Uniqueness and Metaepistemology

Daniel Greco and Brian Hedden

How slack are requirements of rationality? Given a body of evidence, is there just one perfectly rational doxastic state to be in given that evidence? Defenders of *Uniqueness* answer 'yes' while defenders of **Permissivism** answer 'no.' For Permissivists, rationality gives you a certain amount of leeway in what to believe, with the result that rational people with the same evidence can nonetheless disagree with each other. By contrast, Uniqueness theorists hold that if two agents share the same total evidence and are rational in evaluating that evidence, they must have the same beliefs in response to that evidence.[1]

The debate over Uniqueness and Permissivism is important in its own right, but it also has implications for a number of other epistemological disputes. First, it has taken on a central role in debates about the epistemology of disagreement and the role of higher-order evidence.[2] Second, it bears on whether requirements of rationality are diachronic, concerning how your beliefs at different times ought to be related to each other, or only synchronic, concerning only how your beliefs are at particular times, and

---

[1]Uniqueness and Permissivism are compatible with any number of views about how to conceive of rational doxastic states - they could be sets of binary (on/off) beliefs, or credence functions, or pairs consisting of a credence function and a set of binary beliefs, or all manner of other possibilities. What Uniqueness says is that given a body of total evidence, there's a unique doxastic state (whatever structure that might have) that it's rational to be in, while Permissivism denies this.

[2]See Matheson (2009), Christensen (2010), Ballantyne and Coffman (2011, 2012), and Schoenfield (2014).

in particular how they relate to your evidence at that time.[3] Our concern in this paper, however, is not with these related issues, but with the Uniqueness vs. Permissivism debate itself.

On the face of it, Permissivism seems to have a lot going for it. First and foremost, it has a certain intuitive appeal. Here is Rosen (2001, p.71):

> It should be obvious that reasonable people can disagree, even when confronted with a single body of evidence. When a jury or a court is divided in a difficult case, the mere fact of disagreement does not mean that someone is being unreasonable. Paleontologists disagree about what killed the dinosaurs. And while it is possible that most of the parties to this dispute are irrational, this need not be the case.

Second, it seems that multiple 'epistemic values' are at play when you are evaluating a given body of evidence. You want to favor simpler hypotheses, infer to the best explanation, and project natural predicates (or properties), among other things. But sometimes these epistemic values conflict. The simplest hypothesis isn't always the most explanatory one, for instance. And when these epistemic values conflict, there doesn't seem to be any privileged way of weighing them up. Different ways of weighing these epistemic values can each be rationally permissible, but they can license different bodies of belief (or different credence functions) given the same body of total evidence.[4]

Defenders of Uniqueness are not convinced. Against Rosen, they can respond that the jurors, or the paleontologists, don't really have the same *total* evidence, which will include background information and experience not included in the presentations of the defense and prosecution, or in the latest scientific journals.[5] Perhaps once they share the same total evidence, rational disagreement will no longer be possible. Moreover, it

---

[3]See Meacham (2010) and Hedden (Forthcoming) for discussion.

[4]Schoenfield (2014) talks of different agents who have different *epistemic standards*, which license different sets of beliefs even given a fixed body of evidence. Kelly (2013) argues that different agents can assign different weights to the values of gaining true beliefs vs. avoiding false beliefs and that because differing Jamesian trade-offs can each be rational, Permissivism is true.

[5]See Goldman (2009), Hedden (Forthcoming).

is important to distinguish being irrational in the sense of falling short of the rational ideal from being irrational in the sense of being less rational than (most of) the rest of us. The jurors (and the paleontologists) may disagree without being irrational in latter sense - they needn't be on a par with conspiracy theorists wearing tinfoil hats - but that doesn't mean that they can disagree without any of them falling short of the rational ideal. Indeed, on a natural spelling-out of the case, they are fallible, cognitively limited agents, and hence fall short of ideal rationality in all sorts of ways, some of them predictable and systematic, according to much recent work in behavioral economics.[6]

Against the 'competing epistemic values' argument, defenders of Uniqueness have available at least two different defenses. First, they might hold that it is indeterminate how the competing epistemic values are to be weighed against each other, but it is nonetheless determinately the case that there is a uniquely rational way of doing so. As such, it is determinately the case that there is always a uniquely rational response to a given body of evidence, but sometimes indeterminate what that uniquely rational response is.[7] Hence it can sometimes be indeterminate whether a given person's beliefs are irrational. But its being indeterminate whether they are rationally required or rationally impermissible (as they might be on a version of Uniqueness involving indeterminacy) is importantly different from its being determinately the case that they are rationally permissible (as on a Permissivist picture).

Second, they might hold that competing epistemic values motivate the claim that rational doxastic states are more coarse-grained than initially supposed. For instance, in a Bayesian framework, the Permissivist might hold that, given a body of total evidence, there is a set S of probability functions, each of which uses a different, but rationally permissible, weighting of the competing epistemic values to evaluate the evidence. But where the Permissivist might say that each member of S is rationally permissible,

---

[6]See Ariely (2008) for a survey and discussion.
[7]Compare Christensen (2007, p.192, fn. 8).

but you have to pick one, the defender of Uniqueness can hold that you ought to be in a coarse-grained credal state (sometimes called 'mushy' or 'imprecise' credences) represented by the set S itself.[8]

So intuitive case judgments and the observation of competing epistemic values are inconclusive considerations. A standoff threatens to ensue. Our aim in this paper is to break the deadlock by bringing metaepistemological considerations to bear on the debate. We explore two strands of thought about the roles that epistemically evaluative talk—attributions of rationality, justification, and the like—plays in our lives, and argue that each of these roles is inconsistent with Permissivism. The first, inspired by Craig (1990) and recently taken up by Dogramaci (2012, 2013), treats epistemically evaluative language as serving to guide our practices of deference to the opinions of others. Roughly, saying that someone knows whether P or has a rational or justified belief about whether P serves to pick out that person as worthy of being deferred to on the question of whether P, and perhaps on related matters as well. Epistemically evaluative language earns its keep in part by helping us to identify those who are likely to be reliable testifiers on certain subjects. The second metaepistemological insight originated in metaethics with Gibbard (2003) and was recently brought into epistemology by Schafer (Forthcoming).[9] It focuses on our need to formulate plans about what to do and what to believe in a variety of situations. Regarding a belief that P as rational or justified in a given situation is closely tied to planning, for the contingency of finding oneself in that situation, to believe that P.

We think that both Craig and Dogramaci, on the one hand, and Gibbard and Schafer, on the other, have identified central roles played by epistemic evaluations and

---

[8]See Kelly (2013) and Joyce (2010). Joyce endorses this move while Kelly opposes it.

[9]See also Schoenfield (ming), who argues that different conceptions of the relationship between doxastic plans and evaluations of rationality correspond to different views about the significance of higher-order evidence, and Greco (Forthcoming), who doesn't explicitly use the language of "planning", but is naturally interpreted as appealing to considerations about doxastic planning to address questions about which epistemological debates are "merely verbal."

helped to explain why we have such epistemic concepts and rely so heavily on them. But each line of thought supports Uniqueness over Permissivism, or so we shall argue.

First, however, two notes of clarification are in order. Much recent literature contrasts *coherence* requirements on the one hand, with requirements of *reason* on the other.[10] In broad strokes, coherence requirements concern how one's attitudes "hang together." Being coherent doesn't require having any *particular* attitude, but only avoiding incoherent *combinations* of attitudes. For example, you might coherently believe that it will rain tomorrow, or that it won't, but you can't coherently believe both that it will and will not rain tomorrow. That is, the requirement of non-contradiction is a paradigm example of a coherence requirement on belief. Other, more controversial requirements include requirements of probabilistic coherence and enkratic requirements. Probabilsitic coherence doesn't tell you which credences to have, only that whichever ones you do have should satisfy the probability axioms. Enkratic requirements say that whatever beliefs you hold, you should hold them non-akratically—you shouldn't both believe P, and believe that you shouldn't believe P. Requirements of reason, on the other hand, don't concern how your attitudes fit with one another, but instead concern how your attitudes fit with the reasons for them, and they *can* require you to have particular attitudes. For instance, if you've seen a weather report that strongly predicts rain, then even though you might disregard it and expect a clear day without thereby being incoherent, that would be *unreasonable*. All else equal, in such a case, reason would require you to expect it to rain, even though coherence would not. A closely related but more contentious way of putting the distinction is that requirements of reason are *substantive*, while coherence requirements are *merely formal*.[11]

---

[10]Much of the discussion traces back to debates between John Broome (1999, 2007) on the one hand, and Niko Kolodny (2005, 2007, 2008) on the other.

[11]None of this is straightforward. E.g., see Titelbaum (2015) for an argument that what most writers would think of as coherence requirements may nevertheless require agents to hold certain particular attitudes.

When writers draw this distinction, requirements of rationality are typically put on the side of coherence requirements. For instance, when John Broome (2005) asks whether we have reason to do as rationality requires, or more simply, when Niko Kolodny asks "Why Be Rational?" (2005), both presuppose that requirements of rationality are *not* requirements of reason. Arguably, if requirements of rationality *were* requirements of reason, both their questions would be self-answering, in the affirmative—they would boil down to the question of whether we have reason to do what we have reason to do. It's only when talk of rationality is understood as concerning coherence requirements only, and *not* requirements of reason, that the question of what reasons (if any) we have to be rational can be raised.

