

EPISTEMIC PERMISSIVENESS

Roger White
New York University

A rational person doesn't believe just anything. There are limits on what it is rational to believe. How wide are these limits? That's the main question that interests me here. But a secondary question immediately arises: What factors impose these limits? A first stab is to say that one's *evidence* determines what it is epistemically permissible for one to believe. Many will claim that there are further, non-evidentiary factors relevant to the epistemic rationality of belief. I will be ignoring the details of alternative answers in order to focus on the question of what kind of rational constraints one's evidence puts on belief. Our main question concerns how far epistemic *permission* and *obligation* can come apart.¹ Suppose I am epistemically permitted to believe P, i.e., it would not be irrational for me to believe it. Am I thereby obliged to believe P, or are other options rationally available to me?² Might I be equally rational in remaining agnostic about P, or even believing not-P? Or could even a slightly stronger or weaker degree of confidence be just as reasonable?

1. Examples and Motivations for Permissive Epistemology

Following Feldman (forthcoming) we can call the negative answer the Uniqueness thesis.

Uniqueness: Given one's total evidence, there is a unique rational doxastic attitude that one can take to any proposition.³

I call positions that depart from Uniqueness "permissive" as they entail that epistemic rationality permits a range of alternative doxastic attitudes. Many philosophers hold permissive epistemologies. Rejection of Uniqueness is prominent in Bas van Fraassen's (1984, 1989, 2002) epistemology. Proponents of Conservatism (Harman 1986, Lycan 1988) suggest that you are *prima facie* justified in maintaining your beliefs until you have a reason to abandon them.

So had you drawn a different conclusion from the same evidence, you would be fully rational in continuing to hold it, at least until challenged. According to Richard Foley (1987), a belief is epistemically rational if upon sufficient reflection you would think that holding that belief was an effective means to achieving the goal of believing just what is true. This condition fails to single out a unique set of beliefs for any given evidence. It seems that the method of seeking Reflective Equilibrium (Rawls 1999, Goodman 1955) need not terminate in a unique set of convictions. So if that is the ultimate standard of rationality, Uniqueness is false. According to the Subjectivist version of Bayesianism, the only rational constraint on one's initial degrees of belief is conformity to the probability calculus. While some Bayesians are uncomfortable with this degree of permissiveness, the constraints that they are willing to add fall far short of Uniqueness.⁴ Any coherence theory of rational belief would appear to be at odds with Uniqueness. For however we understand coherence there will be multiple ways of achieving it that involve different attitudes to various propositions.

Doubts about Uniqueness may arise from reflection on persistent and widespread disagreement among apparently rational inquirers.⁵ According to Gideon Rosen (2001)

It should be obvious that reasonable people can disagree, even when confronted with the same body of evidence. When a jury or a court is divided in a difficult case, the mere fact of disagreement does not mean that someone is being unreasonable. (p. 71)

In a criminal trial we have a large, complex body of evidence, some of which appears to link the defendant with the crime, some which suggests that he had nothing to do with it. Figuring out what to believe is a matter of weighing various considerations as we try to fit all the pieces together. Difficult cases like this tend to produce sharp disagreement even among the most diligent inquirers. Some may be tempted to give up and conclude that there is no telling who's right. But many of us retain our conviction that having earnestly wrestled with the evidence our conclusions are correct. Now a proponent of Uniqueness will have to insist that those who disagree have failed to respond to the evidence rationally. For if the parties in the dispute have the same evidence, and one's total evidence uniquely determines what one can rationally believe, then they should all be in agreement.⁶ But this can be an awkward position to maintain. Why should I trust my own estimate of the probative force of the evidence, when equally intelligent people judge it differently? One initially appealing way to escape this worry is to join Rosen in insisting that those who disagree needn't display any failure of rationality.

Examples and motivations for permissivism can be multiplied. There are many important differences among permissive epistemologies. For instance, many permissive views will identify some other non-evidential factor which

“takes up the slack” in fully determining what one should believe. And these extra factors will vary among different views. Differences such as these may crucially affect how a permissivist can best respond to the objections that I present. But in order to get out the basic worries that I have without too much fuss, I will ignore these distinctions and treat these permissive views as a generic position. I think there is much to be said for these and other challenges to Uniqueness. But my purpose here will be to raise trouble for permissive views by displaying some difficulties we get into by denying Uniqueness.

