The Foundations of Epistemic Decision Theory*

Jason Konek
University of Bristol
jason.konek@bristol.ac.uk

Ben Levinstein
University of Oxford
ben@levinstein.org

Abstract

According to *accuracy-first epistemology*, accuracy is *the* fundamental epistemic good. Epistemic norms — Probabilism, Conditionalization, the Principal Principle, etc. — have their binding force in virtue of helping to secure this good. To make this idea precise, accuracy-firsters invoke Epistemic Decision Theory (EpDT) to determine which epistemic policies are the best means toward the end of accuracy. Hilary Greaves and others have recently challenged the tenability of this programme. Their arguments purport to show that EpDT encourages obviously epistemically irrational behavior. We develop firmer conceptual foundations for EpDT. First, we detail a theory of praxic and epistemic good. Then we show that, in light of their very different good-making features, EpDT will evaluate epistemic *states* and epistemic *acts* according to different criteria. So, in general, rational preference over states and acts won't agree. Finally, we argue that based on direction-of-fit considerations, it's preferences over the former that matter for normative epistemology, and that EpDT, properly spelt out, arrives at the correct verdicts in a range of putative problem cases.

Keywords. Epistemic Decision Theory, Accuracy, Direction of Fit, Epistemic Rationality

1 Introduction

Credences have a range of epistemically laudable properties. They are more or less *specific* and *informative*. They encode more or less *simple* and *unified* explanations of *prima facie* diverse phenomena. They are more or less *appropriate* or *justified* in light of our evidence. And importantly, they are closer or further from *the truth*, i.e., more or less *accurate*. According to *accuracy-first epistemology*, this final virtue — accuracy — is *the* fundamental epistemic good. It is the primary source of epistemic value. The higher your credence in truths and the lower your credence in falsehoods, the better off you are all epistemic things considered.

Norms of epistemic rationality, on this view, have their binding force in virtue of the following fact: they are good means toward the end of securing accuracy. Obeying them in some way helps in the pursuit of accurate credences. To spell this out, accuracy-firsters co-opt the resources of practical decision theory. Just as decision-theoretic norms explain why certain practical policies — economic policies, environmental policies, etc. — are bad means to practical ends, and hence irrational, they also explain why certain *epistemic* policies are bad means to the epistemic end of accuracy, and hence irrational. For example, Pettigrew (2013, 2014a,b) uses standard decision-theoretic norms — Dominance, Chance Dominance, Maximize Expected Utility, Maximin — together with an appropriate measure of accuracy, to explain why violating various epistemic norms — Probabilism, the Principal Principle, Conditionalization, the Principle of Indifference — is bad epistemic policy. Violating them is a bad means to the epistemic end of accuracy.

Despite these promising beginnings, recent challenges have cast doubt on the tenability of this project. Hilary Greaves (2013) has forcefully argued that any accuracy-first approach sanctions epistemically irrational behavior in problem cases, the most vexing of which is the following:[1]

---

[1] Caie (2013) and Berker (2013) raise related concerns. Our treatment of Greaves' problem cases extends naturally to theirs as well.

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 10 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 $x = 0$ that there is now a child before her, they will come out to play. If she forms degree of belief $x = 1$ 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 $(1 - 0.5x)$. Emily's epistemic decision is the choice of credences in the propositions C_0 that there is now a child before her, and, for each $j = 1, \ldots, 10$ the proposition C_j that the j^{th} summerhouse child will be outdoors in a few minutes' time.

In this case, Emily is offered an epistemic bribe. If she can get herself to deny the manifest and have credence 0 in C_0, she can guarantee herself perfect accuracy in propositions $C_1, \ldots, C_{10}$. So, although her credence in C_0 would be maximally inaccurate, her *overall* level of accuracy would be highest if she took the bribe.[2]

Any plausible way of spelling out an accuracy-first epistemology, Greaves thinks, will sanction taking this epistemic bribe. The reason: any reasonable accuracy measure will assess the accuracy of one's credal state *globally.* It will take the accuracy of all of your credences considered individually, weigh them up in some sensible way, and deliver a 'summary statistic' that captures how accurate they are as a whole. As a result, the accuracy measure will be open, so to speak, to sacrificing accuracy in a relatively small number of propositions in exchange for gaining accuracy in a large number of other propositions. Our intuitive notion of epistemic rationality, however, does not sanction taking this epistemic bribe. It does not sanction lowering one's credence in C_0 to 0 *when the child is standing right there in plain view.*

By entangling Emily's epistemic choices with the external state of the world, IMPS brings forth an issue lying at the foundation of any decision theory, practical or epistemic. Compare: Savage-style unconditional expected utility theory (SAVAGE) — the 'standard model' in practical decision theory since Savage's seminal *The Foundations of Statistics*

---
[2]Taking the bribe in this particular case leads to more accuracy only on some measures. However, for any measure $\mathcal{I}$ of accuracy, we can formulate a case such that bribe-taking leads to the most accuracy according to $\mathcal{I}$.

(1954) — comes with a tacit warning: only apply if probabilities of states are independent of acts (IND). Determining what to do when states and acts are entangled—and thereby doing away with IND— motivated the move from SAVAGE to a Jeffrey-style evidential decision theory (EDT), and eventually to causal decision theory (CDT). In the process, we learned something deep about the nature of practical rationality.

In the early formulations of EpDT, accuracy-firsters have likewise presupposed IND. They only consider cases in which probabilities of states do not depend on which credences you adopt. Doing away with IND, Greaves argues, teaches us something deep about the nature of *epistemic* rationality. It teaches us that epistemic rationality is non-consequentialist. Or at least: the most popular brand of epistemic consequentialism — accuracy-first epistemology — cannot capture our intuitive notion of epistemic rationality.

We agree that cases like IMPS require accuracy-firsters to establish a firmer conceptual foundation for epistemic decision theory (EpDT). But, *pace* Greaves, we do not think that they show epistemic rationality to be non-consequentialist. By drawing the proper moral from developments in practical decision theory over the last 60 years (progressing from SAVAGE to EDT to CDT), we hope to show why accuracy-first epistemology, spelt out correctly, handles Greaves' problem cases in just the right way.

Part of our task will be clearing up just what it is that needs to be explained in Greaves' problem cases. To this end, we ought to make an important observation right away, to which we will return at various points. When an agent adopts a credence function c, there are two very different ways to evaluate her in terms of her overall accuracy. First, we can evaluate how closely the epistemic *state* she occupies conforms to the world. We can evaluate how close her credences for various propositions are, at any particular time, to the actual truth-values of those propositions. Second, we can evaluate how much accuracy her coming-to-occupy c (which we denote $\bar{c}$) produced. That is, we can evaluate the epistemic *action* of *adopting c* as her credence function. Cases like IMPS bring out this distinction: If Emily has credence 0 in C_0 and credence 1 in $C_1, \ldots, C_{10}$, she knows the credence function that assigns 1 to C_0 and $C_1, \ldots, C_{10}$ is more accurate than her own. However, *were she to adopt that state*, she would end up less accurate than she currently is.

EpDT ought to evaluate epistemic states and epis-

temic actions by different criteria. Epistemic actions, like all actions, are properly assessed in terms of their causal impact on the world. They are valuable to the extent that they *make* the world fit our desires; to the extent that they *cause* the world to be good (desirable). Epistemic states, on the other hand, are assessed in terms of their fit *to the world*. They are valuable to the extent that they encode an accurate picture of the world, not to the extent that they causally influence the world so as to make it fit that picture.

Accuracy-first epistemology will yield evaluations both of epistemic states and epistemic acts. But the deliverances of epistemic rationality, we will argue, track evaluations of epistemic *states*, not acts. The reason: epistemic states, rather than acts, have the epistemically interesting direction of fit, *viz.*, mind-to-world.

In Imps, then, accuracy-first epistemology says: The *action* of taking the bribe does the best job of getting Emily what she wants (if all she wants is accuracy). Nevertheless, she would be epistemically irrational for occupying a *state* that assigned credence 0 to C_0.

Here's the plan. In §2, we discuss the theory of the praxic and epistemic *good*. §3 introduces our own theory of praxic and epistemic *preference*, which specifies how to evaluate epistemic actions and states, respectively, in light of their very different good-making features. In §4 we apply the results of our discussion to diagnose the apparent epistemic dilemmas posed by Imps and related cases. Finally, §5 discusses EpDT's recommendations and their implications for epistemic rationality.

2 A Theory of the Good: Praxic and Epistemic

Rational agents, on our view, line up their preferences over options—acts, epistemic states—with their unconditional best estimates of the value—prudential value, epistemic value—of those options. This *general theory of preference* is the common core of practical and epistemic decision theory. The key to spelling out the correct practical decision theory, we will argue, is to pin down the correct theory of *prudential value* or *praxic good* (see §3). The key to spelling out an accuracy-first epistemology is to pin down the correct theory of *epistemic value* or *epistemic good*. In the remainder of §2, we will argue that both causal decision theory's account of praxic good, and the accuracy-firster's account of epistemic good are independently motivated. They

fall naturally out of the direction-of-fit metaphor, once it is properly unpacked.

2.1 Direction of Fit

It's a common adage that beliefs have a mind-to-world direction of fit, while desires have a world-to-mind direction of fit. You might understand this *descriptively*, *e.g.*, as a causal claim:

> A belief that p tends to go out of existence in the presence of a perception with the content that not p, whereas a desire that p tends to endure, disposing the subject in that state to bring it about that p. (Smith, 1994, p. 115)

or perhaps a claim about higher-order attitudes:

> The thetic/telic difference [difference in direction of fit between beliefs and desires] is a difference in the structure of a controlling conditional intention [a higher-order intention]. … The controlling background intention in the case of belief is … [the intention] not to believe that p, given that (or: in the circumstance that) *not p* … in the telic [desire] case, the intention is that it be the case that p, given the telic attitude toward p. (Humberstone, 1992, pp. 75–6)

It would be better, though, to understand the direction of fit metaphor *evaluatively*, as Anscombe does (*cf.* Sobel and Copp 2001).[3] To illustrate Anscombe's position, imagine that a man writes a shopping list and goes to the store. As he shops, a detective hired to follow him writes down everything that she thinks the man is buying. What is the difference between the shopping list (which reflects the man's *desires*) and the detective's records (which reflects her *beliefs*)?

> It is precisely this: if the list and the things that the man actually buys do not agree … then the *mistake* is not in the list but in the man's *performance* (if his wife were to say: 'Look, it says butter and you have bought margarine', he would hardly reply: 'What a mistake! we must put that

---

[3] Sobel and Copp (2001) explore whether the best theory of direction-of-fit could provide an account of belief and desire. While we *do* think Anscombe's proposal best explicates the direction of fit metaphor, we do *not* endorse using it for such a purpose.

right' and alter the word on the list to 'margarine'); whereas if the detective's record and what the man actually buys do not agree, then the *mistake* is in the record. (Anscombe, 1957, p. 56; emphasis ours)

When you desire to buy butter and you put margarine in the basket, your *action* is bad (mistaken), or lacking value (*prudential* value). Your action fails to make the world bend to your will. It fails to causally influence the world in a way that satisfies your desire. And exerting the right sort of causal influence — making good (desired) outcomes come about — is what gives actions (prudential) value.

