

Chapter Twelve

Can Evidence Be Permissive?

Roger White defends a negative answer to this question. He begins with what is known as the Uniqueness principle: the total evidence relevant to proposition P permits one and only one attitude towards P. Philosophers who deny Uniqueness are permissivists. They hold that there are possible cases in which the total evidence regarding P allows for more than just one attitude towards P; it might, for example, allow believing P and suspending judgment about P. According to what White calls strong permissivism, one and the same body of evidence might even permit believing P and permit believing \sim P. Using a variety of thought experiments involving belief-inducing pills, White argues that strong permissivism should be rejected. In his response to White, Thomas Kelly distinguishes between *intrapersonal* and *interpersonal* slack. If my total evidence permits two different attitudes towards P, that's intrapersonal slack; if you and I have the same total evidence and our shared evidence permits one attitude for you and another for me, that's interpersonal slack. Suppose you value collecting truths more than avoiding error; I value error avoidance more than truth collection. In that case, our shared evidence might make believing P permissible for you and suspending judgment about P permissible for me. Further avenues towards forms of interpersonal slack are open to subjective Bayesians. Hence, Kelly argues, there is a significant gap between principles against intrapersonal slack and principles against interpersonal slack, a gap that he thinks is not easy to bridge.

Evidence Can Be Permissive

Thomas Kelly

Roger White'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.

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 US presidential election, it is 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 none 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.¹ 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.

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" (1897). 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's 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's 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's 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. Joseph 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.² 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, unpublished, 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 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.³ 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's 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*.

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, p. 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.⁴

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.⁵ 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 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 *reductio ad absurdum* arguments. We are invited to suppose (for purposes of *reductio*) 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, namely that there could be such cases.⁶

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 intra-personal 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:

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 non-existence 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.⁷

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 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

In addition to considerations having to do with arbitrariness, Roger 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. 314).

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 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.)⁸

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. 315)." I think that that's completely fair. In particular, I think that (for example) 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 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 the 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 twentieth 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.⁹ 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.

Acknowledgments

This paper grew out of two APA sessions that took place seven years apart: in 2005, I served as a commentator on Roger's seminal paper "Epistemic Permissiveness," and in 2012 I took part in an invited symposium (along with Roger and E.J. Coffman) on "The Uniqueness Thesis." I'd like to thank the audiences present on those occasions for their feedback and Michael Titelbaum for organizing the symposium and inviting me to participate in it. I also discussed these issues in two seminars at Princeton University, one of which was co-taught with Bas van Fraassen; thanks to Bas and the students for much helpful discussion. Finally, special thanks to Roger for the intellectual stimulation provided by his work on this topic and our exchanges about it over the years. Although I've been thinking about his challenging arguments for a long time now, I certainly don't pretend to have gotten close to the bottom of them here.

Notes

- 1 Recent critiques include Elga (2010) and (somewhat ironically, if 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).
- 2 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).
- 3 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's 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, namely 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).
- 4 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 uniqueness is one on which a number of us have apparently independently converged, including Lee (manuscript), Meacham (manuscript), and Rachiele (manuscript).
- 5 See, for example, Feldman (2006), Christensen (2007), Kelly (2010), Ballantyne and Coffman (2011, 2012), Matheson (2011), and Cohen (2013). On the epistemology of disagreement, see especially Christensen (2009), Kelly (2005), and the essays collected in Feldman and Warfield (2010).
- 6 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, namely *if one were in a permissive case, then at least in principle one could know that one was* seems extremely plausible.

- 7 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 stand 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).
- 8 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).
- 9 See White (2009). For a critique, see Meacham (manuscript).

