

My contribution to a debate with Roger White for Blackwell's *Contemporary Debates in Epistemology*, edited by Matthias Steup and John Turri. (Although written as a response to White, the paper is intended to be intelligible when read on its own.)

How to be an Epistemic Permissivist¹

Thomas Kelly
Princeton University

Roger's official statement of the thesis that he defends reads as follows:

Uniqueness: If an agent whose total evidence is E is fully rational in taking doxastic attitude D to P, then necessarily, any subject with total evidence E who takes a different attitude to P is less than fully rational.

Following Roger, I'll call someone who denies Uniqueness a *Permissivist*. In what follows, I'll argue against Uniqueness and defend Permissivism.

1. The Strength of Uniqueness

At an intuitive level, one immediate attraction of Permissivism is this: Uniqueness is an extremely strong thesis. We can think of Uniqueness as one possible answer to the following question: How much *slack* exists between the evidence and what it's reasonable to believe given the evidence? In these terms, the friend of Uniqueness thinks that there is never any slack, ever. On the other hand, the Permissivist thinks that in at least some possible cases, there is at least a little bit of slack. As this suggests, a Permissivist might very well think that there are many cases in which there is no slack at all, where there is one and only one response to the evidence that's the fully rational response.

¹This paper grew out of two APA sessions that took place seven years apart: in 2005, I served as a

I mention this possibility—that Uniqueness is false, even though there are many non-permissive cases—in part because of my conviction that this is where the truth lies. Suppose that I pull a coin out of my pocket at random in order to flip it. I invite you to consider the proposition that *the coin will land heads rather than tails*. How much credence should you invest in this proposition? Here it's quite natural to think that, given plausible assumptions about your evidence, you should divide your credence evenly between this proposition and its negation, and that if you did anything other than that, you would be responding less than perfectly to your evidence. This natural verdict is one that a Permissivist can embrace. (Although of course, not every Permissivist will embrace it). Moreover, a Permissivist might clear-headedly hold that the great majority of cases are non-permissive, in the way that this one at least initially appears to be.

One respect in which Permissivism is a very modest thesis then, is that it's compatible with there being relatively few permissive cases. Another respect in which it's a very modest thesis is that the Permissivist might think that what permissive cases there are, aren't all that permissive. At this point, it will be helpful to describe a realistic example that (unlike the coin case) seems to be a good candidate for a permissive case, at least as far as pre-theoretical intuition is concerned.

Suppose that six months before the U.S. presidential election, it's quite unclear whether the Democratic or the Republican nominee will win. (Although it is clear that one or the other will.) I possess a large body of information that I take to bear on this question. Some of this information makes it more likely that the Democrat will win, while some of it makes that outcome less likely. On balance, I regard it as somewhat more likely that the Democrat will win than not, so I invest somewhat more credence in

that proposition than in its negation. If I met someone who had exactly my evidence but was *extremely confident* that the Democrat will win, then I would regard this person as less reasonable than I am. (Perhaps he's in the grips of wishful thinking, or alternatively, pessimistic despair, and that accounts for why he's so confident). Similarly, if I met someone who had exactly my evidence but thought that the Republican was going to win, it would be natural for me to think that this person had made some kind of mistake in responding to our shared evidence. Suppose, however, that you and I agree on the basis of our common evidence that the Democrat is more likely than not to be elected. We similarly agree that although this outcome is more likely than the alternative, it's far from a sure thing. The only difference between us is this: you're a bit more cautious about the Democrat's prospects, and so give a bit less credence to the proposition that the Democrat will win than I do. Here there seems little pressure for me to conclude that you are less reasonable than I am. Moreover, the natural verdict about the case is that it's consistent with everything that's been stipulated so far that you and I might both be fully reasonable in our opinions about the election, despite the fact that those opinions are not identical. But if adding that further detail to the story does not render the story incoherent, then Uniqueness is false.

Again, someone might deny Uniqueness while thinking that what permissive cases there are resemble this one in relevant respects. So Uniqueness seems very strong. How strong is it exactly? Perhaps it matters here how we think about the psychological states to which it is taken to apply. To my mind, uniqueness seems most plausible when we think about belief in a maximally coarse-grained way, so that there are only three options with respect to a given proposition that one has considered: belief, disbelief, or