Desires seem to have a world-to-mind direction of fit, then, in just this sense: the means to satisfying them, *viz.*, actions, are *better* (more valuable) to the extent they *make the world conform to those desires.* They are better to the extent that they *causally influence* the world in the right way, so that those desires are satisfied.

In contrast, when you *believe* there's butter in the basket, but there's not, your *belief* is bad (mistaken) and thereby lacks *epistemic* value. Your belief fails to accurately represent the world. And accurately representing the world, or 'getting close to the truth', is what gives beliefs (epistemic) value.

So beliefs have a mind-to-world direction of fit in the following sense: they are *better* (more epistemically valuable) to the extent that *they* conform to the world. They are better to the extent that they accurately represent the world. Unlike actions, they are *not* valuable in virtue of causally influencing the world, so as to make themselves accurate. Of course, rational inquirers are part of the causal system that they hope to investigate. As such, they may, by adopting some belief or other, influence the world in any number of ways. But — and this is the crucial point — influencing the world in (epistemically) good or bad ways is not what *makes* them epistemically valuable. What *makes* them epistemically valuable — the primary source of all-epistemic-things-consider value — is just accuracy.

To take an example, if God believes there is now light when there is not, then God's belief is not *epistemically* valuable on this view. It lacks any peculiarly *epistemic* virtue. This is so even if God's belief *causes* there to be light. Such a causally efficacious belief is valuable *once true*. But it is not valuable *in virtue of* causally influencing the world in some way or other.

On Anscombe's view, then, the direction of fit adage is best understood as encapsulating a theory of the good. In particular, it is best understood as encapsulating a theory of praxic and epistemic good, respectively. A theory of praxic good specifies which factors conspire to make *actions* (the means to satisfying desires) prudentially good or valuable, and how they do so. A theory of epistemic good specifies which factors make *beliefs* (or doxastic states more generally) epistemically good or valuable. In a bit more detail, our theories say:

> *Praxic Good.* An action A is prudentially valuable at a world w, relative to a state of desire D, to the extent that A makes w satisfy D, by causally influencing it in the right way.

> *Epistemic Good.* A doxastic state B is epistemically valuable at a world w to the extent that B is close to the truth (accurate) at w.

2.2 Praxic Good

To make this more precise, we will focus our attention on an agent whose state of belief or opinion is given by a credence function c defined on a finite algebra Ω, and whose non-instrumental desires are given by a utility function u defined on the atoms w of Ω (the finest-grained possibilities that the agent can distinguish between).[4] Let $\mathcal{W}$ be the set of all such atoms or 'possible worlds'. An agent's credence $c(X)$, roughly speaking, measures the strength of her confidence in X, where $c(X) = 0$ and $c(X) = 1$ represent minimal and maximal confidence, respectively. An agent's utility $u(w)$ measures the strength of her desire that w be true.[5]

In addition to unconditional opinions, captured by c, we will suppose that our agent has various conditional opinions. Her confidence in X on the indicative supposition that Y is given by her credence for X conditional on Y, $c(X|Y)$.[6] So, for example, if she is next to certain that someone else killed Kennedy if Oswald in fact did not, then $c(\textit{Someone else killed K}|\textit{Oswald did not kill K}) \approx 1$. In contrast, her confidence in X on the subjunctive supposition that Y is given by her credence for X *imaged* on Y, $c(X\|Y)$. Roughly speaking, $c(X\|Y)$

---

[4]More carefully, Ω is a finite set of propositions closed under negation and countable disjunction.

[5]For ease, we'll assume that all credence functions under consideration are probability functions, though this restriction is unnecessary.

[6]When $c(Y) > 0$, $c(X|Y)$ is just $c(X\&Y)/c(Y)$. For a theory of conditional probability that allows $c(X|Y)$ to be defined when $c(Y) = 0$, see Rényi 1955 and Popper 1959.

shifts the credence spread over $\neg Y$-worlds to Y-worlds in proportion to their estimated similarity to the actual world (*cf.* Lewis 1986, p. 310). So if our agent is next to certain that no one else would have killed Kennedy had Oswald not done so, then $c(\textit{Someone else killed K} \| \textit{Oswald did not kill K}) \approx 0$.[7]

Typically, $c(\cdot|A)$ will reflect our agent's views about A's *evidential* import and $c(\cdot\|A)$ will reflect her views about A's *causal* powers.[8] In particular, $c(X|A) > c(X)$ only if she thinks that learning A increases the degree to which her total evidence confirms X. And $c(X\|A) > c(X)$ only if she thinks that A has a positive incremental causal impact on X. That is, $c(X\|A) > c(X)$ only if she thinks that (i) A causally promotes X, and (ii) A increases the degree to which the totality of causally relevant factors promote X.[9]

To illustrate how these two types of views might come apart, imagine that you wake up on the roof of an abandoned building. You cannot remember who you are or where you are from. Your identity is a mystery. You look down at your hands — a rifle. Frantically, you search for a clue; any hint as to why you are in this mess. Then you notice: there is a gathering in the square below; a terrible despot is about to take the stage. You grab the binoculars at your feet and scan the buildings surrounding the square. Three government snipers, maybe more. But you doubt you are part of their team. They're in body armour, and you're in ratty jeans and a t-shirt.

"Who am I?" you mutter. Maybe you are a lone vigilante who has been planning to end the despot's reign of terror singlehandedly. Or maybe you are just a patsy, placed on the building by the government to take the fall after the snipers complete their mission.

The despot takes the stage. The snipers lift their rifles. "Should I shoot too?" you wonder. On the one hand, it would be *bad news* to learn you took the shot. If you actually have the nerve to shoot, you think to yourself, then you are probably a slightly scatterbrained vigilante, rather than a patsy. But the despot's security team is notoriously adept at sniffing out scatterbrained vigilantes. They almost certainly have sussed out your plan (if you have one)

and put extra security measures in place to help foil attempts on the despot's life. So pulling the trigger, in your view, provides evidence that the despot will *survive*, rather than die (in virtue of providing good evidence that there is extra security in place). This is reflected in the fact that $c(\textit{Death}|\textit{Shoot}) < c(\textit{Death})$. If you in fact have the nerve to take the shot, then in your best estimate, the despot is *less* likely to end up dead, all things considered.

"On the other hand," you think to yourself, as you lift your rifle and notice how steady your hand is, "taking the shot will make a positive difference." Whether the despot's security team put extra protective measures in place or not, having one more steady-handed marksman (you) taking a shot raises the chance that the despot will meet his end, if only by a small amount. Perhaps, for example, you think that each of the three government snipers has a 90% chance of hitting the despot if there is no extra security in place, and a 20% chance even if there is. And you think that *you* have a 75% chance of hitting the despot if there is no extra security in place, and only a 1% chance if there is. Then adding your shot to the mix raises the chances that *someone* will hit the despot from about 48% to 49% if there is extra security, and from 99.9% to 99.98% if not.[10]

Shooting raises the chances in this way because it has a positive incremental causal impact on the despot's death. Shooting *promotes* that end, and moreover increases the degree to which the *totality* of causally relevant factors promote that end. (It is not swamped by other causal factors.)

These opinions about the causal structure of the world are reflected in your subjunctive conditional (imaged) credences. In particular, the fact that you think shooting has a positive incremental impact on the despot's death is reflected in the fact that $c(\textit{Death}\|\textit{Shoot}) > c(\textit{Death})$. If you were to shoot, the despot would be more likely to die, in your best estimate, than he currently is.

One measure of *how large* an incremental causal impact your shooting has is the 'imaged Bayes factor': $c(\textit{Death}\|\textit{Shoot})/c(\textit{Death})$. When $c(\textit{Death}\|\textit{Shoot})/c(\textit{Death})$ is greater than 1, shooting has a positive incremental impact on the despot's death, according to this metric. When $c(\textit{Death}\|\textit{Shoot})/c(\textit{Death})$ is less than 1, it has negative incremental impact; it causally inhibits the despot's death. When $c(\textit{Death}\|\textit{Shoot})/c(\textit{Death})$ equals 1, it has no incremental impact whatsoever.

Of course, $c(X\|A)/c(X)$ is neither perfect nor the

---

[7] Lewisian imaging is, of course, only one way to model subjunctive supposition in Bayesian epistemology. We assume it henceforth for illustrative purposes. But nothing substantive hangs on this assumption.

[8] See, for example, Hájek and Joyce 2008, Joyce 1999, §5.4, 2000, pp. S10–11, and 2002, p. 74.

[9] Imaged credences will not so straightforwardly reflect one's causal opinions in cases of preemption and trumping. See Lewis 1986, 2000.

[10] Assuming that your respective chances of success are independent.

only measure of *A*'s incremental causal impact on *X*. But it will help illuminate why causal decision theory provides the wrong sorts of tools for building up an accuracy-first epistemology, so we'll be using it as our official measure in what follows. And the morals we draw at the end of the day are entirely general. They do not depend on this particular choice of measure.

With a measure of incremental causal impact in hand, we can fill in our schematic theory of praxic good.

> *Praxic Good.* An action *A* is prudentially valuable at a world *w*, relative to a state of desire *D*, to the extent that *A* makes *w* satisfy *D*.

An action *A* makes the world *w* satisfy *D* to the extent that (i) *A* helps to make *w* true (false), by having a positive (negative) incremental impact on *w*, and (ii) *w* is desirable (undesirable), according to *D*. On the view we've settled on, *A*'s impact on *w* is measured by $c(w\|A)/c(w)$. And *w*'s degree of desirability is measured by $u(w)$, our agent's utility for *w*. Given this, we ought to fill in our schematic theory of praxic good as follows:

> *Praxic Good**. An action *A*'s prudential value at a world *w*, relative to credences *c* and utilities *u*, is given by:
> $$\mathcal{V}_A(w) = [c(w\|A)/c(w)] \cdot u(w).$$

According to our measure, $\mathcal{V}_A$, if *w* is a positively desirable state of the world, so that $u(w) > 0$, then *A*'s prudential value, $\mathcal{V}_A(w)$, increases as $c(w\|A)/c(w)$ increases. It is more valuable the more it does to help make *w* true. If *w* is an undesirable state of the world, so that $u(w) < 0$, then *A*'s value increases as $c(w\|A)/c(w)$ decreases (approaches zero). It is more valuable the more it does to help make *w* false.

Suppose, for example, that the world w^* in which you take the shot and the despot is killed, despite the extra security, is a highly desirable one. Perhaps $u(w^*) = 10$. And suppose that $c(w^*\|Shoot) = 0.4 > 0.3 = c(w^*)$, reflecting the fact that shooting helps to make w^* true. Then shooting has high prudential value in w^*:

$$\mathcal{V}_{Shoot}(w^*) = [c(w^*\|Shoot)/c(w^*)] \cdot u(w^*)$$
$$= [.4/.3] \cdot 10$$
$$\approx 13.3$$

The reason: shooting has a positive incremental causal impact on w^*: $c(w^*\|Shoot)/c(w^*) = .4/.3 \approx 1.33$. It helps to *make* w^* true. And w^* is a desirable state of the world. The upshot: shooting makes the world, w^*, satisfy your desires to a high degree. According to the theory of praxic good on offer, this is exactly what gives an action prudential value.

2.3 Epistemic Good

According to accuracy-first epistemology, what makes a credal state epistemically valuable at a world — the primary source of its all-epistemic-things-considered value — is its accuracy, or closeness to the truth at that world.