References

- Ballantyne, N. and Coffman, E.J. (2011) Uniqueness, evidence, and rationality. *Philosophers' Imprint* 11 (18).
- Ballantyne, N. and Coffman, E.J. (2012) Conciliationism and uniqueness. *Australasian Journal of Philosophy*: 90 (4): 657–670.
- Christensen, D. (2007) Epistemology of disagreement: the good news. *Philosophical Review* 116 (2): 187–217.
- Christensen, D. (2009) Disagreement as evidence: the epistemology of controversy. *Philosophy Compass* 4 (5): 756–767.
- Cohen, S. (2013) A tentative defense of the equal weight view. In *The Epistemology of Disagreement: New Essays*, ed. D. Christensen and J. Lackey. Oxford: Oxford University Press.
- Elga, A. (2010) Subjective probabilities should be sharp. *Philosophers' Imprint* 10 (5).
- Feldman, R. (2006) Epistemological puzzles about disagreement. In *Epistemology Futures*, ed. S. Hetherington, pp. 216–236. Oxford: Oxford University Press.
- Feldman, R. and Warfield, T. eds. (2010) *Disagreement*. Oxford: Oxford University Press.
- James, W. (1897) The will to believe. In *The Will to Believe and Other Essays in Popular Philosophy*. New York: Longmans, Green, and Co.
- Joyce, J.M. (1998) A nonpragmatic vindication of probabilism. *Philosophy of Science* 65 (4): 575–603.
- Joyce, J.M. (2010) A defense of imprecise credences in inference and decision making. *Philosophical Perspectives* 24 (1): 281–323.
- Kelly, T. (2003) Epistemic rationality as instrumental rationality: a critique. *Philosophy and Phenomenological Research* 66 (3): 612–640.
- Kelly, T. (2005) The epistemic significance of disagreement. In *Oxford Studies in Epistemology*, Vol. 1, ed. J. Hawthorne and T. Gendler, pp. 167–196. Oxford: Oxford University Press.
- Kelly, T. (2010) Peer disagreement and higher order evidence. In *Disagreement*, ed. R. Feldman and T. Warfield, pp. 111–174. Oxford: Oxford University Press.
- Lee, M. (manuscript) Sharpening up the Uniqueness thesis.
- Matheson, J. (2011) The case for rational Uniqueness. *Logic and Episteme* 2 (3): 359–373.
- Meacham, C.J.G. (manuscript) Impermissive Bayesianism.
- Rachiele, J. (unpublished).
- Sorensen, R.A. (1988) *Blindspots*. Oxford: Oxford University Press.
- White, R. (2005) Epistemic permissiveness. *Philosophical Perspectives* 19: 445–459.
- White, R. (2009) Evidential symmetry and mushy credence. In *Oxford Studies in Epistemology*, pp. 445–459. Oxford: Oxford University Press.

Further Reading

- Conee, E. (2010) Rational disagreement defended. In *Disagreement*, ed. R. Feldman and T. Warfield, pp. 69–90. Oxford: Oxford University Press.
- Goldman, A. (2010) Epistemic relativism and reasonable disagreement. In *Disagreement*, ed. R. Feldman and T. Warfield, pp. 187–215. Oxford: Oxford University Press.

Evidence Cannot Be Permissive

Roger White

1 Permissivism and Disagreement

In figuring out what to believe we look to our evidence. What sort of guidance does the evidence give? How stringent are the constraints it puts on rational belief? To sharpen our question, consider Alice and Bob who happen each to have the same body of evidence E pertaining to some matter P. They are independent inquirers unaware of each other's opinions. Suppose further that Alice and Bob are *ideally rational agents*, and so their beliefs are rationally impeccable. Might their opinions on whether P differ? If so, by how much? Is it possible, given the description above, that Alice is convinced that P while Bob thinks that not-P? If not, might it still be that Alice is a little more confident that P than Bob is?

Proponents of the Uniqueness thesis will say that Alice and Bob must be of exactly the same opinion.

Uniqueness: There is just one rationally permissible doxastic attitude one can take, given a particular body of evidence.

Or more precisely,

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.¹

Those who deny Uniqueness are *permissivists* about epistemic rationality.² According to permissivism, the evidence leaves us with some leeway as to what to believe. (Permissivists may disagree about how much leeway is allowed in general and in particular cases.)

It will be useful here to have labels for a stronger and weaker permissive thesis.

Strong permissivism: There are cases in which it is rationally permissible to believe P, but it is also rationally permissible to believe not-P instead, given the very same evidence.

Moderate permissivism: There are cases in which there is more than one rationally permissible *degree of confidence* one can have in P, given the same evidence.

Call a case in which an agent's evidence E does not determine a uniquely rational attitude to take to a proposition P a *permissive case*. A permissivist believes that there are some permissive cases, but needn't insist that all cases are permissive, or that they are all strongly permissive.

The issue of permissivism is importantly related to the epistemological problem of *peer disagreement*. But the two issues need to be distinguished. Consider now Carol and Dave who likewise have the same evidence (and they know this), but are not ideally rational (or at least don't take themselves to be). While recognizing their fallibility they share a mutual respect for one another. Carol, for instance, judges Dave to be no more or less likely than herself in general to arrive at rational, and indeed correct, conclusions in response to evidence. As it happens, Carol and Dave reach opposite conclusions: Carol believes that P while Dave thinks that not-P. When Carol learns that Dave thinks differently, how should this affect her opinion about P? Once both are aware of their difference of opinion can they rationally maintain their respective views, or must they somehow converge? This is the central issue in the epistemology of disagreement.³

While the issues are distinct, how we answer the question of permissivism would appear to have consequences for the problem of disagreement. The exact relation is open to debate, but on the face of it permissivism allows you to "stick to your guns" in the face of disagreement. If either opinion is rationally permissible, why should it bother Carol that Dave thinks differently? If Uniqueness holds then the matter becomes more pressing. Given their disagreement and shared evidence it follows from Uniqueness that at least one of Carol and Dave is failing to respond rationally to the evidence. Unless Carol is entitled to pin the blame entirely on Dave, she has reason to suspect that her opinion is not fully rational. And that would appear to give her reason to doubt her opinion.