suspension of judgment. On the other hand, as we begin to think about belief in an increasingly fine-grained way, the more counterintuitive Uniqueness becomes. Consider a thought experiment. Suppose that when we meet the Alpha Centaurians, they differ from us in only one important respect: they routinely take up doxastic attitudes towards propositions that are extremely fine-grained compared to our own. So, for example, the Alpha Centaurians really do have psychological states such as *believing to degree .5436497 that the Democrat will win*, or *believing to degree .5122894 that it will rain tomorrow*. I assume that this is a perfectly coherent possibility. (We might even have empirical evidence that they have such attitudes; it shows up in their betting behavior, and so on.) The friend of Uniqueness might insist that, for any possible evidential situation, the evidence in that situation singles out some one, exact degree of belief that it is uniquely reasonable for the Alpha Centaurians to have, any slight deviation from which already counts as a deviation from perfect rationality. Moreover, this will be so no matter how fine-grained we make the propositional attitudes of the Alpha Centaurians. But as one cuts up the psychology more and more finely, Uniqueness looks increasingly counterintuitive. Even if we are inclined to think that the epistemic facts (i.e., facts about what it's reasonable to believe, given the evidence) are *sharp* and not fuzzy, could there really be no limit to their sharpness? At some point, one wants to say, there must be a range of (presumably adjacent) mutually exclusive attitudes, any one of which would be reasonable to hold, and no one of which is any more reasonable than any other within the range.

What should the friend of Uniqueness say about this? I think that the best move for her at this point is to appeal to so-called “mushy credence”. It's not really that there is

some range of permissible options. Rather, the uniquely reasonable thing for the Alpha Centaurians to do is to go vague over the ostensibly permissible range. On this way of thinking about it, one way of falling short of perfect reasonableness is to have overly precise degrees of belief: that amounts to treating your evidence as though it carries information that it doesn't carry. (And if the Alpha Centaurians are constitutionally incapable of having these coarser attitudes, then they are constitutionally incapable of full rationality.)

Although natural, the appeal to mushy credence in order to defuse the challenge carries risks, inasmuch as whether the mushy credence picture is ultimately viable is currently the subject of intense debate.² I don't propose to enter into that debate here. Instead, I'll simply note that it seems that the friend of Uniqueness has strong incentive to hope that this vigorously contested issue is resolved in one way rather than another.

2. A Jamesian argument for Permissivism

What has been said so far concerns only the intuitive (im)plausibility of Uniqueness. But even if it would be surprising if Uniqueness turned out to be true, perhaps that's where the arguments lead. In this section, I'll sketch one argumentative route by which someone might arrive at the conclusion that Uniqueness is false. For reasons that I'll explain, I think that someone who arrives at the conclusion that Uniqueness is false in this way should not feel especially threatened by the kinds of arguments offered by Roger.

² Recent critiques include Elga (2010) and (somewhat ironically, in my sense of the dialectic with respect to the permissiveness question is on the right track) White (2009). A recent defense of mushy credence is Joyce (2010).

How then might the Permissivist be thinking about things? Consider first a point emphasized by William James in his classic essay “The Will to Believe”. James noted that philosophers often talk about the importance of attaining truth and avoiding error, but that such talk tends to mask certain complexities. On the one hand, there is the goal of *not believing what is false*, a goal that can be successfully achieved with respect to a given issue by suspending judgment on that issue. On the other hand, there is the goal of *believing what is true*, for which suspending judgment is obviously insufficient.

Moreover, as James also emphasized, these two cognitive *desiderata* can pull in opposite directions. In general, the more value one gives to not believing what’s false about some issue, the more it behooves one to be relatively cautious or conservative in forming beliefs about that issue. That is, the more weight one gives to not believing something false, the more it makes sense to hold out until there is a great deal of evidence that *p* is true before taking up the belief that *p*. On the other hand, the more one values not missing out on believing the truth, the more it makes sense to take a somewhat more liberal attitude about how much evidence one expects before taking up the relevant belief. That is, to the extent that one is concerned to avoid *not believing p when p is in fact true*, one shouldn’t wait until there is overwhelming evidence in favor of *p* before taking up the corresponding belief.

My suggestion is that James’ observation is potentially highly relevant to our assessment of Uniqueness. Suppose that the evidence that you and I have that bears on some hypothesis *H* is *E*. Although it’s clear enough that *E* supports *H* over not-*H*, it’s not as though *E* is overwhelming evidence that *H* is true. Indeed, let’s suppose that this is a marginal case, in that *E* is just *barely sufficient* to justify believing *H*: if *E* were any less

supportive than it is, believing H on its basis would be positively *unreasonable*.

Recognizing that E suffices to justify belief in H, I take up the belief in response. I notice, however, that you don't take up the same belief, despite having the same evidence. Let's further stipulate that it's not as though you are dogmatically averse to believing H, or anything like that: in fact, if the evidence for H grows any stronger, than you too will become an H-believer in response.

