. when we don't know what the relevant chances are? Assuming with Ramsey that a subjective MEUP, read descriptively, is 'a useful approximation to the truth', it can tell us what we *will* do in that situation; but that will not tell us what we *should* do. Yet nor, it seems, can an objective MEUP tell us this when we don't know the chances it invokes to tell us what decisions to make. How, then, can any probabilistic decision theory tell us how to act under uncertainty?

To answer that question we must first ask if there *is* a right way of deciding whether an end E justifies a prospective means M when we have no idea what

difference, if any, doing M will make to E's chance of coming about. We certainly want there to be a right way, since we quite often have to make decisions in this situation, decisions that may be serious, as when people who know no first aid must decide how to keep someone going after a sudden heart attack (say) until they can get medical advice. How can they tell what to do, when even doing nothing is only one of the potentially fatal (but also potentially life-saving) options open to us? The answer, I'm afraid, is that they can't, if they really know *nothing* about which option gives the patient his or her best chance of survival. (This is why we should all learn something about first aid.) For here, if not in the Big Bang, *ex nihilo nihil fit*: nothing comes of nothing. When we know none of the relevant empirical facts, theories of decision-making can no more tell us how to act than theories of inference can tell us what to believe; in neither case can any amount of rationality compensate for empirical ignorance. In epistemology, in short, there is no such thing as a fact-free lunch, which is why it is no objection to an objective MEUP that it fails to offer one.

An objective MEUP can, however, offer a useful *low*-fact diet, if not a no-fact one. For it rarely if ever requires us to know precise values of either utilities or chances for it to be able to tell us how to act. All it requires us to know is that their values lie in intervals, which are often quite long, within which all values will prescribe the same decision; and that much we often do know, as two examples will serve to show.

First, suppose you offer to bet me at 2:1 that a coin toss won't land heads. That is, you offer to pay me twice as much if it does land heads as I agree to pay you if it doesn't. Suppose too that, being neither averse nor attracted to the bet for its own sake, I will take the bet, if I do, purely as a means to winning it. Then if the ratio of the amounts to be won or lost measures the relative utilities for me of those outcomes, all I need to know, for an objective MEUP to tell me to take the bet, is that the coin toss's chance of landing heads is greater than 1/3, a fact I can easily know without knowing its precise value.

Now take the more serious case of smoking and cancer. Nearly all cancer patients will agree that, even if treatment stops cancer shortening their lives, the pain and anxiety it causes them and their families makes getting it more than twice as bad as quitting smoking, whatever many as-yet healthy smokers think. This, and the overwhelming statistical evidence that quitting reduces smokers' chances of getting cancer by over 50 per cent, is all an objective MEUP needs in order to advise them to quit smoking, however imprecise their knowledge of the relevant chances and utilities may be.

(Many smokers will, of course, fail to heed this advice, for one of two reasons. First, they may reject it by denying that, for them, quitting smoking will reduce their chance of getting cancer, or that getting it will be so much worse for them than quitting; in short, a smoker's relevant credences and/or subjective utilities may differ so much from the corresponding chances and objective utilities that they decide to go on smoking – a decision whose conformity to a

subjective MEUP makes neither right nor rational. And second, smokers who apparently decide to quit may still not do so, because they can't: their addiction stops them. These, however, I take to be cases where 'the spirit is willing but the flesh is weak' and hence covered by the proposal in the first section to 'ignore complications raised by weakness of the will and other mental or bodily limitations'. We may therefore credit these addicted smokers with decisions, if not actions, that conform to both subjective and objective MEUPs.)

So far so good, but the 'reliabilist' view of knowledge (see, e.g., Armstrong 1973: Chs 11–13; Goldman 1986) that naturally accompanies an objective MEUP lets the latter do even better. For on a reliabilist view, what justifies our beliefs is our acquiring them by processes (e.g. those of our senses) that have a high enough chance of giving us true beliefs. (How high is high enough will vary with context, and doesn't matter for present purposes, provided it can be less than 1.) So in particular, if we acquire beliefs in *chances* by processes that give those beliefs a high enough chance of being true, then on this view they will be justified, even if they are false. And that is how our rules of inference from observed frequencies (e.g. of coin tosses landing heads, or of smokers getting cancer) to chances work: by giving their conclusions a high chance of being true.[3]

Moreover, as the conclusions these inferences justify are always that a given chance lies in a given interval, they are, as we have just seen, precisely what an objective MEUP needs. And the wider the intervals the chances can lie in for this MEUP to prescribe a decision, the less frequency evidence it takes to give those intervals a high enough chance of containing those chances. Thus, in the coin tossing example above, I need far less evidence to give the chance of heads a high chance of being greater than 1/3 than of being greater than 2/3. The evidence may still, of course, be less than I need; and then not even an objective MEUP will justify whatever decision I make. But, then, nor, I have argued, will anything else; in particular, a subjective MEUP will not justify it.

