

Epistemic Decision Theory

HILARY GREAVES

Faculty of Philosophy, University of Oxford

hilary.greaves@philosophy.ox.ac.uk

I explore the prospects for modelling epistemic rationality (in the probabilist setting) via an epistemic decision theory, in a consequentialist spirit. Previous work has focused on cases in which the truth-values of the propositions in which the agent is selecting credences do not depend, either causally or merely evidentially, on the agent's choice of credences. Relaxing that restriction leads to a proliferation of puzzle cases and theories to deal with them, including epistemic analogues of evidential and causal decision theory, and of the Newcomb Problem and 'Psychopath Button' Problem. A variant of causal epistemic decision theory deals well with most cases. However, there is a recalcitrant class of problem cases for which no epistemic decision theory seems able to match our intuitive judgements of epistemic rationality. This lends both precision and credence to the view that there is a fundamental mismatch between epistemic consequentialism and the intuitive notion of epistemic rationality; the implications for understanding the latter are briefly discussed.

1. Beyond pure observation: Some puzzle cases

In most epistemological situations, the agent is a *pure observer*, in the following two senses. (1) What she believes does not causally influence the truth of the propositions that her beliefs are about. (2) While one generally hopes that the agent is more likely to believe that P if P is true than if P is false, still the fact that S believes that P on the basis of evidence E is not itself *additional* evidence in favour of, or against, P.

Interesting epistemological puzzle cases arise when the agent is *not* merely an observer: when the truth of the proposition believed *does* depend, in some stronger way than in the usual cases, on the extent to which the agent believes either that very proposition, or some other proposition. Here are a few such cases.

Promotion:

Alice is up for promotion. Her boss, however, is a deeply insecure type: he is more likely to promote Alice if she comes across as lacking in confidence. Furthermore, Alice is useless at play-acting, so she will come across that way iff she really does have a low degree of belief

that she's going to get the promotion. Specifically, the chance¹ of her getting the promotion will be $(1 - x)$, where x is whatever degree of belief she chooses to have in the proposition P that she will be promoted. What credence in P is it *epistemically* rational for Alice to have?

Presumably, in the Promotion case, there is a unique rationally permitted degree of belief in P : Alice must adopt credence $\frac{1}{2}$ in P , because only in this case will her credences match her beliefs about the chances once she has updated on the proposition that she will adopt that very credence in P . (Here and in the following, we assume that the agent is aware of the specification of his or her case.)

Leap:

Bob stands on the brink of a chasm, summoning up the courage to try and leap across it. Confidence helps him in such situations: specifically, for any value of x between 0 and 1, if Bob attempted to leap across the chasm while having degree of belief x that he would succeed, his chance of success would then be x . What credence in success is it epistemically rational for Bob to have?

One feels pulled in two directions. On the one hand: adopting an extremal credence (0 or 1) will lead to a perfect match between one's credence and the truth, whereas a non-extremal credence will lead to only imperfect match. But on the other: whatever credence one adopts (extremal or otherwise), one's credences will match the chances: they will be the right credences to have given the then-chances. Is any degree of belief in success epistemically rationally permissible, or only an extremal credence?

Embezzlement:

One of Charlie's colleagues is accused of embezzling funds. Charlie happens to have conclusive evidence that her colleague is guilty. She is to be interviewed by the disciplinary tribunal. But Charlie's colleague has had an opportunity to randomize the content of several

¹ Here and throughout the paper, I employ a relatively permissive notion of chance. On a more restrictive notion it might turn out, for instance, that for every degree of belief Alice might have, whether or not Alice would get the promotion if she entered the interview with that degree of belief was already fixed by the present microphysical state of the universe, hence the relevant chance would be either 0 or 1. Some of my analyses (specifically, some of those based on causal epistemic decision theory) would change if I used this more restrictive notion, but *probably* not in ways that would ultimately be of interest for the purposes of this paper.

otherwise informative files (files, let us say, that the tribunal will want to examine if Charlie gives a damning testimony). Further, in so far as the colleague thinks that Charlie believes him guilty, he will have done so. Specifically, if x is the colleague's prediction for Charlie's degree of belief that he's guilty, then there is a chance x that he has set in motion a process by which each proposition originally in the files is replaced by its own negation if a fair coin lands Heads, and is left unaltered if the coin lands Tails. The colleague is a very reliable predictor of Charlie's doxastic states. After such randomization (if any occurred), Charlie has now read the files; they (now) purport to testify to the truth of n propositions P_1, \dots, P_n . Charlie's credence in each of the propositions P_i , conditional on the proposition that the files have been randomized, is $\frac{1}{2}$; her credence in each P_i conditional on the proposition that the files have not been randomized is 1. What credence is it epistemically rational for Charlie to have in the proposition G that her colleague is guilty and in the propositions P_i that the files purport to testify to the truth of?

One again feels pulled in two directions. On the one hand: Charlie has conclusive evidence that her colleague is guilty, so, presumably, she should retain her degree of belief 1 in the proposition G that he is guilty. But if she does, then her colleague will almost certainly have predicted that she would do this, hence will almost certainly have randomized the files, in which case Charlie should adopt credence $\frac{1}{2}$ in each P_i ; this is the best she can do, but she knows that her degree of belief is then bound to be 'one half away from the truth' for each P_i , as the truth-value can only be 1 or 0. If, on the other hand, Charlie moves to degree of belief 0 in G , then her colleague will almost certainly have predicted *that*, and so he will almost certainly have left the files alone; in that case, by trusting the testimony of the files vis-à-vis the P_i s, Charlie can do very well in terms of getting the truth-values of the P_i s correct. Agents who disrespect their evidence concerning G in cases like this will tend to have more beliefs that are closer to the truth—does epistemic rationality, therefore, recommend joining them?

Arrogance:

Dennis is wondering whether or not he is arrogant. He takes a low (resp. a high) degree of belief that one is arrogant to be evidence for the proposition that one in fact is (resp. is not) arrogant: specifically, his credence in the proposition A that he is arrogant, conditional on the proposition that he will end up with credence x in A , is $1 - x$, for

all $x \in [0,1]$. What credence is it epistemically rational for Dennis to have in the proposition A that he is arrogant?

Again there seems to be a unique rationally permitted credence: Dennis must have credence $\frac{1}{2}$ in A, but for a slightly different reason than in Promotion: only if his final credence in A is $\frac{1}{2}$ can that final credence equal his *conditional initial* credence in A, conditional on the proposition that his final credence takes that particular value.

Imps:

Emily is taking a walk through the Garden of Epistemic Imps. A child plays on the grass in front of her. In a nearby summerhouse are n further children, each of whom may or may not come out to play in a minute. They are able to read Emily's mind, and their algorithm for deciding whether to play outdoors is as follows. If she forms degree of belief o that there is now a child before her, they will come out to play. If she forms degree of belief 1 that there is a child before her, they will roll a fair die, and come out to play iff the outcome is an even number. More generally, the summerhouse children will play with chance $(1 - \frac{1}{2}q(C_o))$, where $q(C_o)$ is the degree of belief Emily adopts in the proposition (C_o) that there is now a child before her. Emily's epistemic decision is the choice of credences in the proposition C_o that there is now a child before her, and, for each $j = 1, \dots, n$, the proposition C_j that the j th summerhouse child will be outdoors in a few minutes' time.

Again one is torn. On the one hand: Emily has conclusive evidence that there is now a child before her, so presumably she should retain her degree of belief 1 in the proposition C_o that indeed there is. In that case, there will be a chance of $\frac{1}{2}$ of each summerhouse child coming out to play, so she should have credence $\frac{1}{2}$ in each C_j ; this is the best she can do, but she knows that her degree of belief is then bound to be 'one half away from the truth' for each C_j , as the truth-value can only be 1 or 0 . On the other hand, if Emily can just persuade herself to ignore her evidence for C_o , and adopt (at the other extreme) credence 0 in C_o , then, by adopting degree of belief 1 in each C_j ($j = 1, \dots, 10$), she can guarantee a perfect match to the remaining truths. Is it epistemically rational to accept this 'epistemic bribe'?

One might have strong intuitions about what epistemic rationality recommends for each of these cases. Straw polls, however, suggest that not all of the intuitions in question are universally shared. And whether they are shared or not, it would be illuminating to have a

more precise theoretical framework within which to reason about such cases. Enter the epistemic consequentialists.

2. Epistemic consequentialism

Epistemic consequentialism is the analogue of prudential or ethical consequentialism in the epistemic domain. Epistemic consequentialists recognise a notion of *epistemic value*, analogous to utility or ethical value: a state of affairs is one of high epistemic value for a given agent just in case it is a state of affairs in which there is a good degree of fit between that agent's beliefs and the truth. Where prudential (respectively, ethical) consequentialists evaluate acts such as carrying umbrellas (resp. lying) for prudential rationality (resp. moral rectitude), the epistemic consequentialist evaluates 'epistemic acts'—acts such as believing or 'accepting' particular propositions, or adopting particular credence functions—for epistemic rationality or irrationality. Such acts count as epistemically rational to the extent to which they do, or could reasonably be expected to, bring about states of high epistemic value. (The issue of whether or not they are under voluntary control is orthogonal to that of their epistemic rationality.)

As in the prudential and ethical domains, there is a compelling independent motivation for epistemic consequentialism (independent, that is, of consequentialism's promise to help with theorizing about problem cases such as those above): it would seem paradoxical for epistemic rationality to forbid an epistemic act that would lead, or would reasonably be expected to lead, to more epistemic good than any permitted act.

In the prudential and ethical domains, consequentialist ideas are nicely precisified by decision theory, in which the notion of good being pursued is captured by a utility function (or value function) that assigns higher numbers to preferred (or better) outcomes. This suggests the project of developing an analogous epistemic decision theory, to precisify epistemic consequentialism. This would be a theory in which an epistemic utility function makes quantitative the extent to which the agent has 'achieved a state of high epistemic value', and according to which, in any given epistemic predicament, the epistemically rational (epistemic) 'act' is the one that 'tends to maximize epistemic utility'.