2. Objections to Extreme Permissivism

Let's begin by considering the following radical rejection of Uniqueness:

Extreme Permissivism: There are possible cases in which you rationally believe P, yet it is consistent with your being fully rational and possessing your current evidence that you believe not-P instead.

My main objection to Extreme Permissivism will take some work to develop. I will begin with a quick-and-dirty preliminary one. I am on a jury with the task of deciding whether the defendant Smith is guilty. Prior to considering any evidence it would be quite unreasonable to have any opinion as to Smith's guilt or innocence. Perhaps there are some propositions that one can rationally believe without evidence, but the guilt or innocence of a criminal defendant is not one of them. How could evaluation of the evidence render me rational in believing that Smith is guilty, if it is not reasonable to believe this already? Surely only if the total evidence *supports Smith's guilt*. And likewise only evidence supporting his innocence could *make it* rational to believe that he is innocent. Evidence that has no bearing on whether Smith is guilty is no better than no evidence at all.

But the evidence cannot support both Smith's innocence *and* his guilt. Whatever is evidence *for* P is evidence for the falsity of not-P and hence is evidence *against* not-P. Of course, certain elements of or aspects of the total body of evidence might suggest that Smith is guilty, while others suggest the opposite. But it is incoherent to suppose that a whole body of evidence could count both for and against a hypothesis. So then it is impossible that my examination of the evidence makes it rational for me to believe that Smith is guilty but also rational to believe instead that he is innocent. And since neither view was rational apart from any evidence, the proposed radical departure from Uniqueness cannot be right.

While I think it is right, the argument above might seem a little too quick. The following considerations go into more depth. First consider the following situation. Once again, prior to being presented with the evidence I haven't a clue as to whether Smith is guilty. But I have a couple of magical belief-inducing

pills. One, when swallowed, will give me a true belief on the matter, the other a false belief. I have no idea which one is the Truth pill, so it is clearly irresponsible to take either, given that my purpose is to discern the truth. But I pick one at random and take it, finding myself with the conviction that Smith is guilty, without being aware of any reason to think so. Now perhaps we can rationally retain beliefs after we have lost track of the reasons that led us first to hold them. But in this case I can recall that my belief was not based on evidence pointing to Smith's guilt, but is rather the result of an arbitrarily chosen pill. If I am rational, recognition of this fact should surely undermine my conviction in Smith's guilt. Prior to taking the pill my credence that I would soon hold a true belief should have been about $1/2$. I cannot coherently maintain a conviction that Smith is guilty while being entirely uncertain as to whether my belief that Smith is guilty is true. And it would be perverse to reason backward as follows: Even though I had a only fifty-fifty chance of taking the Truth Pill, since I have formed the belief that Smith is guilty, and indeed he is guilty, I now see that I must have got lucky in choosing the right pill! If it is not already obvious enough that this would be absurd, imagine the following. I take lots of these pills, arbitrarily forming hundreds of beliefs on matters about which I had no clue beforehand. If I retain the resulting beliefs then I will have to conclude that by some extraordinary coincidence I managed to pick a Truth pill each time. And if this were reasonable, then it should make no difference if the pills were selected from a collection 99% of which were Falsity Pills. Surely instead the only reasonable response to reflection on my pill popping is to slip back into agnosticism about Smith's guilt.

Now let's suppose instead that I have not taken a pill, but that under the influence of a permissive epistemologist I come to believe that Extreme Permissivism applies to the matter of Smith's guilt given the available evidence. That is, I believe that upon considering the evidence in court one could rationally conclude that Smith is guilty, but there is an alternative path that one's reasoning could take arriving instead at the rational conclusion that he is innocent. Supposing this is so, is there any advantage, from the point of view of pursuing the truth, in carefully weighing the evidence to draw a conclusion, rather than just taking a belief-inducing pill? Surely I have no better chance of forming a true belief either way. If my permissive assumption is correct, carefully weighing the evidence in an impeccably rational manner will not determine what I end up believing; for by hypothesis, the evidence does not determine a unique rational conclusion. So whatever I do end up believing upon rational deliberation will depend, if not on blind chance, on some arbitrary factor having no bearing on the matter in question.

It might be suggested that rationally evaluating the evidence is a fairly reliable means of coming to the correct conclusion as to whether P, even if that evidence does not determine that a particular conclusion is rational. But it is very hard to see how it could. Of course often one's total evidence does favor a conclusion of P over not-P, and its doing so is a fairly reliable guide to the truth.