In these circumstances, is there any chance that your refraining from believing H is reasonable, given that my believing H is reasonable? As someone who believes H, am I committed to thinking that you're guilty of making some kind of mistake, that you've misjudged the probative force of our shared evidence? Before attempting to answer these questions, let's add one further detail to the story. With respect to the question at hand, you're a bit more concerned than I am to avoid believing what's false, while I'm a bit more concerned than you are to not miss out on believing what's true in virtue of suspending judgment. That is, there is a subtle difference in our cognitive goals, or rather, in the relative weights that we give to the two cognitive goals with respect to the question at hand.

Once this further stipulation is added, your not believing H on the basis of evidence that is only marginally sufficient to justify such belief seems eminently reasonable. As an H-believer, if I learned that we differed in our cognitive goals in this way, I would be disinclined to conclude that the manner in which you are responding to our shared evidence is unreasonable, even though it differs from my own. In fact, I might even think that if you were responding to the evidence in any other way than you are, then *that* would be unreasonable, given your cognitive goals. Moreover, notice that making such a

judgment has no tendency to make me insecure in my conviction that I am also responding to the evidence in a reasonable way, given my cognitive goals. The upshot: subtly different ways of responding to the same body of evidence seem equally reasonable, given corresponding differences in the weights that we give to our shared cognitive goals.

Notice that this route to rejecting Uniqueness does not depend on thinking that “anything goes” with respect to the relative weights that can be permissibly assigned to the two cognitive goals, or even that there is much in the way of permissible variation here at all. So long as there are at least some possible cases in which it is reasonable for different individuals to give at least somewhat different weights to the goals, then this can affect how much evidence they should hold out for before they take up the relevant belief. There will then be possible bodies of evidence that fall within the relevant margin, bodies of evidence relative to which belief is a perfectly reasonable response on the part of the person who is somewhat more concerned to believe the truth, and relative to which suspension of judgment is a perfectly reasonable response on the part of the person who is somewhat more concerned to avoid believing what is false.

It might be objected that this route to rejecting Uniqueness depends on thinking about belief as an all-or-nothing matter, as opposed to a matter of degree. According to this line of thought, the “James point” only comes into play when one combines a fine-grained notion of evidence with a coarse-grained picture of belief. For once that combination is in place, then it seems like the following kind of threshold question is appropriate: How much evidence does one need that p is true, before it becomes appropriate to believe p ? (Presumably, just a little bit of evidence that p is true isn’t

enough.) And once questions about *where* the evidential threshold is located are put in play, it becomes natural to ask why the threshold is where it is, as opposed to someplace higher or lower. It is at this point that James' observation seems to become relevant, inasmuch as it is natural to think that one of the factors that can make a difference to where the threshold is located is the relative weight given to the two cognitive goals. Intuitively, as more relative weight is given to not believing what's false, that tends to exert some upward pressure on the threshold. (More evidence will be required, before it makes sense to take up the belief.) On the other hand, as more weight is given to not missing out on the truth by suspending judgment, that tends to exert some downward pressure on the threshold.

The suggestion of the objector is that (i) James' observation about the potentially competing cognitive goals only gets traction against the background of this threshold picture, but that (ii) we can and should dispense with the threshold picture by doing epistemology in terms of credences or degrees of belief as opposed to all-or-nothing beliefs. Once we think in terms of more fine-grained doxastic states, there is no longer any question about where the threshold is, or which factors play a role in determining where it lies, because there is no need for a threshold at all. The only rule is: proportion your credence to the strength of your evidence. When one's evidence for *p* is very weak, one should invest very little credence in *p*; as one's evidence for *p* grows stronger, one's credence should rise accordingly. Thus, there is never any question about how much evidence one needs before belief (as opposed to suspension of judgment) is appropriate.

This is a tempting line of thought. In fact, for most of the time that I have been thinking about these issues, I believed that it was correct. I now think that it is mistaken.