The scare quotes enclose phrases that we have thus far deliberately left vague: epistemic decision theory comes in several varieties,

according to the way in which each is made precise. There are three main choices to be made:

1. *Identification of the epistemic acts*: Epistemic decision theories (EDTs) have been developed fairly extensively for evaluating acts of *acceptance*. Such theories take the epistemic acts under evaluation to be acts of accepting, rejecting, or suspending judgment about particular propositions, and they have the agent weigh the epistemic risk of accepting falsehoods against the prospect of the epistemic gain of accepting truths. Typically, the agent also has given credences in the propositions in question; the issue of whether or not an agent with given credences should *accept* a given proposition depends on the extent to which her epistemic values recommend taking such risks. (EDTs of this type have been developed by, e.g., Levi (1967) and Maher (1993).)

Whatever one thinks of the notion of ‘acceptance’ that is in play here,² such epistemic decision theories are at best half of the epistemological story, since they presuppose an account of which *credences* it is rational to have. In the present paper, we undertake this prior project: we work in a purely probabilist setting, in which (as in the puzzle cases presented above) there is no notion of acceptance/rejection or of full belief/disbelief in play at all: the epistemic acts are acts of adopting particular *degrees* of belief, or credence functions.

2. *Choice of a ‘thick’ or ‘thin’ notion of epistemic utility*: There are several good-making features of an epistemic state. One would like one’s credences to do a good job of matching the truth, but one would also like them to be informative, to be a rational response to one’s evidence, to obey the Principal Principle, and so on. A ‘thick’ utility function would encode all such epistemic norms directly into the epistemic utility function, so that given, for instance, two agents with identical degrees of belief, one of whom had arrived at those degrees of belief by a long process of empirical investigation and conditionalization on acquired evidence, the other of whom had simply jumped to the degrees of belief in question arbitrarily, the first could count as having higher epistemic utility than the second. We will not take this course: we will employ a ‘thin’ notion of epistemic utility, according to which the epistemic utility of a given agent at a given time depends *only* on (i) the agent’s degrees

² For scepticism about this notion, see e.g. Stalnaker 2002.

of belief at that time, and (ii) the truth of the propositions in which she has credences.

The point of taking this course is that we hope for an *explanatory* theory. Epistemic norms that are encoded directly into the epistemic utility function are *presupposed*; while we must presuppose some norm in order to get the project off the ground, we hope to presuppose relatively little, instead *explaining* such norms as updating by conditionalization, and having informative credences, by showing how they tend to promote the more fundamental norm that is directly encoded.³

3. *Subjective vs objective consequentialism*: As in the ethical domain, there is a subjective and an objective notion of epistemic ‘ought’. Ethically, one objectively-ought to perform the action that will in fact have the best consequences; but since agents typically cannot tell for sure which act this is, we also recognize a subjective ‘ought’. An agent subjectively-ought to perform (roughly) the action that *seems most likely* to lead to the best consequences, given the agent’s beliefs at the time of action: more precisely, an agent subjectively-ought to perform that action that maximizes *expected* utility or value, where the expectation value is taken with respect to the agent’s own credences. In the epistemic case: objectively, there is little to say beyond that agents ought to have credence 1 in all and only true propositions. Our project is to develop *subjective* epistemic consequentialism.

It remains ‘only’ to make the technical details precise: to say something about the mathematical form of the epistemic utility function, about the expression for expected utility, and about the precise sense in which one should ‘maximize expected utility’. These details, however, turn out to contain all the devils: those who have conflicting intuitions about what counts as epistemically rational in one or more of the above cases are suffering conflict over what exactly ‘maximizing expected epistemic utility’ amounts to. The merit of casting the dispute into the framework of epistemic decision theory is that we will then have a precise framework within which to articulate and evaluate the various options. Furthermore, we may thereby avoid reinventing some wheels: lessons already hard-won in the context of practical decision theory can be transferred, *mutatis mutandis*, to the epistemic case. One can, for instance, distinguish ‘evidential’ and ‘causal’

³ Thus our project has much in common with that of Wedgwood (2002).

variants of epistemic decision theory; the Embezzlement case discussed above is fairly straightforwardly an epistemic analogue of the Newcomb problem that teases apart causal and evidential decision theory in the practical case; if one has (as we do) *arguments* for causal decision theory in the practical case, the analogous arguments in the epistemic context may teach us something about epistemic rationality. How far this project of ‘piggy-backing’ on the work of the practical decision theorists can be pushed will depend, of course, on just how similar or different practical and epistemic rationality turn out to be; but that, too, is an interesting question that is nicely made precise by framing the issues in decision-theoretic terms.

One objection must be disposed of before we proceed further. Here are two data. (1) I do not have any reason, epistemic or otherwise, to count the blades of grass on my lawn—I just do not care how many there are, and neither do I have any interest in having true beliefs about that matter. (2) Nevertheless, if I did count them, I would then acquire an epistemic reason to form certain beliefs about how many blades there are—I would have acquired evidence, and epistemic rationality would require me to respect that evidence. Epistemic decision theorists would account for (2) by noting that I have higher epistemic utility if my beliefs about grass conform more closely to the truth. But (1) seems then to show that the fact that doing X would take me to a state of higher epistemic utility fails to imply that I have any reason to do X. Does this show that *I do not, after all, have* the epistemic goal of having true beliefs about the number of blades of grass on my lawn; and does this in turn show that epistemic rationality just is not the sort of goal-seeking activity that epistemic consequentialists in general, and epistemic decision theorists in particular, take it to be? (This line of thought is pressed by Thomas Kelly (2003).)

This objection can be avoided by restricting the domain of applicability of epistemic decision theory: EDT can be a tool only for evaluating the (epistemic) rationality of *epistemic* acts. Forming certain beliefs, *given* that one has acquired certain evidence, is an epistemic act; taking steps to acquire such evidence is a non-epistemic act. What the above line of argument shows is that the theory will deliver false verdicts (such as the verdict that I have an epistemic reason to count blades of grass) if we try to use it to evaluate the rationality of *non-epistemic* acts (such as the act of counting grass-blades). This point is worth noting well, as it means that EDT cannot fulfil one of the ambitions that some have had for it: it cannot be used to give a justification for experimentation (cf. e.g. Oddie 1997). But if it can

disentangle the issues of epistemic rationality involved in puzzle cases such as those above, that is quite motivation enough.

The structure of the paper is as follows. Section 3 sketches an epistemic decision theory analogous to the (practical) decision theory of Savage, and briefly reviews existing work in this setting. The weakness of Savage-style EDT is that it can handle only cases of pure observation: like Savage's original decision theory, it does not give a good treatment of cases in which states depend on acts. The obvious moral is that epistemic decision theory must tread the path that practical decision theory has trodden since Savage in response to this problem: the investigation of evidential decision theory, causal decision theory, and so on; section 4 reviews this path. Section 5 develops evidential EDT, and argues that, of the puzzle cases surveyed in section 1, evidential EDT deals adequately with Promotion and Arrogance, arguably deals adequately with Leap, but deals badly with Embezzlement and Imps. Section 6 develops causal EDT, and argues that it deals well with Promotion and Embezzlement, arguably deals adequately with Leap, but deals badly with Arrogance and Imps. Section 8 reviews the 'deliberational decision theory' that Frank Arntzenius has suggested in response to Andy Egan's counterexamples to causal practical decision theory, develops an epistemic analogue, and argues that deliberational causal decision theory retains the successes of section 7 while also treating the Arrogance problem adequately. Section 9 summarizes, comments on the fact that *none* of our EDTs adequately handles the Imps problem, and assesses the degree of promise of EDT in the light of this fact. The result is inconclusive: the problem presented by the Imps case runs deep and gives a strong reason for doubting EDT, but we also have strong arguments in its favour and, as yet, no better theory.

3. Savage-style EDT for pure observers

Before moving on to more complicated cases, we first develop a simple epistemic decision theory that adequately handles cases of pure observation, to fix ideas and to serve as a starting point for further refinement. Such an EDT can be based on a naive application of the framework of Savage 1972; this is the theory discussed in e.g. Oddie 1997, Greaves and Wallace 2006, Leitgeb and Pettigrew 2010a.

We work with a set S of states of the world. This set is, as in the case of practical decision theory, a set of mutually exclusive and jointly

exhaustive propositions; the agent is uncertain as to which state obtains. Let $\mathcal{P}(S)$ be the set of probability distributions on S . We suppose that at any given time, a given rational agent has some particular credence function,⁴ represented by some particular $Cr \in \mathcal{P}(S)$. Her *epistemic utility*, if she has credence function Cr and the world is in fact in state s , is given by her *epistemic utility function*, $U : S \times \mathcal{P}(S) \rightarrow \mathbb{R}$.⁵ We take an *epistemic act* to be a function from states to credence functions. We assume that epistemically rational agents maximize expected epistemic utility (EEU), where the EEU of an epistemic act $a : S \rightarrow \mathcal{P}(S)$ with respect to credence function p is given by

$$(i) \quad EEU_p(a) = \sum_{s \in S} p(s)U(s, a(s))$$

At first sight, the range of epistemic decision problems appears to be rather limited. Three sorts of problems suggest themselves:

- (i) Suppose you currently have credence function p . You receive no new information. Should you retain your current credence function, or jump to a new one?
- (ii) Suppose you currently have credence function p , and you are about to receive new information from some known set $\{\mathbf{E} = E_1, \dots, E_n\}$ of mutually exclusive and jointly exhaustive propositions. Which newly available epistemic act $a : \{\mathbf{E} \rightarrow \mathcal{P}(S)\}$ should you perform?
- (iii) Suppose you currently have credence function p , and you receive new information in the form of a new constraint that your posterior credence function must meet: for example, something happens to constrain your posterior credences in the elements of some given partition $\{E_1, \dots, E_n\}$ to take

⁴ This particularly idealized assumption can be relaxed without wholesale alteration of the structure of the theory: for instance, one can allow that agents have imprecise credences, represented by a *class* of probability functions, rather than a single probability function.