But this cannot be how 'rationality' is understood in the debate over Uniqueness in epistemology. If all parties to *this* debate accepted that rational requirements were coherence requirements, and not requirements of reason, then uniqueness would be a non-starter. Given almost any conception of evidence, and of coherence requirements, it's uncontroversial that two agents could share a body of evidence, have different beliefs, and yet both be coherent—Uniqueness would turn out to be obviously false.[12] For example, if we hold the Williamsonian view that one's evidence is one's knowledge, (Williamson, 2000a) then two agents might have the same knowledge but differ in their beliefs that do *not* constitute knowledge, without either one violating any coherence requirement. Or if we think of evidence as consisting in experiences, (Dougherty and Rysiew, 2013) then again, two agents might have the same experiences, yet form different, coherent bodies of belief in response to those experiences.

For these reasons, when epistemologists ask whether a body of evidence fixes a

---

[12]Harman (2002) is a possible exception. He seems to hold a view according to which an agent's evidence consists in her body of beliefs. On this view, it's hard to know how to frame the debate over uniqueness—two agents who have different beliefs *thereby* have different evidence. So it's hard to ask questions about whether multiple bodies of belief might be permissible given a body of evidence—the very question that the uniqueness debate asks seems to presuppose a distinction between evidence and belief.

unique rational doxastic state, we suggest they should be interpreted as using "rational" to refer not (or at least not *only*) to coherence requirements, but instead (or at least, in addition) to requirements of reason. Indeed, we suspect the whole Uniqueness debate could be conducted more clearly without talking about rationality at all, but instead by just asking whether, given a body of evidence, there is always a unique maximally *reasonable* doxastic attitude to adopt in response to it, or whether there is sometimes a range of such attitudes. Nevertheless, because epistemologists who've discussed this issue have typically done so using the language of rationality, rather than reasonableness, we will follow suit. While the preceding might seem like a merely terminological point, we think it's an important one—some of the intuitive implausibility of Uniqueness, we suspect, can be alleviated once we recognize that the Uniqueness theorist can admit that there is *an* interpretation of "rationality" on which Uniqueness is obviously false—the interpretation where rationality concerns coherence, and not reasonableness—but that this isn't the interpretation of "rationality" that she's employing.

Second note of clarification: Uniqueness, as we understand it, is a thesis about the rationality of belief states, not a thesis about the rationality of transitions between belief states. Uniqueness is the thesis that given a body of total evidence, there is a unique most rational belief state to be in.[13] It is not the thesis that from a given starting body of evidence, there is a unique most rational belief state to transition into. The former thesis could be true even if the latter is false. Here is an example to illustrate. Suppose, following Williamson (2000), that your evidence consists of all and only the propositions that you know. Moreover, suppose that given a body of total evidence, your credences ought to equal the evidential probabilities, i.e. the

---

[13]Perhaps it would be better to say that given a body of total evidence, there is *at most one* most rational belief state to be in, if some bodies of evidence are such that there is no rational response to them. This might be the case, for instance, if your evidence includes the proposition that most of your credences are irrational. We won't pursue this issue here.

probabilities of hypotheses given your evidence (see Williamson 2000, Ch 9). If such evidential probabilities exist, then this Williamsonian picture entails Uniqueness. But for all that has been said, it might be that there is more than one perfectly rational belief state into which one could transition. For instance, suppose there is some proposition P which you don't currently believe but which is such that, were you to come to believe P, that belief would constitute knowledge. In such a case, there may be more than one perfectly rational belief states into which you could transition - one in which you have the belief that P and have credences which match your new evidential probabilities (probabilities given your old evidence E plus your new evidence P), and one in which you still lack the belief that P and have credences which match your old evidential probabilities (probabilities given just your old evidence E). The fact that Uniqueness does not entail that there is only one rational belief state into which you could transition will be important in Section 4.

1 Deference

Craig (1990) sets himself the task of investigating the nature of knowledge not by looking at particular case judgments and attempting to devise an analysis in terms of necessary and sufficient conditions, but rather by looking at the role that the concept of knowledge plays in our lives. Why do we make such widespread use of this concept? What human needs does it answer to, and what would we lose if it were to disappear from our thoughts and language? Now, our concern is with epistemically evaluative terms like 'rational' and 'justified' (and the concepts to which they refer) rather than knowledge, but we can apply Craig's methodology here as well and ask what is the utility of epistemic evaluation in general. Moreover, as we shall see, Craig's guiding thought about the role of the concept of knowledge is also promising as an account of

a principal role of epistemically evaluative concepts more broadly.

We are first and foremost concerned with gaining true beliefs, since having true beliefs helps us navigate our environments and achieve our ends. But as social animals, it would be beneficial if we could acquire true beliefs not only through our own perceptual faculties and reasoning abilities, but also by gathering information from others. But whose testimony should we trust? To whom should we defer in forming our beliefs? Here is where epistemically evaluative terms come in. They serve to help categorize potential informants into those that should be relied upon and those that should not. Thus, Craig proposes that 'the concept of knowledge is used to flag approved sources of information' (11). In a similar vein, Dogramaci (2012, p.524) emphasizes the role of epistemic evaluations in '[e]xpanding the pool of accessible evidence' by identifying good testifiers.

Let us take this idea on board, so that a major role of epistemically evaluative concepts is to identify whom to defer to. Then, judging that someone has a rational belief about whether p involves, perhaps among other things, a commitment to deferring to that person's view about whether p, unless of course you think you have some relevant evidence that that person lacks.[14] Note that as we are understanding it, deferring to someones belief involves adopting that belief yourself, not merely respecting it or allowing it to govern a group's decision. We can encapsulate this thought by saying that agents ought to satisfy the following conditional:

Deference

If agent S_1 judges that S_2's belief that P is rational and that S_1 does not have relevant evidence that S_2 lacks, then S_1 defers to S_2's belief that P.

---

[14]There are interesting questions, which we will not pursue, regarding how deference works if it can be indeterminate whether a given belief state is rationally permissible. Can you rationally judge that someone's belief is rational while knowing that is it indeterminate whether that belief is rational? And, supposing that you rationally judge someone's belief to be neither determinately irrational nor determinately rational, ought you to defer to this belief? Is it then indeterminate whether you ought so to defer?

Permissivism sits uncomfortably with the role of epistemic evaluations in identifying appropriate targets of deference. Again, Permissivism says that for at least some bodies of total evidence, there is more than one rational belief state to have in response to that evidence.

There are two problems in particular with combining Permissivism with **Deference**. The first is that it threatens to yield inconsistent deferential commitments. The second is that it is difficult to see why you should defer to the beliefs of someone with epistemic standards that are rationally permissible but different from your own.

Start with the first problem. In a nutshell, if two agents have the same total evidence but different beliefs about whether P, then you cannot defer to each's belief about whether P on pain of inconsistency. Let's take this a bit more slowly. Suppose you know that one agent has credence n in p while another has credence m in P, and that they have the same total evidence. Then, you cannot defer to each one's credence in P; you cannot simultaneously adopt credence n in P and credence m in P, where $n \neq m$. Hence, if judging that an agent's credence is rational involves a commitment to deferring that agent's credence, then you cannot judge each agent's credence in p to be rational.

Of course, we did say that judging that an agent's credence in P is rational involves a commitment to deferring to her credence in P *unless you think you have some relevant evidence that she lacks.* So in the above case, knowing each agent's credence in P and judging each's credence to be rational only involves a commitment to deferring to each if they also know each other's credence in P. The Permissivist might then say that once each one knows about the other's credence in P, they should converge upon the same credence in P (thereby appealing to something like the Equal Weight View in the epistemology of peer disagreement; see Elga (2007)). But this is a concession to the Uniqueness theorist. For insofar as one is aware of what one's initial judgment was, this

fact about one's initial judgment is part of one's evidence. Then, adopting the Equal Weight View (or something like it) amounts to conceding that once two agents share the same *total* evidence, *including evidence about which judgment each initially arrived at*, they must in fact have the same credence in the relevant proposition on pain of irrationality. (Note, though, that the Equal Weight View, even on a version where two disagreeing agents are required to completely converge in their opinions, does not entail Uniqueness. It only entails that when two agents share the same total evidence *and that evidence includes facts about what each's initial judgment was*, there is a unique doxastic state that it is rational for either to be in. But if each of two agents is unaware of what judgment either one initially arrived at and also unaware of their disagreement, then neither one's evidence would include facts about either's initial judgment or about their disagreement. Then, they could share the same total evidence without the Equal Weight View kicking in and requiring that they converge in their opinions. If these two agents could rationally have different belief states, then Permissivism would be true.)

Summing up, if we combine the view that differing evaluations of the same evidence can each be rational with the idea that judging an opinion to be rational involves a commitment to deferring to that opinion, then we get inconsistent deferential commitments.