Rachiele (unpublished) argues compellingly that “the James point” holds even in theoretical frameworks that employ credences rather than all-or-nothing beliefs. For even if we do our theorizing in terms of credences, there will still be different dimensions relative to which we can evaluate the accuracy of those credences. Thus, one natural goal is that of minimizing the gradational inaccuracy of one’s credences.³ Relative to this goal, one set of credences is more accurate than another just in case it has a lower mean gradational inaccuracy. Another desideratum is that of lowering the *variance* in the gradational inaccuracy of one’s credences. Even if one set of credences is superior to a second set in having lower mean gradational inaccuracy, the second set might be superior with respect to the variance property. Significantly, neither of these cognitive desiderata seems to be lexically prior to the other (Rachiele pp.11-12). Although these two accuracy related desiderata are complementary, the fact that they are distinct means that trade-offs will sometimes be necessary. (Compare: although the goals of believing truths and not believing falsehoods are complementary—doing well with respect to one is generally helpful with respect to the other—the fact that they are different goals creates the need for trade-offs; the optimal strategy for the achievement of one is not the optimal strategy for the achievement of the other.) On the plausible assumption that different individuals might reasonably differ, at least marginally, in how they resolve these trade-offs, different patterns of belief revision might be appropriate relative to the different resolutions. The upshot is that, to the extent that it works at all, the Jamesian route to

³ If we measure credences with real numbers, we can measure the gradational inaccuracy of a credence by taking the absolute value of the difference between that credence and the actual truth value of the target proposition (where “the actual truth value of the target proposition” = 1 just in case the proposition is true, and 0 just in case the proposition is false). For a useful discussion of gradational accuracy, see Joyce (1998).

vindicating a permissive epistemology sketched in this section works just as well in a framework that employs credences instead of all or nothing beliefs.

3. Interpersonal versus intrapersonal slack

The Permissivist should not rest her case on this Jamesian line of thought.⁴ But even if it ultimately fails to undermine Uniqueness, I believe that there is an important lesson to be learned from it. The lesson concerns the need to distinguish sharply between statements of Uniqueness that have what I will call *interpersonal import* from those that do not.

As noted above, someone who is impressed with James' point might think that the following kind of case is possible: if you are somewhat more concerned than I am to avoid believing what's false about whether *p*, and I am somewhat more concerned than you are to not miss out on believing the truth about *p* by suspending judgment, then there are possible bodies of evidence *E* such that:

- (1) The uniquely reasonable response for you is to suspend judgment about whether *p*, and
- (2) The uniquely reasonable response for me is to believe *p*.

⁴ Having just noted why I am unconvinced by one natural objection, let me mention what I take to be a better (even if more idiosyncratic) reason for skepticism. James' point seems to depend upon thinking about epistemic rationality in a particular way. Specifically, it seems to depend on thinking that epistemic rationality is really a special case of instrumental or means-end rationality, viz. instrumental rationality in the service of one's cognitive goals, goals such as believing what's true and not believing what's false. This is an extremely natural way of thinking about epistemic rationality, and I believe that it is widely accepted within contemporary epistemology (even if many of those who accept it do so only implicitly). Nevertheless, I think that there are good reasons to be skeptical of the general picture. On this, see my "Epistemic Rationality as Instrumental Rationality: A Critique" (2003).

Generalizing this, one might arrive at a view that is *permissive across individuals* but that is *impermissive with respect to the range of options open to any particular individual*. Someone who holds a view of this kind is prepared to countenance *interpersonal* slack (different individuals possessing the same evidence might believe differently, and each be reasonable in believing as they do) but deny the existence of *intrapersonal* slack (for any given individual, there is a uniquely reasonable thing for *her* to believe given her evidence). Roger's official statement of Uniqueness is clearly inconsistent with this kind of view; in this sense, it has interpersonal import. But other principles in the near neighborhood might lack such import. As a possible example, consider Roger's statement of Uniqueness in his seminal 2005 paper on the topic, which I will call Uniqueness*:

Uniqueness*: Given one's total evidence, there is a unique rational doxastic attitude that one can take to any proposition (2005: 445).

On what I take to be its most natural reading—at least, its most natural reading when it is read in isolation--this principle says the following: *there is no slack for a single subject*. (Once you specify what her evidence is, that locks in what it is reasonable for her to believe.) But the principle is silent on whether some other individual with the same total evidence might take up a different attitude towards the same proposition that's fully reasonable. It thus lacks interpersonal import. When read in this way, Uniqueness* is significantly weaker than Uniqueness, which explicitly rules out the possibility of interpersonal slack.⁵

⁵ Although I have yet to see the point appear in print, the importance of distinguishing between principles that have interpersonal import and principles that lack such import in discussions of

Although principles that lack interpersonal import raise philosophically interesting questions in their own right, I believe that there are good reasons to think that the issue that philosophers have been concerned with in the literature on this topic concerns the truth of principles that *do* have interpersonal import, like Uniqueness. First, many philosophers (including Roger in his contribution to this volume) have suggested that there are important connections between this debate and the debate over the epistemic significance of disagreement.⁶ And it is hard to see why a principle that did not have any interpersonal import would be thought relevant to the latter debate.