⁵ Note that at this point, our EDT is structurally *disanalogous* to Savage's PDT. For Savage, the set of consequences (i.e. the domain of the utility function) is another primitive, bearing no special relationship to the sets of states or credence functions; the agent may value anything she likes. For our 'Savage-like' epistemic decision theory, however, we identify the set of consequences with the Cartesian product $S \times \mathcal{P}(S)$. The difference reflects a difference between practical and epistemic rationality: practical utility may depend on anything whatsoever, but we require that 'how well our agent is doing, epistemically' depends only on (i) which state the world is in fact in and (ii) what her credence function is.

some particular values $\{p_1, \dots, p_n\}$. Which full posterior credence function, from those meeting this constraint, should you adopt?

We accept as a constraint on a reasonable epistemic decision theory that the answer to the first question should be ‘stick with p ’, for every $p \in \mathcal{P}(\mathcal{S})$. That is, it must be the case, for every $p \in \mathcal{P}(\mathcal{S})$, that the expected epistemic utility *calculated with respect to p* of having credence function p is higher than the expected epistemic utility *calculated with respect to p* of having any distinct credence function q . In other words, the epistemic utility function must be a ‘proper scoring rule’. Throughout the present paper, we will assume that the epistemic utility functions of the agents we deal with are proper scoring rules.⁶

The second question was analysed in Greaves and Wallace 2006. Greaves and Wallace show that, if the epistemic utility function is indeed a proper scoring rule, then *conditionalization* on the new piece of information received is the unique epistemic act that maximizes expected epistemic utility. In this sense, epistemic decision theory provides a justification for the classical Bayesian rule of updating by conditionalization.

The third question has been examined in Leitgeb and Pettigrew 2010b. Leitgeb and Pettigrew argue that the epistemic decision-theoretic perspective actually *undermines* the usual answer to this question, namely that in such circumstances, one should update by Jeffrey conditionalization, but that it suggests an alternative updating policy which they endorse for such cases.

So far, so good. This simple Savage-style EDT, however, does not perform very well once we go beyond pure observation: Savage EDT is incapable of capturing the subtle and interesting features of the puzzle cases introduced in section 1. Two points are worthy of note:

- (i) *If* the state partition is simply that over which the agent is selecting credences, since the epistemic utility function is a

⁶ Examples of proper scoring rules include:

The logarithmic rule, $U(p, s) = \log p(s)$

The quadratic rule, $U(p, s) = - \sum_{s' \in \mathcal{S}} (\delta_{ss'} - p(s))^2$

The spherical rule, $U(p, s) = \frac{p(s)}{\sqrt{\sum_{s' \in \mathcal{S}} p(s')^2}}$

An example of an improper scoring rule is the otherwise plausible-looking linear rule, $U(s, p) = p(s)$.

proper scoring rule and no new information is to be acquired, the theory will simply recommend retaining one's initial credences, whatever they happen to be. This epistemic decision theory will therefore not capture the sense in which, for example, any credence in P other than $\frac{1}{2}$ in the Promotion case is epistemically deficient.

- (ii) The predictions of 'Savage' EDT depend on the state partition. For example, in the Promotion case, if the states are propositions that specify *how the chance of promotion depends on the agent's choice of epistemic act*, the theory does indeed recover the intuitively correct result that only credence $\frac{1}{2}$ in P is rationally permitted (cf. Sect. 7).

Point 1 is not in itself a decisive objection to the theory: it could be that EDT is true, but is not the *whole* truth about epistemic rationality. In that case, one would expect just this sort of division of labour: perhaps the job of EDT is to tell us what is epistemically rational *given* certain initial credences, but other principles must come into play to explain which initial credences are rationally permitted. Other things being equal, an EDT that delegates less of the work to supplementary principles, and succeeds in explaining a larger part of epistemic rationality itself, is to be preferred; but perhaps we should not be surprised if such an ambitious EDT is simply not available.

Point 2 is more problematic. If Savage EDT says that EEU must be maximized relative to *every* state partition, it is an incoherent theory. The way out of this problem *may* be to supplement the injunction to maximize EEU with a principle identifying the correct state partition; indeed, this is the course taken by causal EDT (Sect. 7). But other ways out are available too; each amounts to the replacement of our naive theory with a different theory.

The *reason* that section 1's puzzle cases are problematic for our naive EDT are familiar: they are all cases in which states depend, in one way or another, on acts, and, as in practical decision theory, a naive application of the Savage formula runs into trouble in cases with this feature. Practical decision theory has moved on in response to this observation: successively more refined theories have been developed, to take account of a wider and wider range of decision problems. We will therefore follow the obvious research strategy: we will recall those developments in practical decision theory, and develop their analogues for our epistemic case.

4. A brief history of practical decision theory

A (very!) brief history of the foundations of practical decision theory is as follows. (For a fuller survey, see, for example, Joyce 1999.)

- (i) We first have Savage's (1972) theory, in naive form: the rational act is that with highest expected utility, where the expectation value is given by the formula

$$(2) \quad EU_{Cr}(a) = \sum_{s \in \mathcal{S}} Cr(s)U(s \& a)$$

- (ii) This theory dealt well with simple cases, but could not in general deal adequately with decision problems in which states depended on acts. For example:

Deterrence Problem: You park your car in a dodgy neighbourhood. A hooligan approaches you, and tells you that he will probably smash your windscreen while you're gone unless you pay him \$10 now; if you do pay, he will probably leave your car alone. The acts are {pay, don't pay}. What should you do?

In this case, Savage is supposed to be committed to a Dominance argument in favour of not paying the \$10, since it is better not to pay than to pay regardless of whether or not the hooligan smashes the window; but this is the wrong answer.

- (iii) Jeffrey (1965) developed an alternative theory, in which acts, states, and consequences are all propositions about which the agent has credences, and the expected utility of an act is calculated with respect to her *conditional* credences in states—conditional on the proposition that she performs the act in question:

$$(3) \quad EU_{Cr}(a) = \sum_{s \in \mathcal{S}} Cr(s|a)U(s \& a)$$

(In the light of the Newcomb Problem, this theory later comes to be called *evidential decision theory*.)

- (iv) Jeffrey's theory gives the right answers in cases like the Deterrence Problem. However, Jeffrey's theory arguably gives the wrong answer in many cases in which causal and

probabilistic dependence come apart—for example, the Newcomb Problem:

Newcomb Problem: There are two boxes, one transparent and one opaque. The transparent box contains \$1,000. The opaque box contains either nothing, or \$1,000,000. Your decision is whether to take only the opaque box, or both boxes.

There is a reliable Predictor of your decisions. The contents of the opaque box depend on what the Predictor predicted your decision would be. Iff she predicted you would take both boxes, she put nothing in the opaque box; iff she predicted you would take only the opaque box, she put \$1,000,000 in that box.

In this case, Jeffrey's theory predicts that you should take only the opaque box. But this, as is *generally* (if not universally) accepted, is the wrong answer.

- (v) The Newcomb Problem motivates the development of *causal decision theory*. In one version of this theory (Lewis 1981), the problems in Savage's theory are fixed by adding a stipulation about which partition of possible worlds may count as a state partition: we require the states to be 'causal dependency hypotheses'.
- (vi) Causal decision theory deals well with all of the cases discussed so far; the mundane cases, the Deterrence Problem, and the Newcomb Problem. However, it seems to give the wrong answer in cases in which there is some act *a* that maximizes causally expected utility with respect to the agent's *initial credences*, but not with respect to the credences that she moves to if she becomes certain that she will perform *a*—for example, Andy Egan's (2007) Psychopath Button problem:

Psychopath Button: Paul has a button. Iff he presses it, all psychopaths will die. Paul's decision is whether or not to press the button. Paul would like to kill all psychopaths as long as he is not one himself, but he strongly prefers not to die. Paul currently has very low credence that he is a psychopath. However, he also has high credence that only a psychopath would press the button.

In this case, causal decision theory predicts that Paul should press the button, if his current credence that he is a psychopath

is sufficiently low. However, Paul is in a position to predict that once he has made this decision and updated on the proposition that he has decided to press, he will regret that decision: with respect to his *updated* credences, not pressing will have a higher causally expected utility than pressing. Vulnerability to this sort of decision instability seems to indicate irrationality.

- (vii) Frank Arntzenius (2008) has recently suggested that a ‘deliberational decision theory’ along lines developed by Brian Skyrms (1990) can solve problems like ‘Psychopath Button’. According to deliberational decision theory, a rational agent should not necessarily perform the act that has highest expected utility according to her *initial* credences; rather, she should allow her credences to develop according to a specified dynamical rule (which rule involves the expected utilities of the acts under consideration), and she should perform the mixed act with probabilities equal to her equilibrium credences. This theory gives intuitively reasonable judgements in all problem cases considered to date.

This progress in the foundations of practical decision theory can serve as a guide for the development of epistemic decision theory. In sections 5–8 we will develop evidential and causal epistemic decision theory and their ‘deliberative’ variants, and we will examine how each of these theories treats our puzzle cases of section 1. To limit our task to a manageable size, here we will consider only decision problems in which the agent is to receive no new information between the time of deliberation and the time at which she adopts her chosen posterior credence function. Thus, we will be considering only epistemic acts that are ‘constant’, in the sense that the agent adopts the same posterior credence function regardless of which state is actual.⁷

5. Evidential epistemic decision theory (evidential EDT)

To go beyond Savage-style EDT, we must distinguish between the ‘target partition’ over which the agent is selecting credences, and the state partition: these may, but need not, be identical.

⁷ To avoid terminological confusion, note that this is distinct from Savage’s sense of ‘constant’: in Savage’s terminology, a ‘constant’ act is one that delivers the same Consequence (hence, the same utility), regardless of which state obtains.