Evidence can be misleading—i.e. point us to the wrong conclusion—but this is not common. So there is no problem in this sort of non-permissive case in seeing how rational assessment of evidence is a reliable means to the truth. The probative force of the evidence is a reliable guide to the truth, and a rational inquirer is sensitive to this force, forming his beliefs accordingly. But in any case in which the evidence does favor one conclusion over another, it is *not* equally rational to draw either conclusion. A rational person when confronted with evidence favoring P does not believe not-P. So a case—if there could be such—in which a rational person can believe P or believe not-P instead, must be one in which the evidence favors neither conclusion.

Even if it is granted that a rational person needn't suspend judgment in such a situation, just how rational evaluation of the evidence could reliably lead us to the truth in such a case is entirely mysterious. It would have to be by virtue of some property of the evidence whose reliable link to the truth is inaccessible to the inquirer. For if an inquirer is aware that the evidence has feature F, which is reliably linked to the truth of P, then surely it would be unreasonable to believe not-P. It is hard to imagine what such a truth-conducive feature could be, let alone how it could act on an inquirer's mind directing him to the truth. Furthermore, it is not enough for forming beliefs in response to evidence to have a rational advantage over pill-popping, that it just so happens that the former is a reliable route to the truth without the inquirer recognizing this. For all I can tell, the Truth pill might be a little stickier, so that randomly fetching one from the bag and taking it will usually yield a true belief. But for all I can tell it might be quite the opposite: I may have a tendency to pick the Falsity pill. Only if I could reasonably suppose that arbitrarily selecting a pill will most likely lead to the truth would this be a way to form beliefs that can rationally survive recognition of having been formed this way. We certainly have no reason to expect that rational assessment of the evidence reliably leads to the truth in a way that is inaccessible to an inquirer, even if we allow for this as a possibility. So forming beliefs in response to evidence that does not determine a rational conclusion seems no better than taking belief-inducing pills.

Perhaps the following will drive home the point further. I might examine all the evidence without drawing a conclusion as to whether Smith is guilty. But I do draw the epistemological conclusion that either verdict could be rationally held given this evidence. Perhaps if I were to judge that a particular verdict was rationally obligatory, I couldn't help but draw that conclusion. But as the evidence doesn't strike me as pointing clearly in one direction—if it did, I couldn't judge it to be equally rational to believe the opposite—I am psychologically capable of suspending judgment. But if I really do judge that believing P in this situation would be rational, as would believing not-P, then there should be nothing wrong with my bringing it about that I have some belief or other on the matter. But then it surely cannot matter how I go about choosing which belief to hold, whether by choosing a belief that I'd like to hold, or flipping a coin, or whatever.

Now I have argued that reflection on my having taken a pill to believe that Smith is guilty should undermine all my conviction on the matter. So likewise, if I have concluded on the basis of the evidence that Smith is guilty, my conviction should be undermined if I really think that a belief in Smith's innocence is also rationally permissible in the light of this evidence. For if I believe this, then I should judge myself no more likely to have arrived at the truth than a random pill-popper. So we have reached the conclusion that I cannot rationally accept the extreme permissivist thesis with respect to one of my own beliefs. That is, believing P is not rationally compatible with believing that one could just as rationally have believed not-P given the same evidence.

To be careful, we should note that there are some logical loopholes available to the permissivist here, even granting my argument thus far. For instance, I might maintain that it is rationally permissible for *you* either to believe P or to believe not-P, given your evidence, but deny that the same liberal standards apply to me. But this would be obviously *ad hoc* and unmotivated. It is hard to imagine any reason I could have to accept a permissive view of rationality that exempts myself. Another move would be to suggest that Extreme Permissivism is true but epistemically destructive if known: In some cases one can rationally believe P or believe not-P given one's evidence, but *only* as long as one mistakenly thinks that one's current attitude to P is rationally obligatory given one's evidence. This strikes me as an odd position, for it seems natural to suppose that a belief can always rationally survive learning the epistemic value of one's evidence. That is, if it is rational to believe P given evidence E, then it is rational to believe P given E & E', where E' correctly states what attitudes to P are rationally permissible given E. But in any event, this position makes it very difficult for us to rationally believe Extreme Permissivism, for the knowledge that it applies in a particular case is self-undermining: By the arguments above, if I believe that both a belief in P and a belief in not-P are each rationally permitted by my evidence, then I should conclude that neither belief is rationally permissible. Hence if I am to believe anything, then for each belief I should hold that it is the only rational attitude that I can take given my evidence. Furthermore, if Extreme Permissivism were plausible with respect to many matters, then surely it could plausibly be applied to itself. That is, Extreme Permissivists should hold that the usual evidence available to philosophers does not determine whether one should believe or disbelieve Extreme Permissivism. After all, plenty of apparently reasonable people deny Extreme Permissivism. But then by the arguments above, Extreme Permissivism cannot rationally be believed on our evidence.