2 The Case against Permissivism

I will mostly focus on strong permissivism. Many of the points I make can be adapted to make a case against moderate permissivism also. In the interests of space I will not pursue these developments in detail here.

2.1 Evidential support

Here are three premises that make up an argument against strong permissivism.⁴ First, we commonly appeal to relations of *evidential support* among propositions. We can ask, for instance, whether certain meteorological data support the hypothesis that human activity is a major cause of global warming. It is by virtue of these evidential support relations that an agent is rational in her beliefs. In short, it is rational for an agent to believe P just in case her total evidence E *supports* P.

Second, as I'll argue in what follows, these evidential support relations hold *necessarily*. If E supports P then necessarily E supports P. Someone might question the necessity claim as follows. That the gas gauge reads Full supports the conclusion that the tank is full. But it need not. Suppose we know that the gauge is stuck on Full, or even that the wiring is switched so that it tends to read Full only when the tank is empty. In these cases the gauge's reading Full seems to support no conclusion or the

opposite conclusion. However, I think the right way to think about the matter is that these are just cases in which our *total evidence* is different. Our evidence concerning the gas tank involves a lot more than just how the gauge reads. We have a lot of information about the purpose of gauges and their reliability. It is not clear that the reading of the gauge all by itself does much if anything to support the tank's being full. By varying the additional evidence, of course, we can vary which conclusion is supported. But when we consider the totality of what we have to go on in assessing the state of the gas tank – that it reads Full, that gas gauges are typically reliable, and so on – it is hard to make sense of the idea that all of *that* information might have supported a different conclusion.

The point can be supported by two further considerations. If evidential support were contingent it would be unclear how we can assess what our evidence supports. I am trying to assess whether it will rain today. I survey the evidence and consider which conclusion if any it supports. Suppose it could be that while the weather reports, color of the sky, and so on *actually* support the conclusion that it will rain, that very evidence *might* have supported a different conclusion, say that it will be sunny all day. How am I to tell whether I'm in a possible world in which this evidence does support rain rather than one in which it doesn't? Will some further evidence indicate this? But then is it a further contingent matter what this extra evidence supports? Surely the fact of the matter is as follows: In principle I can examine the total evidence and recognize that it supports the conclusion of rain. I can recognize this without further investigation because the evidence *necessarily* supports this conclusion. There may be worlds in which the very same evidence is *misleading* – worlds in which we have all the same evidence and yet it doesn't rain (the actual world might even be one of these). But there are no worlds in which it supports a different conclusion than it actually supports.

Suppose again that this were not so. Our total evidence E does in fact support the standard thesis C of climate change. But now suppose that that very evidence need not have supported this conclusion. If this were so then it would seem that industry lobbyists could in principle manipulate the world such that E supports not-C instead. Hence they could make it the case that we can rationally *believe* that carbon emissions do no harm (and rationally act on this!) without having done anything to prevent disastrous climate change. This seems absurd.

Third, evidential support is *unidirectional*. It cannot be that E supports P but also that it supports not-P. Whatever is evidence for P is evidence against not-P. If it could be that the evidence supports both P and not-P then apparently one could rationally hold both contradictory opinions at once. But that can't be right. Putting these points together we have a simple argument against strong permissivism.

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.

The argument is valid. There are, as usual, sophisticated ways in which one might deny one or more of these premises. I will not try to explore all of these here.⁵ My point is

just that avoiding this conclusion appears to require a departure from very natural ways of thinking about evidence and rationality.

2.2 Arbitrariness

As a jury member it is my responsibility to do my best to arrive at the truth about the defendant's guilt or innocence. The responsible way to go about this is to examine the available evidence and seek to proportion my belief to it. There are lazier ways of going about it. I have two magical belief-inducing pills, one is marked GUILTY, the other NOT GUILTY. Ingesting a pill will give me the labeled belief. Clearly it is irresponsible to take such a pill. Why? Arbitrarily taking a pill while I have no clue as to whether the defendant is guilty gives me only a 50 percent chance of arriving at the correct verdict. If instead I base my belief on the evidence I have a much better chance of getting at the truth. There is no guarantee of course – evidence can be misleading – but I have a much better chance. Indeed, if there is evidence available strongly supporting one verdict, then it is highly probable that it supports the correct verdict.