Let \mathcal{C} be the target partition. Then $\mathcal{P}(\mathcal{C})$, the set of probability distributions over \mathcal{C} , is canonically isomorphic to the set \mathcal{A} of (constant) epistemic acts. As in evidential practical decision theory, we will regard acts as propositions; we will write $q \in \mathcal{P}(\mathcal{C})$ for a probability function itself, and $\bar{q} \in \mathcal{A}$ for the proposition that the corresponding act is performed, that is, the proposition that the agent adopts q as her ‘final’ credence function in the decision scenario under consideration. (Thus, for $c \in \mathcal{C}$, ‘ $q(c) = 0.4$ ’ makes sense, but ‘ $\bar{q}(c) = 0.4$ ’ does not; and \bar{q} but not q is a proposition.) The agent has *initial* credences about which act she will perform, as well as concerning which elements of \mathcal{C} obtain, and these may be correlated: thus her initial credences are over the product set $\mathcal{C} \times \mathcal{A}$.⁸

Evidential EDT, like its practical counterpart, assigns a utility (relative to credence function $p \in \mathcal{P}(\mathcal{C} \times \mathcal{A})$) to every proposition. We will assume that epistemic utility is determined by closeness to the truth *concerning \mathcal{C} alone*. This utility is then determined by the following two rules:

- (i) A rule assigning a utility to every proposition of the form $c \& \bar{q}$, where $c \in \mathcal{C}$, $\bar{q} \in \mathcal{A}$. As mentioned above, we will assume that this is a proper scoring rule, in the sense that for fixed $p \in \mathcal{P}(\mathcal{C})$ and variable $q \in \mathcal{P}(\mathcal{C})$, the quantity $\sum_{c \in \mathcal{C}} p(c)U(c \& \bar{q})$ is maximized at $q = p$.
- (ii) The expected utility formula: for any proposition t and any (‘state’) partition \mathcal{S} ,

$$U_p(t) = \sum_{s \in \mathcal{S}} p(s|t)U(s \& t)$$

The evidentially expected utility $EvEEU_{Cr}(\bar{q})$ of an act $\bar{q} \in \mathcal{A}$ with respect to credence function Cr is just its utility: thus we have

$$(4) \quad EvEEU_{Cr}(\bar{q}) = \sum_{s \in \mathcal{S}} Cr(s|\bar{q})U(s \& \bar{q})$$

Evidential EDT then asserts that, in any epistemic decision situation, the agent is rationally required to perform an act that maximizes evidentially expected epistemic utility.

This theory, like Jeffrey’s own, is *partition invariant*: the evidential EEU of a given epistemic act is independent of the choice of state

⁸ It would be nice if the ‘target partition’ could be identified with the domain of the agent’s *initial* credence function, as well as that of the credence function being selected. I do not see how to obtain this feature because $\mathcal{A} \simeq \mathcal{P}(\mathcal{C})$ necessarily has higher cardinality than \mathcal{C} .

partition \mathcal{S} , and hence we can simply choose whichever partition is most convenient, without worrying that we are thereby smuggling in illicit assumptions.

To facilitate definite calculations, we will sometimes assume a *particular* proper scoring rule.⁹ To do this, we select a number of propositions $d \subset \mathcal{C}$ that are of epistemic interest in the decision problem under consideration; let \mathcal{D} be the ‘target set’ of all these propositions. Relative to a choice of \mathcal{D} , we will then impose the simple quadratic scoring rule

$$(5) \quad U(c\&\bar{q}) = - \sum_{d \in \mathcal{D}} (\chi_c(d) - q(d))^2$$

where χ_c is the ‘characteristic function of c ’ (so $\chi_c(d) = 1$ if $c \in d$, and 0 otherwise). Note that $U(c\&\bar{q})$ increases (becomes less negative) as $q(d)$ for a true proposition increases, or as $q(d)$ for a false proposition d decreases. In addition, and unlike the simpler scoring rule $U(c\&\bar{q}) = \sum_{d \in \mathcal{D}} |\chi_c - q(d)|$, (6) is *proper* (cf. Sect. 3, esp. n. 6).

We now wish to apply this theory to the puzzle cases stated in section 1, and see which credence function it recommends that our agent adopt in each of those cases.

Evidential EDT on Promotion: The target partition \mathcal{C} is just $\{P, \neg P\}$, where P is the proposition that Alice will be promoted. We use the same set as state partition \mathcal{S} . The relevant features of Alice’s credence function are exhausted by the constraint that for all possible final credence functions q , $Cr(P|\bar{q}) = 1 - q(P)$. (As in evidential PDT, the agent’s initial credences in her own *acts* are irrelevant to the decision theory.)

We seek the act $\bar{q} \in \mathcal{A}$ that maximizes the value of

$$\begin{aligned} EvEEU_{Cr}(\bar{q}) &= \sum_{s=P, \neg P} Cr(s|\bar{q})U(\bar{q}\&s) \\ &= (1 - q(P))U(\bar{q}\&P) + q(P)U(\bar{q}\&\neg P) \end{aligned}$$

With the quadratic scoring rule (5) (for $\mathcal{D} = \{P\}$), this becomes

$$EvEEU_{Cr}(\bar{q}) = -(1 - q(P))^3 - q(P)^3$$

⁹ This move raises a legitimate concern that the conclusions we draw may be artefacts of our particular choice of scoring rule. If they are, they must be regarded as conclusions concerning (e.g.) evidential EDT *under such-and-such a scoring rule*, rather than blanket conclusions on evidential EDT simpliciter. It would be preferable to proceed in a more abstract way, and we do this when it is straightforward to do so.

which is maximized at $q(P) = \frac{1}{2}$. That is, our evidential EDT delivers the intuitively correct result for our ‘Promotion’ problem case.

Evidential EDT on Leap: The target partition is given by $\mathcal{C} = \{L, \neg L\}$, where L is the proposition that Bob succeeds in leaping across the chasm. We again use the same set as state partition \mathcal{S} . The agent’s initial credence function Cr is such that $\forall q \in \mathcal{P}(\mathcal{C}), Cr(L|\bar{q}) = q(L)$. We seek the $\bar{q} \in \mathcal{A}$ that maximizes the value of

$$\begin{aligned} EvEEUCr(\bar{q}) &= \sum_{s=L, \neg L} Cr(s|\bar{q})U(\bar{q}\&s) \\ &= q(L)U(\bar{q}\&L) + (1 - q(L))U(\bar{q}\&\neg L) \end{aligned}$$

With the quadratic scoring rule (5) (for $\mathcal{D} = \{L\}$), this becomes

$$EvEEUCr(\bar{q}) = q(L)(-(1 - q(L))^2) + (1 - q(L))(-q(L)^2)$$

which is maximized at $q(L) = 0$ or 1 .

Thus, evidential EDT captures the first of the intuitions we canvassed in section 1 in response to this case, namely that only extremal credences are rationally permitted, since these and no others lead with certainty to perfect match between one’s credences and the truth.

Evidential EDT on Embezzlement: The target partition \mathcal{C} is generated by the set $\{G, P_1, \dots, P_n\}$, where G is the proposition that Charlie’s colleague is guilty, and the P_i are the propositions about which the files purport to be informative. That is, \mathcal{C} contains all maximal conjunctions of literals based on the atoms G, P_1, \dots, P_n , up to logical equivalence (so, heuristically, an element of \mathcal{C} can be thought of as an assignment of truth-values to all these propositions; a typical element of \mathcal{C} is $G \& P_1 \& \neg P_2 \& \neg P_3 \& \dots \& P_n$). We use as state partition $\mathcal{S} = \{R, \neg R\}$, where R is the proposition that the files have been randomised. The relevant features of the agent’s initial credence function Cr are given by

$$\begin{aligned} Cr(R|\bar{q}) &= q(G) \\ Cr(G) &= 1 \\ Cr(P_i|R) &= \frac{1}{2}, i = 1, \dots, n \end{aligned}$$

$$Cr(P_i|\neg R) = 1, i = 1, \dots, n$$

We have

$$\begin{aligned} EvEEU_{Cr}(\bar{q}) &= \sum_{s=R, \neg R} Cr(s|\bar{q})U(\bar{q}\&s) \\ &= q(G)U(\bar{q}\&R) + (1 - q(G))U(\bar{q}\&\neg R) \end{aligned}$$

With the quadratic scoring rule (5) (with $\mathcal{D} = \{G, P, \dots, P_n\}$), this becomes

$$\begin{aligned} EvEEU_{Cr}(\bar{q}) &= -(1 - q(G))^2 - q(G) \sum_{i=1}^n \left(q(P_i)^2 - q(P_i) + \frac{1}{2} \right) \\ &\quad - (1 - q(G)) \sum_{i=1}^n (1 - q(P_i))^2 \end{aligned}$$

For $n < 4$, this quantity is maximized at $q(G) = 1, q(P_i) = \frac{1}{2}$ (for $i = 1, \dots, n$; this is unproblematic. However, for $n > 4$, the advantage of perfect accuracy regarding the larger number of propositions P_1, \dots, P_n takes over, and evidentially expected epistemic utility is maximized at $q(G) = 0, q(P_i) = 1 (i = 1, \dots, n)$. The latter is intuitively the *incorrect* answer: intuitively, since Charlie has conclusive evidence that her colleague is guilty, she is at least rationally *permitted*, and arguably is rationally required, to retain credence 1 in G.

Evidential EDT on Arrogance: The target partition \mathcal{C} is $\{A, \neg A\}$, where A is the proposition that Dennis is arrogant. We use the same set as state partition. The relevant feature of the agent’s initial credence function is that for all $\bar{q} \in \mathcal{A}$, $Cr(A|\bar{q}) = 1 - q(A)$.

Evidential EDT is insensitive to the difference between the Arrogance and Promotion cases. It thus recommends adopting credence $\frac{1}{2}$ in A. Again, this is intuitively the correct answer.

Evidential EDT on Imps: The target partition is generated by the set $\{C_0, C_1, \dots, C_n\}$. We use as state partition $\mathcal{S} = \{D, \neg D\}$, where D is the proposition that the summerhouse children will roll their die. The relevant features of Emily’s initial credence function are given by

$$\begin{aligned} Cr(D|\bar{q}) &= q(C_0) \\ Cr(C_0) &= 1 \end{aligned}$$

$$Cr(C_i|D) = \frac{1}{2} \text{ for } i = 1, \dots, n$$

$$Cr(C_i|\neg D) = 1 \text{ for } i = 1, \dots, n$$

Evidential EDT is insensitive to the difference between this case and Embezzlement. It thus recommends adopting a credence function q with $q(C_0) = 0$, and $q(C_j) = 1$ for $j = 1, \dots, n$. As in the Embezzlement case, this is intuitively the *incorrect* answer.