Now turn to the second problem, that **Deference** is unmotivated on a Permissivist picture. The source of the problem is that on a Permissivist picture, you might judge another agent's credence in p to be rational while knowing that if you had that agent's very same evidence, you yourself would arrive at a different credence in p (and still be rational).[15] But if that is the case, why on earth should you defer to the other agent

---

[15]We're presupposing that if permissivism is true, it is sometimes possible to recognize that a certain case is a permissive one. Cohen (2013) argues that that even if Uniqueness is false, it is nonetheless impossible to rationally take oneself to be in a permissive case. More exactly, he defends Doxastic Uniqueness, the claim that '[a] subject cannot rationally believe there are two (or more) rational credences for h on e, while rationally holding either.' Now, many Permissivists will not want to accept Doxastic Uniqueness, for instance if they hold that Permissive cases are the norm rather than

on the question of whether p?

Suppose (as mentioned above) that Permissivism were true because there are different 'epistemic values' and no privileged way of trading them off against each other. Rationally evaluating evidence requires favoring simpler hypotheses over more complex ones, favoring more explanatory hypotheses, favoring hypotheses that involve natural properties rather than grue-like ones, and so on. But sometimes these epistemic values conflict, and different ways of trading them off against each other can be equally rational.

But then, judging some agent's credence to be rational shouldn't motivate you to adopt that credence even if you know that agent has strictly more evidence than you. For you might judge her credence to be rational while knowing that she differs from you in the weights she assigns to various competing epistemic values. For instance, you might know that she puts greater emphasis on a hypothesis' simplicity than you and less weight on its explanatory power. But then, to defer to her credence would be in effect to adopt her epistemic standards in place of your own, even though you disavow her standards. In Dogramaci's apt phrase, if an agent has different epistemic standards, you are not in a position to treat her as an epistemic surrogate.[16]

_____

the exception (this may be what leads Kelly (2013, fn. 6) to regard this move as unpromising). But for present purposes, the important thing to note is that adopting Doxastic Uniqueness will not by itself save Permissivism from the objections we are raising. For the problems with combining **Deference** with Permissivism involve cases in which you take *someone else* (such as an expert) to be in a permissive case. So in order to resist our arguments, the Permissivist would need to hold not only Doxastic Uniqueness, but also the stronger claim that you cannot rationally take someone else to be in a permissive case.

[16]This point can be made precise in a Bayesian apparatus. A Bayesian version of **Deference** can be construed as a constraint on your conditional credences. In particular, your credence in any H, conditional on the proposition that agent A is an expert with credence n in H, should be n. And we can think of an agent's epistemic standards as being represented by her prior probability function (or priors, for short).

Here, then, is the problem: If Permissivism is true, then it is possible for you to know that some expert A has different priors from you. From the supposition that agent A is an expert with credence n in H, you can reverse engineer which propositions she might have as her total evidence, such that you are certain that A is an expert with credence n in H just in case the disjunction of those evidence propositions E_i is true. But your credence in H, conditional on the disjunction of the E_i, differs from your credence in H, conditional on A's being an expert with credence n in H. But this is inconsistent.

Of course, the Permissivist could simply restrict **Deference** so that it only applies to experts who have the same epistemic standards as you.[17] But to say this is to break the link between deference and judgments of rationality. Once we have restricted the application of **Deference** to only experts who share your epistemic standards, it is no longer the case that, in general, judging that someone's doxastic state is rational comes with a commitment to defer to it. As far as deference goes, agents with rational doxastic states based on different epistemic standards from yours are treated just like agents with irrational doxastic states; in neither case are you in any way committed to deferring to them.[18]

By contrast, Uniqueness vindicates the deference-guiding role of rational evaluations. If you believe that your own epistemic standards are the uniquely rational ones, then you believe of anyone you deem rational that she shares your epistemic standards. Therefore, you can treat that person as an epistemic surrogate, confident in the belief that the conclusions she draws from her evidence are those that you yourself would draw from the same evidence.[19] And if you have doubts about whether your own epis-

---

Hence, you cannot obey the probabilistic version of **Deference** on pain of incoherence. The problem is that, since Permissivism entails that an expert can have different priors from you (and you could presumably know this), there will be experts that you cannot treat as epistemic surrogates and hence are not ones that you can legitimately defer to.

[17]Schoenfield (2014) adopts this proposal.

[18]It might be suggested that judgments of *irrationality* are what do the heavy-lifting on a Permissivist view. Judging someone's belief to be irrational commits you to *not* deferring to that person on that question. But what does the Permissivist then say about what happens when you judge someone's belief to be rational? If it commits you to deferring, then this proposal amounts to the one already considered and rejected. If, by contrast, judging someone's belief to be rational permits one to defer or to not defer, then this is a genuinely new proposal. However, we think it is unattractive. Here's why. As noted in footnote 16 above, deferring to beliefs you know to be based on epistemic standards different from your own threatens to yield incoherence in your own beliefs. You could avoid such incoherence by actually adopting the deferee's epistemic standards as your own, but Permissivists (e.g. Meacham (2013)) have been at pains to avoid the result that it's permissible to arbitrarily switch your epistemic standards, which is one of the main objections to Permissivism (White (2005)). For this reason, we think that Permissivists should say not only that you ought not defer to beliefs you judge to be irrational, but also that you ought not defer to beliefs you judge to be rational but based on epistemic standards different from your own. So again, beliefs you judge to be irrational are treated on a par with beliefs you judge to be rational but based on different epistemic standards, and therefore judgments of rationality or irrationality are doing no work in guiding our practices of deference.

[19]With some idealizing assumptions, it is possible to prove that the probabilistic version of **Def-**

temic standards are the uniquely rational ones, it is also clear that you should defer to judgments you take to be rational, for they are (in your view) based either on your own epistemic standards (if your epistemic standards are in fact the rational ones), or on epistemic standards superior to your own (if your epistemic standards are not the rational ones).

Summing up, we started with the guiding thought that part of the role of epistemically evaluative concepts is to flag which judgments to defer to in forming your own beliefs, so that regarding someone's belief (or credence) as rational involves a commitment to adopting it as your own, unless you have evidence that that person lacks. Uniqueness vindicates this guiding thought, whereas Permissivism clashes with it. So much the worse for Permissivism.[20]

---

erence mentioned in footnote 16 above follows from the axioms of the probability calculus on a Uniqueness picture. See Weisberg (2007) and Briggs (2009). The requisite idealizing assumptions are (i) that if you're rational, then you're certain that you're rational and also certain about what your own priors are (and hence what the uniquely rational priors are) and (ii) that the possible bodies of total evidence one might possess are mutually exclusive. This latter assumption amounts to a kind of luminosity assumption for evidence, namely that when your evidence includes some proposition, your evidence entails that that proposition is part of your evidence, and when your evidence doesn't include some proposition, your evidence entails that that proposition isn't part of your evidence. See Williamson (2000b, Ch. 7) for how dropping this luminosity assumption yields counterexamples to such a probabilistic deference principle, and see Pettigrew and Titelbaum (2014) for discussion of difficulties with formulating a precise deference principle if we drop the assumption that ideally rational agents are certain of what rationality requires of them.

[20]Dogramaci (2012, forthcoming) adds a twist to the deferential role of epistemic evaluations. Rather than serving merely to passively identifying whom to defer to, our epistemic evaluations serve also to promote coordination by influencing others to adopt the same epistemic standards that we have. He argues that '[i]nstances of the simple use [of epistemic evaluations] serve to pressure *others* to conform to the belief-forming rules *of the evaluator*' (522). When I pressure you to follow my rules, and you pressure me to follows yours, 'together we push toward an equilibrium in which we follow *shared* rules' (ibid). Coordinating on a set of shared belief-forming rules makes testimony more trustworthy, because when you follow the same rules as I do but have different evidence, I can indirectly make use of your evidence because 'I can trust that *you will* draw the same conclusion from an evidence basis that *I would*' (524). This suggests that my use of epistemic evaluations involves a commitment to there being a uniquely correct set of belief-forming rules, namely my own. In more recent work, Dogramaci endorses this line of argument Dogramaci and Horowitz (Ms.).

2 Planning

In the previous section we focused on a role that epistemic evaluations can play in inter-personal relations—flagging appropriate targets of deference—and argued that Uniqueness must be true in order for them to play that role. But it's hard to think that the function of epistemic evaluations is *exhausted* by their interpersonal role. Couldn't Robinson Crusoe, who's never in a position to interact with others, let alone defer to their opinions, have reason to make some epistemic evaluations?