> *Epistemic Good.* A credal state *c* is epistemically valuable at a world *w* to the extent that *c* is close to the truth (accurate) at *w*.

In order for this idea to be useful in a formal decision theory, we'll need a more precise way of quantifying accuracy. The appropriate mathematical tools for this task are *epistemic scoring rules*, which can be thought of as *in*accuracy scores.[11]

Let $\mathrm{Prob}(\Omega)$ be the set of probability functions over the algebra Ω. An *inaccuracy score* is a function $\mathcal{I} : \mathrm{Prob}(\Omega) \times \mathcal{W} \rightarrow \mathbb{R}_{\geq 0}$ that measures how close a credence function *c* is to the truth if *w* is actual. If $\mathcal{I}(c, w) = 0$, then *c* is minimally inaccurate (maximally close to the truth) at *w*. Inaccuracy increases as $\mathcal{I}(c, w)$ grows larger.

Reasonable inaccuracy scores satisfy a range of constraints (*cf.* Joyce 1998, 2009, and Predd et al. 2009). For example, moving credences uniformly closer to the truth should always improve accuracy; if $|b(X) - w(X)| \leq |c(X) - w(X)|$ for all *X*, and $|b(Y) - w(Y)| < |c(Y) - w(Y)|$ for some *Y*, then $\mathcal{I}(b, w) < \mathcal{I}(c, w)$. Instead of detailing these constraints, though, we'll simply focus on one particularly attractive inaccuracy measure: the Brier score.[12]

Let *w* be a world, $X \in \Omega$, and χ_w be *w*'s characteristic function, *viz.*, $\chi_w(X) = 1$ if *X* is true at *w* and $= 0$ if *X* is false at *w*. The Brier Score is:

$$\mathrm{BS}(c, w) = \frac{1}{|\Omega|} \sum_{X \in \Omega} (\chi_w(X) - c(X))^2$$

---

[11]We use *in*accuracy instead of accuracy for technical convenience. Accuracy is simply negative inaccuracy.

[12]So long as the scoring rule is strictly proper, nothing we say below will hinge on this choice.

That is, the Brier Score identifies c's inaccuracy with its mean-squared divergence from truth-values at w.

With a more precise notion of inaccuracy in hand, we can fill in our schematic theory of epistemic good.

> *Epistemic Good**. A credal state c's epistemic value at a world w, $\mathcal{V}_c(w)$, is given by $-\mathcal{I}(c, w)$.

3   Rational Preference: Praxic and Epistemic

Our next task is to detail and defend a theory of epistemic preference. Such a theory specifies when an agent with credences c and evidence E ought to prefer one credal state b to another b^*. It specifies when, in view of E, she ought to see b as a preferable state to occupy to b^* (whether she or anyone else is currently, or will come to be in that state). When an agent (weakly) prefers b to b^*, we write $b \unrhd b^*$.

Our strategy is as follows. We will explore two ways of generalising SAVAGE that yield EDT and CDT as special cases. The first is Joyce's (1999; 2000; 2002). The second is our own. Both generalisations illuminate what is at issue between EDT and CDT, in a way that tells us something about their suitability for furnishing a theory of epistemic preference. But our generalisation provides *positive* advice too. It tells us how to use our theory of the epistemic good to arrive at the correct theory of epistemic preference.

On Savage's (1954) model, a decision-maker uses her credences about which *state of the world* is actual to choose between *actions* that produce more or less desirable *outcomes*. For expositional ease, we follow Jeffrey (1983) in thinking of states of the world and actions as propositions: elements of the partitions $S = \{S_1, ..., S_n\}$ and $\mathcal{A} = \{A_1, ..., A_m\}$, respectively.[13] The states S_i are the loci of her uncertainty. The actions A_i are, to a first approximation, the propositions whose truth-values she can (more or less directly) control.[14] The outcome of performing action A in state of world S, $o[A, S]$, is

the conjunction $A\&S$.[15] Importantly, for the agent's decision problem to be well-posed, outcomes must be grained finely enough to reflect everything that she cares about. Formally, this means: for any $A \in \mathcal{A}$ and $S \in S$, we have $u(w) = u(w')$ for all $w, w' \in o[A, S]$.

According to Savage, an agent should evaluate her options as follows:

> *Theory of (Praxic) Preference*: An agent ought to weakly prefer act A to B, $A \succeq B$,
>
> iff
>
> $$Est_c(A) \geq Est_c(B)$$

Typically, an agent's best estimate of A's utility (or any other quantity) is given by its expected value: $Est_c(A) = \sum_i c(S_i) \cdot u(o[A, S_i])$. (When harmless, we will talk directly of expectations. But, in certain pathological evidential circumstances, of the sort we examine in §4, estimates and expected values come apart.)

Savage also insisted — though this is not explicit in his formalism — that to properly apply the theory, probabilities of states must be independent of acts (Savage 1954, p. 73). Supposing that you perform act A should not change the credence that you assign to state S, for any $A \in \mathcal{A}$ and $S \in S$. Otherwise, SAVAGE would countenance absurdities such as this. When you face the following decision problem every evening:

	Eat Heartily	Go Hungry
Leave Oven Off	*Satisfied & Don't Pay for Gas*	*Unsatisfied & Don't Pay for Gas*
Turn Oven On	*Satisfied & Pay for Gas*	*Unsatisfied & Pay for Gas*

you ought to prefer (and choose) to leave your oven off. The reason: leaving your oven off *dominates* turning your oven on, relative to this partition of states of the world. It has a better outcome in every

---

[13] In the case of EpDT, we set $\mathcal{A} = \{\overline{c} \mid c$ a credence function over $\Omega\}$.

[14] A more sophisticated account of actions is necessary if we hope to countenance *epistemic* acts as actions proper. After all, whether we come to have one credal state or another is beyond our direct control. So a proposition describing the adoption of a credal state — an epistemic act — does not count as an action proper, according to Jeffrey's proposal. But this need not worry us. Whether or not epistemic acts count as actions proper, they are *evaluable* as such. Epistemic acts *cause* you to

occupy some new doxastic state; a state which represents the world. But they are not themselves representational. So they are not evaluable directly on the basis of their accuracy. Rather, like actions proper, they are valuable to the extent that they have a more or less desirable incremental causal impact on the world. And this, as we will see in §3-5, is the real reason they are not appropriate loci of evaluation for a theory of epistemic rationality.

[15] For Savage, outcomes are disjunctions of Jeffrey outcomes $A_i\&S_j$. This is unimportant for our purposes. All that matters is this: if $w \in A\&S$, then $w \in o[A, S]$. So we have: if $w \in A\&S$, then $u(w) = u(o[A, S])$.

state. So its unconditional expected utility is higher, whatever your credences are.

Properly understood, then, SAVAGE comes with the following caveat:

IND SAVAGE only applies if probabilities of states are independent of acts.

The reason SAVAGE comes with this caveat, Joyce argues, is that it evaluates actions from the wrong epistemic perspective. SAVAGE enjoins agents to evaluate actions from the perspective of their *unconditional* credences. But actions ought to be evaluated *on the supposition that they are performed*. They ought to be evaluated not from the perspective of one's unconditional credences c, but from the perspective of c updated on A. To do otherwise is to ignore relevant information. Let $c(\cdot \mid\mid\mid A)$ go proxy for the appropriately updated credence function, whatever it is. All decision theorists should agree, then:

Joycean General Theory of Preference: An agent ought to weakly prefer act A to B, $A \succeq B$,

iff

$$\sum_i c(S_i \mid\mid\mid A) \cdot u(o[A, S_i]) \geq \\ \sum_i c(S_i \mid\mid\mid B) \cdot u(o[B, S_i]).$$

What evidential and causal decision theorists will disagree on is this: what sort of supposition is appropriate for evaluating actions.

According to EDT, $c(\cdot \mid\mid\mid \cdot) = c(\cdot\mid\cdot)$. That is, the sort of supposition appropriate for evaluating actions is *indicative* supposition. Your indicative-conditional opinions reflect your views about which outcomes are likely to occur if you do, in fact, perform one act or another. And that's precisely the information that you ought to take into account in decision-making, according to EDT. So actions ought to be evaluated from the perspective of your conditional credences, $c(\cdot\mid\cdot)$, which capture your indicative-conditional opinions.

In contrast, CDT says: $c(\cdot \mid\mid\mid \cdot) = c(\cdot\|\cdot)$. That is, the sort of supposition appropriate for evaluating actions is *subjunctive* supposition. Your subjunctive-conditional opinions reflect your views about what sort of *causal influence* your actions will have. And *that* is the information that you ought to take into account in decision-making, according to CDT. So actions ought to be evaluated from the perspective of your imaged credences, $c(\cdot\|\cdot)$, which capture your subjunctive-conditional opinions.

These assumptions, together with Joyce's General Theory of Preference, yield the following:

EDT*'s Theory of (Praxic) Preference*: An agent ought to weakly prefer act A to B, $A \succeq B$,

iff

$$\sum_i c(S_i \mid A) \cdot u(o[A, S_i]) \geq \\ \sum_i c(S_i \mid A) \cdot u(o[B, S_i]).$$

CDT*'s Theory of (Praxic) Preference*: An agent ought to weakly prefer act A to B, $A \succeq B$,

iff

$$\sum_i c(S_i \| A) \cdot u(o[A, S_i]) \geq \\ \sum_i c(S_i \| A) \cdot u(o[B, S_i]).$$

Neither of these theories require that probabilities of states be independent of acts. Further, when states *are* independent of acts, they reduce to SAVAGE.[16]

Joyce's General Theory of Preference illuminates what is at issue between SAVAGE, EDT and CDT in a way that tells us something about their suitability for furnishing a theory of epistemic preference. They disagree about which epistemic perspective to adopt when evaluating actions. For the purposes of building out an accuracy-first epistemology, the important question is this: which theory (if any) — SAVAGE, EDT or CDT — identifies the right perspective for evaluating *epistemic states*, rather than *actions*?

Our theory of epistemic value seems to gesture toward an answer. Compare: the fact that actions are valuable or good to the extent that they *make* the world desirable *suggests* that we ought to evaluate actions from a perspective that reflects your causal opinions. It suggests that CDT's epistemic perspective is the right one for evaluating actions. But epistemic *states* are good (epistemically valuable) to the extent that *they* conform to the world. They are valuable in virtue of encoding an accurate picture of the world. They are *not* valuable in virtue of causally influencing the world, so as to *make* themselves accurate. So, it seems, you should *not* evaluate epistemic states from a perspective that reflects your views about the extent to which they will do *exactly that*, *viz.*, causally influence the world

---

[16]More carefully, when states are *evidentially* independent of acts, so that $c(S_i \mid A) = c(S_i)$, EDT reduces to SAVAGE: $\sum_i c(S_i \mid A) \cdot u(o[A, S_i]) = \sum_i c(S_i) \cdot u(o[A, S_i])$. When states are *causally* independent of acts, so that $c(S_i \| A) = c(S_i)$, CDT reduces to SAVAGE: $\sum_i c(S_i \| A) \cdot u(o[A, S_i]) = \sum_i c(S_i) \cdot u(o[A, S_i])$.

in good (accuracy-conducive) ways. You should instead evaluate them from a perspective that reflects your best estimates about the way the world is. You should evaluate them from the perspective of your *unconditional credences*. So SAVAGE's epistemic perspective is the right one for evaluating doxastic states.