6. Interlude: On the status of the ‘initial’ credence functions

The astute reader may by now be wondering about the precise sense in which my agents’ ‘initial’ credence functions are ‘initial’.¹⁰ An interpretational problem here arises from the ‘simplifying’ policy decision taken in this paper, namely to consider cases that do not involve the acquisition of any new evidence. In cases with that feature, for any time at which the ‘initial’ credence function might otherwise be possessed, an *ideally* rational agent (who, presumably, deliberates at infinite speed) would already have decided which ‘final’ credence function to move to, and done so. In contrast to ‘new evidence’ cases, in ‘no new evidence’ cases there is therefore no time at which the agent actually holds her so-called ‘initial’ credences. What, then, is the status of the ‘initial’ credences that play an essential role in the epistemic decision theories discussed in this paper? To put the question another way: What distinguishes ‘the’ initial credence function from any other member of the space of ‘possible’ initial credence functions, given that it is not the feature of corresponding to the agent’s actual credences at some specified time?

There are four possible ways of dealing with this issue. The first involves stipulating that we are *not* theorizing about ideally rational agents: we are theorizing about agents who are ideally rational in all respects *except* that they deliberate at finite speed, so that there really is some time at which the ‘initial’ credences are actually held. This stipulation would indeed render the status of the credences in question unmysterious, but since it has no independent motivation, it would be somewhat *ad hoc*.

The second possible response is to hold that indeed there cannot be anything distinguishing one ‘initial’ credence function from another *in cases that do not involve the acquisition of new evidence*. On this view,

¹⁰ I am grateful to Ralph Wedgwood for pressing this question.

epistemic decision theory is near-impotent in such no-new-evidence cases; it issues the synchronic requirement that one's 'final' credence function maximizes expected epistemic utility relative to itself, but that is an extremely weak constraint. It does not follow, though, that epistemic decision theory is impotent *simpliciter*: it follows only that, to see the theory properly at work, we need to consider cases involving the acquisition of new evidence. This line of thought is very natural, in particular, in the light of the experience of Savage-style EDT: we saw above (Sect. 3) that most of the interesting results in the Savage framework concern such diachronic norms as classical and Jeffrey conditionalization, and arise only in 'new evidence' cases. The suggestion would therefore be to replace the cases discussed in this paper with 'new evidence' ones.

This suggestion, however, seems unsatisfactory. Puzzle cases like Embezzlement *do* appear to be epistemic analogues of the Newcomb Problem; if epistemic decision theory has anything going for it at all, it should be able to deliver a genuinely decision-theoretic treatment of these cases.

According to the third possible response, we are again theorizing about only one credence function; thus, on this proposal the terminology of 'initial' and 'final' credence functions is again misleading. The proposal is: the role of the 'initial credence function' is to specify *some but not all* of the one and only actual ('final') credence function; the role of the epistemic decision theory is then to specify what values the remaining degrees of freedom in the one and only credence function ought to have, *given* that the 'initial' degrees of freedom have the stated values.

This account works fairly well for the Promotion, Leap, and Arrogance cases: in each of *those* cases, the point of the 'initial' credence function is to fix certain *conditional* credences, leaving open the values of the corresponding unconditional credences, and the epistemic decision theory is then wheeled in to determine the latter. But it does not *seem* to work for the Embezzlement or Imps case. As we saw above, there the epistemic decision theory recommends adopting a final credence function that directly contradicts the 'initial' one—since it disagrees with the 'initial' credence function on the credence assigned to, respectively, the propositions G or C_o —in which case, clearly the 'initial credence function' is not a partial specification of the 'final' one. (See, however, Sect. 9.)

We will, therefore, opt for a fourth response. On this fourth account, initial and final credence functions are (or at least can be) genuinely

distinct credence functions, but there is no time at which the ‘initial’ credence function is actually possessed. While the final credence function is characterized (i.e. singled out from the myriad other mathematically possible credence functions) by correspondence to the agent’s actual credences, the ‘initial’ credence function is characterized instead by the stipulation that it ‘respects the specification of the case’. In other words, the agent’s awareness of certain facts—the facts given in the case-specification—gives rise to rationality constraints on her credences; the ‘initial’ credence function is one respecting these rationality constraints. For example, in the Promotion case, Alice’s awareness that the chance of P is equal to one minus her final credence in P gives rise to rationality constraints on her conditional credences $Cr(P|\bar{q})$. (The constraints, of course, need not single out a *unique* credence function; thus ‘the’ initial credence function will typically be a large class of credence functions, agreeing (however) on all features that are relevant to the evaluation of ‘final’ credence functions.)

In making this response, we must note well that it amounts to wheeling in a constraint of epistemic rationality that is not grounded in epistemic decision theory, in order to get the epistemic decision theory itself off the ground; thus our epistemic decision theory turns out to be less than fully ‘ambitious’ in the sense discussed at the end of section 3.

7. Causal epistemic decision theory

Returning to the main thread: recall (from Sect. 5) that evidential EDT does not give the right answers to all of our puzzle cases. Its treatment of Promotion and Arrogance is satisfactory. Its treatment of Leap may be satisfactory. Its treatment of Embezzlement and Imps, however, are not: since, in those cases, you have conclusive evidence that your colleague is guilty of embezzlement (respectively, that there is a child in front of you), epistemic rationality *requires* you to retain credence 1 in this proposition, however your beliefs about other propositions may fall.

Let us focus first on the Embezzlement problem.¹¹ On reflection, it is perhaps unsurprising that evidential EDT gives the wrong answer

¹¹ One may be sceptical of this tactic’s chances of success: the fix we propose here to obtain a satisfactory epistemic-decision-theoretic treatment of Embezzlement predictably will not work for the Imps case, and, since our existing epistemic decision theory seems to fail on these two cases for the same reason, one might well expect that the true epistemic decision theory will fix both problems in the same way as one another. This scepticism is healthy. We will return to it in Sect. 9; we shelve it in the meantime only for clarity of exposition.

here: this is an epistemic analogue of the Newcomb Problem,¹² to which evidential *practical* decision theory (arguably) gives the wrong answer.

This observation suggests a solution to the problem. Those who agree that evidential PDT gets the Newcomb Problem wrong usually opt instead for causal (practical) decision theory. Our natural next move, then, is to develop *causal epistemic decision theory* (causal EDT), hoping that this theory will retain the successes of evidential EDT while improving on its failures.

We formulate causal EDT as follows. As in section 5, the target partition \mathcal{C} contains the propositions in which the agent is choosing credences; $\mathcal{A} \simeq \mathcal{P}(\mathcal{C})$ is then the set of acts.

The next move is the identification of a set \mathcal{S}_C of causal states. Here we follow Lewis (1981), who formulates causal (practical) decision theory by retaining the Savage formula (2) for expected utility—that is, expected utility is to be calculated with respect to *unconditional* credences in states, rather than the conditional credences used by evidential decision theory. (Unconditional credences in acts, and conditional credences $Cr(state|act)$, play no role in causal decision theory.) Lewis avoids the incoherence problem by adding a constraint on which partitions of logical space may serve as the state partition \mathcal{S} . Specifically, Lewis's prescription is that states must be causal states, or 'causal dependency hypotheses': propositions that specify how every outcome that the agent cares about depends causally on her choice of act. Sometimes a choice of act causally determines a chance distribution over outcomes, but not a unique outcome; thus we take the causal states to be conjunctions of causal counterfactuals, of the form $s_j \equiv \bigwedge_{a_i \in \mathcal{A}} (a_i \square \rightarrow Ch_{ij})$. Here, i indexes acts, j indexes causal states, and each Ch_{ij} is a chance distribution over the target partition \mathcal{C} . The agent's *initial* credence function is defined on (at least) the set of such causal states.

¹² Since (as stipulated in the description of the case) the colleague is in fact guilty, any epistemic act \bar{q} with $q(G) = 1$ leads to higher epistemic utility than any epistemic act \bar{q}' that agrees with \bar{q} on the P_i but that has the property $q'(G) < 1$, whether or not the files have been randomized and whatever values the $q(P_i)$ take. In this sense credence 1 in G dominates any lesser credence in G (despite being evidentially correlated with unfortunate circumstances), as two-boxing dominates one-boxing (but is evidentially correlated with unfortunate circumstances) in the Newcomb case.

The *causally expected epistemic utility* of an act $\bar{q} \in \mathcal{A}$, with respect to probability function $Cr \in \mathcal{P}(\mathcal{S}_C)$, is given by

$$(6) \text{ CausEEU}_{Cr}(\bar{q}) = \sum_{s \in \mathcal{S}_C} Cr(s)U(s\&\bar{q})$$

The epistemic utility $U(s\&\bar{q})$ of a conjunction of causal state $s \in \mathcal{S}_C$ and act $\bar{q} \in \mathcal{A}$ is in turn given by the objective expectation value

$$(7) U(s\&\bar{q}) = \sum_{c \in \mathcal{C}} Ch_{qs}(c)U(c\&\bar{q})$$

Thus, combining (6) and (7),

$$(8) \text{ CausEEU}_{Cr}(\bar{q}) = \sum_{c \in \mathcal{C}} \left(\sum_{s \in \mathcal{S}_C} Cr(s)Ch_{qs}(c) \right) U(c\&\bar{q})$$

Causal epistemic decision theory then asserts that, in any epistemic decision situation, the agent is rationally required to perform an epistemic act that maximizes causally expected epistemic utility.

Five remarks are in order.