Crusoe certainly needs to think about what to do. On a picture defended by Gibbard (2003), this is enough for Crusoe to have a use for normative concepts—including the concept of epistemic rationality. Suppose Crusoe is trying to decide how to spend his day tomorrow. He is hungry and will starve soon if he doesn't eat. He could try to hunt big game. If he succeeds, he'll have enough food for a long time, and will be able to focus on getting off the island. But if he fails, he'll be even hungrier and weaker than he is now. He could also forage for nuts and berries. If he forages, he'll almost certainly find *some* food, but there's little chance of a windfall. How will he put this question to himself? It's extremely natural to do so using normative language: "what should I do tomorrow?" And if he ultimately decides to forage, it's natural to express that conclusion using normative language too: "hunting would be too risky, so I should forage instead." The connection between normative language and planning isn't limited to questions about what to *do*, however—it also extends to questions about what to *believe*. Suppose Crusoe is planning for the contingency of finding berries. He might ask himself directly: "what should I do if I find berries?" But he might instead first ask himself what to believe, and develop his action plans accordingly: "if I find berries, how confident should I be that they are poisonous? How confident should I be that they are edible? What sort of evidence would favor one hypothesis over the other?" He might systematize his plans for what to believe under what circumstances using the

15

language of rationality: "in circumstances C, the rational thing to think would be P; in circumstances C', the rational thing to think would be Q..." etc.

Must Crusoe use epistemically normative language to systematize his commitments? Couldn't he express the same commitments using by using conditionals?—e.g., "if I see berries, they'll be poisonous, if I see nuts, they'll be safe..." This would work in some cases, but it's not obviously general enough. What if, e.g., Crusoe thinks the rational attitude to have upon seeing berries is a credence of 0.5 that they're poisonous? He might try to express this judgment to himself without using epistemically normative language as follows: "if I see berries, there's a 0.5 chance they'll be poisonous." But this just raises the question of how to understand the species of chance involved in this judgment. If it must be objective, physical chance, then the proposal fails—Crusoe might think a 0.5 credence that the berries are poisonous is rational, while being entirely agnostic about the physical chance that they're poisonous.[21] The proposal also fails if the chance judgment is understood as involving purely subjective credence—Crusoe might think that a 0.5 credence is rational, while not expecting that he will in fact have such a credence when he sees berries. The proposal will work if "there's a 0.5 chance they'll be poisonous" is interpreted as "the reasonable credence to have that they're poisonous will be 0.5", but then the proposal isn't an alternative to ours; Crusoe will be still be using (covertly) epistemically normative language to systematize his plans.

So epistemically normative language plays an important, potentially ineliminable role in doxastic planning. If we are ambitious metaepistemological expressivists, we might try to explain all uses of epistemological thought and talk as ultimately boiling down to a species of contingency planning. But for the purposes of the present paper, that's unnecessary—it's enough that *one* central function of epistemological thought

---

[21]To fix ideas, we might imagine that for some reason Crusoe thinks it's already determined whether any berries he sees will be poisonous, and so the objective, physical chance of encountering poisonous berries is either 0, or 1, but he doesn't know which. This might hold if, e.g., all berries on his island are of the same species, perhaps poisonous, perhaps not.

and talk is its connection to planning. But just how should we understand the connection between planning what to believe on the one hand, and epistemic evaluations on the other?

Once we ask this question, we face a dialectic very similar to the one explored in the previous section. If Uniqueness is true, then we are in a position to endorse a simple, straightforward thesis about the connection between planning and rationality. If Permissivism is true, however, it is much harder to say anything systematic and plausible about this connection.

If Uniqueness is true, then the following thesis about the connection between epistemic evaluations and doxastic planning is a natural one to endorse:

Planning

If S judges that it would be rational to hold doxastic attitude B given evidence E, then S plans to hold B, given evidence E.[22,23]

But, for much the same reasons the permissivist cannot endorse **Deference**, she cannot endorse **Planning**. If she ever judges of any particular body of evidence E that it is a permissive one—that various, conflicting attitudes B would be rational given E—and **Planning** is true, then she will find herself with inconsistent plans concerning which attitudes to adopt upon getting evidence E.

If the Permissivist still wants to accept *some* connection between epistemic evaluation and planning, what can she say? The following thesis seems like a natural fallback:

---

[22]Depending on our metaepistemological commitments, we might endorse **Planning** as part of a broadly reductive theory of what *constitutes* epistemic evaluative judgment, or instead as a merely normative thesis about how two metaphysically distinct sorts of mental state—epistemic evaluations on the one hand, and plans on the other—*ought* to be related to one another. Gibbard (2003) opts for the former sort of thesis—he ultimately wants to reduce normative judgments to planning states—but this more ambitious view is not necessary for our project.

[23]Talk of plans for belief might seem to involve commitment to an implausible form of doxastic voluntarism, but it needn't be interpreted that way. For our purposes, it's enough that planning to believe that P in circumstances C involves some sort of commitment to believing that P in C—it needn't be the case that this commitment can be adopted voluntarily, for purely practical reasons. See Schafer (Forthcoming, §5).

Permissive Planning

If S judges that it would be rational to hold doxastic attitude B given evidence E, then S does not plan *not* to hold B, given evidence E—her plans don't rule out holding B.

In the special case where subjects' judgments conform to Uniqueness—where they never judge more than one attitude to be rationally permissible given a fixed body of evidence—**Planning** and **Permissive Planning** coincide. But when subjects' judgments do *not* conform to Uniqueness, **Planning** generates inconsistent plans, while **Permissive Planning** merely generates permissive plans.

How much better is that? Suppose Phil is a paleontologist who thinks that while an asteroid impact was a major factor in the Cretaceous extinction, volcanic eruptions also played a significant role. But Phil is a permissivist, and he holds that paleontologists who give full credit to an asteroid are also rational in their beliefs. Phil plans to continue holding his hybrid theory until new evidence comes in.

This last stipulation, however, doesn't sit nicely with **Permissive Planning**. After all, if Phil holds that both the asteroid theory and the hybrid theory are rationally permissible, then if **Permissive Planning** holds, he cannot plan to accept one but not the other—his plans would have to rule out neither. The problem is that if we accept that permissive cases are plentiful and not too hard to recognize (as most Permissivists will want to do),[24] then if we also accept **Permissive Planning**, we'll find ourselves endorsing strict limits on just how committal subjects' doxastic plans can be—we'll have to say that in permissive cases, subjects' plans cannot rule out any of the permissible attitudes.[25] And that seems contrary to the spirit of Permissivism—the permissivist should want to allow not only that permissive cases exist, but that in such cases, it's

---

[24]As noted above, an important exception is Cohen (2013), who holds that while Permissivism is true, it is impossible to ever rationally take yourself to be in a permissive case.

[25]Again, whether this "cannot" claim is a metaphysical one or a normative one will depend on how we were understanding **Permissive Planning** in the first place. See footnote 21.

permissible to plan to hold any of the permissible attitudes, *to the exclusion of the others.*[26] So while **Permissive Planning** is more congenial to the permissivist than the unmodified **Planning**, it still doesn't give her what she'd like.

Suppose Phil repudiates his Permissivism, and adopts Uniqueness. How will he think about his predicament then? First, he'll probably become agnostic about which theory of the cretaceous extinction, if either, it is rational to believe. But this won't paralyze him, or leave him unable to form plans for what to believe—he will still let his plans for how much credence to invest in such theories be guided by his tentative, probabilistic judgments about how likely they are to be rational given his evidence. While he won't necessarily be fully rational—if Uniqueness is true, it's *hard* to be fully rational—he will have in place a framework for letting his doxastic planning be guided by his epistemic evaluations that doesn't lead to absurd results.

So far we haven't found a Permissivist-friendly account of how epistemic evaluations and doxastic plans relate to one another. But perhaps this is because we've been conceiving of plans in the wrong way. As noted, many Permissivists hold that the reason a body of evidence doesn't determine a unique permissible response is because

---

[26]What about Buridan's ass cases in the practical realm, where you regard the reasons for each of two options - going left and going right, say - as equally balanced? We'll talk about such cases in Section 4, but some preliminary remarks are appropriate at this point. Don't such cases force us to come up with an account of the relationship between rationality and planning on which it can be permissible to plan to perform one action to the exclusion of another—to plan to go left, rather than right—without regarding the excluded action as irrational? That is (roughly) the conclusion Gibbard (2003, pp.54-5) draws, in distinguishing between plans that involve "plumping" from indifference, and plans that derive from genuine preference.

But Gibbard's isn't the only way to conceive of Buridan's ass cases. It strikes us as an equally natural description to hold that a rational agent, for the contingency situation in which she has not yet started down either path (where this includes not having made any mental commitment to either path), will neither plan to go left nor plan to go right. However, for the contingency in which she has started down the left path, say, she will plan to go left. In other words, in planning for the contingency where she has not already tipped the balance in favor of one of the options, she leaves both options open. On this description, no distinction between genuine planning from preference, and mere "plumping" from indifference need be drawn.

Even if Gibbard's treatment of Buridan's ass cases is preferable to ours, however, we doubt that an attractive version of Permissivism can be founded on an analogy to such cases; an attractive version of permissivism will hold that one can regard alternatives as permissible even when one is not indifferent between them.

epistemic obligation is determined, not just by evidence, but also by some further factor—perhaps a set of epistemic standards or values.[27] The Permissivist might then hold that given a body of evidence *and* a set of epistemic standards, there is a unique rational response. It's just that there is no single rationally obligatory set of epistemic standards.

If the Permissivist takes this position, then it's natural for her to endorse the following modified version of **Planning**:

> **Planning***
>
> If S judges that it would be rational to hold doxastic attitude B given evidence E and epistemic standards S, then S plans to hold B, given evidence E and standards S.