This still seems to leave room for a coherent permissive position. Call a case in which my evidence rationally permits me to believe P or to believe not-P, a *permissive case*. We could maintain that there are possible (and even actual) permissive cases, but that these are very rare, and one cannot tell, for any case, that it is a permissive one. So my attitude with regard to each case considered individually should be that it is not permissive, while I may maintain that some

rare cases are permissive.⁷ (There is an obvious parallel here with the so-called Preface Paradox, where I believe regarding each of my beliefs individually, that it is true, while maintaining that at least some of my beliefs are false). But while this position may be coherent and escape the objections thus far, I doubt that anyone holds such a view, as it is hard to see what could motivate it. Typical permissive epistemologies such as Subjectivist Bayesianism entail that permissive cases are quite pervasive, and hence cannot easily avoid the problems I have raised.

3. Alternative Permissible Standards?

Here is an objection that one might make to the main argument that I've put forth.⁸ A crucial step in my case was the claim that if you believe P but maintain that you could have rationally believed not-P on the same evidence, then you should take the means by which you arrived at your belief in P to be no better than arbitrarily popping a belief-inducing pill. This might be a point of resistance for a permissivist. He might insist that he has come to his conviction in P by means that are very conducive to forming true beliefs, namely by assessing the evidence according to the appropriate epistemic standards. This is not at all like randomly popping a belief-inducing pill. He might claim that such standards are not permissive ones presenting subjects with alternative options for belief, but rather they state uniquely what is to be believed given certain evidence. Even before assessing the evidence, if someone accepts such standards of evidence assessment, then he should judge it very likely that he will arrive at the truth by their application. For the standards will sanction only those beliefs that according to the standards, are well supported by the evidence, and hence likely to be true. So far this doesn't seem like a permissive position. For given his acceptance of standards that sanction only a belief in P given his evidence, then he shouldn't just believe not-P instead. However, the permissivist that I'm imagining might maintain he could have adopted a different set of standards that sanction very different attitudes given the same evidence. This possibility, he might suggest, is consistent with his being fully rational and possessing his current evidence; and were he to accept standards sanctioning a belief in not-P, he would be rational in doing so. In this way the subtle permissivist tries to have his cake and eat it too. On the one hand he can proudly maintain that he is responsibly forming beliefs in a manner that he takes to be reliable in attaining the truth. On the other, he denies that he is bound by rationality to form just his current opinions on his evidence. Likewise, he needn't accuse those who share his evidence but differ in their conclusions of any irrationality, provided their views are in accord with their own alternative epistemic standards.

I think we ought to be suspicious of this position for similar reasons as before. First note that if our permissivist takes his own standards to be a reliable guide to the truth, then since the alternative standards deliver very different conclusions, he must judge them to be rather unreliable. So he should

judge himself very lucky to have adopted truth-conducive standards, since with full rationality he could have followed ones that would lead him wildly into error. We see here that the same worries about arbitrariness just arise at a different position. How have I come to hold the epistemic standards which lead me from my evidence to conclude that P? According to this permissivist it was not by virtue of being rational, since it is consistent with my being rational that I adhere to rather different standards that would have me believe not-P instead. But then it seems that my applying the correct standards and hence arriving at the right conclusion is just a matter of dumb luck, much like popping a pill. And hence I ought to doubt that I really have been lucky enough to do so.