- (i) For the purposes of causal EDT itself, as noted above, it is not necessary that initial credences be defined on any algebra larger than that generated by the causal states: the causally expected epistemic utility of a given act depends on no features of the initial credence function other than the marginal credences $Cr(s)$ in causal states. However, we wish to consider such issues as whether and when causal EDT recommends retaining one’s initial credences over the target partition \mathcal{C} , and when causal and evidential EDT coincide. For this latter purpose, we need initial credences to be defined, in the first instance, over the product set $\mathcal{S}_C \times \mathcal{A}$; any such probability distribution can be extended, using the chance distributions Ch_s and the Principal Principle,¹³ to a probability distribution over $\mathcal{S}_C \times \mathcal{A} \times \mathcal{C}$.
- (ii) We now seek sufficient conditions for coincidence between the recommendations of causal and evidential EDT. Comparing (8) and (4), we see that a sufficient condition is that, for all $c \in \mathcal{C}$ and $\bar{q} \in \mathcal{A}$, $\sum_{s \in \mathcal{S}_C} Cr(s)Ch_{qs}(c) = Cr(c|\bar{q})$. This

¹³ If we wish to allow for the possibility that the agent’s credence function may violate the Principal Principle, we can easily do so, by simply taking Cr to be defined directly on $\mathcal{C} \times \mathcal{S}_C \times \mathcal{A}$. Consideration of this possibility lies outside the scope of this paper.

condition obtains in Promotion, Leap, and Imps, but not in Embezzlement or Arrogance.

- (iii) It is worth highlighting a special case of this sufficient condition, arising when there is one particular causal state s in which the agent has unit credence. In that case, (8) and (4) coincide if, for this particular state s , $Cr(c|\bar{q}) = Ch_{\bar{q}s}(c)$ for all c, \bar{q} .
- (iv) In the special case in which the chance distributions $Ch_{\bar{q}s}$ depend only on causal state s , and not on act \bar{q} , we may write simply Ch_s for those chance distributions. The coefficients of $U(c\&\bar{q})$ on the right-hand side of (8) are then independent of \bar{q} . It then follows from the fact that U is a proper scoring rule that the credences q that causal EDT recommends adopting over \mathcal{C} are those given precisely by those coefficients: that is, the recommended credence function q is the one given by

$$\forall c \in \mathcal{C}, q(c) = \sum_{s \in \mathcal{S}_C} Cr(s)Ch_s(c)$$

Further, since we have stipulated that initial credences obey the Principal Principle, we must have, for all $c \in \mathcal{C}$,

$$Cr(c) = \sum_{s \in \mathcal{S}_C} Cr(s)Ch_s(c)$$

thus, under these conditions causal EDT will always recommend retaining one's initial credences in the 'target' propositions.

Of our puzzle cases, this condition obtains in Embezzlement and Arrogance, but not in Promotion, Leap, or Imps.

- (v) In particular, if the causal state partition \mathcal{S}_C coincides with the target partition \mathcal{C} , then the conditions analysed in the previous remark must obtain.

We now run through causal EDT's more detailed treatment of each puzzle case in turn.

Causal EDT on Promotion: In the Promotion problem, our agent is certain which causal state obtains: that is, she knows exactly how the chances of her getting promoted depend on her choice of epistemic act (since it is precisely that information that is given in the

specification of the case). The analysis of the decision problem thus consists solely in optimizing the *objective* expected utility (7) of each act, given the causal state with chance function

$$Ch_{\bar{q}}(P) = 1 - q(P)$$

But these chances match the conditional credences $Cr(P|\bar{q})$; hence, causal EDT's analysis of this problem is mathematically identical to evidential EDT's. Therefore, causal EDT also sanctions the (intuitively correct) verdict that evidential EDT gave us for the Promotion case.

Causal EDT on Leap: The agent is again certain which causal state obtains: we have, for all $\bar{q} \in \mathcal{A}$,

$$Ch_{\bar{q}}(L) = q(L).$$

Again, these chances are numerically equal to the conditional credences $Cr(L|\bar{q})$ that are crucial to evidential EDT, and so the mathematical analysis of this decision problem is the same for causal EDT as it is for evidential EDT; causal EDT, like evidential EDT, recommends credence 0 or 1 in success.

Causal EDT on Embezzlement: This is our first case in which causal and evidential EDT come apart. They come apart here because, in the Embezzlement case, Charlie's adopting high credence that her colleague is guilty is *evidence for* the proposition that the files have been randomized, but does not *cause* them to be randomized.

The causal states can be taken to be $\mathcal{S}_C = \{R, \neg R\}$, where R is the proposition that the colleague has randomized the files. Charlie is uncertain which causal state obtains. Each chance distribution Ch_{ij} actually depends only on the causal state s_j , and not at all on the act. The recommended act will therefore be given by

$$q(c) = \sum_{s=R, \neg R} Cr(s)Ch_s(c)$$

This yields, for the propositions G, P_1, \dots, P_n ,

$$\begin{aligned} q(G) &= 1 \\ q(P_i) &= Cr(R) \cdot \frac{1}{2} + (1 - Cr(R)) \cdot 1 \\ &= 1 - \frac{Cr(R)}{2} \end{aligned}$$

According to Causal EDT, therefore, Charlie should retain credence 1 in G.

We noted above that in this puzzle case, causal EDT is guaranteed to recommend retaining one's initial credences in the propositions G, P_1, \dots, P_n . But the specification of the Embezzlement case does not *directly* constrain the initial credences $Cr(R), Cr(\neg R)$ in causal states, so we are as yet in the dark as to what those initial credences are. A further constraint does arise, however, if Charlie *foresees* that she will obey the dictates of causal EDT. In that case, she foresees that she will retain credence 1 in G; since the case specifies that the rogue colleague is a very reliable predictor of Charlie's mental states, Charlie should therefore have $Cr(R) \approx 1$, in which case the recommended act has $q(P_i) \approx \frac{1}{2}$ for $i = 1, \dots, n$.

Unlike evidential EDT, causal EDT thus delivers the intuitively correct verdict for the Embezzlement case. This is perhaps unsurprising (to two-boxers), since this is an epistemic analogue of the Newcomb Problem.

Causal EDT on Arrogance: Again causal and evidential connections, hence the recommendations of causal and evidential EDT, come apart: Dennis's adopting a high degree of belief that he is arrogant is *evidence* against his being arrogant, but does not *cause* him to be humble.

The causal states are $\mathcal{S}_C = \{A, \neg A\}$: either Dennis is arrogant or he is not and, *ex hypothesi*, nothing he does now will *causally* affect which is the case. Thus this is a case in which the partition \mathcal{S}_C of causal states is identical with the target partition \mathcal{C} , and in which (as per the last of the five remarks above) causal EDT will recommend retaining your initial credences over \mathcal{C} , whatever those happen to be.

There is, however, a further complication. Let $x \in [0,1]$ be the agent's initial credence in A, $Cr(A)$. By the specification of the case, we must have, $\forall \bar{q} \in \mathcal{A}$, $Cr(A|\bar{q}) = 1 - q(A)$. We therefore *cannot* have $Cr(\overline{Cr}) = 1$ —that is, the agent cannot *initially* have credence 1 that he will retain his initial credences over $\mathcal{C} \equiv \{A, \neg A\}$ —unless $Cr(A) = \frac{1}{2}$.

Suppose, though, that the agent comes to believe that he will obey the dictates of causal EDT, and notices that causal EDT recommends performing the act \bar{q} with $q(A) = Cr(A)$. *Updating by*

conditionalization on the proposition that he will perform this act takes our agent to a new credence function Cr' , such that

$$Cr'(A) = Cr(A|\bar{q}) = 1 - Cr(A)$$

What happens, then, to agents who *do* happen to start off with credence other than $\frac{1}{2}$ in A , and who obey causal EDT (but, for consistency, do not initially have credence 1 that they will)? Let us illustrate this by picking an example. Suppose an agent begins with credence $p(A) = 0.3$ that he is arrogant, and decides (as causal EDT would recommend) to perform the epistemic act $\bar{q} \in \mathcal{A}$ of retaining this credence in A : that is, he decides to perform \bar{q} such that $q(A) = 0.3$. This is unproblematic until he updates on the proposition that that is the act he will perform; but updating by conditionalization on the proposition that *he will adopt q with $q(A) = 0.3$* changes his credence in A to 0.7. Further, *the epistemic act that maximizes causally expected epistemic utility with respect to his thus-updated credences is not the same as the act that maximizes causally expected epistemic utility with respect to his initial credences*. A self-aware agent trying to obey the dictates of Causal EDT can thus suffer from ‘epistemic decision instability’: the act that is judged optimal by the theory relative to his *initial* credences may be judged sub-optimal by the same theory relative to the credences the agent has once he has updated on the proposition that he will perform the act in question.

To put it another way: taking the epistemic decision that is recommended by causal EDT and updating on the proposition that one will act in accordance with this decision leads to the agent’s regretting that very decision. I take it that vulnerability to this sort of epistemic decision instability is an indicator of epistemic irrationality. If so, we can conclude from this case alone that causal EDT does not correctly capture epistemic rationality. We will return to this problem in section 8.

Causal EDT on Imps: Again there is a single causal state s in which the agent has initial credence 1. The corresponding chance function is: $Ch_{\bar{q}}(C_o) = 1$, $Ch_{\bar{q}}(C_j) = 1 - \frac{q(C_o)}{2}$: that is, the causal state of which the agent is certain specifies that the chance of C_o (the proposition that there is now a child before Emily) is 1, and that the chance of each C_j (the proposition that the j th summerhouse child will come

out to play) is the above decreasing function of Emily's final credence in C_o . As in the Promotion and Leap cases, this chance function coincides with the conditional credences $Cr(C_o|\bar{q}), Cr(C_j|\bar{q})$, so that causal EDT agrees with evidential EDT on this case.

The crucial difference between the Embezzlement and Imps cases is that in the latter, adopting high credence in C_o is not merely evidentially correlated with less extremal chances of the C_j , but actually *causes* the chances of the C_j to take less extremal values. That is why causal EDT, in conflict with pretheoretic intuitions on epistemic rationality, joins evidential EDT in recommending that Emily accept the 'epistemic bribe' that we noted in section 1 in connection with this case.