My coin tossing bet shows too how reliabilism enables an objective MEUP to let decisions be justified which it also says are wrong. For suppose, in the example, that I accept your bet because a freak run of heads makes me infer, justifiably but falsely, that the coin's chance of landing heads when tossed is greater than 1/3. An objective MEUP can now admit that my decision, although made wrong by the falsity of my belief that the coin toss's chance of landing heads is greater than 1/3, was nevertheless justified by the statistical inference from the frequency data that gave this false belief of mine a high chance of being true. Similarly in the smoking and cancer case. An improbably low frequency of cancer in a population of smokers might justify their false belief, that quitting would not raise their chance of escaping cancer, and hence their objectively wrong decision to go on smoking.

That, in brief, is how an objective MEUP can prescribe and/or justify many, although of course not all, of the decisions we make under uncertainty about the relevant chances. But what, then, of the decisions under uncertainty that it

cannot justify, let alone prescribe? What, for example, can an objective MEUP tell the people above who know too little first aid to justify any decision about how to keep a heart attack victim alive until they can get medical advice? I say it can tell them nothing, for the reasons I have given, except, as I have also said, that anyone who is likely to have to cope unaided with such situations *should* learn enough first aid beforehand. And similarly for all decisions made in situations of such uncertainty: all an objective MEUP can tell us is not to let these situations arise. Yet that is perhaps its most valuable prescription: to learn enough about the risks we are likely to encounter to enable us, if and when we encounter them, to make decisions that are objectively right, or at least objectively justified.[4]

Notes

1 Note that I do not mean 'doing M' to imply that it takes *action*, as opposed to inaction, to bring about M rather than ~M. It does not, since MEUP ignores the action–inaction distinction, which often does not apply or does not matter, as indeed it does not in the case of quitting or going on smoking. Where it does matter, as in the distinction between killing and letting die, MEUP can allow for it by letting it affect the relevant utilities.

2 In Mellor (2005b) I argue that a prescriptive MEUP also needs objective utilities, which may differ from what we *think* they are, i.e. from subjective utilities. But as this essay is chiefly about the kinds of probability, chances and credences that distinguish risk from uncertainty, I shall not repeat that argument here. Readers need only bear in mind that by 'utilities' in what follows I shall, unless I say otherwise, mean *objective* utilities, however, if at all, these differ from subjective ones. In particular, I shall now use the phrase 'utilities for us' introduced in the first section to mean utilities which, despite being objective (since they may differ from what we think they are), may be different for different people.

3 This statement is too simple as it stands; for a more detailed and accurate statement of how reliabilism can justify beliefs about chances, see Mellor (2005a: Ch. 8.III). The statement is also made contentious by its appeal to reliabilism and the existence of chances; but that is no objection to it as a statement of how reliabilism can assist an objective MEUP which takes chances for granted.

4 I am indebted to the editor Tim Lewens and to Wlodek Rabinowicz for invaluable comments on and corrections to earlier drafts of this essay.

References

Armstrong, D.M. (1973) *Belief, Truth and Knowledge*, Cambridge: Cambridge University Press.

Bayes, T. (1763) 'An essay towards solving a problem in the doctrine of chances', in Swinburne, R.G. (ed.) *Bayes's Theorem*, Oxford: Oxford University Press.

British Medical Association (2002) *Illustrated Medical Dictionary*, London: Dorling Kindersley.

Crane, T. and Mellor, D.H. (1990) 'There is no question of physicalism', *Mind*, 99: 185–206.

Davidson, D. (1963) 'Actions, reasons, and causes', *Journal of Philosophy*, 60: 685–99.

de Finetti, B. (1937) 'Foresight: its logical laws, its subjective sources', in Kyburg, H.E. Jr and Smokler, H.E. (eds) *Studies in Subjective Probability*, New York: Wiley.
Goldman, A.I. (1986) *Epistemology and Cognition*, Cambridge, MA: Harvard University Press.
Jeffrey, R.C. (1983) *The Logic of Decision*, Chicago: University of Chicago Press.
Kyburg, H.E. Jr and Smokler, H.E. (eds) (1964) *Studies in Subjective Probability*, New York: Wiley.
Lewis, D.K. (1980) 'A subjectivist's guide to objective chance', in his *Philosophical Papers Volume II*, Oxford: Oxford University Press.
Luce, R.D. and Raiffa, H. (1957) *Games and Decisions*, New York: Wiley.
Mellor, D.H. (1995) *The Facts of Causation*, London: Routledge.
——(2005a) Probability: *A Philosophical Introduction*, London: Routledge.
——(2005b) 'What does subjective decision theory tell us?', in Lillehammer, H. and Mellor, D.H. (eds), *Ramsey's Legacy*, Oxford: Oxford University Press.
Ramsey, F.P. (1926) 'Truth and probability', in his *Philosophical Papers*, Cambridge: Cambridge University Press.