If this is right, then the correct theory of epistemic preference is:

> *Theory of Epistemic Preference*: An agent ought to weakly prefer credal state p to q, $p \unrhd q$,
>
> iff
>
> $$Est_c(\mathcal{I}(p)) \leq Est_c(\mathcal{I}(q))$$

Here, $Est_c(\mathcal{I}(p))$ is c's estimate of p's inaccuracy. In normal cases, where estimates and expected values coincide, $Est_c(\mathcal{I}(p)) = \sum_w c(w) \cdot \mathcal{I}(p, w)$. So normally, an agent ought to prefer credal state p to q just in case she assigns lower unconditional expected inaccuracy to p than to q.

Despite our work so far, we still lacking a general theory of preference that does more than *suggest* what the right perspective is for evaluating epistemic states. It would be nice to have a theory that allows you to simply *plug in* a theory of the good (praxic or epistemic), and have the preferred perspective *fall out*. We will now provide such a theory.

3.1 The General Theory of Rational Preference

On our view, rational agents line up their praxic preferences — preferences over acts — with their best estimates of the prudential value or goodness of those acts. They also line up their epistemic preferences — preferences over doxastic states — with their best estimates of the epistemic value of those states. In particular:

> *Our General Theory of Preference*: An agent ought to weakly prefer act A to B, $A \succeq B$
>
> iff
>
> $$Est_c(\mathcal{V}(A)) \geq Est_c(\mathcal{V}(B))$$
>
> She ought to weakly prefer credal state p to q, $p \unrhd q$,
>
> iff
>
> $$Est_c(\mathcal{V}(p)) \geq Est_c(\mathcal{V}(q))$$

where $Est_c(\mathcal{V}(A))$ is the agent's estimate of A's prudential value, and $Est_c(\mathcal{V}(p))$ is her estimate of p's epistemic value.

The idea here is that an agent's best estimates of the value of various options *rationalize*, or *justify*, or provide *good reasons* for having the preferences that line up with (are represented by) those estimates. And she ought to prefer what she has most reason to prefer. Treating estimates as explanatorily basic, in this way, is nothing new. It has a long history in Bayesian epistemology and decision theory; a history which includes, *e.g.*, de Finetti, Jeffrey, and Joyce. Jeffrey (1986), for example, treats estimation as *the* basic concept in Bayesian epistemology, and defines probability in terms of it.

When an agent's best estimates and expectations coincide, which they typically will (save for in certain pathological evidential circumstances; see §4), our General Theory says that she should prefer an action A to B just in case $\sum_w c(w) \cdot \mathcal{V}_A(w) \geq \sum_w c(w) \cdot \mathcal{V}_B(w)$, and should prefer a credal state p to q just in case $\sum_w c(w) \cdot \mathcal{V}_p(w) \geq \sum_w c(w) \cdot \mathcal{V}_q(w)$.

According to *our* generalisation, what is at issue between SAVAGE, EDT and CDT is this: SAVAGE employs the wrong theory of praxic good.[17] EDT and CDT aim to rectify this, but disagree about what the right theory of the good is. They disagree, in the first instance, about which *quantity* to estimate, for the purposes of evaluating actions (in virtue of disagreeing about which quantity measures praxic goodness). The crux of their dispute is thus not about which epistemic perspective to estimate quantities (utility) from.

CDT agrees with our theory from §2.2:

> CDT*'s Theory of Praxic Good*. An action A's prudential value at a world w, relative to credences c and utilities u, is given by:
>
> $$\mathcal{V}_A(w) = [c(w\|A)/c(w)] \cdot u(w).$$

Actions are good to the extent that they *make* the world desirable.[18] EDT, in contrast, says:

---

[17]Better: SAVAGE identifies a mere constraint on the correct theory of praxic good, *viz.*, $\mathcal{V}_A(w) = u(w) = u(o[A, S])$ if act-state independence holds. It only partially specifies the correct theory of praxic good.

[18]It is worth noting that this account requires prudential value (or praxic good) to be measured on a ratio scale, and hence to have a theoretically significant zero point. This is as it should be. There are two factors which jointly conspire to make an action A prudentially valuable at a world w, according to CDT's Theory of Praxic Good: how large (or small) of an incremental

EDT*'s Theory of Praxic Good.* An action A's prudential value at a world w, relative to credences c and utilities u, is given by:

$$\mathcal{V}_A(w) = [c(w|A)/c(w)] \cdot u(w).$$

Actions are good to the extent that they provide good *evidence* (incremental evidential support) that the world is in a desirable state.[19]

_____

causal impact A has on w, and how desirable (or undesirable) w is. And one of those factors — incremental causal impact — *does* have a natural and theoretically significant zero point. An action A can have no smaller an incremental impact on any proposition X than it does on a contradiction $\perp$. (A can't help to make $\perp$ true *at all*, since $\perp$ is *necessarily* false.) So incremental causal impact, and in turn prudential value, *is* plausibly measured on a ratio scale.

[19]Both theories of praxic good — EDT's and CDT's, respectively — appear at first glance to be "doubly subjective" in an objectionable sort of way. Not only is an action A's prudential value at a world w (for an agent S) a function of one subjective quantity, *viz.*, the utility or degree of desirability of w according to S, $u(w)$. It is also a function of a second subjective quantity: the Bayes factor, $c(w|A)/c(w)$, and imaged Bayes factor, $c(w||A)/c(w)$, according to EDT and CDT, respectively.

This might strike you as odd. Informally, on our view, rational decision-making is a matter of choosing the action that you expect to have the most desirable incremental evidential/causal impact. But the formal story doesn't seem to match up. It says that rational decision-making proceeds *not* by estimating incremental evidential/causal impact and desirability *per se.* Rather, it proceeds by estimating some other quantity: a quantity whose value at a world w depends not only on w's degree of desirability, but also on *your own credence* that w is actual.

This oddness, however, is an artefact of presentation. For example, we might have cast CDT's theory of praxic good as follows (*mutatis mutandis* for EDT):

CDT*'s Theory of Praxic Good.* An action A's prudential value at a world w (for an agent S) is given by:

$$\mathcal{V}_A^*(w) = C_A(w) \cdot u(w).$$

where $C_A : \mathcal{W} \to \mathbb{R}$ is a random variable which maps each world w to a number $C_A(w)$ which measures A's incremental objective causal impact on w, and $u : \mathcal{W} \to \mathbb{R}$ is S's utility function, which maps each world w to a number $u(w)$ which measures w's degree of desirability according to S.

This presentation would have made it clear that rational decision-making *does* proceed by estimating incremental causal impact and desirability directly. But it would also have immediately collapsed into our current proposal without furnishing much additional insight. To see this, recall that on the Bayesian view, if an agent has credences c, then her best estimate of A's incremental causal impact, conditional on being in w, is given by the imaged Bayes factor, $c(w||A)/c(w)$, *i.e.*,

$$\sum_x c(C_A = x|w) \cdot x = c(w||A)/c(w)$$

where '$C_A = x$' is the set of worlds w' such that $C_A(w') = x$. Given this, estimating $\mathcal{V}_A$ — prudential value framed as a "doubly subjective" quantity — and $\mathcal{V}_A^*$, respectively, come to the same thing:

These theories of praxic good, together with our general theory of preference, yield CDT and EDT:[20]

CDT*'s Theory of (Praxic) Preference*: An agent ought to weakly prefer act A to B iff

$$\sum_w c(w) \cdot \mathcal{V}_A(w)$$
$$= \sum_i \sum_{w \in S_i} c(w) \cdot \mathcal{V}_A(w)$$
$$= \sum_i \sum_{w \in S_i} c(w) \cdot [[c(w\|A)/c(w)] \cdot u(w)]$$
$$= \sum_i \sum_{w \in S_i} c(w\|A) \cdot u(w)$$
$$= \sum_i c(S_i\|A) \cdot u(o[A, S_i])$$
$$\geq \sum_i c(S_i\|B) \cdot u(o[B, S_i])$$
$$= \sum_w c(w) \cdot \mathcal{V}_B(w)$$

for any partition of states of the world $S = \{S_1, ..., S_n\}$.

EDT*'s Theory of (Praxic) Preference*: An agent ought to weakly prefer act A to B iff

$$\sum_w c(w) \cdot \mathcal{V}_A(w)$$
$$= \sum_i \sum_{w \in S_i} c(w) \cdot \mathcal{V}_A(w)$$
$$= \sum_i \sum_{w \in S_i} c(w) \cdot [[c(w|A)/c(w)] \cdot u(w)]$$
$$= \sum_i \sum_{w \in S_i} c(w|A) \cdot u(w)$$
$$= \sum_i c(S_i|A) \cdot u(o[A, S_i])$$
$$\geq \sum_i c(S_i|B) \cdot u(o[B, S_i])$$
$$= \sum_w c(w) \cdot \mathcal{V}_B(w)$$

for any partition of states of the world $S = \{S_1, ..., S_n\}$.

This puts us in a better position to explain why CDT provides the right epistemic perspective for evaluating actions. Before, we said: the fact that actions are valuable or good to the extent that they *make* the world desirable *suggests* that you ought to evaluate actions from a perspective that reflects your causal opinions. Now we can say: the right theory of praxic good, *viz.*, $\mathcal{V}_A(w) = [c(w\|A)/c(w)] \cdot u(w)$ *entails* that CDT's epistemic perspective is the right one for evaluating actions, given our general theory of preference. It *entails* that you ought to evaluate actions by $\sum_i c(S_i\|A) \cdot u(o[A, S_i])$.

_____

$$
\begin{aligned}
Exp_c(\mathcal{V}_A^*) &= \sum_w c(w) \cdot C_A(w) \cdot u(w) \\
&= \sum_w \sum_x c(w \& C_A = x) \cdot x \cdot u(w) \\
&= \sum_w c(w) \cdot \left[\sum_x c(C_A = x|w) \cdot x \right] \cdot u(w) \\
&= \sum_w c(w) \cdot [c(w\|A)/c(w)] \cdot u(w) \\
&= Exp_c(\mathcal{V}_A)
\end{aligned}
$$

[20]We simply assume regularity, for ease of exposition; $c(w) > 0$ for all $w \in \mathcal{W}$ and that estimates and expected values coincide.