8. Deliberational decision theory

Both causal and evidential EDT give a clearly incorrect answer to two of our puzzle cases: causal EDT gets the Arrogance problem wrong, evidential EDT gets the Embezzlement problem wrong, and both get the Imps problem wrong. The purpose of the present section is to show that refining causal EDT to generate a 'deliberational causal EDT', following the work of Arntzenius (2008) in practical decision theory, enables us to solve the 'decision instability' problem that plagued causal EDT's treatment of the Arrogance problem while retaining the successes of that EDT.

Since the relevant work in practical decision theory (PDT) is less familiar, we first provide a brief review. Arntzenius's suggestion is that, just as the Newcomb Problem drove us from evidential to causal PDT, similarly Egan's 'Psychopath Button' problem should drive us from our existing causal PDT to *deliberational* causal PDT.

8.1 Deliberational causal practical decision theory

According to deliberational causal PDT, if initial credences are unconstrained, then (*contra* unrefined causal PDT) one should not necessarily choose the action that maximizes causally expected utility with respect to one's initial credences. Rather, one should allow those initial credences to evolve under a 'deliberational dynamics'. The basic idea of such a deliberational dynamics is as follows: credences are, as in evidential decision theory, defined on an algebra that includes both states and acts. Starting from arbitrary initial credences, one calculates the causally expected utility of each available act. One

then revises one's credences about which act one will in fact perform, in a (precisely specified) way that tends to increase one's credence that one will perform the act that currently has the highest expected utility. In 'Psychopath Button cases' (by definition), this process of revision naturally alters one's credences about which state obtains, thus one updates the latter also. One then recalculates the expected utilities of the various available acts, using one's updated credences in states. This process continues until equilibrium is reached. One performs the mixed act corresponding to one's equilibrium credences in acts.

The formal development of deliberational causal PDT is as follows. As before, let \mathcal{A} be the set of pure acts; let $\mathcal{P}(\mathcal{A})$ be the set of probability distributions over \mathcal{A} , representing mixed acts. The agent's credence function at any given time, as in our version of unrefined causal PDT, is defined over the larger space $\mathcal{S}_C \times \mathcal{A}$, but of course any such credence function also delivers marginal credences over \mathcal{A} . A *deliberational dynamics* is (in discrete time) a mapping $h : \mathcal{P}(\mathcal{S}_C \times \mathcal{A}) \rightarrow \mathcal{P}(\mathcal{S}_C \times \mathcal{A})$. h is formed from two parts:

- (i) A map $f : \mathcal{P}(\mathcal{A}) \rightarrow \mathcal{P}(\mathcal{A})$. To build a deliberational *causal* decision theory, we choose a map f that is suitably related to the notion of *causally* expected utility. Specifically, the dynamics we are interested in are those with the following ('causal EU-seeking') features:

(DCDT-1) For all $p \in \mathcal{P}(\mathcal{A})$ and all $a \in \mathcal{A}$, $f(p)(a) > p(a)$ only if $\text{CausEU}_p(a) > \sum_{a' \in \mathcal{A}} p(a') \text{CausEU}_p(a')$.

(At each timestep, the agent increases her credence only in acts whose causally expected utility is greater than that of the 'status quo', i.e. of the mixed act corresponding to her *current* credences in acts.)

(DCDT-2) Let $\mathcal{A}_{p+} \subset \mathcal{A}$ contain just those pure acts whose causally expected epistemic utility relative to p is greater than that of the ('status quo') mixed act p . Then, $\sum_{a \in \mathcal{A}_{p+}} f(p)(a) > \sum_{a \in \mathcal{A}_{p+}} p(a)$. (At each timestep, the agent increases the sum of her credences in all acts that have greater causally expected utility than the status quo.)

- (ii) A rule $g : \mathcal{P}(\mathcal{S}_C \times \mathcal{A}) \times \mathcal{P}(\mathcal{A}) \rightarrow \mathcal{P}(\mathcal{S}_C \times \mathcal{A})$ for updating credences in states in the light of one's revised credences in acts. The rule g must preserve the credences over \mathcal{A} supplied by its second argument. (g may, but need not, be Jeffrey conditionalization.)

Given such ingredients f and g , h is then given by the condition that for all $x \in \mathcal{P}(\mathcal{S}_C \times \mathcal{A})$, $h(x) = g(x, f(x|_{\mathcal{A}}))$.

According to deliberational causal PDT, the agent is rationally required to perform a mixed act corresponding to (the restriction to \mathcal{A} of) some equilibrium of the deliberational dynamics.¹⁴

As shown by Arntzenius, at any equilibrium of any such dynamics in the Psychopath Button case, the agent has credence strictly between 0 and 1 that he will push the button. If the agent's credence that he will push is lower (resp. higher) than its equilibrium value, his credence that he is a psychopath will be low (resp. high), and so pushing the button has higher (resp. lower) causally expected utility with respect to his credences than does not pushing. Relative to his equilibrium credences, the two available acts have equal causally expected utility to one another. Thus deliberational causal decision theory recommends a non-pure act for the Psychopath Button case.

This verdict is plausible. Hence, *extensionally*, deliberational causal PDT seems to perform well. It is, however, worth pausing to note a feature of the theory that is conceptually odd.

The feature in question concerns the interpretation of mixed acts. As Arntzenius also notes, if these acts were really acts of setting in motion some chancy process that will result in one doing various things with various chances, we could treat that setting-in-motion itself as a pure act. But in that case the notion of mixed acts would

¹⁴ The existence of at least one equilibrium is guaranteed, under natural continuity assumptions, by a fixed-point argument. The equilibrium will not in general be unique: indeed, in mundane cases (such as those handled unproblematically by a Savage-style decision theory) *any* credence distribution over states, combined with credence 1 that one will perform the pure act that maximizes expected utility with respect to those credences in states, constitutes an equilibrium of any allowed deliberational dynamics.

In stipulating that the rational requirement is (merely) to correspond to *some* equilibrium, we depart slightly from Arntzenius' version of deliberational decision theory: Arntzenius suggests insisting that 'an ideally rational person is always in an equilibrium state such that there is no other equilibrium state that has higher [expected] utility' (Arntzenius 2008, p. 295). This move, however, would wrongly classify agents who believe there is only £2000 in their bank accounts and plan accordingly as irrational, on the grounds that the state of believing that there was £100,000 in their accounts and planning accordingly would be another equilibrium and would have higher expected utility.

be eliminable from the theory; meanwhile, it does not seem to be so eliminable. Arntzenius suggests taking instead the endpoint of deliberation to be *credences* about which (pure) acts one will perform. But this is odd: one would expect the endpoint of deliberation to be a *decision*.

This strongly suggests that deliberational decision theory is not entirely correct. The theory's natural way of solving the 'decision instability' problem, however, suggests that it may contain *some* truth; thus investigation of its epistemic analogue is called for.

8.2 Deliberational causal epistemic decision theory

We claim that applying the analogous refinement to causal *epistemic* decision theory produces an EDT that gives a treatment of the Arrogance problem that is at least as satisfactory, and possibly *more* satisfactory, than deliberational PDT's solution of the Psychopath Button problem, while retaining the existing successes of causal EDT.

As in ordinary causal EDT, and as in deliberational causal PDT, let the agent's credence function be defined on the space $\mathcal{S}_C \times \mathcal{A}$. The 'analogous refinement', exactly as in deliberational causal PDT, is the replacement of the prescription 'maximize causally expected epistemic utility' with the prescription 'perform a mixed act corresponding to some equilibrium of a deliberational dynamics $h: \mathcal{P}(\mathcal{S}_C \times \mathcal{A}) \rightarrow \mathcal{P}(\mathcal{S}_C \times \mathcal{A})$ such that the corresponding maps $f: \mathcal{P}(\mathcal{A}) \rightarrow \mathcal{P}(\mathcal{A})$ and $g: \mathcal{P}(\mathcal{S}_C \times \mathcal{A}) \times \mathcal{P}(\mathcal{A}) \rightarrow \mathcal{P}(\mathcal{S}_C \times \mathcal{A})$ satisfy **DCDT-1** and **DCDT-2**'.

Three general observations will help us to draw out the implications of this theory for our puzzle cases (the analogous remarks apply equally to deliberational causal PDT):

- (i) If a given credence function p on $\mathcal{S}_C \times \mathcal{A}$ is an equilibrium of an allowed deliberational dynamics, then all pure acts $a \in \mathcal{A}$ that are in the support of p have equal causally expected epistemic utility (with respect to p) to one another, and no pure act outside that support has higher causally expected epistemic utility (with respect to p) than those inside.
- (ii) If there is only one (non-null) causal state (as in Promotion, Leap, and Imps), the causally expected epistemic utility of each act is invariant under the deliberational dynamics. The

equilibrium points are then exactly those that assign credence 1 to the act, or acts, that have highest causally expected epistemic utility. If there is a unique optimum (as in Promotion and Imps), then deliberational causal EDT agrees with unrefined causal EDT: the act in question is rationally required. If there is more than one optimum pure act (as in Leap) then, whereas unrefined causal EDT required that one of these pure acts be performed, deliberational causal EDT also permits mixtures thereof.

- (iii) If there is more than one non-null causal state and acts are evidentially relevant to causal states, then deliberational causal EDT can be more restrictive than unrefined causal EDT, via the restriction that the credences with respect to which pure acts are ultimately evaluated must be equilibrium credences. This happens in the Embezzlement and Arrogance cases.

We now make more detailed comments on the three cases—Leap, Embezzlement, and Arrogance—on which the deliberational and unrefined causal EDT issue different verdicts.

Leap: There is one causal state, but there are two pure acts that have joint highest causally expected epistemic utility: the acts of adopting credence 0 and 1 in the proposition L that Bob's leap is successful. *Deliberational* causal EDT will permit these two pure acts, but will also permit also any convex combination thereof. That is, for any $x \in [0,1]$, the credence function given by

$$Cr(\text{final credence in L is 1}) = x$$

$$Cr(\text{final credence in L is 0}) = 1 - x$$

$$Cr(L) = x$$

is an equilibrium of any allowed deliberational dynamics, and the corresponding mixed act is rationally permitted according to deliberational causal EDT.