To make the situation more vivid, let's fill in the details with a make-believe story about how we come to adhere to certain epistemic standards. Suppose it is just a matter of education. I follow standards S because I was inculcated with them at MIT. But had I attended Berkeley, I would have been inculcated with standards S' instead. Given my total evidence as input, S and S' deliver the conclusions P, and not-P respectively. (The story may be far-fetched, but nothing hinges on the details. And if my being rational is compatible with adopting either set of standards, then there is *some* further factor which led me to hold S rather than S'.) Now I can imagine myself in a counterfactual situation before graduate school where my sole motive for study is to answer the question whether P. I have all the available relevant evidence, I'm just not sure yet what to make of it. Now I learn that if I attend MIT I will inevitably inherit standards S from my mentors, which given the evidence will lead me to believe P. Attending Berkeley will result in my adhering to standards S' and hence arriving at the conclusion not-P. Now surely the prospect of several years of graduate school will seem rather pointless no matter how passionately curious I am as to whether P. Indeed my prospects for answering whether P to my satisfaction should seem dim. I might as well choose a grad school to attend and hence opinions to hold by a preference for Massachusetts weather, or by flipping a coin. Once I have filled out the enrollment form for MIT say, I will know that unless something gets in the way, in a few years I will be of the opinion that P. If I am a permissivist, I should take it that I will soon *rationaly* believe P. If this is so, why shouldn't I just believe it now and save myself the time and trouble? But of course it would be absurd to form an opinion on the matter by an arbitrary choice when I don't even know what to make of the evidence. If this is the sorry state I find myself in with respect to answering whether P before I begin my inquiry, then I should judge myself no better off having arrived at a conclusion, if I judge that my adoption of epistemic standards was such an arbitrary matter.⁹

4. Moderate Permissivism and Practical Deliberation

I have been considering an extreme rejection of Uniqueness that might appeal only to the most permissive epistemologists. We can back off from this

level of permissiveness without going all the way to Uniqueness. Many will find it easier to accept that rationality allows for a more restricted range of differences of opinion. Perhaps there are cases in which a reasonable assessment of the evidence rules out a belief that not-P, but does not dictate whether one should believe that P or suspend judgment. Or more cautiously still, there may at least be some leeway in the *degree* of conviction that is rationally permissible given one's evidence.

These moderately permissive positions still run into trouble. Suppose that having carefully considered the evidence, my conviction in Smith's guilt is strong enough to warrant voting Guilty. Surely this is what I should do—to let him go free given that I am (rationally) confident that he's guilty would be grossly irresponsible. If I had more substantial doubts it would be appropriate to vote Not Guilty, even if I strongly suspected that he was guilty. Now let's try to suppose that the range of rationally permissible degrees of confidence in Smith's guilt is just broad enough to include my own conviction as well as a degree of doubt that would make a vote of Not Guilty appropriate.¹⁰ It appears that if this were so I should have no qualms about letting Smith go free, even though I'm sure he's a murderer.¹¹ For suppose that before considering the evidence I am persuaded that there is this range of rationally permissible degrees of conviction. If I am correct in really thinking this, then there should be nothing wrong with arbitrarily choosing a verdict without bothering to look at the evidence. For there should be nothing wrong with my arbitrarily choosing a degree of conviction (induced perhaps by a magic pill) that is within the rationally permissible range. And given such a degree of conviction, I should vote accordingly, whatever it is. I can do no better by examining the evidence, forming a view and voting in accordance with it—to suppose otherwise involves denying the permissive view. For assuming that rationality and evidence do not determine a unique degree of conviction, then even if I am perfectly rational, there is no predicting what my degree of conviction will be upon evaluating the evidence. And hence there is no predicting how I will vote, whether Smith is guilty or not. Justice is no more likely to be done given that I examine the evidence than if I just vote however I like.

Now if this is an appropriate way to proceed before considering the evidence, then I can reason similarly once I have viewed it and have become convinced that Smith is guilty:

Given the same evidence I could no less rationally have had significant doubts about Smith's guilt, given which I could responsibly vote Not Guilty. That I happen to lack these doubts is either a result of chance or some arbitrary causal factor. So I would be doing no better with respect to the goal of securing a just verdict by voting with my convictions. For what ultimately matters for justice is the correspondence between the verdict and whether Smith *is* guilty. What I happen believe and how strongly I believe it makes no difference to the justice of the verdict. My rationally forming an opinion in the light of the evidence is just

the *means* by which I try to attain a just outcome. But however good I am at assessing the probative force of the evidence, it is entirely open whether my degree of rational conviction will fall within the Guilty or Not Guilty range. As long as my vote corresponds to an attitude that is within the bounds of rationally permissible conviction, I have done as well as I possibly can given the available evidence to attain a just verdict. So although I'm sure that Smith is guilty, I'll vote Not Guilty just because I happen to like the guy.