It also puts us in a better position to explain why you should *not* evaluate credal states from a perspective that reflects your views about how occupying those states will causally influence the world. The reason: credal states are *not* valuable in virtue of causally influencing the world, so as to *make* themselves accurate. You should *not* measure the epistemic value of credal state p at world w by $\mathcal{V}_p(w) = -[c(w\|\overline{p})/c(w)] \cdot \mathcal{I}(p,w)$, where $\overline{p}$ is the epistemic act of adopting credal state p. Instead, credal states are good to the extent that *they* conform to the world. According to the *correct* theory of epistemic good, $\mathcal{V}_p(w) = -\mathcal{I}(p,w)$. Together with our general theory of preference, this entails:

> *Theory of Epistemic Preference*: An agent ought to weakly prefer credal state p to q, $p \trianglerighteq q$,
>
> iff
>
> $\sum_w c(w) \cdot \mathcal{I}(p,w) \leq \sum_w c(w) \cdot \mathcal{I}(q,w).$

Or to put matters fully generally:[21]

> *Theory of Epistemic Preference*: An agent ought to weakly prefer credal state p to q, $p \trianglerighteq q$,
>
> iff
>
> $Est_c(\mathcal{I}(p)) \leq Est_c(\mathcal{I}(q))$

---

[21]Greaves objects to this way of spelling out EpDT on the grounds that

> The predictions of 'Savage' EpDT depend on the state partition. … If Savage EpDT 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 [expected epistemic utility] with a principle identifying the correct state partition; indeed, this is the course taken by causal EpDT.… But other ways out are available too; each amounts to the replacement of our naïve theory with a different theory. (2013, p.12, minor changes)

Two points in response. Firstly, *no* plausible epistemic decision theory — Savage-style, evidential, causal, or deliberational — is partition invariant. To see why, suppose that you are certain of X. Then the expected epistemic utility of your credence function restricted to $\{X, \neg X\}$ is maximal (your expected inaccuracy is 0). But this, of course, will not hold in general, for any partition whatsoever. If $0 < c(Y) < 1$ for some Y, then the expected epistemic utility of your credence function restricted to $\{Y, \neg Y\}$, according to any reasonable measure, is non-maximal (your expected inaccuracy is greater than 0).

Secondly, on the accuracy-first approach, epistemic norms have their binding force in virtue of a being a good means to the end of epistemically valuable *total* doxastic states. Total doxastic states are the relevant loci of epistemic evaluation. So there *is* a correct state partition, *viz.*, the set of atoms w of Ω. The reason: the atoms w are exactly the propositions determine the truth-values for *every* proposition in Ω, and hence fix the overall inaccuracy (epistemic value) of any total credal state $c : \Omega \to \mathbb{R}$.

We will now use our theory of epistemic preference to explain why accuracy-first epistemology does not sanction epistemic bribe-taking. Before we proceed, though, it is worth contrasting this theory of preference with a closely related one. Suppose that what you care about all things considered is just accuracy — not money, prestige, or fame. Then our theory of praxic preference says:

> *Theory of Preference over Epistemic Acts*: An agent ought to weakly prefer $\overline{p}$ to $\overline{q}$, $\overline{p} \succeq \overline{q}$,
>
> iff
>
> $\begin{aligned} &\sum_w c(w) \cdot \mathcal{V}_{\overline{p}}(w) \\ =&\sum_w c(w) \cdot [[c(w\|\overline{p})/c(w)] \cdot u(w)] \\ =&-\sum_w c(w) \cdot [[c(w\|\overline{p})/c(w)] \cdot \mathcal{I}(p,w)] \\ =&-\sum_w c(w\|\overline{p}) \cdot \mathcal{I}(p,w) \\ \geq&-\sum_w c(w\|\overline{q}) \cdot \mathcal{I}(q,w) \\ =&\sum_w c(w) \cdot \mathcal{V}_{\overline{q}}(w) \end{aligned}$
>
> iff
>
> $\begin{aligned} &\sum_w c(w\|\overline{p}) \cdot \mathcal{I}(p,w) \\ \leq&\sum_w c(w\|\overline{q}) \cdot \mathcal{I}(q,w). \end{aligned}$

Epistemic *acts*, like actions more generally, are good to the extent that they *make* the world desirable. Epistemic *states* are not. As a result, a rational agent's preferences over epistemic acts and states will *not*, in general, coincide. Greaves' concerns about accuracy-first epistemology result from running these very different sorts of evaluations — evaluations of epistemic states and acts — together. Carefully separating them out is the key to seeing that accuracy-first epistemology does *not* sanction epistemic bribe-taking.

4 Leap, Promotion and Imps

4.1 Estimates and Expectations

Before analyzing the cases presented above, we'll need two additional principles relating rational preference and chance, since both play an important role in the cases under discussion. The first is the familiar:

PRINCIPAL PRINCIPLE An agent with evidence E ought to have a credence function $c : \Omega \to \mathbb{R}$ such that $c(X|\phi_{ch}) = ch(X|E)$, for all $X \in \Omega$ and all ch with $c(\phi_{ch}) > 0$ in the set of possible ur-chance functions C, where ϕ_{ch} is the proposition that ch is the true chance function.

While variants like the New Principle improve on the Principal Principle, we'll be using the latter primarily

for expositional ease, since no added nuance is needed in what follows. Furthermore, since some chance-credence norm like PP is nearly universally accepted in the literature, we won't argue for it here.[22]

For reference below, we note that an agent with credence function c who follows the PRINCIPAL PRINCIPLE can calculate expected inaccuracy with any of the formulæ below:

$$\sum_w c(w) \cdot \mathcal{I}(p, w)$$
$$= \sum_w \sum_{ch \in C} [c(w|\phi_{ch})c(\phi_{ch})] \cdot \mathcal{I}(p, w)$$
$$= \sum_{ch \in C} c(\phi_{ch}) \cdot \left[\sum_w c(w|\phi_{ch}) \cdot \mathcal{I}(p, w) \right]$$
$$= \sum_{ch \in C} c(\phi_{ch}) \cdot \left[\sum_w ch_E(w) \cdot \mathcal{I}(p, w) \right]$$

The second principle relates chance and epistemic preference:

DEFERENCE TO CHANCE If an agent with credences c and evidence E is such that:

(i) $\sum_w c(w) \cdot \mathcal{I}(p, w) = x$

but she is also certain that the chance function ch is such that:

(ii) $\sum_w ch_E(w) \cdot \mathcal{I}(p, w) = y$, with $x \neq y$

in which case she violates the PRINCIPAL PRINCIPLE, then nonetheless:

(iii) $Est_c(\mathcal{I}(p)) = y$.

That is, she ought to line up her own best estimate of p's inaccuracy with what she knows p's objective expected inaccuracy to be.

It's always the case, on our theory of epistemic preference, that an agent ought to prefer the credal state that in her best estimate is least inaccurate. Normally, her best estimate of a state's inaccuracy *just is* her expected value of its inaccuracy. However, in special pathological cases like the ones we'll be

considering below, the agent may have knowingly diverged from chance in order to secure herself lower overall inaccuracy. For instance, she may know that $ch_E(w) = 1$ while her own credence $c(w) = 0$. In that case, she should recognize that the expected inaccuracy of an epistemic state p as calculated in the traditional way (i.e., $\sum_w c(w) \cdot \mathcal{I}(p, w)$) is not actually her best estimate of p's inaccuracy. Instead, since she knows the salient objective probabilities, she ought to defer to those objective probabilities, and line up her best estimate with the expected value *chance conditional on her evidence* assigns (i.e., $\sum_w ch_E(w) \cdot \mathcal{I}(p, w)$).[23]

The reason to defer to chance in this way, by satisfying DtC, is exactly the reason given by Pettigrew (2013) for satisfying the PRINCIPAL PRINCIPLE: the set of estimates that result from substituting chance's expectations for your own when the two diverge *chance-dominates* your original set of expectations. Every possible ur-chance function conditional on your evidence expects the amended set of estimates to be more accurate.[24]

With these additions in hand, we can devote the remainder of the section to the analysis of cases above.

4.2 Imps

While we already gave a preliminary diagnosis of what's happening in IMPS in the Introduction, we'll return to it here for more in depth analysis in light of the theory developed above.

Let's first consider what Emily's evidence E includes. From the description, we know E entails both C_0 and $ch(C_j|c(C_0) = x) = (1 - 0.5x)$. We'll also assume— here and throughout—that Emily's credences are *luminous*. That is, she can tell what her credence function c is. Therefore, E also includes $c(C_0) = x$.

To understand the full range of Emily's epistemic

[22]One may wonder why an accuracy-firster would endorse the PRINCIPAL PRINCIPLE. For an extended discussion see Pettigrew 2013.

[23]Objection: It's easy to cook up a case like IMPS in which an agent can attain overall more accuracy if she intentionally forms inaccurate credences about what the chances themselves are. In such a situation, DEFERENCE TO CHANCE will get the wrong answer. Reply: We don't claim DEFERENCE TO CHANCE is a *fully* general principle for forming estimates. Instead, it improves the usual identification of estimates with expected values to correctly cover the class of cases we'll be studying.

[24]This is a straightforward consequence of the fact that (i) DtC only applies when every possible ur-chance function conditional on the evidence E agrees that p's expected inaccuracy is y, and (ii) every possible chance function agrees that moving any set of estimates uniformly closer to its own expectations — *e.g.*, by moving $Est_c(\mathcal{I}(p))$ closer to y and keeping all other estimates fixed — decreases the expected inaccuracy of that set of estimates (so long as inaccuracy is measure by a strictly proper scoring rule).

options, we'll evaluate her epistemic *doppelgängers*. The idea here is that we put agents with different credence functions in the same case and then evaluate the epistemic state and behaviour of each.

Here, we'll let Em_x be the Emily doppelgänger who adopts credal state c_x, which is such that $c_x(C_0) = x$ and $c_x(C_j) = 1 - 0.5x$ for all $j \geq 1$. So, for example,

- $c_{.8}$: $c_{.8}(C_0) = .8$ and $c_{.8}(C_j) = .6$.

- $c_{.1}$: $c_{.1}(C_0) = .1$ and $c_{.1}(C_j) = .95$.

- c_0: $c_0(C_0) = 0$ and $c_0(C_j) = 1$.

By considering each Em_x we can now determine what EpDT's verdicts are for any epistemic state Emily might be in and for any epistemic action she may have taken.

Regarding epistemic *states*, the question is: How should Em_x evaluate her own credences? Should she prefer her own credal state over the alternatives in light of E? Or should she prefer some other credal state?

The answer: Em_1 ought to prefer her own credal state c_1 to all alternatives b. Since c_1 satisfies PP,

$$Est_{c_1}(\mathcal{I}(b)) \quad = \sum_w c_1(w) \cdot \mathcal{I}(b,w)$$
$$= \sum_{ch' \in C} c_1(\phi_{ch'}) \cdot [\sum_w ch'_E(w) \cdot \mathcal{I}(b,w)].$$

Note also that, since she's sure her credences match the chances — $c_1(\phi_{ch}) > 0$ only if $ch_E(C_0) = c_1(C_0) = 1$ and $ch_E(C_j) = c_1(C_j) = 1/2$, for $j \geq 1$ — we have:

$$Est_{c_1}(\mathcal{I}(b)) \quad = \sum_w c_1(w) \cdot \mathcal{I}(b,w)$$
$$= \sum_{ch' \in C} c_1(\phi_{ch'}) \cdot [\sum_w ch'_E(w) \cdot \mathcal{I}(b,w)]$$
$$= \sum_w ch_E(w) \cdot \mathcal{I}(b,w).$$

Finally, recall that we're identifying $\mathcal{I}$ with the Brier score. The Brier score is a 'strictly proper' scoring rule, *i.e.*, $\sum_w p(w) \cdot \mathcal{I}(p,w) < \sum_w p(w) \cdot \mathcal{I}(q,w)$, for any probabilistically coherent credence function p and any $q \neq p$.[25] So we have:

[25] The Brier score is also *separable* (*cf.* Joyce 2009, p. 271). Separability guarantees that the inaccuracy of c_1's credences over the C_i, as well as the inaccuracy of all alternatives b, can be assessed independently of what probabilities they assign to propositions other than the C_i. Further, since c_1 and b agree on all other probabilities (*ex hypothesi*), their comparative accuracy over the C_i is all that matters to comparative accuracy *tout court*. So we restrict attention to the propositions of interest—the C_i—in what follows.

$$Est_{c_1}(\mathcal{I}(c_1)) \quad = \sum_w c_1(w) \cdot \mathcal{I}(c_1,w)$$
$$= \sum_w ch_E(w) \cdot \mathcal{I}(ch_E,w)$$
$$< \sum_w ch_E(w) \cdot \mathcal{I}(b,w)$$
$$= \sum_w c_1(w) \cdot \mathcal{I}(b,w)$$
$$= Est_{c_1}(\mathcal{I}(b))$$

for all $b \neq c_1 = ch_E$. Hence, by our theory of epistemic preference, Em_1 ought to prefer her own credal state c_1 to any alternative credal state b, *i.e.*, $c_1 \rhd b$.