Now suppose that permissivism is true and that the present case is a permissive one and that I know this about the case. (Perhaps the Epistemology Oracle has revealed to me that this is one of those cases.) What reason do I have to form my belief by an examination of the evidence rather than just popping a pill? If either conclusion can be rationally held given the evidence, why not just randomly pick one? In a non-permissive case where the evidence directs us to a particular conclusion, following the evidence is a reliable means of pursuing the truth. It is hard to see how this could be so in a permissive case (if there could be such). In such a case forming a belief that is rationally permitted by the evidence leaves it underdetermined what my conclusion will be. Whatever ends up causing me to accept one conclusion over the other will involve some non-evidentiary factor. And how could some factor other than *evidence* increase my chances of arriving at a true conclusion? Think of it this way. Suppose a hundred fully rational agents are given evidence E. Now of course in a non-permissive case we can predict that they will all arrive at the same conclusion and they will likely be right. But in a permissive case surely we can't expect this. If either conclusion can be rationally held it would be natural to expect around a 50–50 split of opinions. In this case only about half of the inquirers will be correct in their conclusions. If I am one of these inquirers, how can I sensibly expect to be one of the lucky ones? So it appears that in a permissive case I have no good reason to form my beliefs by examining the evidence rather than just popping a pill. But this is absurd. Surely it is always wiser to let the evidence guide us in inquiry.

Here is a different way of thinking about it. Suppose I have no idea whether the defendant is guilty, but I take one of the pills, the GUILTY one say, and find myself with the conviction that he is guilty. Could I reasonably maintain my belief while recognizing that I formed it just by popping a pill? Surely not. Prior to ingesting the pill I must have only 50 percent confidence that the result will be a true belief. Once I believe that the defendant is guilty and recognize that this is what I believe, I can't coherently doubt that my belief on the matter is true. But wouldn't it be strange to conclude, "What a stroke of luck! I now believe that the defendant is guilty, and he *is* guilty. So I must have gotten lucky in selecting the right pill." If this doesn't seem bizarre enough, imagine taking *many* such pills by random selection from a bag. Could

you sensibly think that you are lucky enough each time to pick the right pill? It seems that it can't be rational to maintain my belief in these circumstances. But now according to permissivism, it may be that whichever pill I take my resulting belief is a rationally permissible response to the evidence. So it seems that a permissivist is committed to the surprising conclusion that I can rationally maintain my conviction that I have remarkable luck at randomly forming true beliefs.

It might be noted that as a result of taking the pill, my belief, while rationally permissible, will not be well-founded. While the resulting opinion will conform to the evidence, it will not be appropriately *based* on this evidence. Such a belief may be *propositionally* but not *doxastically* justified, and to this extent it will be defective. Might this explain why a permissivist should prefer reasoned inquiry over pill popping? I can't see that it does. In inquiry my first concern is to arrive at a *true* conclusion regarding the defendant's guilt. And it is not clear why I should be so concerned with having my beliefs appropriately based unless this is conducive to the goal of getting things right. Consider again a *non*-permissive case. Suppose that the Epistemology Oracle reveals to me that the belief I have just formed was not in fact based on my evidence, but rather formed in some random way (whether or not it is supported by the evidence). In this case it would be a matter of luck if my belief were supported by the evidence. And unsupported beliefs are less likely to be true. So if I realize that my belief was not based on the evidence I should doubt that my belief has evidential support and hence doubt that it is true. But here the worry stems not from the lack of appropriate basing as such, but rather from the doubt raised as to whether my belief is supported by the evidence. The Oracle may continue, "Luckily for you in this case your conclusion that P is the uniquely rational conclusion to draw from your evidence E, even though your possession of this evidence isn't what caused you to first believe P." Having heard this I can sensibly maintain my belief in P since if it is the only rational conclusion to draw from my evidence then it is likely to be true. The matter is different if permissivism is correct. If this is a permissive case then there isn't the risk that a randomly formed belief will be rationally impermissible given the evidence. Whichever doxastic attitude I end up with – believing P, disbelieving P, or suspending judgment about P – will be fine as far as the evidence goes. Nevertheless, the knowledge that my resulting belief is a rationally permissible response to the evidence can do nothing to suggest that my belief is true, since in this case rationality does not distinguish between the true and the false: both are deemed to be rational conclusions from the evidence.

In any event, we can sidestep worries about well-foundedness by building it into the belief pills. There are more sophisticated pills that not only induce a belief that P, but cause me to base my belief on the evidence in the appropriate way. In an alleged permissive case there are possible worlds in which I examine the evidence and thereby come to believe P basing it on this evidence. There are other worlds where I examine this same evidence and come to the conclusion that not-P, basing it on this same evidence. A P-belief pill just ensures that one of the former set of worlds is actualized. If I take one of these pills then I will have a well-founded rational belief and there should be no reason to abandon it.