That is, the Permissivist can hold that doxastic plans should be contingent not just on what evidence is received, but also on the epistemic standards the subject is using. The Permissivist's plans, then, would differ from those of the Uniqueness advocate in that the Permissivist would plan to hold different beliefs given different epistemic standards, even holding fixed her evidence. The Uniqueness advocate, on the other hand, would plan to stick with the same beliefs given the same evidence, even for the contingency of having different standards.[28]

If this is going to give the Permissivist a viable story about the relationship between epistemic evaluations and doxastic planning, and how her plans differ from those of the

---

[27]See, e.g., Kelly (2013), Schoenfield (2014).

[28]Of course, an agent might gain evidence that bears on how much she should favor simplicity as opposed to explanatory power, for instance. But this is evidence is still evidence; it's just higher-order rather than first-order evidence. So planning for a contingency in which you have gained higher-order evidence which supports a different trade-off between simplicity and explanatory power (say) is not planning for a contingency in which you have the *same* evidence but different standards. Rather, we should think of your current standards as giving a verdict on how you should respond to both first-order and higher-order evidence. Therefore, as we are thinking of epistemic standards, your epistemic standards do not recommend abandoning them in favor of a alternative set of standards in the face of higher-order evidence. Indeed, there are good arguments to the effect that a given set of epistemic standards *cannot* recommend abandoning those very standards in the face of evidence, whether first-order or higher-order. See Lewis (1971) for related discussion.

Uniqueness advocate, we need to be able to make sense of the idea of planning for the contingency of having different epistemic standards. It's not immediately obvious that we can do so, however. One way of understanding such planning involves seeing it as planning for the contingency of finding out what sort of standards you've had in the past. Suppose you now have the cautious epistemic standards—while you've observed a fair number of green emeralds, it seems to you only slightly likely that *all* emeralds are green, and it strikes you as hasty to jump to that conclusion on the basis of your present evidence. Now you find out that these judgments are very out of character for you—in the past, you've been happy to hastily generalize concerning white swans, black ravens, etc. Finding out about these judgments surprises you—they each strike you as hasty *now*—but you're convinced that for most of your life, they wouldn't have. Should the discovery that you're, by character, a hasty generalizer, make any difference to your current plans? It's hard to see why; from your current perspective, it's not all that likely that all emeralds are green. The fact that, in some sense, you had standards that permit hasty generalizations just doesn't seem relevant to what you should *now* think. After all, in the first place, you're trying to figure out whether all emeralds are in fact green—not trying to make an "in character" judgment concerning whether all emeralds are green.[29]

Is there any other way the Permissivist might make sense of planning for the contingency of having different plans? We might take the example of Ulysses as a jumping off point. Ulysses plans for a future in which he will have different plans. His plan for that future is *not* to do what he'll plan to do then—swim over to the sirens—but instead to stick with his current plan of avoiding them. And he acts on this plan by binding

---

[29]Many writers have argued that diachronic coherence requirements aren't genuinely normative—see Christensen (2000), Kolodny (2007), Moss (Forthcoming), Hedden (Forthcoming). The present point is along similar lines; to the extent that, in forming doxastic plans, one is aiming at accuracy, coherence with one's past self will seem relevant only insofar as it is instrumental to accuracy. In cases where coherence and accuracy seem to come apart, at least from a planning perspective, accuracy will win out.

himself to the mast. How might this suggest a model for the Permissivist? Sometimes, like Ulysses, we plan for circumstances in which we'll have different plans. But we don't always, and plausibly shouldn't always, engage in self-binding behavior. Right now I plan to pick up chocolate ice cream, because that's what I happen to like, but when I contemplate a contingency in which I plan to pick up vanilla, I don't recoil in horror; it doesn't seem worthwhile or desirable to bind myself to the chocolate ice cream plan, even under contingencies in which I no longer prefer chocolate. So it looks like we can make sense of both (a) planning for contingencies under which one has different plans, and (b) either planning, for those contingencies, to act on one's *current* plans (this involves self-binding), or planning, for those contingencies, to act on the plans one *would* have (this involves no self-binding).

The Permissivist, then, can offer an account of what's involved in holding an attitude, taking it to be rational, but also taking it to be permissible to hold a different attitude, given the same evidence. Somebody with this set of judgments will plan, for at least some of her doxastic attitudes, *not* to stick with those attitudes under contingencies in which she has different plans (i.e. different epistemic standards) but the same evidence. She will hold the attitudes, but will not bind herself to holding them under contingencies in which her plans change. This seems like an *intelligible* way for the Permissivist to explain what sort of planning behavior is entailed by permissive epistemological judgments. But is it an *attractive* one?

We think not. We want our theory of epistemic rationality to mesh with our theory of instrumental rationality. After all, part of the point of being epistemically rational is that doing so is a good way of achieving our aims. This isn't to say that epistemic rationality is a species of instrumental rationality (seeKelly (2003) for a compelling critique), and it is not to deny that beliefs can be epistemically rational but instrumentally irrational, or *vice versa* (more on this shortly). Rather, it's just to say that

the Permissivist should be loath to accept that epistemic and instrumental rationality often, or even typically, pull in opposite directions. But failure to self-bind to one's current beliefs—at least in the absence of receiving new evidence—seems instrumentally irrational, as the following example should bring out. Suppose you currently believe that the Democrats will win the 2016 United States presidential election. Furthermore, let's suppose that there are no side benefits or costs of holding this belief—e.g., nobody will reward you or punish you based on your belief. Insofar as it is an (in)auspicious one to hold, it is because of the costs and benefits of acting on it.[30] And you're likely to act on it; you have the opportunity to make bets on political events, and frequently do so.

Might you rationally take a permissive attitude towards your belief? On the present suggestion, that would amount to believing that the Democrats will win, but failing to be ready to bind yourself to continuing to hold this belief under circumstances where your belief (but not your evidence) changes. But it's hard to see how this won't amount to a straightforward failure of instrumental rationality. All else equal (and we may imagine that all else *is* equal in this case) you prefer more money to less, and by your current lights (i.e., given your belief that they'll win), you will win money if you believe that the Democrats will win and you'll lose money if you don't; believing that they'll win will lead you to bet on them, while believing that they won't will lead you to make contrary bets. So if you believe they'll win, you must prefer to continue believing this (at least if we restrict attention to scenarios in which you get no new evidence).

To drive this point home, imagine that we flesh out the story as follows. Suppose you know that in one week's time, you will be asked to bet on which party will win the 2016 presidential election. You now believe that the Democrats will win. You are also now in a position to determine whether or not, in one week's time, you will

---

[30]We might say that its only value to you is its "guidance value", to use the terminology of Gibbard (2007).

believe that the Democrats will win, or will believe that they won't. Perhaps you can take a pill that will make you think that the Republicans will win. By your current lights, it certainly looks as if you should prefer *not* to take the pill—if you take the pill, you are likely to lose money. Of course, you might not care whether or not you take the pill—on the present view, this is what would be involved in taking a permissive attitude towards my belief. But such an attitude seems instrumentally irrational. Quite generally if you believe that P, then you'll think that you're more likely to satisfy your desires by acting on the belief that P, than on the belief that $\sim P$. So as long as you're contemplating situations under which neither your desires nor your evidence changes (as in this example), you should prefer that you continue to believe that P. Diachronic dutch book arguments are controversial (Christensen, 1991), but that's not what we're offering. The point of the present example is that, if one isn't willing to bind oneself to one's beliefs, then one will have preferences that, at a single time, are practically irrational. One will, e.g., prefer more money to less, and believe that by self-binding, one will get more money rather than less, but nevertheless not prefer to self-bind.

The thrust of the foregoing considerations is that it would be instrumentally irrational to plan to act on beliefs or credences that you think are likely to be less accurate than others you might have arrived at. After all, you're less likely to satisfy your preferences by acting on inaccurate beliefs or credences than on accurate ones. But plausibly, rational epistemic standards must regard themselves as best, in the sense that according to those epistemic standards, the expected accuracy of beliefs or credences based on those very epistemic standards is greater than the expected accuracy of beliefs or credences based on other epistemic standards. In the terminology of Lewis (1971), rational epistemic standards (or what he calls 'inductive methods') must be *immodest*. As Lewis argues, modest epistemic standards would not be worthy of your trust. By definition, a modest set of epistemic standards would say that in some cases some par-

ticular alternative set of epistemic standards would yield more accurate credences, and so you should switch to those alternative standards.[31] You could not fully trust a set of epistemic standards that sometimes recommends trusting some alternative standards instead. As he writes, 'It is as if *Consumer Bulletin* were to advise you that *Consumer Reports* was a best buy whereas *Consumer Bulletin* itself was not acceptable; you could not possibly trust *Consumer Bulletin* completely thereafter' (p. 56). We can bolster this thought by appeal to epistemic utility theory. In a Bayesian framework, we can think of a set of epistemic standards as a prior probability function, where the credences recommended by those epistemic standards given evidence E are the result of taking that prior probability function (or those 'priors,' for short) and conditionalizing it on E. Greaves and Wallace (2006) show that, from the point of view of your own credences, the expected accuracy of the credences resulting from conditionalizing your own priors on total evidence E is greater than the expected accuracy of the credences resulting from conditionalizing some other set of priors on E.[32] What this means is that you should take your own epistemic standards to yield (expectedly) more accurate beliefs or credences than alternative epistemic standards. So given that it's instrumentally rational to base your actions on the most accurate beliefs or credences around, it would be instrumentally irrational to plan in certain circumstances to hold beliefs or credences based on alternative epistemic standards. So much the worse for **Planning***, which recommends doing just that.