On the other hand, if $x \neq 1$, Em_x ought to prefer ch_E — which is such that $ch_E(C_0) = 1$ and $ch_E(C_j) = 1 - 0.5x$ — to her own credal state, c_x. Since she's certain that the true chance function ch is such that:

$$\sum_w ch_E(w) \cdot \mathcal{I}(ch_E,w) < \sum_w ch_E(w) \cdot \mathcal{I}(b,w)$$

for all $b \neq ch_E$ (including $b = c_x$), by Deference to Chance (DtC), we have: $Est_{c_x}(\mathcal{I}(ch_E)) < Est_{c_x}(\mathcal{I}(b))$. Hence, by our theory of epistemic preference, $ch_E \rhd b$.

Less formally, since Emily ought to prefer to be in whatever credal state is, in her best estimate, most valuable, *i.e.*, most accurate; and since she treats chance's best estimates as her own (by DtC); and moreover since she is certain that the true chance function ch_E estimates itself to be most accurate, she ought to prefer ch_E — a credal state which assigns 1 to the proposition that there is a child before her — to any other credal state b, including her own, c_x.

Regarding epistemic *acts*, the question is: Assuming she cares only about accuracy, how should Em_x evaluate the *epistemic act* that she performed? How should she evaluate the *action $\overline{c_x}$* of adopting credal state c_x? Should she prefer it over the alternative epistemic acts she might have performed?

The answer: Em_0, and indeed all Em_x, ought to prefer epistemic act $\overline{c_0}$ to the alternatives $\overline{c_y}$. First reason:

$$\sum_w c_x(w\|\overline{c_y}) \cdot \mathcal{I}(c_y,w) = \sum_w ch_E(w|\overline{c_y}) \cdot \mathcal{I}(c_y,w)$$

Em_0 thinks that, were she to raise her credence in C_0 (the proposition that there is now a child before her) from 0 to y, for some $y > 0$, the j^{th} summerhouse child would be less likely to come outdoors (C_j would be less likely). Indeed, she's sure that C_j would be exactly this likely: $ch_E(C_j|\overline{c_y}) = 1 - 0.5y$ ($< c_0(C_j) = 1$). So $c_0(C_j\|\overline{c_y}) = ch_E(C_j|\overline{c_y}) = 1 - 0.5y$. More generally, $c_x(X\|\overline{c_y}) = ch_E(X|\overline{c_y})$. Hence:

$$\sum_w c_x(w\|\overline{c_y}) \cdot \mathcal{I}(c_y,w) = \sum_w ch_E(w|\overline{c_y}) \cdot \mathcal{I}(c_y,w)$$

13

From this, we can derive:[26]

$$\sum_w c_x(w \| \overline{c_y}) \cdot \mathcal{I}(c_y, w)$$

$$= \sum_w ch_E(w | \overline{c_y}) \cdot \mathcal{I}(c_y, w)$$

$$= \sum_w ch_E(w | \overline{c_y}) \cdot \sum_{i=0}^{10} (\chi_w(C_i) - c_y(C_i))^2$$

$$= (1-y)^2 + \sum_{k=0}^{10} \binom{10}{k} \left(1 - \frac{y}{2}\right)^k \left(\frac{y}{2}\right)^{10-k}$$

$$\bullet \left[k\left(\frac{y}{2}\right)^2 + (10-k)\left(1 - \frac{y}{2}\right)^2 \right]$$

$$= -\frac{3y^2}{2} + 3y + 1$$

which is uniquely minimized at $y = 0$. So, by our theory of preference over epistemic acts, $\overline{c_0} \succ \overline{c_y}$.

Less formally, since Emily ought to prefer to perform whatever epistemic action is, in her best estimate, most valuable — *i.e.*, will produce the most accuracy (assuming that she cares exclusively about accuracy) — and she is certain that dropping her credence that there's a child before her down to 0 will impact the chances in the best-possible way — *i.e.*, that way that produces the highest objective expected accuracy — she ought to prefer that epistemic action to any other.

Before elaborating on the dissonance between the evaluation of states and the evaluation of acts in this case, consider Greaves' own commentary on the case. We discuss the disparate intuitions that we label [1] and [2] below:

> ...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_0 that indeed there is** [1]. In that case, there will be a chance of $1/2$ of each summerhouse child coming out to play, so she should have credence $1/2$ 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_i, 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_0, and adopt (at the other extreme) credence 0 in C_0, then, by adopting degree of belief 1 in each C_j (j = 1, …, 10), she can guarantee a perfect match to the remaining truths**. Is it epistemically rational to accept this 'epistemic bribe'? (2013, p. 4; emphasis ours)

The two different intuitions here track perfectly the two different modes of analysis that EpDT provides. Intuition [1] is driven by Emily's evaluations of credal *states*, not epistemic *acts*. *Every* Emily doppelgänger Em_x should prefer to have a credal state that assigns probability 1 to C_0. Em_1 should prefer her own credal state c_1 to all alternatives b. Em_x should prefer the chance function conditional on her evidence (i.e., $ch_E(C_0) = 1$ and $ch_E(C_j) = 1 - .5x$) to all her alternatives, including her own credal state c_x.

Intuition [2] is driven by Emily's evaluations of *epistemic acts*, not *states*. *Every* Emily doppelgänger Em_x, if she cares exclusively about accuracy, should prefer to perform the *act* of adopting credal state c_0. This is the act that *causes* the world to satisfy her desires to the greatest degree possible. $\overline{c_0}$ influences the truth-values of the C_j in just the right way, so as to make them as close as possible to her credences.

4.3 Leap

Another telling case Greaves (2013) provides is the following:

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 S would then be x.

---

[26]The following observations should clarify the fourth line of this derivation:

1. $ch_E(w | \overline{c_y}) > 0$ only if $\chi_w(C_0) = 1$.
   - In turn, $ch_E(w | \overline{c_y}) > 0$ only if $ch_E(w | \overline{c_y}) \cdot \mathcal{I}(c_y, w) = ch_E(w | \overline{c_y}) \cdot [(1-y)^2 + \mathcal{I}(c_y|_{j\geq 1}, w)]$, where $c_y|_{j\geq 1}$ is the restriction of c_y to $\{C_j | j \geq 1\}$.

2. Given (1), $\sum_w ch_E(w | \overline{c_y}) \cdot \mathcal{I}(c_y, w) = (1-y)^2 + \sum_{k=0}^{10} ch_E(\# = k | \overline{c_y}) \cdot \mathcal{I}(c_y|_{j\geq 1}, \# = k)$, where $\# = k$ is the proposition that exactly k of summerhouse children 1 through 10 come outdoors.

3. $ch_E(\# = k | \overline{c_y}) = \binom{10}{k} \left(1 - \frac{y}{2}\right)^k \left(\frac{y}{2}\right)^{10-k}$

4. $\mathcal{I}(c_y|_{j\geq 1}, \# = k) = k\left(\frac{y}{2}\right)^2 + (10-k)\left(1 - \frac{y}{2}\right)^2$

In this case, Bob's evidence E includes $ch(S|c(S) = x) = x$, and because of our luminosity assumption, E also includes $c(S) = x$.

So we can see what EpDT says for any possible credence function Bob might have, we'll let B_x be the Bob-doppelgänger who adopts credal state c_x with $c_x(S) = x$.

Regarding epistemic *states*, the question is: How should B_x evaluate his own credences? Should he prefer his own credal state over the alternatives in light of E? Or should he prefer some other credal state?

The answer: Every B_x ought to prefer his own credal state c_x to all alternatives c_y. Since c_x satisfies PP, and is certain that $c_x = ch_E$, we have:

$$
\begin{aligned}
Est_{c_x}(\mathcal{I}(c_x)) \quad &= \textstyle\sum_w c_x(w) \cdot \mathcal{I}(c_x, w) \\
&= \textstyle\sum_w ch_E(w) \cdot \mathcal{I}(ch_E, w) \\
&< \textstyle\sum_w ch_E(w) \cdot \mathcal{I}(c_y, w) \\
&= \textstyle\sum_w c_x(w) \cdot \mathcal{I}(c_y, w) \\
&= Est_{c_x}(\mathcal{I}(c_y))
\end{aligned}
$$

for all $c_y \neq c_x = ch_E$. Hence, by our theory of epistemic preference, $c_x \vartriangleright c_y$.

Informally: since Bob ought to prefer to be in whatever credal state is, in his best estimate, most valuable — *i.e.*, most accurate — and since he estimates his own credences (which he is sure agree with the chances) to be most accurate, he ought to prefer his own credal state to any other.

Regarding epistemic *acts*, the question is: Assuming he cares only about accuracy, how should B_x evaluate the *epistemic act* that he performed? How should he evaluate the *action* $\overline{c_x}$ of adopting credal state c_x? Should he prefer it over the alternative epistemic acts he might have performed?

The answer: Every B_x ought to prefer epistemic acts $\overline{c_0}$ and $\overline{c_1}$ to all alternatives $\overline{c_x}$. The reason:

$$
\begin{aligned}
&\sum_w c_x(w \| \overline{c_y}) \cdot \mathcal{I}(c_y, w) \\
&= \sum_w ch_E(w | \overline{c_y}) \cdot \mathcal{I}(c_y, w) \\
&= y(1 - y)^2 + (1 - y)(0 - y)^2 \\
&= y(1 - y)
\end{aligned}
$$

which is minimized only at $y = 0$ and $y = 1$.

Informally: since Bob ought to prefer to perform whatever epistemic action is, in his best estimate, most valuable — *i.e.*, will produce the most accuracy — and he is certain that either dropping his credence

that he'll clear the chasm to 0, or raising it to 1, will impact the chances in the best possible way — *i.e.*, that way that produces the highest objective expected accuracy — he ought to prefer either of those epistemic actions to any other.

Again, Greaves identifies two dissonant intuitions:

> 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** [1]. **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** [2]. Is any degree of belief in success epistemically rationally permissible, or only an extremal credence? (2013, p. 2; emphasis ours)

Intuition [1] is driven by Bob's evaluations of *epistemic acts*, not *states*. *Every* Bob doppelgänger B_x, if he cares exclusively about accuracy, should prefer to perform the *act* of adopting credal state c_0 or c_1. These are the acts that *cause* the world to satisfy his desires to the greatest degree possible. $\overline{c_0}$ and $\overline{c_1}$ influence the truth-value of S in just the right way, so as to make it as close as possible to his credence. As a result, $\overline{c_0}$ and $\overline{c_1}$ are, in his best estimate, more prudentially valuable than all alternatives $\overline{c_x}$.