Embezzlement: There are two causal states (R and $\neg R$), and acts are evidentially relevant to states. If Cr is an equilibrium credence function then it must not only have $Cr(\text{final credence in G is 1}) = 1$, but also $Cr(R) = 1$. Thus moving from unrefined to deliberational causal

EDT removes the need for the qualification ‘if she foresees that she will obey the dictates of causal EDT’ in section 7: deliberational causal EDT requires without such qualification that the agent must adopt not only credence 1 in G , but also credence $\frac{1}{2}$ in each P_i . (This, of course, is because deliberational decision theory effectively *requires* agents to foresee that they will indeed obey the decision theory.) It is not clear whether this is an advantage of the deliberational theory, a disadvantage, or neither.

Arrogance: Again there are two causal states (here A and $\neg A$), and acts are evidentially relevant to states. Again there is a unique equilibrium credence function q , given this time by¹⁵

$$q(A) = \frac{1}{2}$$

$$q\left(\text{final credence in } A \text{ is } \frac{1}{2}\right) = 1$$

Thus the move to a deliberational variant of causal EDT removes the ‘decision instability’ problem faced by unrefined causal EDT.

Why say, as I did at the start of section 8.2, that this may be *more* satisfactory than deliberational causal PDT’s solution to Psychopath Button? The point is that there is an interesting disanalogy between deliberational causal PDT’s verdict on Psychopath Button and deliberational causal EDT’s verdict on *Arrogance*: the latter, while using the deliberational dynamics (and associated contemplation of mixed acts) to rule out pure acts that lead to decision instability, *nevertheless recommends a pure act*. Heuristically, this is possible because the set of acts in the *epistemic* case already comes with a suitable continuum structure—that furnished by the $[0,1]$ target space of degrees of belief—whereas in the case of PDT the action set {push, don’t push} was discrete. There is thus an enticing prospect, in the *epistemic* case, that the theory may not actually require *performing* any of the conceptually problematic ‘mixed acts’; and if so there is a prospect of a variant theory that *requires*

¹⁵ Proof: In order to be an equilibrium of a dynamics that complies with DCDT-2, it must be that no pure act has strictly higher *CausEEU* than the status quo. But, since U is a proper scoring rule, we know that retaining current credence in A has strictly higher *CausEEU* than any distinct act. Hence, the status quo must assign credence 1 to the retention of one’s current credence in A . Formally: if Cr is an equilibrium credence function then $Cr(\overline{Cr|C}) = 1$. But in that case $Cr(A|\overline{Cr|C}) = Cr(A)$, and we know that for all $\bar{q} \in \mathcal{A}$, $Cr(A|\bar{q}) = 1 - q(A)$; hence we have $1 - Cr(A) = Cr(A)$, which is satisfied only if $Cr(A) = \frac{1}{2}$.

performing some pure act that is permitted by the deliberational theory developed herein, and thereby avoids the need to interpret mixed acts. (In Leap, mixed acts were *permitted*, but not *required*.) Whether this can be established under any interestingly general conditions (as opposed to for particular examples, as here) is an open question.

9. Conclusions

Previous work in epistemic decision theory (EDT) has, to my knowledge, explored only the simplest sorts of case, in which decision-theoretic ‘states’ are causally and probabilistically independent of epistemic acts, so that a Savage-style decision theory suffices.¹⁶ In this paper, we have explored epistemic analogues of the main developments in the foundations of practical decision theory. For the most part, the similarities between the practical and epistemic cases have been more striking than the differences: thus we have been able to develop evidential, causal, and deliberational causal epistemic decision theory, in direct analogue to their practical counterparts, apparently running against the current of thought according to which epistemic rationality is a *very* different beast to practical rationality. The most successful of these theories, deliberational causal epistemic decision theory, deals well with all the puzzle cases we have considered *except* the Imps case.

There is, however, a crucial difference, and one that *may* be fatal for epistemic decision theory: causal practical decision theory would issue the correct verdict on a practical analogue of our Imps case, but we have *no* epistemic decision theory that deals adequately with this case. We have, that is, no theory that recovers the obviously correct result that an agent (epistemically-) should retain credence 1 in propositions for which she has conclusive evidence, even in the face of ‘epistemic bribes’. The *reason* this is difficult to recover via a decision-theoretic approach is, heuristically, that a decision-theoretic utility function always assesses epistemic utility *globally*, and hence will always be open to the move of increasing *overall* expected epistemic utility by making a sacrifice of a relatively small number of propositions; our intuitive notion of epistemic rationality, meanwhile, does not seem to exhibit this willingness.

¹⁶ Caie forthcoming, which also considers act-state dependence, was brought to my attention while the present paper was under review.

There are three options from here. Firstly, one could allow an otherwise well-performing theory to revise the problematic intuitions concerning these epistemic tradeoffs: that is, one could conclude that it is, initial intuition notwithstanding, epistemically rational to accept epistemic bribes. Secondly, one could seek some way of further refining epistemic decision theory, in such a way as to retain the theory's existing successes but deliver the desired verdict also for the Imps case. Or, thirdly, one could conclude that this is indeed a fatal objection to epistemic decision theory, and abandon the approach altogether. The first seems implausible. The third may be correct, but is a last resort: we should try every possible way of making epistemic decision theory succeed before giving up on it.

Here, then, is a suggestion for how the second might be pursued. We might simply *impose* any rationality constraints that are generated by the agent's awareness of the specification of her case, by way of side-constraints. That is, expected epistemic utility is to be maximized *subject to* the requirement that the 'final' credence function respects any such side-constraints. Such a modified epistemic decision theory is of course guaranteed not to violate any constraints that intuitively seem to be imposed by 'the evidence' (such as the constraint $q(C_0) = 1$ in the Imps case).

One might worry that this move is ad hoc. The degree of ad hocery, however, is significantly reduced by the fact that it is not merely a matter of intuition, outside of epistemic decision theory, insisting that certain constraints on the 'final' credence function are requirements of epistemic rationality (so that it counts as a failure of the epistemic decision theory if the latter recommends 'final' credence functions violating those constraints). We also saw, in section 6 above, that the decision theory *itself* has already needed to acknowledge these rationality constraints, in giving an account of the status of the decision theory's 'initial' credence function. The proposal currently under consideration is therefore to make a second use of a resource to which our theories have already helped themselves once. (It amounts, in effect, to adopting the 'third possible response' discussed in Sect. 6 to the worry about the status of epistemic decision theory's 'initial' credence functions.)

The fact does remain, however, that there is no epistemic-consequentialist rationale for the side-constraints in question. Illumination may be gained by considering the question, thrown into sharp relief by the successes and apparent failures of (side-constraint-free) epistemic

decision theory: *Why do we have* a notion of epistemic rationality that forbids epistemic bribes?

We can flesh out this question by returning to the ‘ultimate argument’ for consequentialism mentioned in section 1: it is paradoxical to hold, of an action that leads to higher good (or expected good), that it is wrong, and this thought is no less compelling in the epistemic than in the ethical case. It is thus *prima facie* mysterious why we would have a non-consequentialist intuitive standard of epistemic rationality. We could, of course, come up with an extensionally adequate and *formally* consequentialist theory by building the desired factors into the theory of the good (i.e. abandoning a ‘thin’ in favour of a ‘thicker’ epistemic utility function); but, just as in the ethical case, this move does not seem to yield genuine explanatory power (it is merely ‘gimmicky representation’). In the ethical case, we might try to resolve the paradox via explorations of the consequences of *publicly advocating* adherence to certain rules, of internalizing such rules and/or of adopting them as a decision procedure, rather than simply of conforming to them; it is an open question how much success any such tactic could yield in the epistemic case.¹⁷

References

- Arntzenius, F. 2008: ‘No Regrets, or: Edith Piaf Revamps Decision Theory’. *Erkenntnis*, 68, pp. 277–97.
- Caie, M. 2013: ‘Rational Probabilistic Incoherence’. *The Philosophical Review*, 112, pp. 527–75.
- Egan, A. 2007: ‘Some Counterexamples to Causal Decision Theory’. *The Philosophical Review*, 116, pp. 93–114.
- Greaves, H. and D. Wallace 2006: ‘Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility’. *Mind*, 115, pp. 607–32.
- Jeffrey, R. 1965: *The Logic of Decision*. Chicago: University of Chicago Press.
- Joyce, J. 1999: *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.

¹⁷ I am grateful to Frank Arntzenius, Grant Reaber, and audiences at Bristol, Oxford, Leeds, the London School of Economics, and Nottingham, for useful discussions. Special thanks are due to Ralph Wedgwood, for extensive and very helpful comments on an earlier version of this paper.

- Kelly, T. 2003: 'Epistemic Rationality as Instrumental Rationality: A Critique'. *Philosophy and Phenomenological Research*, 66, pp. 612–40.
- Leitgeb, H. and R. Pettigrew 2010a: 'An Objective Justification of Bayesianism I: Measuring Inaccuracy'. *Philosophy of Science*, 77, pp. 201–35.
- 2010b: 'An Objective Justification of Bayesianism II: The Consequences of Minimising Inaccuracy'. *Philosophy of Science*, 77, pp. 236–72.
- Levi, I. 1967: *Gambling with Truth*. New York: Knopf.
- Lewis, D. 1981: 'Causal Decision Theory'. *Australasian Journal of Philosophy*, 59, p. 5.
- Maher, P. 1993: *Betting on Theories*. Cambridge: Cambridge University Press.
- Oddie, G. 1997: 'Conditionalization, Cogency, and Cognitive Value'. *British Journal for the Philosophy of Science*, 48, pp. 533–41.
- Savage, L. J. 1972: *The Foundations of Statistics*, second edition. New York: Dover.
- Skyrms, B. 1990: *The Dynamics of Rational Deliberation*. Harvard: Harvard University Press.
- Stalnaker, R. 2002: 'Epistemic Consequentialism'. *Proceedings of the Aristotelian Society Supplementary Volume*, 76, pp. 153–68.
- Wedgwood, R. 2002: 'The Aim of Belief'. *Philosophical Perspectives*, 16, pp. 267–97.