More importantly, certain views in epistemology that everyone would be inclined to treat as paradigms of “permissive” views seem to be consistent with uniqueness principles that lack interpersonal import. Consider, for example, a subjective Bayesian who thinks that the only rational constraints on one’s doxastic corpus are the following: (i) One’s initial probability distribution must be coherent (beyond that, “anything goes”), and (ii) one must update one’s credences by conditionalization upon gaining new information. The subjective Bayesian should presumably count not only as a Permissivist, but as an “Extreme Permissivist” in Roger’s sense. For she thinks that even if you and I have exactly the same evidence, I might be extremely confident that the Democrat is going to win the election, and you might be extremely confident that the Republican is going to win (while both being perfectly reasonable). Nevertheless, the subjective Bayesian might very well accept Uniqueness*, given a reading of that

uniqueness is one on which a number of us have apparently independently converged, including Lee (manuscript), Meacham (manuscript) and Rachiele (manuscript).

⁶ See, e.g., Feldman (2006), Christensen (2007), Kelly (2010), Ballantyne and Coffman (2011, forthcoming), Cohen (forthcoming) and Matheson (2011). On the epistemology of disagreement, see especially Christensen (2009), Kelly (2005) and the essays collected in Feldman and Warfield (2010).

principle on which it lacks interpersonal import. Given the total evidence that I have, there really is one place that I should be, and if I were anywhere else, I would be less than fully reasonable. What the subjective Bayesian will deny is that it follows from this that *you* are less than fully reasonable, if you are somewhere else.

What I have argued for thus far in this section is the following. First, there is a significant gap between statements of uniqueness that have interpersonal import and those that lack such import: the former are significantly stronger than the latter, as witnessed by the fact that there are positions in contemporary epistemology with actual, flesh-and-blood proponents that are inconsistent with the former and consistent with latter. Second, the debate in the literature on this topic is really about whether the stronger principles are true. Notably, however, many of the kinds of considerations that friends of Uniqueness offer in its favor actually seem best suited to establishing the *weaker* principles, principles that lack interpersonal import. For example, both Roger's "arbitrariness argument" and his "arbitrary switching" cases invite us to consider how things look from the perspective of a *single subject*, and whether we can make good sense of the possibility that such a subject might be faced with a choice between incompatible but perfectly rational options with respect to his or her beliefs. As I understand them, these arguments have the form of *reducio ad absurdum* arguments. We are invited to suppose (for purposes of *reducio*) that a particular subject is in a permissive case and knows that she is. Roger then proceeds to ingeniously draw out the many apparent absurdities that seem to follow from these suppositions. For example, the subject might decide to switch her opinions randomly back and forth between the ostensibly permissible options, by popping a pill, or some other mechanism that has nothing to do

with the truth, and then rationally maintain her latest opinion in the full knowledge that this is how she had arrived at it. We are then invited to conclude that this shows that there is something absurd about the original supposition, viz. that there could be such cases.⁷

However, I don't think that arguments of this general form could possibly establish anything as strong as Uniqueness, a principle that has interpersonal as well as intrapersonal import. This is because a theorist might very well *agree* with the conclusion that there is something incoherent or absurd about the supposition that a person could be in a situation in which she had rationally permissible doxastic options, while holding that some other person (say, someone with a different prior probability distribution) might reasonably believe something else on the basis of the same evidence. The kind of subjective Bayesian described above is an example of such a theorist. Notice that this possible combination of view is no mere occupier of logical space, something cooked up in order to avoid having to accept Uniqueness; rather, it follows immediately from independently motivated positions in epistemology that have prominent defenders.

Of course, that isn't the end of the story. If the kind of arbitrariness arguments put forward by Roger do suffice to establish that there is no intrapersonal slack, then one might attempt to argue from that intermediate conclusion or lemma to the stronger conclusion that there is no interpersonal slack, either. For example, suppose that the following bridge principle could be established:

⁷ It is sometimes objected to this style of argument that it will inevitably fall short of showing that there are no permissive cases; rather, at best it shows that even if there are permissive cases, one could never know that one was in one. In effect, the objection is that *I'm in a permissive case* might be a "blind spot proposition", in the sense of Sorensen (1988). Like Roger, I doubt that this objection ultimately has much force, inasmuch as the assumption needed to close the gap in the argument, viz. *if one were in a permissive case, then at least in principle one could know that one was* seems extremely plausible.