But of course it can't be right to try to let someone go free when you are rationally persuaded that he is a murderer. Perhaps it is obligatory to vote Not Guilty, if I know that one might have rational doubts about Smith's guilt. For any evidence that it is rational to have doubts may itself cast enough doubt on Smith's guilt to make a Guilty vote inappropriate. But in this case I should not retain my conviction that he is guilty. Holding this moderately permissive view either leads to the absurd consequence that it is appropriate to vote against one's convictions, or it is self-undermining.

Here is a related worry. Instead of immediately inducing a degree of conviction, a magic pill might ensure that I rationally respond to the evidence to arrive at that conviction. We are supposing for the sake of argument that there is a certain range of rationally permissible degrees of confidence in Smith's guilt, given the evidence. So among the possible worlds in which I respond to the same evidence in a perfectly rational manner, there are some in which I'm sure he's guilty and others in which I have doubts. The pill simply ensures that one of the latter worlds obtains. Taking such a pill at random is not essentially different from evaluating the evidence without it. In either case, some factor other than the evidence or my ability to rationally respond to it causes me to form a degree of conviction—why not take charge of these arbitrary influencing factors as I choose? But since which pill I take determines how sure I will be, which in turn will determine how I vote, taking a pill is tantamount to arbitrarily choosing how to vote with no evidence at all. But this cannot be the right thing to do.

5. Responsiveness to New Evidence

A rational person is responsive to new evidence. My confidence in a proposition given certain evidence should depend partly on my confidence prior to obtaining this evidence, and partly on the strength of the evidence.¹² The problem I want to raise is that to the extent that we reject Uniqueness, it is hard to see why one's convictions ought always to be responsive to confirming evidence. Suppose that you and I share our total evidence *E*. My subjective probability for *P* is *x*, and yours is lower at *y*. We each now obtain additional evidence *E'*, which supports *P*. My confidence in *P* rises to *x'* and yours to *y'*, which happens to be equal to the *x* that I held prior to obtaining *E'*. We have each updated our convictions appropriately in response to the new evidence. But

now let's suppose that we were each fully rational in holding our different degrees of belief x and y given just evidence E . Although I've been a little more confident than you, I would have been no less rational in sharing your doubts. Why then shouldn't I just keep my confidence in P at x , if it suits me? After all, that is where I would have ended up had I enjoyed your lower degree of confidence prior to obtaining E' . If that degree of doubt would have been rational for me back then, what could be wrong with the corresponding lower credence now? Perhaps in between obtaining E' and responding to it, I take one of those pills shifting my confidence down to yours at y . Then recognizing the significance of E' , my confidence rises appropriately to x . There should be nothing wrong with taking such a pill given that the resulting conviction is no less rational. But the net effect is that I remain entirely unmoved by E' , even though I correctly recognize that it confirms P .

This can't be right. Imagine a juror who is on the fence, undecided as to whether Smith is guilty. He receives new evidence and correctly recognizes that it confirms Smith's guilt. But our juror remains unmoved, not the slightest bit more inclined toward a guilty verdict. Such a person is not epistemically competent to serve on a jury.

6. Rationality and First-person Deliberation

The arguments of Sections 4 and 5 suggest that a permissive account of rationality introduces a kind of arbitrariness to our beliefs that can infect both practical and theoretical deliberations. We can roughly pose the general challenge in terms of outright belief as follows. If my current beliefs are not rationally obligatory for me, why should I take propositions that I *actually* believe as a basis for action and reasoning, rather than some others that I don't believe, but would be rational in believing? Why should *my* beliefs be privileged in my practical and theoretical deliberations, over equally rational alternative beliefs? Suppose that due to a bump on the head, I lose my belief in P and forget that I ever held it. The Epistemology Oracle informs me that believing P is just one of a range of epistemically permissible options for me given my evidence. In figuring out what to believe now, surely the fact that *I used to believe P* is entirely irrelevant; I might just as well start over again and form a new rational belief. Why then should it be any different if I still happen to believe P ?

Here is a possible response. To vote Not Guilty when I'm rightly sure that Smith committed murder amounts to deliberately trying to bring about an injustice. This is surely wrong regardless of whether a different degree of conviction would be equally rational. Similarly, to reason from propositions one takes to be false, is to purposely undermine the aim of inquiry, namely to reach the truth. And taking pills that change one's beliefs is not a rational option, even if the resulting beliefs are such that one could have rationally held them given one's evidence. For taking such a pill amounts to attempting to deceive one's

self. To induce a belief that one takes to be false violates the aim of believing only what is true.