But consider just how odd this situation is. I have no idea whether P is true and have no idea whether the belief I'm about to form by taking a pill will be true. I now choose a P-pill and hence know that soon I will believe P, but presumably I still have no idea whether my belief will be true (I know of course that I will soon *think* that my belief

is true). As the pill takes effect I can think to myself, "What do you know? P, and I believe P, so by a stroke of luck I picked the pill that gave me a true belief. I'm glad I didn't take the other pill because if I had I would be mistaken, even though I would still be rational." But now I can take a not-P-pill. This will revert my mind back into a state of considering E and lead me to rationally base my conclusion that not-P on E. At this point I will rationally think, "So, it turns out I was wrong in thinking that P. But that's okay. I'm glad I took this not-P-pill and corrected myself." And of course I can switch back again. Each time I toggle my beliefs in this manner I am relieved to find that my resulting opinion is true. The absurdity of this should make us wonder whether permissive cases are possible.

2.3 Reflection argument

The following seems to be a plausible general principle:

Reflection: If I know that tomorrow I will come to rationally believe P on the basis of new evidence without having lost any of my old evidence, then it is rational for me to believe P now.⁶

Consider now a different case. Instead of a P-pill and a not-P-pill we have a Truth pill and a Falsity pill. The Truth pill gives me a true belief regarding P and the Falsity pill gives me a false belief. (In either case the pill I take will cause me to base my belief on the available evidence.) You will flip a fair coin and feed me a pill depending on how the coin lands: Truth if the coin lands *heads*; False if it lands *tails*. (You will not let me know how the coin landed or which pill you have dropped in my mouth.) Suppose I come to believe P. Assuming that this is a permissive case my belief will be rational, and since I will know that P is the conclusion I have reached, I will also rationally believe that I have a true belief concerning P. But the same goes if I take the other pill and believe $\neg P$: from my rational belief in $\neg P$ and my knowledge that I believe it I will rationally conclude that I have a true belief concerning P. So I can know in advance that regardless of which pill I take and whether P is true, upon examining the evidence and taking a pill I will rationally believe that my belief regarding P is true. It follows by reflection that I can rationally believe *now* that my future belief will be true. But I know it will be true only if the coin lands *heads*. It follows from strong permissivism and reflection that I can rationally predict in advance that this fair coin will land heads! But that is absurd.

The strong permissivist of course will just have to deny reflection.⁷ Perhaps that can be done, but let's note how plausible this principle seems. Suppose I know that tomorrow I will believe that P. There are circumstances in which this would be regrettable. Perhaps I know that I soon will have forgotten some important information that would otherwise prevent me from concluding that P. Or perhaps I suspect that I will go mad and believe (what I can now recognize to be) nutty things like P. In such a case I have reason to distrust my future opinion and may want to try to prevent myself from forming such a belief tomorrow. But suppose that none of this is the case and that I'm quite sure that everything will go swimmingly. I will just consider all the evidence and form a rational opinion on the matter. It is hard to see what reason I could have to prevent my future self from forming such an opinion. What better way is there to go

about inquiry than to believe rationally on the evidence? And if I willingly anticipate forming this belief, what is to stop me forming it now?

The situation seems even stranger when we consider the connection with decision making. Before the coin is tossed I will surely have credence $1/2$ that the coin will land heads and hence that I will form a true belief concerning P. So I will naturally be unwilling to take a bet where I win \$1 if my resulting belief is true, but lose \$2 if it is false. But I know that once I take the pill I will rationally form some opinion and hence rationally conclude that my belief on the matter is true. At this point, according to strong permissivism, I will rationally take the bet that my belief concerning P is true. Surely I have no reason to be reluctant to take the pill since it will simply result in me forming a perfectly rational belief on the basis of more evidence than I have to go on now. So it seems that I must *now* approve of the bet that I will take once I've ingested the pill. But I can't now approve of this bet since I'm only 50 percent confident that the coin will land heads. It appears my attitudes are incoherent.

2.4 Arbitrary switching

Fleeing from a tiger escaped from the zoo, you come to a fork in the road. One path leads safely home, the other off a cliff. Given your evidence you rationally believe that path A will lead you home. But you take your case to be a strongly permissive one where a rational agent with your evidence might conclude that B is the safe path to take instead. You don't think that B is the safe path, but you think that you could rationally think so given your evidence. Why not then switch your belief? A quick pill could change your state to one of a rational and well-founded conviction that B is the path that leads home. It is hard to see what could be wrong with taking such a pill. As a permissivist you don't think the worse of someone who has the contrary opinion. So you can't think there would be anything wrong with you epistemically speaking if you had believed that B is the safe path to begin with. Why then be wedded to your actual belief? Of course once you take the pill and rationally believe that B is the safe path you will rationally take that path. But if you can see nothing wrong with taking an action with that result, you might as well just go ahead and take path B. The scenery on that route is rather more enjoyable after all. But that is just bonkers. Insofar as you have considered the evidence and concluded that A is the only path that leads home, you can't sensibly think that it would be okay to take the other path, or even to take a pill which will result in your going the other way. The challenge is to say why this is so in a way that is consistent with strong permissivism.