BRIDGE: If it is currently reasonable for some subject S1 to hold doxastic attitude D1 towards P on the basis of evidence E, and it either is or would be reasonable for some other possible subject S2 to hold a different doxastic attitude D2 towards P on the basis of evidence E, then it is also currently reasonable for S1 to hold doxastic attitude D2 instead of D1 towards P on the basis of evidence E.

If the principle BRIDGE could be established, and if Roger's arguments suffice to show that there is no intrapersonal slack, then we could conclude that there is no interpersonal slack either, by reasoning in the following way:

If there were a case that was interpersonally permissive, then there would be a case that was intrapersonally permissive (by BRIDGE). But Roger's arguments show that there are no intrapersonally permissive cases. Therefore, there are no interpersonally permissive cases, either.

However, the principle BRIDGE is far from obvious. Indeed, many would flatly deny that it is true. In any case, it's the kind of thing for which we should insist on arguments. In the absence of actually looking at what arguments might be offered in its favor, it's difficult to say anything very definitive about the prospects for establishing it (or some sufficiently close principle). So here let me simply record my conviction that the gap between "no intrapersonal slack" and "no interpersonal slack" will not be an easy one to bridge, and that there will be plenty of promising points along the way for the Permissivist to dig in her heels.

Notice, for example, that any reason that might be offered for thinking that conditionalization is the rule that governs belief change over time will cast doubt on intrapersonal slack (given one's initial prior probability distribution, and the evidence that one has accumulated since then, there is some particular probability distribution that one

would have now if one were ideally rational) but *won't* be a reason for thinking that there is no interpersonal slack.

More generally, the fact that there are substantive coherence requirements that constrain permissible combinations of beliefs at the intrapersonal level (what I rationally believe constrains what else I can rationally believe), requirements that do not in general carry over to the interpersonal level (what I rationally believe does not constrain what you can rationally believe, in anything like the same way) generates obstacles for the project of arguing from the putative absence of intrapersonal slack to the nonexistence of interpersonal slack. For example, in Roger's "belief toggling" cases, we are asked to place ourselves in the situation of an agent who can, by means of a pill, swap his current belief that *p* for a belief that not-*p*. (The case is designed to bring out the odd consequences of Extreme Permissiveness, or at least, of taking oneself to be in an extremely permissive case, in which believing either *p* or not-*p* on the basis of one's evidence would be perfectly reasonable.) But it seems that *everyone*—including Extreme Permissivists—will have good reason to deny that one could end up with a fully reasonable belief that not-*p* in this way. After all, the proposition *p* stands in logical and evidential relations to countless other propositions that are potential objects of belief (or disbelief) for me. So if I am currently a *p*-believer who is fully rational, the fact that I am fully rational depends in part on the fact that my belief that *p* perfectly coheres with a large number of other doxastic attitudes that I take towards other propositions. When I contemplate swapping my current belief that *p* for a belief that not-*p*, I should recognize this as a change that is bound to make me less coherent—and therefore, less rational—than I am now. This seems like a good reason to decline to take the pill. But the Extreme

Permissivist can say this, along with everyone else. For it is enough for the truth of Extreme Permissivism if the following is possible: some other person with my evidence is fully rational in believing not-p rather than p. If there is such a person, then she will presumably differ from me a great deal in her doxastic states, inasmuch as her belief that not-p will cohere perfectly well with all of her other doxastic attitudes towards propositions that stand in logical and evidential relations to not-p. Of course, the Extreme Permissivist should also say that the fully rational not-p believer has a good reason to decline to take a belief-toggling pill that will reverse her belief about whether p, inasmuch as such a change is bound to make her less coherent, and therefore, less rational, than she is now.⁸

My advice to the Permissivist then, is that she should resist the slide from

Given that my evidence is E, there is some doxastic attitude D that is the only fully rational doxastic attitude for me to take towards proposition p

to

Given evidence E, there is some doxastic attitude D that is the only fully rational doxastic attitude for anyone to take towards proposition p (including including all of those with different prior probability distributions, or those who assign different weights to the cognitive goals, etc.)

But Uniqueness requires the truth of the latter claim.

4. Evidential Support

⁸ So perhaps we should think of the pills as altering not simply one's doxastic attitude towards the target proposition p, but as altering a large cluster of one's opinions, namely, all of those opinions about propositions that stands in logical or evidential relations to the proposition p. However, it's not obvious that such a change in the case is innocent, or that once the case is changed in this way it elicits the same intuitive responses that the original version was designed to elicit (at least in my case, it doesn't).