What this reply suggests is that I have inappropriately been taking a “third-person” approach in questioning the relevance of *what my beliefs happen to be* if alternative attitudes are permissible. Rather, what I must ask myself in my deliberations is *whether P*. Since I can’t help but answer from my own perspective, my answer will correspond to what I happen to believe. But it is the (apparent) fact *that P* which is relevant, not the fact that *I* happen to believe it. If my aim is to convict Smith if and only if he is guilty, then the question I must answer is *whether he is guilty*. If I answer that yes he is guilty, then of course I should I vote accordingly, for a guilty person should be convicted. From my perspective, it is the fact (as I take it) that he is guilty that is a sufficient reason to convict him, not the fact that I happen to believe that he’s guilty. So whether or not the beliefs that I happen to hold are epistemically obligatory is just not relevant to whether I should act on them or reason from them.

I don’t find this response entirely satisfying. For it seems that at least in some cases my assessment of the epistemic status of a belief can legitimately override my own take on the facts in theoretical and practical deliberation. Suppose that due to irrational fears, I can’t help but think that my plane will crash. Yet I’m well aware that the statistical data make it unreasonable to believe this. It is not that I take myself to have some special insight that can override the regular evidence. I realize that my belief is irrational, but I just can’t manage to abandon it. Admittedly, this is a very odd state to be in. It is crucial to this case that I genuinely expect the plane to crash (I don’t merely fear that it will), while genuinely recognizing the epistemic irrationality of my attitude (I don’t merely recognize that according to typical epistemological theories, it would count as irrational). But I would suggest that it is at least possible to get into such a state, even if it involves severe cognitive dissonance.

It seems to me that in this situation it would be quite reasonable just to force myself to board the plane, or perhaps ask someone to drag me on. Depending how urgent my flight was, it might be irresponsible not to do so. Similarly, it may be appropriate to take a pill to help me believe that the plane is safe. If my faculties were functioning as they should, my recognition of the irrationality of my belief would cause me to abandon this belief and attain the rational conviction that the plane will not crash. As I’m aware of this, but still can’t come to my senses, a magic pill may be useful to nudge my beliefs into order.

Now in a sense, forcing myself on the plane while believing it will crash almost amounts to attempting suicide. I may not intend to die, but I’m acting in a way that I fully expect to have this consequence. Similarly, if I take the pill to induce the belief that the plane is safe, I am trying to adopt a view that I take to be false. Nevertheless, it seems quite appropriate to do so in this case. Instead of acting on what I take to be the facts, I step outside my own perspective, so to

speak, and let my epistemological convictions guide me. Even though I expect the plane to crash, and take the belief that it won't be false, I correctly take the *rationality* of the belief that the plane is safe as the best guide in action. For rationality is the best guide to the truth, and when it conflicts with my own view on things, it is best that rationality wins out.

Now if my epistemic convictions can legitimately overrule my beliefs in a case where they *conflict*, it is hard to see why this should not be appropriate if rationality neither conflicts with my beliefs nor *dictates* what they should be. I should be able to consider my views from the outside, as it were, and recognize that as far as rationality is concerned, they are on a par with various competing views. As there is no better guide to the truth than a rational assessment of the evidence, I should see that there is no advantage to acting on, or reasoning from my actual beliefs rather than other rational alternatives. But as I have argued, this leads to various absurdities.

7. Conclusion

I have been drawing out some of the difficulties we face by being epistemically permissive to various degrees. These arguments should tend to push us toward a less permissive account of rationality, and perhaps all the way to Uniqueness. This will have consequences both for our general theory of rationality and specific issues such as what impact disagreement with peers should have on our convictions. Perhaps whether or not Uniqueness is tolerable, all things considered, is yet to be seen.¹³

Notes

1. I'm following Rosen (2001) in the use of these terms. There is a natural analogue with permissible/obligatory action. But they should not be understood as involving a commitment to a deontological conception of rationality according to which believing rationally is entirely a matter of fulfilling certain epistemic duties.
2. For practical reasons such as clutter avoidance (Harman 1986) it might be wiser to take no attitude to P at all, even though one's evidence adequately supports P. We could refine our question as 'If I am to take any doxastic attitude to P, is believing P the only rational option?'
3. Feldman takes the options to be just belief, disbelief and suspension of belief. I am understanding it as covering degrees of belief, or subjective probabilities. Distracting complications arise here. The relation between outright belief and degrees of belief is controversial. Also, perhaps our convictions do not, and even should not come in precise degrees, but rather cover vague ranges. I will ignore these matters as they are not crucial here.
4. See Earman (1992) for a survey of more and less permissive versions of Bayesianism.