---

[31]As Horowitz (2013) notes, in her excellent discussion of immodesty and Uniqueness, we must be careful about scope. It will generally be the case that you think there exists some set of epistemic standards that would do better than your own in this case, for instance epistemic standards that ignore any evidence that is in fact misleading. But as she rightly notes, 'What's *not* rational, then, is to regard some *particular* [epistemic standards] as more truth-conducive than one's own, while also knowing which belief state that rule recommends in every case.'

[32]A caveat: this only applies where E is logically stronger than your current total evidence. Obviously, if E were something weaker than your current evidence, you'd expect to do at least as well by just sticking with your current credences, which, after all, are based on more evidence. But a plan to have, in circumstances in which your evidence is impoverished, the beliefs that would be supported by a larger body of evidence is a plan which you likely couldn't take yourself to be able to follow.

The Permissivist might grant that taking a permissive attitude towards one's beliefs involves a failure of *instrumental* rationality, while holding that it nevertheless involves no failure of *epistemic* rationality—after all, we are familiar with the idea that, in Pascalian cases, epistemic and instrumental rationality can pull in different directions. Perhaps permissive cases should be understood as simply another sort of case in which this conflict arises. While such a position is available to the Permissivist, it doesn't seem to capture the spirit of Permissivism. If we have Permissivist sympathies, we will want to say that taking a permissive attitude towards one's beliefs needn't involve any failure of rationality at all; the view that permissive cases necessarily present dilemmas in which one can be epistemically rational or instrumentally rational but not both is an unhappy fallback position for the Permissivist.

Note also that the foregoing considerations - that plans not to bind yourself to your actual epistemic standards are instrumentally irrational - also militate against a related Permissivist proposal on which judging a belief to be irrational involves planning not to hold that belief given that evidence, but judging a belief to be rational permits you to plan to hold that belief or to plan not to hold it. (This is analogous to the proposal for a modified deference principle considered and rejected in footnote 16, above.) On a Permissivist picture, it looks as though you should *not* plan to hold beliefs you judge to be rational but based on different epistemic standards, for such plans would be instrumentally irrational.

In this section we started with the idea that there was a connection between epistemic evaluation and doxastic planning. As Uniqueness theorists, we had no trouble making sense of this connection. But we explored how the Permissivist might try to do so as well. If she endorses the simple theses about planning that the Uniqueness theorist can, she ends up with inconsistent plans. It looked as if she needed to endorse both the idea that doxastic plans must be contingent, not just on evidence, but on

something like what standards one has, *and* the idea that one should take a permissive attitude towards one's epistemic standards. We examined two ways of making sense of these commitments. On the first, it amounted to planning to hold different beliefs, upon finding out that one has historically had different standards. On the second, it amounted to not being willing to bind oneself to sticking with one's current plans (even in the absence of new evidence). In both cases, while we get a picture of what permissivist planning might look like, neither is pretty. In the first, the Permissivist is poised to change her beliefs upon receiving evidentially irrelevant information about what she used to think. In the second, the Permissivist has self-defeating plans—she's willing to accept bets on P, but not to guard herself against betting on $\sim P$, even without receiving new information.

3 Proving Too Much?

In this section we'll consider and respond to an objection, and hopefully better clarify our position along the way. The objection runs as follows:

> Everything you've said to motivate the idea that a body of evidence fixes a unique set of rational *doxastic* attitudes could be used to motivate an analogous idea about rational *preference*. We judge some courses of action to be rational, and others irrational, and it's plausible that part of the point of such judgments involves deference and planning—e.g, when we come to believe that some course of action would be rational, we often go on to form an intention to take it. But no uniqueness thesis in the practical realm is plausible. Fix someone's practical situation, and there might be various, equally rational courses of action available to her. In short, if your arguments established uniqueness in the epistemic realm, then analogous arguments would establish uniqueness in the practical realm. But uniqueness is obviously false in the practical realm. So your arguments must be flawed.

Our response to this objection has two parts. First, assuming that uniqueness theses for the practical realm are false, we can explain why analogues of our arguments for Uniqueness in the epistemic realm don't carry over to the practical realm (we restrict the capitalized term 'Uniqueness' to refer only to the thesis for the epistemic realm). Second, we argue that it is not as obvious as it might first appear that uniqueness is false in the practical realm.

Start with the first part of our response. While considerations about the roles played by judgments about the rationality of beliefs support Uniqueness, it is far from clear that the roles played by judgments about the rationality of, say, preferences likewise support uniqueness theses for such practical attitudes. To begin with, while it is plausible that judging a belief to be rational carries with it a commitment to defer to it (except if you think you have relevant evidence that the other person lacks), it doesn't seem that we have a general practice of deference to the preferences of others that we judge rational. We only defer to others' preferences in cases where we know that they share our tastes; but then, facts about their preferences are evidence about what will likely satisfy our own tastes. Now, it is a deep and interesting question why this should be so, and we won't try to fully address it here. One promising approach to answering this question might emphasize that while there is an objective, agent-neutral standard of correctness for beliefs (namely truth), there may not be an objective, agent-neutral standard of correctness for preferences. For instance, it may be that preferences are correct only insofar as they line up with agent-relative facts about goodness, and this is why judging a preference to be rational doesn't involve a commitment to defer to it.

Next, consider the planning role of judgments about the rationality of beliefs. Note that we didn't argue that there is no *consistent* story about planning that the Permissivist can tell. Rather, we argued that the only consistent Permissivist story about planning yields plans that are instrumentally irrational. In particular, planning to

28

sometimes base your belief on epistemic standards that differ from your actual epistemic standards is instrumentally irrational, for by your own lights, the expected accuracy of such beliefs is less than the expected accuracy of beliefs based on your actual epistemic standards. However, planning to have different preferences in a contingency situation in which your tastes and values differ from those you actually have needn't involve any instrumental irrationality. This is the case at least for preferences that are not conditional on their own persistence, to use Parfit (1984, p.151) useful phrase. I only want myself to go to the opera rather than a rock concert in contingencies in which I get more enjoyment out of opera than rock. So it wouldn't be instrumentally irrational to plan to prefer rock over opera (and to act on this preference) for contingencies in which my tastes favor rock over opera. Indeed, binding myself to continuing to hold my actual preference for opera over rock even in circumstances in which my musical tastes differ would seem paradigmatically irrational. So, while we have argued that the plans recommended by **Planning*** would be instrumentally irrational, the plans that would result from an analogue of **Planning*** in the practical realm would not be instrumentally irrational. Therefore, our planning-based argument for Uniqueness doesn't carry over to support a uniqueness thesis for preferences.

Turn now to the second part of our response to the 'proving too much' objection, namely that uniqueness theses in the practical realm aren't obviously false. This will take some work. First, the same qualification we made earlier concerning the distinction between rationality and reasonableness is in order here. Uniqueness theses concerning practical rationality will seem obviously false if by "practical rationality" we mean a domain of purely coherence requirements—e.g., requirements of instrumental rationality, or requirements banning intransitive preferences. Two agents can be faced with the same practical situation, have wildly different intentions, and yet both be instrumentally rational, have transitive preferences, etc. If a uniqueness thesis in the practical domain

are to have any plausibility, they must be theses to the effect that a given practical situation fixes a uniquely *reasonable* set of practical attitudes (intentions, preferences, etc.), where reasonableness is a substantive notion, and not merely a matter of internal coherence among one's practical attitudes.

But even with that qualification, uniqueness theses in the practical realm can look crazy. Take the example discussed earlier, of preferring chocolate to vanilla. Couldn't we fix a practical situation—let someone be faced with a choice of ice creams—without thereby fixing a unique reasonable preference for her to have? Suppose you in fact prefer chocolate. Couldn't you have equally reasonably preferred vanilla, in the same situation? To put the point in a way that seems particularly difficult to accommodate given the arguments we've offered so far—couldn't you regard the choice of vanilla as rational, while nevertheless decisively planning to pick chocolate?

Given the right understanding of a practical situation—the understanding on which it is most closely analogous to a body of evidence in the epistemic realm—we believe the answer to these questions is "no", or at least "not obviously." In a natural version of the case, your preference for chocolate isn't ungrounded or brute. You prefer chocolate to vanilla because you enjoy the taste of chocolate more than the taste of vanilla. While somebody with different tastes could be equally rational in opting for vanilla, it's far from clear that somebody with the *same* tastes as you could be rational in preferring vanilla to chocolate. This is analogous to the idea that, in the epistemic realm, it's no counterexample to Uniqueness to hold that two people *with different evidence* could be equally reasonable in holding conflicting beliefs. That is, once we allow a practical situation to include factors such as personal tastes,[33] some apparently

---

[33]We might make the practical case even closer to the epistemic case by saying that it's not your tastes themselves that rationalize your preference for chocolate over vanilla, but rather your evidence about your tastes. This has some plausibility. After all, if your tastes are such that you enjoy chocolate over vanilla, but you don't realize that and have good evidence that your tastes are the reverse, then plausibly you rationally ought to prefer vanilla over chocolate.

obvious counterexamples to practical uniqueness theses—examples involving gustatory preference—can be accommodated.