Here is what I take to be the most promising line of response to this challenge. It can be developed in more or less sophisticated ways. I will stick to a simple formulation to get across the key idea.⁸ I have asked the strong permissivist, "Assuming that the contrary opinion about the paths can be rationally held with the same evidence, why wouldn't it be rational to take a pill that switches your opinion?" The permissivist has an answer that might seem like a no-brainer: "Because if I do that I'll fall off a cliff! I don't judge someone who thinks otherwise as necessarily any less *rational*. But they are sadly *mistaken*. And of course I don't want to be mistaken on this matter." There are subtle issues here. The permissivist has described her deliberation in a way that seems reasonable. But to answer the question "Why would it be irrational for you to take the pill?" we need to identify what it is about her situation that makes this

choice irrational. This cannot be the *fact* that path B runs off a cliff. It might not even be a fact. Even if B is in fact the safe path but the agent rationally thinks otherwise it would be crazy for her to change her belief on the basis of no further evidence. Instead she might say, "Well I *believe* that taking the pill would result in my falling off the cliff. And I don't want that!" But now it seems we don't have an adequate answer to my question. Sure, you happen to believe that changing beliefs would result in your death since you believe that B leads off a cliff. But the question is, why be wedded to that belief? Upon taking the pill you will no longer think that taking the pill has the result that you run off the cliff. Indeed you will be relieved to conclude that the pill has saved you from that fate. And according to strong permissivism you will be rational in concluding this. In general, it doesn't sound like an adequate answer to the question, "Why not act so as to make it the case that P, rather than Q?" to say "Because it's not the case that P: Q is the case." The answer to the question "Why shouldn't you change your belief from A to B?" can't be "Because I don't believe B; I believe A."

Still, perhaps there does seem to be something right about the agent's original response. In practical deliberation we are focused outward on the world. It does seem appropriate to think, "Doing A will result in my death. So I won't do it." The fact (if it is a fact) that doing A will result in my death can't be the fact in virtue of which it is irrational for me to do A. Nevertheless, it is appropriate in practical deliberation to ask myself whether doing A will lead to my death and to act on the conclusion that I draw. (The relevant question in deliberation is not whether I *believe* that doing A will lead to my death, but whether it *will* lead to my death. For it is my death that I'm worried about, not my psychological states.) So the permissivist may insist that by displaying her apparently sensible line of reasoning she reveals why she is rational in refraining from doing A.

The response still seems unsatisfying. Consider a case involving only moderate permissiveness. You don't know whether either path leads home. Perhaps they both do, perhaps neither does. But on examination of the evidence (carefully attended to as you flee from the tiger!) you are somewhat more confident that path A leads to safety than you are that B does. Naturally you choose A. But as a moderate permissivist you take it to be rationally permissible for you to be somewhat less confident that A is safe and maybe a little more that B is safe such that the balance tips over in favor of B. The question arises again: why not take a pill to effect these small changes in your degrees of confidence in these hypotheses? But here the answer cannot be, "If I did so I would take path B and die!" You don't even believe this. Perhaps you even believe that B *will* lead you home; you are just not as sure of this as you are that A will. Instead you may say, "I'm more *likely* to get home safely if I take path A." But how are we to understand "more likely"? You can't be thinking that the *objective chance* of making it home is greater. The relevant chances of getting home may be around 0 or 1 depending on where the path leads (if you're on the right path you can be sure to get to your home before the tiger catches up). So the only way you could believe that the chance of getting home via A is greater is if you thought that A leads home but B does not. But you don't think that in this case.

A natural way to understand it is in terms of evidential or epistemic probability, which amounts to something like *the evidence more strongly supports the conclusion that A leads home than that B does*. If you believe this to be the case then naturally you will be motivated to take path A. But this does not fit with permissivism. It can't

be fully rational to put more confidence in a hypothesis that enjoys less evidential support.

For a permissivist it looks as though the likelihood here can only be understood as subjective probability, or credence. This is a little odd, however, as it doesn't seem that when I have the thought that *A is more likely to lead me home* I am just reflecting on my own psychology. And anyway, the thought *my credence that A leads home is greater than my credence that B does* can't play the same role in practical deliberation as the thought *only A will lead me home* does. In decision making I am concerned with the consequences of my actions, not with my psychological states. There does not seem to be any fact that you can appeal to in your deliberations in this case to make sense of your reluctance to switch that is consistent with permissivism.⁹ So we have an odd situation. If strong permissivism is plausible then a fortiori moderate permissivism is also. Yet our best response to the arbitrary switching objection works, if it works at all, only in defense of strong permissivism.