Intuition [2] is driven by Emily's evaluations of credal *states*, not epistemic *acts*. *Every* Bob doppelgänger B_x should prefer his own credal state c_x over all the alternatives c_y.

4.4 Promotion

For our third case study, we turn to:

PROMOTION Alice is up for promotion. Her boss, however, is a deeply insecure type: he's more likely to promote Alice if she comes across as lacking in confidence. Furthermore, Alice is useless at play-acting, so she'll come across that way iff she really does have a low degree of belief that she's going to get the promotion. Specifically, the chance 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'll be promoted.

Given the setup, Alice's evidence E includes $ch(P|c(P) = x) = 1 - x$. By luminosity, E also

includes $c(P) = x$.

Let A_x be the Alice-doppelgänger who adopts credal state c_x with $c_x(P) = x$.

Regarding epistemic *states*, the question is: How should A_x evaluate her *credal state*? Should she prefer it over the alternatives in light of E? Or should she prefer some other credal state?

The answer: If $x = .5$, she ought to prefer her own credal state $c_{.5}$ to all alternatives b. But If $x \neq .5$, she ought to prefer ch_E, which is such that $ch_E(P) = 1 - x$. The reason:

$$\sum_w ch_E(w) \cdot \mathcal{I}(c_y, w) = (1 - x)(1 - y)^2 + xy^2$$

which is uniquely minimized at $y = 1 - x$. By DtC, then, we have $Est_{c_x}(\mathcal{I}(c_{1-x})) < Est_{c_x}(\mathcal{I}(c_x))$ if $x \neq .5$. Hence, by our theory of epistemic preference, $c_{1-x} \rhd c_x$.

Informally: since Alice ought to prefer to be in whatever credal state is, in her best estimate, most valuable, *i.e.*, most accurate; and since she treats chance's best estimates as her own (by DtC); and moreover since she is certain that the true chance function ch_E estimates itself to be most accurate, she ought to prefer ch_E — a state which assigns $1-x$ to the proposition that she'll be promoted, rather than x — to any other credal state c_y, including her own, c_x.

Regarding epistemic *acts*, the question is: Assuming all she cares about is accuracy, what should A_x think of the *epistemic act* she ended up performing? That is, what should she think of the *action* $\overline{c_x}$?

The answer: A_x prefers epistemic act $\overline{c_{.5}}$ to any alternative $\overline{c_y}$. The reason:

$$\sum_w ch_E(w | \overline{c_y}) \cdot \mathcal{I}(c_y, w)$$
$$= (1 - y)\mathcal{I}(c_y, w_P) + y \cdot \mathcal{I}(c_y, w_{\neg P})$$
$$= (1 - y)(1 - y)^2 + y \cdot y^2$$

which is uniquely minimized at $y = .5$.

Informally: since Alice ought to prefer to perform whatever epistemic action is, in her best estimate, most valuable — *i.e.*, will produce the most accuracy — and since she is certain that setting her credence that she'll be promoted to exactly $1/2$ will impact the chances in the best-possible way — *i.e.*, that way that produces the highest objective expected accuracy — she ought to prefer that epistemic action to any other.

At first glance, it appears that this case lacks any dissonance. Indeed, Greaves thinks that EpDT's recommendation is clear:

> Presumably, in the Promotion case, there is a unique rationally permitted degree of belief in P: Alice must adopt credence .5 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. (2013, p. 2)

Nevertheless, we maintain that matters are not so simple. If we evaluate Alice's *epistemic acts*, then adopting credence .5 will minimize her expected inaccuracy. Every other option, in her best estimate, *causes* her credence to be less accurate. So, $A_{.5}$ is best from this perspective.

On the other hand, if we evaluate Alice's *credal states*, the problem is more nuanced. If she is in state x, she most prefers to be in state $1 - x$. So, for any $x \neq .5$, the optimal credal state to be in, by Alice's lights, is *not* in fact .5. Which alternative state is optimal varies based on which credal state Alice is in. So, while Greaves does not here identify a pull in two different directions, there in fact is one. From the act point of view, adopting credence .5 is uniquely best, regardless of what credal state c_x Alice occupies. From the state point of view, occupying state c_{1-x} is best.

It is true that $A_{.5}$ (and only $A_{.5}$) is in a conflict-free mental state. $A_{.5}$ should prefer both the *credal state* that she adopted and the *epistemic act* that she performed to all the alternatives. Every other A_x should prefer some other state *and* some other act over her own. It does *not* follow from this, however, that accuracy-first epistemology straightforwardly recommends the *state* $c_{.5}$, since not every A_x most prefers to be in that state.

5   What EpDT Recommends

We've now identified the source of the problem: When the act of adopting a credal state can influence the world, which epistemic acts an agent wants to perform can come apart from which credal states she'd most like to occupy. That is, an agent can prefer to perform epistemic act $\overline{c}$ to an alternative epistemic act $\overline{c'}$ while preferring to be in epistemic state c' to state c.

So, given this dissonance, what does epistemic decision theory recommend in the end? Does it advise agents to go with c or with c'?

In one sense, EpDT equivocates in these cases. It recommends both performing act $\bar{c}$ and occupying state c'. On the face of it, this might seem problematic. These recommendations are not cosatisfiable. You cannot both perform act $\bar{c}$ and occupy state c'.

However, following our discussion of praxic and epistemic good in §2, we nonetheless maintain that EpDT's recommendations concerning which states to occupy are of primary concern to the normative epistemologist, for it is only here that EpDT is returning purely epistemic evaluations and only here that EpDT has implications for purely epistemic rationality. There is a kind of rational dilemma, but not a dilemma of purely epistemic rationality. Instead, these are cases where epistemic rationality and what is ultimately practical rationality come apart. And *this* kind of dilemma is familiar. For example, it might be *epistemically* rational for you to believe that your partner is cheating on you, in light of your evidence, even though it is *practically* rational for you to perform the act of adopting the belief that he is not (if you can).

To see why EpDT's recommendations concerning *states* are the only ones relevant for *epistemic* rationality, recall how we evaluate the epistemic *act* of adopting credence c' from c's point of view:

$$\text{EEU}_c(\overline{c'}) = -\sum_w c(w\|\overline{c'})\mathcal{I}(c',w) \qquad (1)$$

Equation (1) is merely an instance of the more general causal decision-theoretic method of evaluation of *any* action whatsoever. I.e., (1) is a special case of:

$$\text{EU}_c(A) = \sum_w c(w\|A) \cdot u(w) \qquad (2)$$

There are two ways (1) restricts (2) in particular. First, (1) restricts the domain of actions to merely epistemic actions. Second, it identifies $-\mathcal{I}$ with the utility function u. Thus, EpDT evaluates epistemic acts just as causal decision theory does for an agent whose only concern in life is accuracy. That is, if what you care about all things considered is just accuracy—and not money, prestige, or fame—EpDT and causal decision theory will tell you to perform the same epistemic acts.

But the restriction to epistemic acts is arbitrary. (Worse, it verges on *incoherence*.[27]) There is no

reason why we can't evaluate the expected epistemic utility of building the Large Hadron Collider, or reading your sister's diary, or choosing Lucky Charms over Cap'n Crunch for breakfast. Each of these may effect changes in the world as well as in your epistemic state, and they can therefore be evaluated in terms of the expected accuracy they'll deliver in the same way EpDT evaluates epistemic acts.

We can agree that failing to build the LHC, respecting your sister's privacy, and foregoing breakfast altogether are not in themselves epistemically irrational in the sense we're after. Nonetheless they may lead to less accuracy than you could have achieved through other means and are therefore not the optimal acts for an agent whose only concern is how close her credal state is to actual truth-values. The same, we suggest, is true of 'epistemic' acts. Performing 'the' act (supposing that there is such a thing) of adopting some credence function c may not be *epistemically* irrational, even if adopting c leads to less accuracy than you could have achieved through other means.

The reason: while the utility function $-\mathcal{I}$ cares only for your epistemic and not your practical well-being, EpDT nonetheless *evaluates* epistemic acts like any other act, *viz.*, on the basis of the extent to which it *produces* desirable consequences. But this is a *practical* evaluation. It is a practical evaluation whether you care primarily about accuracy or apple pie. And these sorts of evaluations have no bearing on epistemic rationality *per se*. The direction of fit is wrong.

On the other hand, EpDT evaluates epistemic states based solely on how well they fit the world, *i.e.*, based solely on how epistemically good they are. Such verdicts don't tell you to change the world, but merely what states are best to occupy given the way the world is. They provide the sorts of evaluations that a pure observer — one who sees no particular value in influencing the system she is investigating — would use to gauge how successful her inquiry has been. It's then EpDT's recommendations on

---

[27]The very notion of an epistemic act itself is problematic for EpDT. Imagine the following variant of Imps: In order to change from one credence function c to another c', Emily can run one of two cognitive processes $A_{c'}$ and $B_{c'}$. $A_{c'}$ and $B_{c'}$ function like mental switches that Emily can turn on that result in her adoption of c'. Generally, it doesn't matter at all which switch she flips. She simply has two different means of getting herself into state c'. However, in our redux version of Imps, it makes a difference. In particular, if she performs $\overline{c'}$ by initiating $A_{c'}$, then the chance of $C_1,\ldots,C_{10}$ is $1 - \frac{c'(C_0)}{2}$ as before. If she performs $\overline{c'}$ by initiating $B_{c'}$, the chance of $C_1,\ldots,C_{10}$ is $1/2$. Now, EpDT seems to recommend the *act* of $\overline{c_0}$ performed via A_{c_0}, not the act $\overline{c_0}$ *simpliciter*. Indeed, the latter is no longer fine-grained enough to have a place in the space of acts $\mathcal{A}$, since it will lead to different outcomes depending on how it is effected.

which epistemic states to prefer that concern purely epistemic rationality:

STATE-BASED ACCOUNT OF EP RAT. An agent with credal state c is epistemically irrational iff she prefers or, given her evidence, ought to prefer some alternative state b to her own.

In light of our theory of epistemic preference, such an agent is epistemically irrational iff either $Est_c(\mathcal{I}(b)) < Est_c(\mathcal{I}(c))$ or she *should* estimate b to be more accurate given her evidence.

To recap, the state-based account of epistemic rationality yields the following predictions regarding Imps, Leap, and Promotion:

- IMPS: Emily is epistemically irrational if she accepts the 'epistemic bribe'. That is, she is irrational if she drops her credence in the proposition C_0 that there's a child before her down to 0 (or to any credence less than 1), in order to secure more accurate credences about the other children. The reason: if she drops her credence in this way, then she ought to strictly prefer some other credal state to her own, *viz.*, ch_E: the true chance function conditional on her evidence E, which unlike Emily assigns probability 1 to C_0.

- LEAP: Bob is epistemically rational, whether his credence that he'll successfully clear the chasm is 0, 1, or anything in between. The reason: whatever credal state he adopts, he ought to estimate that *that* very state (which he is sure agrees with the chances) is as at least as accurate as any other credal state. In turn, he ought to weakly prefer his own credal state to any other.