This might seem like no defense at all—perhaps talk of tastes and gustatory enjoyment is just elliptical for talk of preferences. When we say you prefer chocolate because you enjoy the taste more, we are not explaining one fact in terms of a different one, but merely stating the same fact—that you prefer chocolate—in two different ways. And if there's no notion of what you enjoy that is distinct from what you prefer, then we're back to the idea that the case is an obvious counterexample to uniqueness—we can't appeal to some way in which the practical situations of the chocolate-chooser and the vanilla-chooser differ that would *explain* why they are both reasonable despite having different preferences.

While this response can seem initially attractive, we think it's ultimately untenable. Perhaps in most ordinary cases, preferences track tastes and enjoyment, but much recent work in psychology turns on the idea that they can come apart.[34] For instance, much addictive and compulsive behavior is naturally understood as involving wanting something that will not be enjoyed when it is achieved. It's possible—indeed, in some cases actual—that the nicotine addict doesn't enjoy the experience of smoking a cigarette any more than we do, even though she strongly desires to have this experience, while we do not. And once we distinguish the sorts of attitudes uniqueness theses might apply to—preferences, desires, intentions, and the like—from facts about an agent's practical situation—understood as including tastes, facts about what an agent enjoys, etc.—then there's room to hold that a practical situation will fix a unique set of reasonable practical attitudes.

Perhaps this response is adequate to deal with *some* apparently obvious counterexamples to uniqueness. But certainly not all. Suppose, rather than being faced with a

---

[34]See, e.g., Berridge (2009).

choice between chocolate (which you like) and vanilla (which you dislike), you are faced with a choice between two identical pints of chocolate ice cream, both of which you would enjoy. Here you have two choices—pick the pint on the left, or the pint on the right—that are clearly equally rational. You might plan to reach for the pint on the left, while recognizing that you could have, just as rationally, planned to pick the pint on the right. Even these Buridan-style cases can be accommodated with a practical analogue of the Uniqueness thesis, or so we'll argue.

Start with the case of preference, as it is easiest for us to accommodate the idea that this sort of example does not threaten a uniqueness thesis concerning reasonable preference. It is intuitively plausible in this case that that one would *not* be reasonable to *prefer* the pint on the left to the pint on the right. Rather, one should have no preference between the two pints.

But what about intentions? Couldn't one form the intention to pick the pint on the left, while recognizing that one could've equally rationally formed the intention to pick the pint on the right? And if the sort of plans for belief we discussed in the previous section are closely analogous to intentions in the practical realm, then we're still faced with a problem—that is, our arguments concerning planning in the epistemic realm would generalize to suggest that a Uniqueness thesis, not just concerning preferences, but also concerning intentions, is plausible in the practical realm. So if Buridan cases threaten uniqueness theses for intentions, we're in trouble.

But there is at least one stance about intentions in Buridan cases that is compatible with a uniqueness thesis for intentions. Recall (Section 1) that, crucially, uniqueness theses are not committed to the claim that from any starting point, there is a uniquely rational state to transition into. It is instead the claim that, given a situation (an evidential situation in the case of beliefs, and a practical situation in the case of preferences or intentions), there is a uniquely rational state to be in right then. We sketched

32

how the latter claim could be true even while the former is false. Suppose that there is an objective evidential probability function, and that your credences ought to match the evidential probabilities (i.e. the result of taking the evidential probability function and conditionalizing it on your total evidence). This means that Uniqueness is true. Suppose also that your evidence consists of all and only your knowledge (E=K; see Williamson 2000). Then, even though Uniqueness is true, so that given a body of evidence, there is a uniquely rational doxastic state to be in, nevertheless there will often be more than one rational doxastic state into which you could transition. For instance, suppose there is a belief that you don't have but which is such that if you were to form that belief, it would constitute knowledge. Then, you could form that belief and change your credences so that they match the new evidential probabilities, or you could refrain from forming that belief and stick with your existing credences, which match the current evidential probabilities. Either way, you will wind up with the uniquely rational doxastic state, given the evidence you then possess.

Here is another, more contentious, case. Consider self-reinforcing beliefs. Suppose that you have to give an important talk, and you believe (plausibly) that self-confident people tend to give good talks, while people lacking in confidence tend to give poorer talks. If you believe that you'll give a great talk, you're likely to do so, while if you believe that you'll give a mediocre talk, then your talk is likely to be mediocre. Suppose further that at the outset, you have no belief one way or the other regarding how your talk will go. As a simplifying assumption, let us assume that you are always aware of what beliefs you have. Relative to your initial evidential state, you ought to suspend judgment on whether you will give a talk. But if you somehow then acquire the belief that you'll give a great talk, then this belief is rational, relative to your new evidential state, for having the belief is evidence for its content. Similarly, if you somehow acquire the belief that you'll give a mediocre talk, then *this* belief is rational, relative to your

new evidential state. All this is quite compatible with Uniqueness, for each body of total evidence uniquely fixes what beliefs you ought to have. It's just that you can affect your evidential state by forming certain beliefs. Note that what has been said doesn't commit us one way or another on the question of whether the transition–formation of such a self-fulfilling belief–would count as rational, irrational, or arational.

We might then say something similar about intentions in a Buridan case. (Indeed, some philosophers think that intentions just *are* self-fulfilling beliefs. Velleman (1989) holds, roughly, that an intention to ϕ is a belief that you will ϕ in virtue of that very belief.) Here is how this would go: Relative to a state in which you lack both the intention to go right and the intention to go left, neither intention is rational. However, relative to a state in which you have the intention to go right, the intention to go right is uniquely rational, and relative to a state in which you have the intention to go left, the intention to go left is uniquely rational. These verdicts about the rationality of states are compatible with different views about the rationality of transitions - whether forming the intention to go right (or left) would be rational, irrational, or arational.

Adopting this position on the rationality of intentions requires that having an intention provide a reason for having that intention. Intentions must be self-supporting in this way. Before going into why this might be, note that the reason an intention provides for itself need not be a strong one. It only needs to be enough to break ties, like in Buridan's ass. So the reason an intention provides for itself could be lexically weaker than other reasons for or against intentions, and the proposal under discussion would still go through. So why might an intention provide a reason, even a very weak one, for itself? One might support this claim by holding that intending to perform some action provides a reason for performing that action, and so it derivately provides a reason for itself. The view that intentions provide reasons to carry them out is controversial (see Bratman (1987)), however. Alternatively, one might hold that having an

intention provides a reason for having that intention because reconsidering an intention carries costs, including both the cognitive costs involved in carrying out the subsequent deliberation and perhaps other costs incurred by having one's intentions fail to fulfill the stabilizing role widely thought to be central to intentions (Bratman (1987), Holton (2004), Broome (2013)). Intentions are important in large part because they serve as fixed points in our reasoning, providing a background against which we can then make further plans as we go. It would be bad for us if our intentions weren't by and large stable, and this might motivate the thought that intentions can provide reasons for themselves.[35] Again, the reason that an intention provides for itself might be very weak, and this allows us to maintain that intentions can often be irrational, namely in cases where the reasons they provide for themselves are overridden by other reasons.

Admittedly, we have not given a watertight case for uniqueness theses in the practical realm. But we hope to have shown that such theses are not as implausible as they might have initially seemed. They can be defended against apparent counterexamples by showing that, with a suitable interpretation of what counts as one's practical situation, the cases aren't counterexamples at all. Once we note that one's practical situation includes facts about one's tastes (or perhaps one's evidence about one's tastes), it no longer seems that one's practical situation will permit either a preference for chocolate over vanilla or vice versa. And once we say that one's practical situation includes facts about what intentions one has (and that intentions provide reasons for themselves), we no longer need to say that in a Buridan case, one's practical situation will permit either the intention to go left, or the intention to go right, or neither. If these defenses of uniqueness theses in the practical realm are on the right track, then our arguments for Uniqueness in the epistemic realm, which has been our main concern, don't prove too

---

[35]One might want to say instead that an intention you have at time t doesn't provide a reason to having that intention at time t, but instead that it provides a reason not to reconsider the intention at time $t + \epsilon$. We won't attempt to adjudicate this issue here.

much after all.