A second point to note here is that it is clearly not always appropriate to appeal to the fact that P as a reason to avoid changing my belief about P. Even if I rationally believe that it is path A that leads home, it is typically a good idea to be open to, or even seek out, further evidence on the matter. In doing so I open myself to the possibility of changing my opinion. For it may turn out that new evidence clearly suggests that it is B that takes me home, and being rational I will change my view accordingly. It would be absurd for me to reason, "Path B runs off a cliff. So anyone who believes otherwise, even rationally, on the basis of further evidence, is unfortunately mistaken and will come to a bad end. I must take whatever steps I can to avoid this myself by shielding myself from any new evidence on the matter."¹⁰

Or consider a variation on this case. I might simply come to think that I was confused in judging that A is the safe path. Without new evidence I might reassess the case and conclude that actually the evidence suggests that it is B that is safe. Perhaps I am so fixated on the plan to take path A home that I initially continue to think that it leads to safety even after I've come to think that the evidence suggests otherwise. It would be perverse to think, "Well, now I see clearly that the only *rational* opinion to hold is that path B will lead me home, given my evidence. However, the fact remains that it is only path A that leads home, so I should avoid changing my opinion even if the alternative opinion is rational." In each case it seems entirely appropriate to let my judgments about *what it would be rational to believe* override my convictions about *what is actually the case*. But if that is so, it is hard to see why, in cases I take to be permissive, I can't give myself permission to change my opinion to one that I judge to be rational even if false by my current lights.

Note further that the anti-permissivist has a good explanation as to why he is always open to obtaining new evidence and conforming his belief and action to it. While of course I take my current opinion that P to be true, I also judge that *conditional on my obtaining further evidence such that my total evidence supports the opposite conclusion*, my current opinion is likely to be false. So I am not concerned that obtaining new evidence will lead me away from the truth (although that is always possible, just as it is possible that I am now mistaken). Compare this with the permissivist's predicament. Suppose I believe P but have the opportunity to obtain substantial new evidence on the matter. I don't yet know what this new evidence will involve, but the

Epistemology Oracle informs me that it will remain rationally permissible to take either of the two contrary opinions on P in the light of the new evidence. Even though I know that I'm perfectly rational, there is no predicting which belief I will end up with if I examine the new evidence since the Oracle informs me that either belief can be rationally held on the new evidence. Upon examining this new evidence it might strike me that I was mistaken and I may (rationally) come to believe $\neg P$. If we are to take the response sketched above to the arbitrary switching problem ("I won't switch my belief from A to B because then my belief will be false"), then I must take it that I have reason to refuse to view the new evidence. For there is a good chance that viewing the new evidence will result in my believing what I now take to be false. I could perhaps try to steel myself against the possibility of changing my opinion. But even if I can do this, it only follows that there is no point in considering the new evidence as it will have no effect on my opinion. But refusing to consider new evidence is absurd. Surely I should always take the opportunity to form a rational opinion on the basis of more evidence than I currently have.

I will leave you with you with one last odd consequence of strong permissivism. You have a hundred beliefs B1, B2, ..., B100, which you take to be permissive cases. I offer you a belief-toggle pill for B1. (The pill replaces B1 with a belief in the negation of its content.) "No thanks," you reply. "While the resulting belief would be rational it would also be false. And I don't want that." I slip it into your coffee anyway. Perhaps you don't thank me for violating your wishes, but you are more than happy with the result. As you see it now, your new opinion is not only rational but true. You couldn't have been more wrong in thinking that my pill would lead you into error. Quite the opposite! Or so you must now rationally think. I offer you a belief-toggle pill for B2. "No thanks," you reply. "While the resulting belief would be rational it would also be false. And I don't want that." I slip it into your coffee anyway ...

At some point shouldn't you start to trust me more? It's not that you have any reason to suppose that I'm an authority on the subjects of B1–B100 that you should defer to. I haven't any clue about these matters. I just enjoy experiments with belief pills. Suppose I feed you belief-toggle pills for the first 60 of these beliefs. You must find the results rather surprising. Of course you are not surprised that you *think* that your new beliefs are true. You knew that would happen. But you are surely surprised to find that they *are* true, which is what you now must take them to be. You are surprised to find your new beliefs true for several reasons. First of all, you had every reason to suppose that your old beliefs B1–B60 were true. You were after all quite rational in holding these beliefs given your evidence. And you haven't lost any of your evidence, so you must still have every reason you ever had to suppose that B1–B60 were true. But it turns out that none of them were! (Or so you must rationally think, on permissivist assumptions.) Still, it has all turned out okay by a remarkably lucky coincidence. Fortunately for you, those 60 of your of 100 beliefs that were mistaken have been corrected by a fist-full of randomly chosen belief-toggle pills! What are the odds of that? The rest of your beliefs you take to be true. So had you taken more pills or less pills, or a different selection (e.g., by counting backwards from B100, or taking the even-numbered pills), then your resulting beliefs would be partially or mostly false. But you must take it that you were lucky enough that the bunch chosen happen to line up perfectly with your prior mistakes. Perhaps you have an Epistemic Guardian Angel looking out for you!