- PROMOTION: Alice is epistemically irrational if she adopts any credence other than $1/2$ that she'll be promoted. The reason: if she adopts some credence $x \neq 1/2$, then she ought to strictly prefer some other credal state to her own, *viz.*, ch_E: the true chance function conditional on her evidence E, which unlike Alice assigns probability $1 - x$ to the proposition that she'll be promoted.[28]

Now, none of this tells Emily, Bob or Alice what *to do* exactly. Epistemic rationality, on our view, tells you when you've landed in a bad spot — a credal state that, from your own perspective, is less epistemically valuable — *i.e.*, accurate — than some alternative state. That is, it tells you where or where not to be, not what to do.[29]

You may see this as a serious drawback. You might object: it is the primary aim of practical decision theory to guide *action*, and any epistemic decision theory worth its salt should yield an account of epistemic rationality that guides action as well. Strictly speaking, we disagree, since what makes an agent epistemically irrational is just her occupation of a bad state regardless of what action she performed to get there.

Nevertheless, there *is* an important sense in which the state-based account of epistemic rationality guides action. By determining which states are preferable to which others in which worlds, we can identify actions that lead to those states. In particular, only the actions that result in states that are weakly preferred to all other states could be performed by an epistemically rational agent, e.g., only $\overline{c_1}$ in IMPS, only $\overline{c_{.5}}$ in PROMOTION, and every $\overline{c_x}$ in LEAP. Thus, although EpDT in the first instance delivers verdicts about which epistemic states are preferable to which others, we can, in a loose sense, call epistemic actions rational or irrational based on the rational status of the states they lead to. Because epistemic actions and states are deterministically coupled, we can thus answer the critic's complaint that EpDT does too little to guide action while simultaneously respecting the direction-of-fit considerations that underwrite our theory of epistemic rationality proper.

You might also object that EpDT's focus on epistemic states makes it ill-suited to do the job that accuracy-firsters set out for it: explaining why epistemic norms have their binding force by showing that they are a good means to the end of accuracy (*cf.* Carr 2015, §5.3). After all, as IMPS illustrates, the state-based account sometimes requires rational agents to have credences that they are certain will turn out to be less accurate than some other

---

[28]IMPS, LEAP and PROMOTION are all cases in which the state of the world depends *causally* on which credences you adopt. But Caie (2013) and Berker (2013) also consider cases in which the state of the world depends *constitutively* on which credences you adopt. Our treatment of IMPS, LEAP and PROMOTION extends naturally to these cases as well. What the state-based account predicts in *any* case involving act-state dependence — causal, constitutive, etc. — is this: An agent with credal

state c is epistemically irrational iff she estimates (or ought to estimate) some alternative state b to be more accurate. Of course, in arriving at these estimates, she must incorporate any information that she has about what her credences are, and how those credences influence the world, whether they do so causally, constitutively, etc. But it is no demand of epistemic rationality, on our account, that she *exploit* these dependency relations to *make* her credences more accurate.

[29]Note that, by identifying the epistemic right with which states to occupy, EpDT is still fully consequentialist, since what is epistemically right is fully determined by the epistemic good.

credences they might have adopted. So, it seems, a state-focused EpDT does not have the resources to show that Probabilism, Conditionalization, etc., are a good means to the end of accurate credences. Whatever its virtues, it fails to furnish the accuracy-firster's preferred explanation of epistemic normativity.

But a state-focused EpDT *can* be used to show that epistemic norms are a good means to the end of accurate credences. Epistemic norms are, in the first instance, a good means to the end of *epistemically valuable credal states*, on our view. And credal states, as we have stressed, are better or worse (more or less valuable) to the extent to that they conform to the world by encoding an accurate picture of it. Accuracy is the principal determinant of epistemic value. As a result, epistemic norms are indeed a good means to securing accuracy. But recall that credal states are *not* valuable in virtue of causally influencing the world, so as to *make* themselves accurate. So epistemic norms will not in general encode sensible policies for securing accuracy *tout court.* They are only a good means to securing accuracy *in the right way*, *viz.*, by conforming one's credences to the world, rather than the other way around.

Consider a practical analogy. You might evaluate strategies for buying your partner a birthday gift on the following basis: how well it produces a match between your gift and your partner's desires. If you do, then STRATEGY 1 will look pretty good:

STRATEGY 1 Whatever your partner wants, in your best estimate, do the following: (i) buy her socks, and (ii) give her a pill that makes her desire socks.

In contrast, you might evaluate gift-buying strategies on the following basis: how well it produces match *in the right way*, *viz.*, by conforming your gift to your partner's desires, rather than the other way around. If you evaluate gift-buying strategies in *this* way, then STRATEGY 1 will look pretty bad, and STRATEGY 2 will look pretty good:

STRATEGY 2 Pay attention to your partner, and buy her the gift that, in your best estimate, she currently desires most.

Unless you know what your partner currently desires, STRATEGY 2 will probably produce less gift-desire match than STRATEGY 1. But that does *not* mean that it's not a good means to the end of good (highly desired) gifts. It is. But it's only a good

means to the end of securing gift-desire match *in the right way*, *viz.*, by conforming your gift to your partner's desires, rather than the other way around.

Similarly, the fact that the state-based account sometimes requires rational agents to leave accuracy on the table does *not* show that its recommended epistemic policies are not a good means to the end of accuracy. They are. They are a good means to securing accuracy because they are a good means to securing epistemic value, and accuracy is the principal determinant of epistemic value. But they are only a good means to securing accurate credences *in the right way*, *viz.*, by conforming one's credences to the world, rather than the other way around. The reason: credal states are simply *not* epistemically valuable in virtue of causally influencing the world, so as to make themselves accurate. So good policies for securing such value will *not* in general be good policies for securing accuracy *tout court.*

Finally, you might object that the state-based account of epistemic rationality countenances *dilemmas of pure epistemic rationality*: circumstances in which there are *no* rationally permissible credal states consistent with your evidence. But surely, you might continue, it is at least *possible* to satisfy the demands of epistemic rationality, whatever your evidence, even if we in fact often fall short of that ideal. To illustrate, consider a modified version of Promotion (*cf.* Caie (2013, pp. 562-6) for a similar case):

PROMOTION* Alice is up for a promotion. Her boss, however, is deeply insecure: he's sure to promote Alice if she comes across as lacking in confidence, and sure not to promote her if she comes across as brimming with confidence. Furthermore, Alice is useless at play-acting, so she'll come across that way iff she really does have a low/high credence in the proposition P that she'll be promoted. Specifically, if her credence in P is less than $1/2$, then she will certainly be promoted: the chance of P is 1. If her credence in P is greater than or equal to $1/2$, then she will certainly not be promoted: the chance of P is 0.

In Promotion*, Alice is epistemically irrational whatever credence she adopts for P, given what she knows about the chances. If her credence that she will be promoted is less than $1/2$, then she ought to strictly prefer having credence 1. Likewise, if her credence is greater than or equal to $1/2$, then she ought to strictly prefer having credence 0. In either case, she is epistemically irrational.

But it is simply *not*, in general, possible to satisfy the demands of rationality. Consider a practical analogue of Promotion*:

SHIFTY DEPOSITS Betsy has two options: push button *A*, or push button *B*. Each button deposits some amount of money into her bank account. The catch: both buttons are equipped with preference-reading sensors. And they're set up to guarantee that Betsy simply can't prefer the option that will cause the best outcome. If she strictly prefers *A* to *B*, then *A* will deposit $50 and *B* will deposit $100. If she strictly prefers *B* to *A*, then *B* will deposit $50 and *A* will deposit $100. And if she is indifferent between the two, or has no preference, then *A* will deposit $10, and *B* will deposit nothing.

In Shifty Deposits, Betsy is practically irrational whatever preference ordering she adopts. If she strictly prefers *A* to *B*, then *B* is sure to cause a better outcome. So she ought to strictly prefer *B* to *A*. Likewise, if she strictly prefers *B* to *A*, then *A* is sure to cause a better outcome. So she ought to strictly prefer *A* to *B*. (*Mutatis mutandis* if she is indifferent between *A* and *B*, or has no preference between the two. In that case, *A* is sure to cause a better outcome. So she ought to prefer it.) In any case, she is practically irrational. She prefers an option that she is sure will cause a worse outcome than some other option.

Epistemic rationality is no different. Just as we are sometimes forced to violate the demands of *practical* rationality, we are also sometimes forced to violate the demands of *epistemic* rationality. The problem is with the world — its possible cruelty knows no bounds — not with the account of rationality on offer. It is no strike against the state-based account that it fails to make the world a nicer place.

References

Anscombe, G. (1957). *Intention*. Oxford: Basil Blackwell.

Berker, S. (2013). Epistemic teleology and the separateness of propositions. *Philosophical Review* (122), 337–393.

Caie, M. (2013). Rational probabilistic incoherence. *Philosophical Review 122*(4), 527–575.

Carr, J. (2015). Epistemic utility theory and the aim of belief. *Ms*.

Greaves, H. (2013, December). Epistemic decision theory. *Mind*.

Hájek, A. and J. M. Joyce (2008). Confirmation. In S. Psillos and M. Curd (Eds.), *The Routledge Companion to Philosophy of Science*, Chapter 11, pp. 115–128. Routledge.

Humberstone, I. L. (1992). Direction of fit. *Mind 101*(401), 59–83.

Jeffrey, R. C. (1983). *The Logic of Decision* (2nd ed.). University of Chicago Press.

Jeffrey, R. C. (1986). Probabilism and induction. *Topoi 5*(1), 51–58.

Joyce, J. M. (1998). A nonpragmatic vindication of probabilism. *Philosophy of Science 65*, 575–603.

Joyce, J. M. (1999). *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.

Joyce, J. M. (2000). Why we still need the logic of decision. *Philosophy of Science 67*(Supplement), S1–S13.

Joyce, J. M. (2002). Levi on causal decision theory and the possibility of predicting one's own actions. *Philosophical Studies 110*, 69–102.

Joyce, J. M. (2009). Accuracy and coherence: Prospects for an alethic epistemology of partial belief. In F. Huber and C. Schmidt-Petri (Eds.), *Degrees of Belief*, Volume 342, pp. 263–97. Springer.

Lewis, D. (1986). Postscripts to 'Causation'. In *Philosophical Papers: Volume II*. Oxford: Oxford University Press.

Lewis, D. (2000). Causation as influence. *Journal of Philosophy 97*, 182–97.

Pettigrew, R. (2013). A new epistemic utility argument for the Principal Principle. *Episteme 10*(01), 19–35.

Pettigrew, R. (2014a). Accuracy and evidence. *Dialectica*, 1–16.

Pettigrew, R. (2014b). Accuracy, risk, and the principle of indifference. *Philosophy and Phenomenological Research*, 1–24.

Popper, K. R. (1959). *The Logic of Scientific Discovery*. New York: Basic Books.

Predd, J. B., R. Seiringer, E. H. Lieb, D. N. Osherson, H. V. Poor, and S. R. Kulkarni (2009, October). Probabilistic Coherence and Proper Scoring Rules. *IEEE Transactions on Information Theory 55*(10), 4786–4792.

Rényi, A. (1955). On a new axiomatic theory of probability. *Acta Mathematica Academiae Scientiarum Hungaricae 6*, 286–335.

Savage, L. (1954). *The Foundations of Statistics.* New York: Dover.

Smith, M. (1994). *The Moral Problem.* Oxford: Blackwell.

Sobel, D. and D. Copp (2001). Against direction of fit accounts of belief and desire. *Analysis 61*(269), 44–53.