In addition to considerations having to do with arbitrariness, Rogers also offers an argument that appeals to the nature of evidential support:

1. Necessarily, it is rational for S to believe P iff S's total evidence supports P.
2. If E supports P then necessarily E supports P.
3. It cannot be that E supports P and E supports not-P.
4. Therefore, if an agent whose total evidence is E is rational in believing P, then it is impossible for an agent with total evidence E to rationally believe not-p (p.x).

Notice that this argument is directed at Extreme Permissiveness, so one could accept it while consistently denying Uniqueness. Nevertheless, it is worth exploring how the argument might be resisted.

One point of potential resistance that will appeal to many is this: the argument relies on the assumption that the relation of evidential support should be understood as a two place relation ("E supports P") as opposed to a three place relation ("E supports P relative to background Z"). It is uncontroversial that whether a particular *piece* of evidence supports a given hypothesis often depends on considerations of background knowledge or theory. In the context of Roger's argument, however, what matters is whether the relation of evidential support should be understood as a two- or three place relation when what is at issue is the bearing of one's *total* evidence on particular hypotheses. Suppose that we take one's total evidence E to include everything that one has learned. Notably, even on this inclusive understanding of what is included in E, orthodox confirmation theorists will insist that the relation of support should be understood as a three place relation, inasmuch as whether evidence E supports P (or the extent to which it supports P) will depend on the agent's initial probability distribution. A philosopher who thinks that the support relation

is a three place relation will thus insist on rewriting the premises of Roger's argument to reflect that fact:

1* Necessarily, it is rational for S to believe P iff S's total evidence supports p relative to S's prior probability distribution.

2* If E supports P relative to a prior probability distribution then necessarily E supports P relative to that prior probability distribution.

3* It cannot be that E supports P relative to a probability distribution and E supports not-P relative to that prior probability distribution.

Once the premises are rewritten in this way, however, even an Extreme Permissivist can happily accept them, for she can then point out that the argument from 1*-3* to 4 is invalid. Rather, what follows from premises 1*-3* is something like the following:

4* If an agent whose total evidence E is rational in believing P given her prior probability distribution, then it is impossible for an agent with total evidence E and the same prior probability distribution to rationally believe not-P.

But this conclusion falls well short of the original conclusion 4 and is consistent with Extreme Permissivism. For even if specifying an agent's total evidence and her prior probability distribution suffices to pin down some doxastic attitude as the uniquely reasonable one, it does not follow that merely specifying her total evidence suffices to do the same. More specifically, an Extreme Permissivist might hold that while an agent with total evidence evidence E might be reasonable in believing (or investing high credence in) P given her prior probability distribution, another agent with the same total evidence might be reasonable in believing (or investing high credence in) not-P given his different

prior probability distribution. (Here again the gap between “no intrapersonal slack” and “no interpersonal slack” is significant.)⁹

Of course, even if evidential support is in fact better understood as a three place relation than a two place relation, it doesn't follow that Extreme Permissivism is true. For it might be that there are substantive rationality constraints on prior probability distributions (that is, constraints beyond that of coherence), constraints that guarantee that it is impossible for an extremely permissive case to arise. Even if *that* is true, however, it doesn't follow that Uniqueness is true, for the rationality constraints might be such as to allow for at least some moderately permissive cases. What would vindicate Uniqueness is if it turned out that there is some uniquely reasonable prior probability distribution, which at least in this context would be tantamount to thinking of the relation of evidential support as a two place rather than a three place relation. So what the Permissivist should claim is this: (i) the relation of evidential support is best understood as a three place relation, and (ii) there is no uniquely rational starting point for all agents. But of course, many philosophers are already committed to thinking exactly this.

I don't imagine that any of this is news to Roger. In fact, he is quite modest in his claims for the argument. As he puts it: “My point is just that avoiding this conclusion [i.e., 4. above] appears to require a departure from very natural ways of thinking about evidence and rationality (p.x). I think that that's completely fair. In particular, I think that (e.g.) understanding the evidential support relation as a three place rather than a two place relation *does* involve a certain “departure from very natural ways of thinking about evidence”, inasmuch as much of our ordinary thought and talk about evidence suggests

⁹ Meacham (manuscript) emphasizes the difference made by thinking about the support relation as three place rather than two place in the course of criticizing a similar argument in White (2005).

the latter understanding when taken at face value. However, the fact that this way of avoiding the conclusion of the argument involves a departure from very natural ways of thinking about evidence and rationality should not be confused with claim that it is an *ad hoc* response to the argument, or even that it should be regarded as a costly one. After all, the fact that many contemporary philosophers think that (i) and (ii) are true is not attributable to their desire to avoid the conclusion of Roger's argument, or any similar argument. Rather, what popularity (i) and (ii) enjoy is largely due to a common perception that these are among the lessons to have emerged from the systematic investigation of the nature of confirmation that has been pursued by philosophers and others in the decades since World War II.