Common wisdom has it that examining the evidence and forming rational beliefs on the basis of this evidence is a good means, indeed the best means, to forming true beliefs and avoiding error. The result so far must shake your confidence in this wisdom. You formed a hundred rational beliefs upon careful examination of the evidence yet 60 percent of them were mistaken! In stark contrast, the method of arbitrarily toggling your beliefs seems to have worked wonderfully. It has a 100 percent track record of leading you to the truth so far. Could this be an accident? Perhaps then you should decide to trust that the pills will have an enlightening effect and take the remaining 40. But there is no way to coherently think so. If you think that the belief-toggle pill will lead you to the truth regarding P, then you must take your current belief that P to be false. You can change that belief to not-P, but then you must think that the pill will change your belief back to P, which you must think is false given that you believe not-P. There is no coherent option except to think that all future belief-toggle pills will have the unfortunate effect of changing your opinion for the worse, even though you think they've done just the opposite so far.

I offer you a belief-toggle pill for B61. "No thanks," you reply. "While the resulting belief would be rational it would also be false. And I don't want that." I slip it into your coffee anyway ...

Acknowledgments

Too many people have helped me think about this topic. But the main ones to thank are Tom Kelly and Miriam Schoenfield. I presented some of this material at the Central APA in Chicago and at MIT. Matthias Steup provided helpful corrections and feedback on a draft.

Notes

- 1 I take the label "Uniqueness" from Feldman (2007) who defends a position along these lines. Other defenses of Uniqueness can be found in Christensen (2007), Sosa (forthcoming), and White (2005).
- 2 Permissivism is clearly the majority view among epistemologists. And many views that are widely held seem to entail it. See Ballantyne and Coffman (2011), Douven (2009), and White (2005) for details. Some recent defenders of permissivism include Ballantyne and Coffman (2011), Bruckner and Bundy (2012), Douven (2009), Meacham (manuscript), Schoenfield (forthcoming), and Titelbaum (2010). Some of these contain helpful critiques of the arguments I make here (derived from my earlier work), which I haven't had the space to address.
- 3 Christensen (2009) provides a useful overview of the disagreement debate.
- 4 Here I am clarifying an argument I've presented before (White, 2005) which was recently challenged by Ballantyne and Coffman (2011).
- 5 The interested reader might look at Schoenfield (forthcoming), and Meacham (manuscript).
- 6 This is a modification of the principle defended by van Fraassen (1984). One source of potential counterexamples involves beliefs expressed with the essential use of indexicals like "now." Rather than get distracted by this issue we can assume that the proposition P that we discuss here involves nothing of the sort.
- 7 See Schoenfield (forthcoming) for a discussion of how to do this.

- 8 See Schoenfield (forthcoming) for a more subtle discussion.
- 9 For an account of probabilistic assertions that might appeal to the permissivist here see Moss (2013).
- 10 This is one version of the puzzle introduced by Kripke (2011).

References

- Ballantyne, N. and Coffman, E.J. (2011) Uniqueness, evidence, and rationality. *Philosophers' Imprint* 11 (18).
- Bruckner, A. and Bundy, A. (2012) On "epistemic permissiveness." *Synthese* 188: 165–177.
- Christensen, D. (2007) Epistemology of disagreement: the good news. *Philosophical Review* 116: 187–217.
- Christensen, D. (2009) Disagreement as evidence: the epistemology of controversy. *Philosophy Compass* 4: 756–767.
- Douven, I. (2009) Uniqueness revisited. *American Philosophical Quarterly* 46: 347–361.
- Feldman, R. (2007) Reasonable religious disagreement. In *Philosophers without God*, ed. L. Antony, pp. 197–214. Oxford: Oxford University Press.
- Kripke, S. (2011) Two paradoxes of knowledge. In *Philosophical Troubles: Collected Papers*, Vol. 1, pp. 27–51. Oxford: Oxford University Press.
- Meacham, C. (manuscript) Impermissive Bayesianism.
- Moss, S. (2013) Epistemology formalized. *Philosophical Review* 122 (1): 1–43.
- Schoenfield, M. (forthcoming) Permission to believe: why permissivism is true and what it tells us about irrelevant influences on belief. *Noûs*.
- Sosa, E. (forthcoming) The epistemology of disagreement. In *The Epistemology of Disagreement: New Essays*, ed. J. Lackey and D. Christensen. Oxford: Oxford University Press.
- Titelbaum, M. (2010) Not enough there: evidence, reasons, and language independence. *Philosophical Perspectives* 24: 477–528.
- van Fraassen, B. (1984) Belief and the will. *Journal of Philosophy* 81: 235–256.
- White, R. (2005) Epistemic permissiveness. *Philosophical Perspectives* 19: 445–459.