Conclusion

We have argued that the use of rational evaluations plays two central roles, and that in order for them to play these two roles, requirements of epistemic rationality must be very strict, in the sense that there is a uniquely rational belief state to be in, given a body of total evidence. First, rational evaluations are tied to deference. They serve to help pick out whom to defer to, so that judging that an agent's beliefs are rational commits you to deferring to that agent's beliefs, unless you know you have some relevant evidence that she lacks. But it only makes sense to defer to an agent's beliefs if she shares your epistemic standards, so that you can treat her as an epistemic surrogate. It must be that the beliefs the agent arrived at are the same ones that would be reached by applying your own epistemic standards to evaluate that agent's body of evidence. In this way, the deference role of rational evaluations requires that for each body of total evidence, there is just one belief state that it is rational to be in. Second, rational evaluations aid us in our contingency planning for what to believe in different circumstances. The most natural way of spelling out this connection between rational evaluations and planning is to say that if you judge that it would be rational to have a certain belief given a body of evidence, then you plan to hold that belief in the contingency in which you have that body of evidence. This works out well if your judgments conform to Uniqueness, so that you never judge more than one different belief state to be rational given a fixed body of total evidence. But if your judgments conform to Permissivism, attempting to spell out a connection between rational evaluations and contingency planning leads to either inconsistent or instrumentally irrational contingency plans. So much the worse for Permissivism.

It may be that Permissivists want a notion of rationality that plays different roles. Perhaps they want an evaluative notion that is more closely tied to consistency, and certainly there are lots of different belief states that are consistent with a given body of evidence. Or perhaps they want rationality to be linked with defensibility by one's own lights, and again, even fixing on a single body of evidence, there are many different belief states that are internally coherent. Permissivists might aso want to preserve the pejorative term 'irrational' for those who are doing worse in forming their beliefs than what is normal for humans and to save the term 'rational' for those who are doing fairly well.

If you want a notion of rationality that answers to these latter sorts of functions, that's fine. Perhaps we can live with many different notions which serve a variety of different functions.[36] But if you want a close connection between rationality, and truth, then it is natural to think that epistemic evaluations are tied to deference and contingency planning. And in order for epistemic evaluations to play their part in deference and planning, Uniqueness must be true. Rational requirements must be strict rather than slack.

References

Ariely, Dan. *Predictably Irrational: The Hidden Forces That Shape Our Decisions.* HarperCollins, 2008.

Ballantyne, Nathan and E. J. Coffman. "Uniqueness, Evidence, and Rationality." *Philosophers' Imprint* 11 (2011).

---

[36]Chalmers (2011) argues that one way to avoid descent into merely verbal disputes is to clearly state what theoretical roles the concept in question is supposed to play, and then to theorize about what would enable to the concept to best play those roles. This is the strategy we have taken here, focusing on the work that epistemically evaluative concepts do in aiding deference and planning.

Ballantyne, Nathan and E. J. Coffman. "Conciliationism and Uniqueness." *Australasian Journal of Philosophy* 90 (2012): 657–670.

Berridge, Kent C. "Wanting and Liking: Observations From the Neuroscience and Psychology Laboratory." *Inquiry* 52 (2009): 378–398.

Bratman, Michael. *Intention, Plans, and Practical Reason.* Center for the Study of Language and Information, 1987.

Briggs, Rachael. "Distorted Reflection." *Philosophical Review* 118 (2009): 59–85.

Broome, John. "Normative Requirements." *Ratio* 12 (1999): 398–419.

Broome, John. "Have We Reason To Do As Rationality Requires?—A Comment on Raz." *Journal of Ethics and Social Philosophy* (2005).

Broome, John. "Wide or Narrow Scope?." *Mind* 116 (2007): 359–370.

Broome, John. *Rationality Through Reasoning.* Wiley-Blackwell, 2013.

Chalmers, David. "Verbal Disputes." *The Philosophical Review* 120 (2011): 515–566.

Christensen, David. "Clever Bookies and Coherent Beliefs." *Philosophical Review* 100 (1991): 229–247.

Christensen, David. "Diachronic Coherence versus Epistemic Impartiality." *Philosophical Review* 109 (2000): 349–371.

Christensen, David. "Epistemology of Disagreement: The Good News." *Philosophical Review* 116 (2007): 187–217.

Christensen, David. "Higher-Order Evidence." *Philosophy and Phenomenological Research* 81 (2010): 185–215.

Cohen, Stewart. "A Defense of the (Almost) Equal Weight View." *The Epistemology of Disagreement: New Essays.* . Oxford University Press, 2013.

Craig, Edward. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis.* Oxford: Clarendon Press, 1990.

Dogramaci, Sinan. "Reverse Engineering Epistemic Evaluations." *Philosophy and Phenomenological Research* 84 (2012): 513–530.

Dogramaci, Sinan. "Communist Conventions for Deductive Reasoning." *Noûs* 47 (2013).

Dogramaci, Sinan and Sophie Horowitz. "An Argument for Uniqueness About Evidential Support." *Unpublished Ms.*

Dougherty, Trent and Patrick Rysiew. "Experience First." *Contemporary Debates in Epistemology.* Ed. Matthias Steup and John Turri. Blackwell, 2013. 2.

Elga, Adam. "Reflection and Disagreement." *Nous* 41 (2007): 478–502.

Gibbard, Allan. *Thinking How to Live.* Cambridge, Ma.: Harvard University Press, 2003.

Gibbard, Allan. "Rational Credence and the Value of Truth." *Oxford Studies in Epistemology Volume II.* Ed. Tamar Szabo Gendler and John Hawthorne. Oxford University Press, 2007.

Goldman, Alvin. "Epistemic Relativism and Reasonable Disagreement." *Disagreement.* Ed. Richard Feldman and Ted Warfield. Oxford University Press, 2009.

Greaves, Hilary and David Wallace. "Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility." *Mind* 115 (2006): 607–632.

Greco, Daniel. "Verbal Debates in Epistemology." *American Philosophical Quarterly* (Forthcoming).

Harman, Gilbert. "Reflections on Knowledge and its Limits." *Philosophical Review* 111 (2002): 417–428.

Hedden, Brian. "Time Slice Rationality." *Mind* (Forthcoming).

Holton, Richard. "Rational Resolve." *Philosophical Review* 113 (2004): 507–535.

Horowitz, Sophie. "Immoderately Rational." *Philosophical Studies* 167 (2013): 1–16.

Joyce, James. "A Defense of Imprecise Credences in Inference and Decision Making." *Philosophical Perspectives* 24 (2010): 281–323.

Kelly, Thomas. "Epistemic rationality as instrumental rationality: A critique." *Philosophy and Phenomenological Research* 66 (2003): 612–640.

Kelly, Thomas. "How to be an Epistemic Permissivist." *The Epistemology of Disagreement: New Essays.* . Oxford University Press, 2013.

Kolodny, Niko. "Why Be Rational?." *Mind* 114 (2005): 509–563.

Kolodny, Niko. "How Does Coherence Matter." *Proceedings of the Aristotelian Society* 107:1 (2007): 229263.

Kolodny, Niko. "Why Be Disposed to Be Coherent." *Ethics* 118 (April 2008): 437–463.

Lewis, David. "Immodest Inductive Methods." *Philosophy of Science* 38(1) (1971): 54–63.

Matheson, Jonathan. "Conciliatory Views of Disagreement and Higher-Order Evidence." *Episteme: A Journal of Social Epistemology* 6 (2009): 269–279.

Meacham, Christopher J. G. "Unravelling the Tangled Web: Continuity, Internalism, Non-Uniqueness and Self-Locating Beliefs." *Oxford Studies in Epistemology, Volume 3*. Ed. Tamar Szabo Gendler and John Hawthorne. Oxford University Press, 2010. 86.

Meacham, Christopher J. G. "Impermissive Bayesianism." *Erkenntnis* (2013): 1–33.

Moss, Sarah. "Time-Slice Epistemology and Action Under Indeterminacy." *Oxford Studies in Epistemology, Volume 5.* . Oxford University Press, Forthcoming.

Parfit, Derek A. *Reasons and Persons.* Oxford University Press, 1984.

Pettigrew, Richard and Michael Titelbaum. "Deference Done Right." *Philosophers' Imprint* 14 (2014): 1–19.

Rosen, Gideon. "Nominalism, Naturalism, Epistemic Relativism." *Nous* 35 (2001): 69–91.

Schafer, Karl. "Doxastic Planning and Epistemic Internalism." *Synthese* (forthcoming): 1–21.

Schoenfield, Miriam. "Permission to Believe: Why Permissivism Is True and What It Tells Us About Irrelevant Influences on Belief." *Noûs* 48 (2014): 193–218.

Schoenfield, Miriam. "Bridging Rationality and Accuracy." *The Journal of Philosophy* (Forthcoming).

Titelbaum, Michael G. "How to Derive a Narrow-Scope Requirement From Wide-Scope Requirements." *Philosophical Studies* 172 (2015): 535–542.

Velleman, David. *Practical Reflection.* Volume 94 . Princeton University Press, 1989.

Weisberg, Jonathan. "Conditionalization, Reflection, and Self-Knowledge." *Philosophical Studies* 135 (2007): 179–97.

White, Roger. "Epistemic Permissiveness." *Philosophical Perspectives* 19 (2005): 445–459.

Williamson, Timothy. *Knowledge and its Limits.* Oxford University Press, 2000.

Williamson, Timothy. "Scepticism and Evidence." *Philosophy and Phenomenological Research* 60(3) (2000): 613–628.