Here the general trajectory of confirmation theory in the 20th century is perhaps significant. Carnap's original vision for an "inductive logic" was that of a system that would assign a unique "degree of confirmation" that would attach to any hypothesis given a particular body of evidence. (The fact that this was a *desideratum* is perhaps a testament to the naturalness of thinking about the relation of evidential support in the way that Roger's argument requires.) But Carnap ultimately abandoned this ambitious vision as unworkable, and he and many of those who followed him in the development of quantitative confirmation theory came to advocate more liberal accounts of confirmation. Thus, for many contemporary philosophers the assumptions about evidential support that are needed to resist Roger's argument are independently motivated: in replying to Roger's argument along the lines suggested here, such philosophers need not say anything that they did not already believe about evidence or rationality. From such a perspective, even if resisting the argument does involve a departure from a very natural

way of thinking about evidence and rationality, whatever theoretical costs are involved in such a departure have already been judged worth paying.

Of course, perhaps those who embraced more liberal views of confirmation did so for bad reasons. Notably, Roger has recently attempted to rehabilitate a version of the Principle of Indifference, a project that many had written off as hopeless.¹⁰ Success in that venture would undoubtedly lend the argument considered in this section a dialectical effectiveness that it currently lacks. For this reason as well as for others, the debate over epistemic permissiveness is surely a long way from over.

¹⁰ See White (2009). For a critique, see Meacham (manuscript).

REFERENCES

- Nathan Ballantyne and E. J. Coffman (2011). "Uniqueness, Evidence, and Rationality". *Philosophers' Imprint* 11 (18).
- Nathan Ballantyne and E. J. Coffman (forthcoming). "Conciliationism and Uniqueness". *Australasian Journal of Philosophy*: 1-14.
- David Christensen (2007). "Epistemology of Disagreement: The Good News". *Philosophical Review* 116 (2): 187-217.
- David Christensen (2009). "Disagreement as Evidence: The Epistemology of Controversy". *Philosophy Compass* 4 (5): 756-767.
- Stewart Cohen (forthcoming). "The Evidential Weight of Disagreement". In David Christensen and Jennifer Lackey (eds.), *The Epistemology of Disagreement: New Essays* (Oxford: Oxford University Press).
- Earl Conee (2010). "Rational Disagreement Defended". In Feldman and Warfield, pp.69-90.
- Adam Elga (2010). "Subjective Probabilities Should Be Sharp". *Philosophers' Imprint* 10 (05).
- Richard Feldman (2006). "Epistemological Puzzles about Disagreement". In Stephen Hetherington (ed.), *Epistemology Futures*. Oxford University Press, pp.216-236.
- Richard Feldman and Ted Warfield, eds. (2010). *Disagreement*, OUP.
- Alvin Goldman (2010). "Epistemic Relativism and Reasonable Disagreement". In Feldman and Warfield, pp. 187-215.
- William James (1897). "The Will to Believe". In *The Will to Believe and Other Essays in Popular Philosophy*. New York: Longmans, Green, and Co.
- James M. Joyce (1998). "A Nonpragmatic Vindication of Probabilism". *Philosophy of Science* 65 (4): 575-603.
- James M. Joyce (2010). "A Defense of Imprecise Credences in Inference and Decision Making". *Philosophical Perspectives* 24 (1): 281-323.
- Thomas Kelly (2003). "Epistemic Rationality as Instrumental Rationality: A Critique". *Philosophy and Phenomenological Research* 66 (3): 612-640.

Thomas Kelly (2005). "The Epistemic Significance of Disagreement". In John Hawthorne and Tamar Gendler (eds.), *Oxford Studies in Epistemology, Volume 1*. OUP pp.167-196.

Thomas Kelly (2010). "Peer Disagreement and Higher Order Evidence". In Feldman and Warfield, pp. 111-174.

Mathew Lee (manuscript). "Getting Clear About the Uniqueness Thesis".

Jonathan Matheson (2011). "The Case for Rational Uniqueness". *Logic and Episteme* 2 (3): 359-373.

Christopher J.G. Meacham (manuscript). "Impermissive Bayesianism".

Joseph Rachiele (manuscript). "Deniers of Uniqueness, Do Not Fear".

Roy A. Sorensen (1988). *Blindspots*. Oxford University Press.

Roger White (2009). "Evidential Symmetry and Mushy Credence". In *Oxford Studies in Epistemology*. Oxford University Press, pp.445-459.