5. See van Inwagen (1996), Kelly (forthcoming), Feldman (forthcoming), and Christensen (unpublished mss.).
6. Whether jurors really do possess the same evidence in a case like this is questionable. Even if they have seen the same data presented in court, different judgments may be due to different background beliefs, which in turn are due to different past experiences. So the relation between epistemic permissiveness and the possibility of rational disagreement is not so simple.
7. I believe it was either Ralph Wedgwood, or Jason Stanley, or both, or someone else who suggested this as a coherent position.
8. Adam Elga and Ralph Wedgwood each raised the kind of worry that I'm responding to here. Apologies if I haven't fully captured their concerns.
9. The issues raised here are related to ones considered by Adam Elga (unpublished ms.). Here, as elsewhere, I have benefited much from discussions with Elga.
10. The case is complicated by the fact that one's vote is supposed to depend on whether it is "beyond reasonable doubt" that the defendant is guilty. Suppose I am convinced beyond a shadow of a doubt that Smith is guilty, correctly judge that my conviction is reasonable, and yet allow that one might reasonably have doubts. Does this count as judging it to be beyond reasonable doubt that he's guilty? Let's avoid this complication by pretending that one's vote should be based just on one's own convictions. In voting one judges: if it were up to me to determine the verdict, this would be my decision.
11. No more qualms that is, than someone who has reasonable doubts about Smith's guilt and votes Not Guilty. Such a person may worry that he is letting a murderer go free, but he can be sure that he has acted responsibly.
12. The standard rule for updating one's degrees of belief is Bayesian Conditionalization. But the problem that I'm raising should apply to any account of rational updating. Thanks to Adam Elga (who in turn gives the credit to Alan Hájek) for suggesting that I state the problem more generally than in the context of Bayesianism.
13. Thanks to Karen Bennett, Paul Boghossian, Liz Camp, David Christensen, Adam Elga, Alan Hájek, Liz Harman, Tom Kelly, Jill North, Karl Schafer, Nico Silins, Ted Sider, Jason Stanley, Michael Strevens, and Ralph Wedgwood for helpful discussion and feedback on this topic. Versions of this material were presented at the Analytic Philosophy Summer School, Venice, July 2004, the Corridor group, an NYU graduate seminar, the Australasian Association of Philosophy Conference July 2005 and the Bellingham summer philosophy conference, August 2005. I'm grateful for the questions and comments I received on these occasions.

References

- Christensen, D. "Epistemology of Disagreement: The Good News" unpublished manuscript.
- Earman, J. 1992. *Bayes or Bust?* Cambridge, Mass.: MIT Press.
- Elga, A. "The Problem of Arbitrary Factors" unpublished manuscript.
- Feldman, R. "Reasonable Religious Disagreements" forthcoming in L. Antony ed. *Philosophers without God*
- Foley, R. 1987. *The Theory of Epistemic Rationality*, Cambridge: Harvard University Press.

- Goodman, N. 1955. *Fact, Fiction, and Forecast*, Cambridge, Mass.: Harvard University Press.
- Harman, G. 1986. *Change in View: Principles of Reasoning*, Cambridge, Mass.: MIT Press.
- Kelly, T. "The Epistemic Significance of Disagreement" forthcoming in J. Hawthorne, and T. Gendler eds. *Oxford Studies in Epistemology* (Oxford University Press)
- Lycan, W. 1988. *Judgment and Justification*, Cambridge, Cambridge University Press.
- Rawls, J. 1999. *A Theory of Justice Revised Edition*, Cambridge, Mass.: Harvard University Press.
- Rosen, G. 2001. "Nominalism, Naturalism, Epistemic Relativism" *Philosophical Perspectives* 15: 69–91.
- van Fraassen, B. 1984. "Belief and the Will" *Journal of Philosophy* 81: 235–256.
- , 1989. *Laws and Symmetry*, Oxford: Clarendon Press.
- , 2002. *The Empirical Stance*, New Haven: Yale University Press.
- van Inwagen, P. 1996. "It is Wrong, Always, Everywhere, and for Anyone, to Believe Anything, Upon Insufficient Evidence" in J. Jordan and D. Howard-Snyder eds. *Faith, Freedom, and Rationality*, Hanham, MD: Rowman and Littlefield, pp. 137–154.